



Interpreting Cohesive Forms in the Context of Discourse Inference

Citation

Kehler, Andrew. 1995. Interpreting Cohesive Forms in the Context of Discourse Inference. Harvard Computer Science Group Technical Report TR-11-95.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:25104741>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story


The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

**Interpreting Cohesive Forms in the
Context of Discourse Inference**

Andrew Kehler

TR-11-95



Center for Research in Computing Technology
Harvard University
Cambridge, Massachusetts

Interpreting Cohesive Forms in the Context of Discourse Inference

A thesis presented

by

Andrew Kehler

to

The Division of Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Computer Science

Harvard University

Cambridge, Massachusetts

June, 1995

©1995 by Andrew Scott Kehler
All rights reserved.

ABSTRACT

In this thesis, we present analyses and algorithms for resolving a variety of cohesive phenomena in natural language, including VP-ellipsis, gapping, event reference, tense, and pronominal reference. Past work has attempted to explain the complicated behavior of these expressions with theories that operate within a single module of language processing. We argue that such approaches cannot be maintained; in particular, the data we present strongly suggest that the nature of the coherence relation operative between clauses needs to be taken into account.

We provide a theory of coherence relations and the discourse inference processes that underly their recognition. We utilize this theory to break the deadlock between *syntactic* and *semantic* approaches to resolving VP-ellipsis. We show that the data exhibits a pattern with respect to our categorization of coherence relations, and present an account which predicts this pattern. We extend our analysis to gapping and event reference, and show that our analyses result in a more independently-motivated and empirically-adequate distinction among types of anaphoric processes than past analyses.

We also present an account of VP-ellipsis resolution that predicts the correct set of *strict* and *sloppy* readings for a number of benchmark examples that are problematic for past approaches. The correct readings can be seen to result from a general distinction between *referring* and *copying* in anaphoric processes. The account also extends to other types of reference, such as event reference and ‘one’-anaphora.

Finally, we utilize our theory of coherence in analyses that break the deadlock between *definite-reference* and *coherence-based* approaches to tense and pronoun interpretation. We present a theory of tense interpretation that interacts with discourse inference processes to predict data that is problematic for both types of approach. We demonstrate that the data commonly cited in the pronoun interpretation literature also exhibits a pattern with respect to coherence relations, and make some preliminary proposals for how such a pattern might result from the properties of the different types of discourse inference we posit.

ACKNOWLEDGEMENTS

Believe it or not, I can actually say that I'm glad that I left a comfortable and rewarding job and returned to school. My three and three-quarter years at Harvard have left me few things to complain about. This is first and foremost due to my having Stuart Shieber as a thesis advisor. Stuart gave me the best of all worlds: freedom to explore my own ideas, with just the right amount of well-placed guidance when I needed it most. I hope that the positive influence that Stuart has had on my work will continue to be manifest in my future endeavors.

Next I want to thank the other two members of my committee, Barbara Grosz and Fernando Pereira. Barbara gave me continual guidance throughout my tenure at Harvard, ranging from research ideas to where to get my Toyota fixed, and just about everything in between. Fernando took an impressive amount of time out of his busy schedule to devote to my research and my thesis, and both are better products as a result.

It was my fellow graduate students who kept me going day-to-day. Big-time thanks go to Stan 'Big Dog' Chen, Jon 'Slacker Boy' Christensen, Allyn Dimock, Josh Goodman, Roni Khardon, Rebecca Hwa, Lillian Lee, Andy Leung, Dong Lin, Karen Lochbaum, Christine Nakatani, Ted Nesson, Dan Roth, Wheeler Ruml, Kathy Ryall, Nadia Shalaby, and John Zao. Special thanks go to my officemates Stan and Josh for putting up with me as I wrote my thesis, and to Jon, the other member of the two-person thesis support group.

I had the pleasure of working with several excellent groups of people during my summers. My experience at IBM T.J. Watson enhanced my skills considerably in statistical NLP and nerf ball (not necessarily in that order), and provided a welcome break from long nights pondering VP-ellipsis. Thanks to everyone there that made it such a fun environment, especially Peter Brown, Stephen Della Pietra, Vincent Della Pietra, Jan Hajic, and Bob Mercer.

Spending the next summer at Sun Microsystems was equally fun. Thanks go to everyone I worked with there, especially Eric Baatz, Gina Levow, Paul Martin, Bill Woods, and Nicole Yankelovich. Extra special thanks go to Paul, who is not only a great guy to work with, but who is also responsible for my addiction to Peet's coffee.

Part of my final summer was spent at Xerox PARC, located in paradise (a.k.a. Palo Alto, CA). Thanks go to all of my friends and colleagues there, especially Mary Dalrymple, Ron Kaplan, John Lamping, John Maxwell, and Vijay Saraswat. Mary deserves special thanks for bringing me into PARC and for being such a great friend and colleague since.

Various other people deserve thanks for one reason or another. Gregory Ward is another of my close colleagues with whom it is always a blast to work. Candy Sidner has gone out of her way to give me encouragement at every opportunity. Shalom

Lappin has remained a steady contact and a great person to talk shop with. Sam Epstein and Susumu Kuno, both in the Harvard Linguistics Department, have always taken the time to be interested in my work. David Magerman has kept me in shape by dragging me to the gym, and kept me poor by dragging me to his poker games.

I am grateful for the financial support that I received to pursue the work reported on in this thesis. The work was supported in part by National Science Foundation Grants IRI-90-09018, IRI-94-04756, CDA-94-01024, IRI-91-57996, IRI-93-50192, and a grant from the Xerox Corporation.

Most of all I want to thank my parents, Frank and Ruby. Getting to this point was only possible because of their continual love and encouragement. Rick, Nancy, and Alex round out the most supportive family one can have; many thanks go to them too. Finally, I want to thank Jill Jarecki, who has always been there when I needed her, and to Katy Jarecki, for just being who she is.

Contents

1	Introduction	1
1.1	The Problem	1
1.2	Discourse Inference	3
1.3	Contributions of the Thesis	5
1.4	Formalism	7
1.5	Thesis Outline	8
2	Discourse Coherence	10
2.1	A Theory of Coherence Relations	11
2.1.1	Resemblance Relations	12
2.1.2	Cause-Effect Relations	14
2.1.3	Contiguity Relations	16
2.2	Discourse Inference	16
2.2.1	Resemblance Relations	19
2.2.2	Cause-Effect Relations	20
2.2.3	Contiguity Relations	20
2.3	Basic Principles of Coherence	21
2.4	Relationship to Past Categorizations of Coherence Relations	24
2.5	Other Approaches to Discourse Coherence	25
2.5.1	Rhetorical Structure Theory	25
2.5.2	Discourse Parsing	26
2.5.3	Intentional-Level Approaches to Coherence	27
2.6	Conclusions	28
3	VP-Ellipsis and Coherence	30
3.1	An Abstract Characterization of the Approaches	31
3.1.1	Syntactic Approaches	31
3.1.2	Semantic Approaches	32
3.2	A Pattern in the Data	33
3.2.1	Voice Alternation	34

3.2.2	Non-VP Antecedents	35
3.2.3	Binding Theory Constraints	36
3.2.4	Subjacency Violations	39
3.3	Syntactic and Semantic Properties of VP-Ellipsis	41
3.4	Interaction with Discourse Inference	44
3.5	Applying the Analysis	46
3.6	Comparison with Past Work in VP-Ellipsis	51
3.6.1	Syntactic Accounts	51
3.6.2	Semantic Accounts	56
3.6.3	VP-Ellipsis and Discourse	57
3.7	Remaining Issues	60
3.7.1	Contiguity Relations	60
3.7.2	Lack of Coherence Relation	61
3.7.3	Situationally-Evoked Antecedents	62
3.8	Conclusion	63
4	A Resolution Algorithm for VP-Ellipsis	65
4.1	The Account	67
4.1.1	Representation	67
4.1.2	Resolution Algorithm	68
4.2	Examples	71
4.2.1	Cascaded Ellipsis	71
4.2.2	Stripping	72
4.2.3	Missing Readings	72
4.2.4	5-Reading Example	75
4.3	Implementation	78
4.4	Extensions to the Analysis	82
4.4.1	Non-Pronominal Referential Elements	82
4.4.2	Extended Parallelism	82
4.4.3	Referring and Copying	83
4.4.4	Ellipsis and Quantification	85
4.5	Comparison with Source-Determined Accounts	87
4.5.1	Source-Determined, Identity-of-Relations Analyses	88
4.5.2	Source-Determined, Non-Identity Analyses	89
4.6	Comparison with Discourse-Determined Accounts	91
4.6.1	Discourse-Determined Accounts	91
4.6.2	Predictions of Discourse-Determined Accounts	93
4.6.3	The Scope of Discourse-Determined Accounts	95
4.6.4	Relationship Between the Two Approaches	97
4.6.5	Remaining Problematic Examples	99

4.7	Conclusions	101
5	Other Cohesive Forms	103
5.1	Gapping	104
5.1.1	Syntactic and Semantic Properties of Gapping	105
5.1.2	The Analysis Applied to Gapping	109
5.1.3	Another Form of Local Ellipsis: Stripping	112
5.1.4	Interaction of Gapping and VP-Ellipsis	112
5.1.5	Comparison to Past Work	114
5.1.6	Conclusions	116
5.2	Event Reference	117
5.2.1	The Sag and Hankamer Dichotomy	117
5.2.2	Syntax and Semantics of Event Reference	118
5.2.3	A New Dichotomy	120
5.2.4	Another Problem for Sag and Hankamer: ‘Do So’ Anaphora	121
5.2.5	Conclusions	126
5.3	Tense Interpretation	127
5.3.1	Background and Past Work	127
5.3.2	Problematic Data	131
5.3.3	The Account	132
5.3.4	Examples	134
5.3.5	Comparison to Definite Reference Approaches	135
5.3.6	Comparison to Coherence-Based Approaches	139
5.3.7	Conclusion	140
5.4	Pronoun Interpretation	142
5.4.1	Past Work	142
5.4.2	Accounting for the Role of Coherence Relations	146
5.4.3	Mixed Approaches	149
5.4.4	Conclusion	152
6	Conclusion	153
6.1	Summary	153
6.2	Future Directions	155

List of Figures

1.1	Syntactic and Semantic Representations for <i>Bill became upset</i>	7
1.2	Syntactic and Semantic Representations for <i>Ross likes his mother</i> . . .	8
1.3	Syntactic and Semantic Representations for <i>Bill, Ross dislikes</i>	8
3.1	Syntactic and Semantic Representations for <i>Ross likes his mother</i> . . .	31
3.2	Syntactic and Semantic Representations for <i>Bill does</i>	31
3.3	Syntactic and Semantic Representations for <i>Bill does [like his mother]</i> . . .	32
3.4	Syntactic and Semantic Representations for <i>Hillary did</i>	43
3.5	Syntactic and Semantic Representations for <i>Bill became upset</i>	46
3.6	Syntactic and Semantic Representations for <i>Hillary did</i>	47
3.7	Syntactic and Semantic Representations for <i>Hillary did [become upset]</i> . . .	47
3.8	Syntactic and Semantic Representations for <i>This problem was looked into by Bill</i>	48
3.9	Syntactic and Semantic Representations for <i>Hillary did</i>	49
3.10	Syntactic and Semantic Representations for <i>he_j did [introduce Bill_j to everyone]</i>	49
3.11	Syntactic and Semantic Representations for <i>This problem was to have been looked into</i>	50
3.12	Syntactic and Semantic Representations for <i>obviously nobody did</i> . . .	50
3.13	Syntactic and Semantic Representations for <i>obviously nobody did</i> after resolution.	51
5.1	Syntactic and Semantic Representations for <i>Bill became upset</i> after fronting.	107
5.2	Syntactic and Semantic Representations for <i>Hillary angry</i>	107
5.3	Syntactic and Semantic Representations for <i>Hillary angry</i> after reconstruction.	108
5.4	Syntactic and Semantic Representations for <i>Hillary Bill</i> after reconstruction.	114
5.5	Syntax and Semantics of <i>Hillary did it</i>	119
5.6	Event times in Hinrichs' system	130

List of Tables

2.1	Resemblance Relations	14
2.2	Cause-Effect Relations	16
3.1	Syntactic and Semantic Properties of Ellipsis and Event Reference . .	41
3.2	Recovery of Missing Information in the Context of Discourse Inference	45

Chapter 1

Introduction

1.1 The Problem

Natural languages possess a variety of expressions that obtain their meaning through reference to entities previously evoked in a discourse. Examples include various types of ellipsis, pronominal reference, and the interpretation of tense with respect to temporal information. Such linguistic forms have been called *cohesive*, because they add to the cohesion of a discourse by establishing ties between the utterances containing them and the utterances that give rise to their antecedents (Halliday and Hasan, 1976).

The pervasiveness of cohesive forms in naturally-occurring discourse requires that natural language understanding systems recover their meaning to be effective. As a result, these forms have received much attention in both computational and theoretical linguistics. Despite this attention, however, the processes required to interpret these forms remain poorly understood, and thus the lack of adequate algorithms for recovering their meaning continues to be a critical bottleneck for a broad range of natural language processing applications.

In many cases, researchers do not even agree on the level of language processing at which the meaning of these expressions should be resolved. Verb phrase (VP) ellipsis, a phenomenon to which we pay special attention in this thesis, provides a paradigmatic example of this disagreement. VP-ellipsis is exemplified in sentence (1).

- (1) Ross likes his mother, and Bill does too.

The stranded auxiliary in the second clause (henceforth, the *target* clause) indicates the deletion of a verb phrase; the meaning of this deleted VP must be determined from the meaning of another clause, in this case, the first clause (henceforth, the *source* clause).

A central facet of any analysis for resolving VP-ellipsis is the level of linguistic representation at which the resolution processes operate. Past theories fall primarily into two categories. Inherent in *syntactic* accounts (Sag, 1976; Williams, 1977; Haïk, 1987; Hellan, 1988; Lappin and McCord, 1990; Hestvik, 1993; Lappin, 1993b; Fiengo and May, 1994, inter alia) is the claim that VP-ellipsis is resolved at some level of syntactic structure. On the other hand, inherent in *semantic* accounts (Dalrymple, Shieber, and Pereira, 1991; Hardt, 1991a; Kehler, 1993a, inter alia) is the claim that VP-ellipsis is resolved at a purely semantic level of representation.

This fundamental question remains a point of contention, particularly because there are data to support each approach. Syntactic treatments have used the unacceptability of examples such as sentences (2)-(4) to support their view.¹

- (2) # The lawyer defended Bill_i, and he_i did too. [defend Bill_i]
(from Lappin (1993b))
- (3) # John_i blamed himself_i, and Bill did too. [blamed him_i]
(from Kitagawa (1991))
- (4) # John read everything which Bill believes the claim that he did. [read ϕ]
(from Haïk (1987))

Sentence (2) is unacceptable assuming the indicated coreference between *he* and *Bill*. This unacceptability is predicted by a syntactic account, because if the elided syntactic structure were present in the target clause, a *Condition C* violation² would result. Because Condition C is a constraint on syntactic representations, a semantic approach would not predict such a violation. Likewise, *Condition A* predicts that sentence (3) does not have the reading where *Bill* blamed *John*, and constraints on *subjacency* predict the unacceptability of sentence (4) in a syntactic theory.

On the other hand, sentences (5)-(7) have been used to support a semantic theory of VP-ellipsis.

- (5) In March, four fireworks manufacturers asked that the decision be reversed, and on Monday the ICC did. [reverse the decision]
(from text of Rosenthal (1988), cited in Dalrymple (1991))
- (6) Harry used to be a great speaker, but he can't anymore, because he lost his voice [speak]. (from Hardt (1993))

¹Phrases shown in brackets after each example indicate the elided material under the intended interpretation. Expressions which are coindexed refer to the same entities.

²In Binding Theory (Chomsky, 1981; Lasnik, 1989), Condition C rules out sentences such as *He_i defended Bill_i* due to the syntactic *c-command* relation between *he* and *Bill*. We discuss this in more depth in Chapter 3.

- (7) I expected Bill_i to win even when he_i didn't. [expect Bill_i to win]
(from Dalrymple (1991))

VP-ellipsis is felicitous in sentence (5) despite the fact that a syntactic structure for the VP *reverse the decision* does not appear in the source clause. Sentence (6) illustrates another type of syntactic-structural mismatch; in this case the antecedent is evoked by a nominalization. Finally, sentence (7) is felicitous despite the fact that Condition C predicts unacceptability under a syntactic account. In contrast to sentences (2)-(4), sentences (5)-(7) support a purely semantic approach to VP-ellipsis resolution as opposed to a syntactic one.

This contradictory data has placed the field in a state of deadlock with respect to the level of representation at which VP-ellipsis is resolved. The data suggest that any theory of VP-ellipsis that operates solely within a single module of language processing will not be empirically adequate. In this thesis, we claim that this situation and similar situations with other cohesive phenomena are the result of two aspects of the state of current research: (i) a prevailing assumption that the resolution process is uniform and articulable within a single module of language processing, and (ii) existing data that appear to be contradictory with respect to this assumption.

This leaves two options for theories of cohesive forms such as VP-ellipsis: (i) a non-unified account could be posited in which different processes apply in different circumstances, or (ii) a unified account could be posited in which other processes are seen to affect the final distribution of the data. In this thesis, we present theories of VP-ellipsis, gapping, event reference, tense interpretation, and pronoun interpretation that fall into the second category. For each of these theories, there is a common external factor at work: the process of establishing that a discourse is coherent, which we will refer to as *discourse inference*. We will show that by providing a unified account of each linguistic form in question, along with a theory of the properties of different discourse inference processes, we can predict data that is beyond the reach of existing theories. In the case of VP-ellipsis, this includes predicting the behavior exhibited by the data presented above. We discuss the notion of discourse inference in the next section.

1.2 Discourse Inference

It is well known that a discourse is more than just an arbitrary set of utterances; the difference between the two is that discourses are *coherent*. Coherence requires that a hearer be able to establish some type of connection between an utterance and the context in which it is uttered.

To illustrate one such type of connection, consider example (8), from Hobbs (1979).

(8) John took a train from Paris to Istanbul. He likes spinach.

As Hobbs notes, passage (8) may not initially strike one as a coherent passage. However, a hearer might suggest an explanation to make it coherent; for instance, he might conjecture that the spinach crop failed in France and Turkey is the closest country where spinach is available. This assumption allows for a type of cause-effect relation to be established between the propositions denoted by the two clauses, in that the information expressed by the second clause could then be understood as the cause of the action expressed by the first clause. The very fact that hearers are driven to such explanations is indicative of the need to establish coherence; we claim in fact that this need constitutes a fundamental part of the human language faculty.

A theory of discourse coherence must articulate the possible connections that can hold between successive utterances in order for them to form a coherent discourse. In this thesis, we utilize a fundamental distinction that was first articulated by the philosopher David Hume (1748), who makes the following general statement concerning these types of connections.³

“Though it be too obvious to escape observation that different ideas are connected together, I do not find that any philosopher has attempted to enumerate or class all the principles of association—a subject, however, that seems worthy of curiosity. To me there appear to be only three principles of connection among ideas, namely *Resemblance*, *Contiguity* in time or place, and *Cause* or *Effect*.”

Two hundred and fifty years have passed since Hume wrote these words; relatively recently several researchers have attempted the enumeration that he speaks of, where such connections are equated with a set of *coherence relations* that hold between utterances in a discourse (Halliday and Hasan, 1976; Hobbs, 1979; Longacre, 1983; Mann and Thompson, 1987; Polanyi, 1988; Hobbs, 1990, inter alia; see Hovy (1990) for a compendium of over 350 relations that have been proposed in the literature). We will argue, however, that many of the proposed relations can be seen as being subsumed by Hume’s three general categories; the view taken here is that Hume’s categories comprise a small set of basic types of cognitive principles that give rise to such relations. In the next section, we discuss how this distinction is leveraged in our analysis of cohesive forms.

³The relevance of this passage to coherence in natural language was noted by Hobbs (1990).

1.3 Contributions of the Thesis

In this thesis, we account for problematic data concerning a wide range of cohesive forms by demonstrating the role of discourse inference mechanisms in establishing coherence relations. The forms we address and the results we achieve are summarized below.

VP-Ellipsis: We address the question of the level of representation at which VP-ellipsis is resolved, as discussed in Section 1.1. We show that the seemingly contradictory data exhibits a pattern which correlates with the type of coherence relation holding between the source and target clauses. We give an account of the properties of VP-ellipsis as well as how they interact with different types of discourse inference mechanisms; this account predicts the pattern in the data noted.

Having determined the levels of linguistic processing at which VP-ellipsis is to be resolved, an algorithm for resolution is articulated. A benchmark that has been used for such algorithms is the generation of the correct set of strict and sloppy ambiguities in various complex cases. Example (9) displays such an ambiguity; Bill may like *Ross's* mother (the strict reading) or *his own* mother (the sloppy reading).

(9) Ross likes his mother, and Bill does too.

We describe how our algorithm makes the correct predictions for a number of examples that are problematic for past approaches. Furthermore, we show that our algorithm extends directly to other similar types of reference, including pronominal event reference and ‘one’-anaphora.

Gapping: Gapping is similar to VP-ellipsis in that material has been elided in the target clause. In this case, however, two bare constituents are extant in the target, neither of which is a stranded auxiliary. The target clause shown in sentence (10) in the result of gapping.

(10) Bill became upset, and Hillary angry.

We show that gapping exhibits a pattern similar to VP-ellipsis; this pattern correlates with the nature of the coherence relation holding between clauses. However, the distribution of gapping will be shown to be most restricted in situations where VP-ellipsis is least restricted, and vice-versa. We articulate the syntactic and referential properties of gapping constructions, and show that the account of discourse inference employed for VP-ellipsis also predicts the gapping data.

Event Reference: Several forms of pronominal event reference may be used to stand proxy for an event that has been introduced in the prior discourse, as exemplified by the form *do it* in sentence (11).

(11) Bill sent out the memo before Hillary had a chance to do it.

Other types of event reference include *do that* and *do so* anaphora. We discuss how resolution of these forms compares to that for VP-ellipsis and gapping, and show that the predictions that our account makes for all of these forms is more empirically adequate than those provided by the distinction between *deep* and *surface* anaphora made by Hankamer and Sag (1976) and Sag and Hankamer (1984).

Tense: The use of simple past tenses in example (12) not only indicates that both events occurred in the past, but also induces a temporal relation in which the second event is understood as occurring after the first.

(12) Bill walked up to the podium. He began his speech.

We present a set of data that is problematic for two types of previous approaches. We provide an account that combines a theory of tense interpretation with the constraints imposed by coherence relations. The resulting account correctly predicts the data, and is shown to have additional advantages over past work.

Object Reference: In sentence (13), the pronoun *he* refers to one of the people that have been previously mentioned.

(13) Bill had a chat with George. He got really angry.

We show that the examples discussed in the pronoun interpretation literature display a pattern with respect to Hume's three types of coherence relation. We suggest ways in which pronoun resolution may be dependent on the properties of the discourse inference mechanisms underlying the establishment of these relations.

Past analyses of each of these phenomena have attempted to account for the data with a theory articulated within a single module of language processing. In each area, we present a set of data that is problematic for such approaches; we show that the nature of the coherence relation operative between clauses must be taken into account to adequately explain this data. We give analyses that combine a unified account of the linguistic form in question with a unified characterization of the role of the discourse inference processes that establish such relations. In each case, we show how the interaction of these two aspects of language processing results in a theory that handles the data cited as problematic for previous accounts.

1.4 Formalism

The analyses in this thesis utilize a model in which the syntactic representations of utterances are coupled with their semantic interpretations; the combination of the two serves as the input to the discourse inference process. Here, we digress to briefly discuss the formalism that we use to describe the syntactic and semantic properties of the forms of interest in this thesis.

We use a formalism related to the version of Categorical Semantics described by Pereira (1990). We have chosen this system both for its straightforward properties and its tight integration of syntax and semantics. It should be noted, however, that the analyses presented in the remainder of the thesis are not dependent on any particular aspects of this formalism beyond these general desiderata.

Figure 1.1 shows the syntactic and semantic representations for the sentence *Bill became upset*.⁴

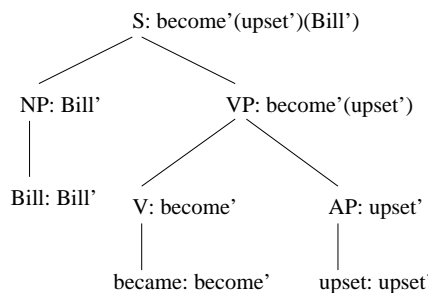


Figure 1.1: Syntactic and Semantic Representations for *Bill became upset*.

The leaf nodes of derivation trees are labeled by *assumptions* which may or may not need to be subsequently *discharged* for a valid derivation. One type of assumption is introduced by lexical items in the syntax (such as by *Bill*, *became*, and *upset* in Figure 1.1); these assumptions are not discharged. Other assumptions express a dependency between one form and another, such as those introduced by pronouns and traces; these assumptions need to be licensed and subsequently discharged by a suitable construction or assignment.

The process of introducing and discharging assumptions for pronominal binding is encoded by two rules: pronoun licensing [**pron-lic**] and pronoun abstraction [**pron-abs**]. The appearance of an unbound pronoun introduces an assumption (us-

⁴In Pereira's paper, derivation trees are shown with the semantic representations and their corresponding types. Along with the constraints imposed by the types of the semantic representations, syntactic analyses are used to control the semantic derivation; a standard Montagovian pairing is used between rules combining constituents in the syntax and rules for combining the meanings of those constituents in the semantics. Because only the syntactic categories and corresponding semantic representations are crucial for our purposes; we will omit the types in our depictions.

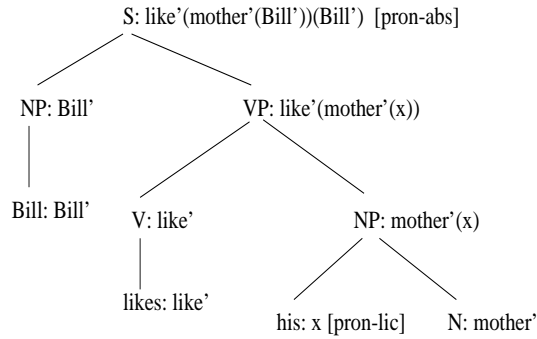


Figure 1.2: Syntactic and Semantic Representations for *Ross likes his mother*.

ing [**pron-lic**]) that is subsequently discharged when the pronoun is bound by an antecedent entity, a trace, or a quantifier (using [**pron-abs**]). The use of these rules is shown in the representations for the sentence *Bill likes his mother* in Figure 1.2.

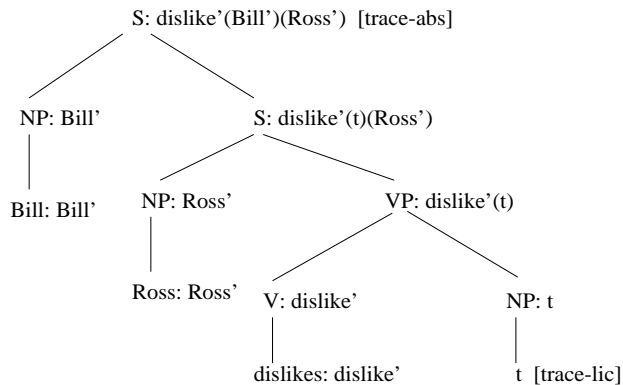


Figure 1.3: Syntactic and Semantic Representations for *Bill, Ross dislikes*.

Similarly, trace assumptions are licensed by the occurrence of a trace in the syntactic representation using the rule [**trace-lic**], and are discharged by a syntactic construction taking a constituent containing a trace using the rule [**trace-abs**]. Figure 1.3 contains a structure resulting from the use of these rules. Further discussion of these topics may be found in Pereira (1990).

1.5 Thesis Outline

In this chapter, we have given an overview of the analyses that we present in the remainder of the thesis, which is organized as follows. In Chapter 2, we detail our account of coherence relations and the discourse inference mechanisms underlying their recognition. Chapter 3 describes the first part of the account of VP-ellipsis, addressing the question of the level of linguistic representation at which resolution

occurs. Chapter 4 articulates the second part of the account, focusing on the resolution mechanism itself and the derivation of strict and sloppy readings. Chapter 5 discusses problems and solutions for the other linguistic forms that we have analyzed, including gapping, event reference, tense interpretation, and pronoun interpretation. We conclude in Chapter 6 and discuss future work.

Chapter 2

Discourse Coherence

It is well-established that the meaning of a discourse is greater than the sum of the meanings of its parts. That is, beyond the meaning of the individual utterances comprising the discourse, there is additional meaning to be gleaned from the fact that the utterances co-occur in a particular configuration. This additional meaning is the result of adopting the inferences required for establishing the way in which the utterances are related under the assumption that the discourse they form is coherent.

This effect is most easily illustrated by considering cases where it is difficult to recover such a relation. We mentioned one such case from Hobbs (1979) in Chapter 1, repeated below as sentence (14).

(14) John took a train from Paris to Istanbul. He likes spinach.

As Hobbs argues, example (14) does not form a coherent passage as it stands. It may become coherent, however, if one believes or finds it plausible to believe that the spinach crop has failed in France and that Turkey is the closest country where spinach is available. In this case, a hearer is assuming information that will allow an *Explanation* relation to be established between the propositions denoted by the utterances, in which the proposition denoted by the second clause is understood as the cause of the event specified by the first clause.

Of course, there are other types of connection that can be established between propositions. We consider example (15).

(15) Bill likes to jog, and George hates broccoli.

As it stands, passage (15) also lacks coherence. In this case, one wants to identify reasonably specific commonalities between the two sentences; the failure to do this here results in the incoherence of the passage. However, if sentence (15) is stated in a context where it was clear that *Bill* refers to Bill Clinton and that *George* refers to George Bush, then the passage becomes much more coherent under the common topic

of *properties of presidents*. Here, a hearer is making assumptions needed to allow for a *Parallel* relation to be established between the propositions.

As a final example, we consider passage (16), from Hobbs (1990).

(16) At 5:00 a train arrived in Chicago. At 6:00 George Bush held a press conference.

Again, without additional assumptions, passage (16) is of questionable coherence. To understand the passage as coherent, hearers may infer that Bush is on the train, or the train arrival is otherwise somehow relevant to him. This assumption allows a hearer to understand the passage as a *Narration*.

In each of these three examples, a hearer is presented with two clauses that independently are readily understood. Interpretation does not stop there however; the hearer is further inclined to assume unstated information necessary to analyze the passage as coherent. The point is summarized nicely in the following passage from Hobbs (1979):

.. the very fact that one is driven to such explanations indicates that some desire for coherence is operating, which is deeper than the notion of a discourse just being “about” some set of entities. (Hobbs, 1979, page 67)

Examples (14)–(16) show that the need to resolve coherence is a central facet of language understanding. Just as naturally as humans attempt to recover the implicit syntactic structure of a string of words they are presented with, they also attempt to recover the implicit “coherence structure” of a discourse they are presented with.

This process has led a number of researchers to posit a set of possible *coherence relations* that can hold between adjacent segments of a discourse. Some of these researchers have also posited interpretation schemes for establishing such relations when applied to discourses. In this chapter, we present an analysis of coherence relations and the discourse inference processes that underlie their resolution.

Despite the centrality of coherence resolution processes in language interpretation, the potential role of such processes is rarely considered by researchers in positing theories of the interpretation of various linguistic expressions. In the remainder of this thesis, we utilize our account of coherence in analyses of the interpretation of the linguistic forms that were introduced in Chapter 1, showing how these analyses predict data that is beyond the scope of existing theories.

2.1 A Theory of Coherence Relations

Examples (14)–(16) illustrate three types of connection between utterances that a hearer might try to establish in analyzing a discourse as coherent. Given suitable

assumptions, a *parallel* relation may be established for example (15), an *explanation* relation established for example (14), and a *narration* relation established for example (16). These three relations are instances of the three types of “connection among ideas” posited by Hume (1748) and discussed in Chapter 1, namely *Resemblance*, *Cause-Effect*, and *Contiguity*.

In this section, we analyze a set of coherence relations as belonging to these three general categories. The three classes are shown to differ systematically in the type of arguments over which the coherence constraints are applied, as well as in the type of inference process underlying this application. We use the theory described in Hobbs (1990) as a starting point; most of the relation definitions are either taken or adapted from definitions given therein. In this section we give the relation definitions; in Section 2.2 we discuss properties of the resolution processes underlying each of the three categories. Differences between our relations and Hobbs’ are discussed in Section 2.4.

2.1.1 Resemblance Relations

Establishing a passage as coherent under a Resemblance relation requires that commonalities and contrasts among corresponding sets of properties and entities be recognized. For each relation, the hearer identifies a relation p_0 that applies over a set of entities a_1, \dots, a_n from the first sentence S_0 , and a corresponding relation p_1 that applies over a corresponding set of entities b_1, \dots, b_n from the second sentence S_1 . Coherence results from these corresponding components being related; in this case a common (or contrasting) relation p subsuming p_0 and p_1 is inferred along with common (or contrasting) properties q_i of the corresponding elements a_i and b_i .

The canonical instance of a Resemblance relation is the *Parallel* relation.

Parallel: Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where for some property vector q , $q_i(a_i)$ and $q_i(b_i)$ for all i .

An example of the Parallel relation is given in sentence (17).

(17) John organized rallies for Clinton, and Fred distributed pamphlets for him.

Here, the parallel relations p_0 and p_1 correspond to the meanings of *organized rallies for* and *distributed pamphlets for* respectively; the common relation p that subsumes these might be a meaning corresponding to *do something to support*. Likewise, the elements a_1 and b_1 correspond to the meanings of *John* and *Bill*, which share the common property q_1 that they are people relevant to the conversation. The elements a_2 and b_2 correspond to the meanings of *Clinton* and *him*, which share a trivial common property q_2 in that the two terms denote the same individual.

Instead of inferring only commonalities among corresponding relations and entities in the utterances, one may infer a point of departure among either of these. This yields two definitions for the *Contrast* relation, one in which the relations expressed by the utterances are contrasted, and the other in which a set of corresponding entities is contrasted.

Contrast (i): Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $\neg p(b_1, b_2, \dots)$ from the assertion of S_1 , where for some property vector q , $q_i(a_i)$ and $\neg q_i(b_i)$ for all i .

Contrast (ii): Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where for some property vector q , $q_i(a_i)$ and $\neg q_i(b_i)$ for some i .

Passage (18) is an example where the first definition of Contrast applies.

(18) John supported Clinton, but Mary opposed him.

The second definition of Contrast is operative in passage (19).

(19) John supported Clinton, but Mary supported Bush.

The *Exemplification* relation holds between a general statement followed by an example of the generalization.

Exemplification: Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where b_i is a member or subset of a_i for some i .

Example (20) illustrates the Exemplification relation.

(20) Young aspiring politicians often support their party's presidential candidate. For instance, John campaigned hard for Clinton in 1992.

The *Generalization* relation is the same as *Exemplification* with reversed clause ordering.

Generalization: Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where a_i is a member or subset of b_i for some i .

An example of Generalization is shown in sentence (21).

(21) John campaigned hard for Clinton in 1992. Young aspiring politicians often support their party's presidential candidate.

The *Elaboration* relation is a limiting case of the *Parallel* relation, where the similar entities a_i and b_i are in fact identical.

Elaboration: Infer $p(a_1, a_2, \dots)$ from the assertions of S_0 and S_1 .

Elaborations are generally restatements, either from a different perspective or in more detail.¹ An example of Elaboration is given in (22).

- (22) A young aspiring politician was arrested in Texas today. John Smith, 34, was nabbed in a Houston law firm while attempting to embezzle funds for his campaign.

These relations are summarized in Table 2.1. Resemblance relations require that points of commonality (parallelism) and departure (contrast) between sets of corresponding entities (the a_i and b_i) and properties (the p_0 and p_1) be determined.

Relation	Constraints	Conjunctions
Parallel	$p_0 = p_1, q_i(a_i)$ and $q_i(b_i)$	and
Contrast	(1) $p_0 = \neg p_1, q_i(a_i)$ and $q_i(b_i)$ (2) $p_0 = p_1, q_i(a_i)$ and $\neg q_i(b_i)$	but
Exemplification	$p_0 = p_1 ; b_i \in a_i$ or $b_i \subset a_i$	for example
Generalization	$p_0 = p_1 ; a_i \in b_i$ or $a_i \subset b_i$	in general
Elaboration	$p_0 = p_1, a_i = b_i$	in other words

Table 2.1: Resemblance Relations

2.1.2 Cause-Effect Relations

Establishing a passage as coherent under a Cause-Effect relation requires that a path of implication be established between the propositions denoted by the utterances.² For each relation, the hearer identifies a proposition P from the first sentence S_0 and a proposition Q from the second sentence S_1 . Coherence results from these two propositions being related; in this case an implication relationship is inferred between the two.

The canonical case of a Cause-Effect relation is *Result*.

Result: Infer P from the assertion of S_0 and Q from the assertion of S_1 , where normally $P \rightarrow Q$.

An example of Result is given in sentence (23).

¹That is, while the corresponding relations and entities are constrained to be the same, the way in which they are *described* in the utterances will generally be different.

²We are using implication in a very loose sense here, meaning “could plausibly follow from”.

(23) Bill is a politician, and therefore he's dishonest.

Here, P corresponds to the meaning of *Bill is a politician*, and Q corresponds to the meaning of *he's dishonest*. The constraint that must hold therefore corresponds to the implication that *being a politician implies being dishonest*.

The *Explanation* relation is *Result* with reversed clause ordering.

Explanation: Infer P from the assertion of S_0 and Q from the assertion of S_1 , where normally $Q \rightarrow P$.

An Explanation relation is often marked by the conjunction *because*, as in sentence (24).

(24) Bill is dishonest because he's a politician.

The *Violated Expectation* relation is used to contrast an actual effect with the expected effect in light of a potential cause.

Violated Expectation: Infer P from the assertion of S_0 and Q from the assertion of S_1 , where normally $P \rightarrow \neg Q$.

Violated Expectation is exemplified in sentence (25).

(25) Bill is a politician, but he's honest.

Finally, *Denial of Preventer* is Violated Expectation with reversed clause ordering.

Denial of Preventer: Infer P from the assertion of S_0 and Q from the assertion of S_1 , where normally $Q \rightarrow \neg P$.

An example of the Denial of Preventer relation is given in example (26).

(26) Bill is honest, even though he's a politician.

These relations are summarized in Table 2.2. Cause-Effect relations all require that some implicational relationship be established between the two propositions.

Relation	Presuppose	Conjunctions
Result	$P \rightarrow Q$	and (as a result) therefore
Explanation	$Q \rightarrow P$	because
Violated Expectation	$P \rightarrow \neg Q$	but
Denial of Preventer	$Q \rightarrow \neg P$	even though despite

Table 2.2: Cause-Effect Relations

2.1.3 Contiguity Relations

There is only one relation in the Contiguity category, namely *Narration*. The Narration relation allows one to express a coherent sequence of events centered around some system of entities. The definition for this relation is given in less formal terms than the other definitions.

Narration: Infer a change of state for a system of entities from S_1 , inferring the initial state for this system from S_0 .

An example of the Narration relation is given in passage (27).

(27) Bill picked up the speech. He began to read.

The Narration relation, combining facets of Hobbs' (1990) *Occasion* and *Ground-Figure* relations, allows for complex situations to be communicated in a multi-utterance discourse. The exact constraints required for Narration are not well understood, a point to which we return in the next section.

2.2 Discourse Inference

The set of coherence relations we have presented establish connections between utterances in a discourse. These relations each impose a set of constraints that must be met in order for the discourse to be coherent under that relation. What is required now is a set of procedures for applying these constraints.

A variety of computational models might be employed for this purpose. The theories presented later in this thesis are not dependent on any particular model, although we do argue for certain features that any such model should have. To afford an additional degree of concreteness to our discussion of discourse inference processes, we consider one possible computational scheme, the abductive interpretation method of Hobbs et al. (1993).

Hobbs et al. (1993) outline a procedure for utterance interpretation and coherence resolution based on the mechanism of logical abduction. In their framework, an utterance is interpreted by “proving” its logical form, where assumptions can be made at a specified cost where necessary. This approach is then scaled up to the multi-utterance level, where world and domain knowledge is used to determine the most plausible coherence relation holding between utterances in a discourse segment. To accomplish this, two axioms are adopted. The first, shown in statement (28), states that a sentence is a discourse segment.³

$$(28) (\forall w, e) s(w, e) \supset \text{Segment}(w, e)$$

The second axiom, shown in statement (29), allows for two smaller segments to be composed into a larger one if a coherence relation can be established between the two.

$$(29) (\forall w_1, w_2, e_1, e_2, e) \text{Segment}(w_1, e_1) \wedge \text{Segment}(w_2, e_2) \\ \wedge \text{CoherenceRel}(e_1, e_2, e) \supset \text{Segment}(w_1 w_2, e)$$

To interpret a coherent text W , one must prove that it is a segment, as expressed by statement (30).

$$(30) (\exists e) \text{Segment}(W, e)$$

To illustrate, we consider interpreting passage (25), repeated below as (31), under the Violated Expectation relation.

$$(31) \text{Bill is a politician, but he's honest.}$$

To understand this passage as a segment, the three premises in axiom (29) need to be established. The first two are established by axiom (28), what remains is to establish a coherence relation. Because Violated Expectation is a coherence relation, we also have the axiom shown in (32).

$$(32) (\forall e_1, e_2) \text{ViolatedExpectation}(e_1, e_2) \supset \text{CoherenceRel}(e_1, e_2, e_2)$$

Violated Expectation requires that the presupposition $\text{cause}(e_1, \neg e_2)$ be abduced, as expressed by the rule shown in formula (33).

$$(33) (\forall e_1, e_2) \text{cause}(e_1, \neg e_2) \supset \text{ViolatedExpectation}(e_1, e_2)$$

³Here, w is a string of words, and e is its assertion or topic.

Let us assume that we have formula (34) in a knowledge base (and is therefore derivable at zero cost), or can otherwise abduce it at a cost corresponding to some measure of plausibility.⁴

$$(34) (\forall x, p_1) \text{Politician}(p_1, x) \supset (\exists h_1) \neg \text{Honest}(h_1, x) \wedge \text{cause}(p_1, \neg h_1)$$

That is, being a politician normally implies being dishonest. Starting from the logical forms of the two clauses in example (31), the abductive proof of the relation (and ultimately the segment) can then be completed; the plausibility measure assigned to the result will be inversely proportionate to the cost assigned to formula (34). Assuming a high plausibility measure, the Violated Expectation interpretation is inferred as a likely relation between the clauses in example (31).⁵

For the analyses in this thesis, we can assume that this abduction strategy is the mechanism that underlies each of our types of discourse inference; again, however, the analysis is compatible with other possible methods. Further, we will not have anything more to say about the process by which a particular coherence relation is chosen over another; the account developed here can be viewed to operate within a framework in which proofs for each possible relation are attempted in parallel.

In the following sections, we discuss properties of the discourse inference mechanisms supporting each of our three types of coherence relation. For the purpose of this discussion, we consider the discourse inference procedure as a two-step process:

1. *Identify* and *retrieve* the arguments to the coherence relation, and
2. *Apply* the constraints of the relation to those arguments.

We consider first the identification and retrieval of arguments. In our account, utterance interpretation results in a syntactic structure annotated with the semantic representation for each constituent, as depicted in Chapter 1. The semantic representations serving as arguments to a coherence relation are identified and retrieved by way of their corresponding nodes in the syntax. The constraints of the relation are then applied using these arguments. As we discuss in the following section, the different types of coherence relation are distinguished in terms of both the level in the syntax from which the arguments are retrieved, and in the nature of the constraints that are imposed on them.

⁴Obviously, the use of the universal quantifier is too strong, but we will ignore this detail for the point made here.

⁵For further details of the system, see Hobbs et al. (1993).

2.2.1 Resemblance Relations

Understanding segments of utterances standing in a Resemblance relation requires the determination of points of commonality (parallelism) and departure (contrast) between sets of corresponding entities and properties within the utterances. For each relation shown in Table 2.1, the hearer understands the relation by inferring $p_0(a_1, \dots, a_n)$ from sentence S_0 and inferring $p_1(b_1, \dots, b_n)$ from sentence S_1 under the listed constraints. The hearer first needs to identify the corresponding relations p_0 and p_1 and the corresponding sets of entities a_i and b_i . The discourse inference process is then reliant on performing comparison and generalization operations on these corresponding representations (Scha and Polanyi, 1988; Hobbs, 1990; Prüst, 1992; Asher, 1993). This process yields the relation p that subsumes both p_0 and p_1 , as well as the properties q_i that apply to the a_i and b_i . We will refer to the inference process underlying the recognition of Resemblance relations as *Common Topic* inference (Kehler, 1994a).

Each of the example passages discussed in Section 2.1.1 is coherent by virtue of the inferences resulting from identifying such parallel elements and properties. Passage (17), for instance, is coherent by virtue of identifying that John and Fred participate in a common relation with respect to Clinton, for instance, *having done something to support* him. This is the p that is inferred by computing the common ancestor of the actual predicates denoted by the sentences. Again, John and Fred are similar by virtue of the fact that they are two people presumably known to the hearer at the time of utterance. For passage (20), by inferring an *Exemplification* relation one infers that John is a young aspiring politician and that he is a Democrat, since *campaigning hard for Clinton* is taken as an instance of *supporting his party's candidate*. None of these relations are expressed explicitly in the text; they result from generalizations that need to be made to meet the constraints of the coherence relation.

Resemblance relations share the characteristic that they require the identification of parallel entities (i.e., the a_i and b_i) and relations (p_0 and p_1) as arguments to the constraints of the relation. The argument identification process therefore requires that the semantics of subsentential constituents be accessed. Because we are assuming that the syntactic structure is available to the discourse inference process, this structure can be used both to guide the identification of parallel elements (e.g., in allowing the utilization of structural cues), and as a point at which to access their semantic representations. After this is accomplished, the comparison and generalization operations can be applied pairwise on those arguments.

2.2.2 Cause-Effect Relations

To understand utterances standing in a *Cause-Effect* relation, a path of inference must be established between the situations (i.e., events or states) described in the participating utterances as a whole. In all four cases summarized in Table 2.2, the hearer is to infer P from sentence S_1 and Q from sentence S_2 under the constraint that the presuppositions listed be abduced (Hobbs et al., 1993).⁶ We will refer to the inference process underlying the recognition of Cause-Effect relations as *Coherent Situation* inference (Kehler, 1994a).

Each of the example passages discussed in Section 2.1.2 is coherent by virtue of generating a presupposition. Beyond what is asserted by the two clauses individually, understanding each of these sentences requires presupposing that *being a politician implies being dishonest*. A hearer of any of these sentences who does not share this belief may appropriately respond by asking *Why do you assume that politicians are dishonest?*, even though the speaker did not assert this belief directly.

Cause-Effect relations share the characteristic that they only require the identification of the sentential-level semantics for the clauses as a whole (the P and Q); unlike Resemblance relations there are no p , a_i , or b_i to be independently identified as arguments to the relations. The argument identification process therefore only requires that this proposition be read off of the top-level sentence node of the syntax; the semantics of subsentential constituents do not need to be accessed. Given the two propositions, the abduction of an implication between them can then be attempted.

2.2.3 Contiguity Relations

Understanding utterances standing in a *Contiguity* relation requires that one identify a coherent sequence of events centered around a common system of entities. Whereas the constraints for the other two types of relation and the type of inferential process underlying their application are relatively well-understood, it is much less clear what constraints need to be satisfied in order for a text to be coherent under a Contiguity relation. The sole Contiguity relation, Narration, is also unlike the other relations in that it remains in a class by itself. Similar relations do not result from negating a proposition denoted by one of the clauses or reversing the clause ordering.

Past treatments of coherence relations (e.g., Halliday and Hasan (1976), Longacre (1983)) have equated Narration with temporal progression, the only constraint being that the events described in the discourse display forward movement in time. Hobbs (1990) shows that this is not enough; he discusses passage (16), repeated below as (35).

⁶Again, we are using implication in a loose sense here, meaning “could plausibly follow from”.

- (35) At 5:00 a train arrived in Chicago. At 6:00 George Bush held a press conference.

As discussed in the introduction to this section, to understand passage (16) as a coherent Narration, hearers generally infer that Bush is on the train, or that the train arrival is otherwise somehow relevant to him. Thus, the temporal progression displayed by this passage is not enough for coherence under Narration; additional inferences are required.

Narration can perhaps best be seen as a mechanism for describing complex situations that require more than one clause to communicate. Where the type of connection for Resemblance and Cause-Effect are manifest in common relations and implicational relationships, respectively, the points of connection between successive utterances in a narration is provided by the system of entities that are central to the passage. This aspect of narratives motivates work in identifying discourse “foci” for problems such as pronoun resolution; we will return to the topic of pronoun resolution and the three categories of coherence relations in Section 5.4. In any case, the characterization of the exact constraints required for *Contiguity* remains an open question.

In this section, we have described properties of the different types of discourse inference processes underlying the establishment of coherence relations in our three classes. The constraints and processes underlying Resemblance and Cause-Effect relations are more well-understood than those for Contiguity, and the analyses presented in this thesis will focus on the first two. In the next section, we describe some basic principles underlying our analysis of coherence relations.

2.3 Basic Principles of Coherence

A common objection to theories of coherence that posit a set of coherence relations is that they are inadequate on scientific grounds, in the sense that no explanatory basis is provided for arriving at the particular set posited. The relations are presented as a laundry-list; one instead desires an explanation for why a particular set of relations is primitive, as opposed to a set containing more, fewer, or different relations. The following passage from Hobbs (1978) summarizes this goal:

“Coherence in discourse can be characterized by means of a small number of coherence relations which are definable in terms of the operations of an inference system.” (Hobbs, 1978, page 3)

Our categorization paves the way for an account of how our coherence relations might be derivable from more basic principles. The two types of discourse inference underlying Resemblance and Cause-Effect relation recognition (i.e, Common Topic

and Coherent Situation inference respectively) can be seen as being based on two familiar operations from artificial intelligence: identifying common ancestors of sets of objects with respect to a semantic hierarchy, and computing implication relationships with respect to a knowledge base. The two types of operation can be seen as corresponding to a sort of reasoning by analogy and reasoning by cause and effect. Taking these two basic operations as indicative of the two classes, other factors then differentiate the relations in each class.

The operation that is central for Common Topic inference is the computation of common ancestors. We consider the possibilities given that we start with a pair of relations p_0 and p_1 , and entities participating in these relations a_i and b_i . In the base case, common ancestors for each pair are identified; this is the definition of the *Parallel* relation. Negating either a common ancestor for the relations, or a common ancestor for a property of a pair of entities results in the two definitions of *Contrast*. Neither *Parallel* nor *Contrast* discriminate clause ordering. The case where the common ancestor is already one of the classes explicitly specified by one of the relations or entities (as opposed to having to generalize from both in a pair) results in *Exemplification* for one clause ordering, and *Generalization* for the other ordering. The limiting case is where both the relation and the entities specify the *same* objects; this is *Elaboration*.

On the other hand, the operation that appears to be central for Coherent Situation inference is (a loose form of) implication. We see a similar pattern in the definitions for the Cause-Effect relations. The base case is where the semantics of the first clause implies that of the second, which is *Result*. Reversing the clause ordering results in *Explanation*. Negating the second proposition in each of these cases leads to *Violated Expectation* and *Denial of Preventer*, respectively.⁷ (These presuppositions would be the same for pure implication.) As we have indicated, the Contiguity relation *Narration* is alone in its class, additional relations in this class are not derivable by reversing clause ordering or changing polarity.

We note that there is no one-to-one relationship between conjunctions and coherence relations; instead the mapping is many-to-many. On the other hand, conjunctions do appear to constrain the type of relation that can be inferred. We might treat

⁷One might ask whether any relations result from negating the proposition denoted by the first clause instead of the second. We have not seen a relation of this sort proposed in the literature, but one might call such a relation “Denial of Cause”. Two versions of this relation could be seen to hold of sentences (36) and (37).

(36) Bill is not a politician, but he is still dishonest.

(37) Bill is dishonest, but he is *not* a politician.

The presuppositions required for sentences (36) and (37) are of the form $\neg P \rightarrow Q$ and $\neg Q \rightarrow P$ respectively.

certain conjunctions as bearing semantic features that interact with the discourse inference process. For example, consider the canonical relations of the three Hume categories: *Parallel* (Resemblance), *Narration* (Contiguity), and *Result* (Cause-Effect); these are the relations that correspond to standard clause ordering without any of the arguments to the relation being negated. Interestingly, these are just the three meanings of “and” that are commonly cited in the literature. They are exemplified by the sentences given in (38).

- (38) a. Bill went to the movies, and Hillary went to the store. (Parallel)
 b. Bill went to the movies, and (then) he came home. (Narration)
 c. Bill went to the movies, and (as a result) Hillary got upset. (Result)

In a feature-based system it is not necessary to treat *and* as being ambiguous. It can instead be seen as having negative values for all features that would indicate a relation other than the canonical one in its class. The three meanings result from the fact that there are three basic categories of coherence relation and that *and* is unspecified with respect to these categories.

Similarly, the conjunction *but* has meanings corresponding to both Resemblance relations and Cause-Effect relations. The difference between these two can be seen in examples (39a-b).

- (39) a. Bill went to the movies, but Hillary went to the store. (Contrast)
 b. Bill went to the movies, but (nevertheless) Hillary got upset. (Violated Expectation)

We have categorized *Violated Expectation* as a Cause-Effect relation instead of lumping it with *Contrast*; however it obviously has features of both *Contrast* and Cause-Effect. Therefore, we might consider *contrast* as a distinct and tangential property (perhaps expressed by *but*), where *Cause-Effect* with *+contrast* yields *Violated Expectation*, whereas *Resemblance* with *+contrast* results in what we are currently calling the *Contrast* relation.

A similar approach to decomposing coherence relations has been taken in work by Sanders, Spooren, and Noordman (1992).⁸ They classify coherence relations corresponding to four primitives: the *basic operation* (*additive* or *causal*), the *source of coherence* (*semantic* or *pragmatic*), the *clause order*, and the *polarity* (positive or negative) of the relation. We have not utilized the notion of source of coherence in our account, which is concerned with whether it is the propositions expressed by

⁸I thank Luuk Lagerwerf for bringing this article to my attention.

the clauses that are related, or the illocutions that are related. With regard to the other features, while they use a different terminology for relations than we have, the resulting categories appear to be similar to those we have outlined.

2.4 Relationship to Past Categorizations of Coherence Relations

In this section, we briefly discuss the relationship between our categorization of coherence relations and others that have been proposed in past work.

The relations and definitions that we have used here are either identical to or variants of those proposed by Hobbs (1990). We have also added a relation, *Denial of Preventer*, to round out the Cause-Effect category, and have collapsed Hobbs' *Occasion* and *Ground-Figure* relations into the single relation *Narration*.

Hobbs presents his account in an attempt to answer the following question: *why do we want to call a sequence of utterances a single discourse rather than simply a sequence of utterances?* Viewing this question in terms of the situations in which discourses take place, he identifies four considerations that may hold when a speaker and hearer engage in a discourse, listed in (40).

- (40) a. The speaker wants to convey a message.
- b. The message is in service of some goal.
- c. The speaker must link what he says to what the listener already knows.
- d. The speaker should ease the listener's difficulties in comprehension.

He claims that these considerations give rise to four classes of coherence relations, respectively: the *Occasion* relation, the *Evaluation* relation, the *Ground-Figure* and *Explanation* relations, and the *Expansion* relations. This categorization suggests that a given relation addresses one of these four needs. We do not find this motivation for grouping relations to be particularly convincing, because we submit that *all four* of these considerations hold in any discourse segment, not just one consideration that would in turn suggest a relation from a particular category. However, Hobbs also suggests another way of classifying coherence relations based more on the types of coherence we extract out of the external world. As he notes, his categorization can be seen to correlate with the Hume categorization: Causality, Figure-Ground, and Similarity, which correspond to the categories *Cause-Effect*, *Contiguity*, and *Resemblance* respectively.

Other similar ways of categorizing coherence relations have also been proposed. Halliday and Hasan (1976) discuss four categories: *Additive*, *Temporal*, *Causal*, and

Adversative. The first three of these are strongly reminiscent of Hume's three categories, Resemblance, Contiguity, and Cause-Effect respectively. The final category, *Adversative*, appears to factor out the relations that involve contrast (e.g., they would include *Contrast*, *Violated Expectation*, and *Denial of Preventer* in this class). Also, Longacre (1983) distinguishes four categories, namely *Conjoining*, *Temporal*, *Implication*, and *Alternation*. Again, the first three categories could be seen as terminological variants of Hume's categories. The final category, *Alternation*, is exemplified by passages conjoined with *or*, which we (and presumably others) categorize with *Resemblance* relations. Finally, Sanders, Spooren, and Noordman (1992) discuss a classification that distinguishes between *Additive* and *Causal* relations; Additive relations can be seen to correspond to Resemblance and Contiguity, and Causal relations to Cause-Effect.

The categorization of relations given here goes beyond previous ones in two respects. First, it presents what is to my knowledge the first categorization that is based on a particular aspect of the formal definitions of such relations (that is, the types of arguments identified and the discourse inference procedure underlying recognition), and not only on subjective judgements about what relations intuitively cluster naturally together. Second, the analyses presented in subsequent chapters are *applications* of the theory, using the distinction to predict a certain behavior with respect to distinct linguistic phenomena. Past work has generally been discussed within the confines of a theory of text coherence itself, without any external interpretation-based application identified.

2.5 Other Approaches to Discourse Coherence

In this section, we briefly discuss several other approaches to coherence that are prominent in the computational linguistics literature.

2.5.1 Rhetorical Structure Theory

Rhetorical Structure Theory (Mann and Thompson, 1987, henceforth RST) was proposed as a descriptive theory of text organization. It, like the approach considered here, argues that text structure is hierarchical; pairs of structural components are related in functional terms by one of twenty-three relations. RST relations connect two non-overlapping text spans, termed the *nucleus* and the *satellite*. The nucleus is the text span that is most central to the purpose of the text; the satellite generally provides less central, supportive information.

RST relation definitions are made up of four fields: *Constraints on Nucleus*, *Constraints on Satellite*, *Constraints on the Combination of Nucleus and Satellite*, and

The Effect. There are no constraints on ordering between nucleus and satellite; either can precede the other (although Mann and Thompson note a strong tendency in some cases for one or the other order). A few relations are multi-nuclear, such as *contrast*, *sequence*, and *joint*. An example relation definition, namely *Evidence*, is given below (from Mann and Thompson (1987)).

Relation name: EVIDENCE

Constraints on nucleus: The hearer might not believe nucleus to a degree satisfactory to the Speaker.

Constraints on satellite: The hearer believes the satellite or will find it credible.

Constraints on the combination of nucleus and satellite: The hearer's comprehending the satellite increases the hearer's belief of the nucleus.

Effect: The hearer's belief of the nucleus is increased.

RST has been popular as a basis for natural language generation systems (Hovy, 1993; Moore and Paris, 1993, *inter alia*). However, it does not serve our interest in using it for *interpretation* very well. First, no formal mechanism for relation identification is provided in RST, the definitions are intended more for human analysts. The definitions are informal and descriptive, and thus do not provide a particularly useful formal machinery to discuss interactions with syntax and semantics.⁹ Second, the relations have a laundry-list feel; there is no external justification for having these relations as opposed to a larger or smaller set of them. While RST may provide a useful descriptive tool for text analysts and developers of natural language generation systems, it does not serve as a useful foundation for our purposes.

2.5.2 Discourse Parsing

The Linguistic Discourse Model (LDM) (Polanyi, 1988; Scha and Polanyi, 1988) is a theory of the incremental building of discourse structure. During the processing of a discourse, a *Discourse Parse Tree* is built on a clause-by-clause basis, in direct analogy with how a sentence is built on a constituent-by-constituent basis.

The basic unit is the *discourse constituent unit* (dcu) of which there are four basic types: the *Sequence*, the *Expansion Unit*, the *Binary Structure*, and the *Interruption*. These types are then divided into sub-types. Individual clauses are the minimal units

⁹In fact, as Rambow (1993) and others point out, the relation definitions themselves are often tautologous. For instance, in the definition of *Evidence* above, the effect simply restates the *a priori* constraint on the combination of the nucleus and satellite.

of discourse, being one-element Sequences. Clause dcus are combined together by “proceeding through the discourse, examining the syntactic encoding form of each clause, its propositional content, and its situation of utterance.”

The other primitive category of the LDM is the *discourse operator*. These “modify the force of discourse constituents in some way.” There are three types of discourse operator: *assigners*, *connectors*, and *discourse PUSH/POP markers*. Assigners include clauses such as *hello* and proper names used as vocatives. Connectors include coordinating and subordinating conjunctions (e.g., *and*, *therefore*, *because*). Discourse PUSH/POP markers include phrases such as *o.k.*, *well*, and *anyway*.

Discourse parsing provides a mechanism for interpretation;¹⁰ the coherence structure results from parsing with a discourse grammar on analogy with a sentence-level grammar. We prefer Hobbs’ system however, as it provides more of an explanatory framework in which the relations are derivable through independent cognitive principles, as opposed to positing a more descriptive discourse-level grammar.

2.5.3 Intentional-Level Approaches to Coherence

The approaches to coherence that we have discussed so far are based on the idea that for a discourse to be coherent, one of a limited set of semantic relations must hold between the propositions specified within each of its segments. This position has been referred to as the “informational” view (Moore and Pollack, 1992). In contrast, other researchers (Grosz and Sidner, 1986, inter alia), have argued for what has been called the “intentional” view of coherence. Following work in speech act theory and plan recognition (e.g., Cohen and Perrault (1979), Allen and Perrault (1980)), these researchers have argued that the role of the utterance in fulfilling the speaker’s intentions is the determining factor of coherence. As Hobbs (1993) notes, the informational-level approach has been applied mostly to single-authored texts, whereas the intentional-level approach has been applied mostly to dialogue. In the intentional view, discourse structure and coherence derive from intention-based relations; in the informational view, they derive from rhetorical or coherence relations.

An example that motivates the intentional view is shown in the excerpt from a task-oriented dialogue shown in passage (41), where the expert *B* is assisting the apprentice *A* in fixing an air compressor (Grosz, 1977).

- (41) B. Tighten the bolt with the ratchet wrench.
- A. What’s a ratchet wrench?
- B. It’s between the wheel puller and the box wrenches.

¹⁰Prüst (1992) discusses an implementation of the theory.

In the third sentence, the expert is not answering the question directly, but instead infers the reason the apprentice is asking it (i.e., so he can identify the wrench) and responds to that goal instead. In short, in determining intentional coherence the hearer must answer the question: *Why is the speaker saying this to me?*

More recently, Moore and Pollack (1992) have argued that both the intentional and informational levels of analysis must co-exist in interpreting a discourse. Their argument utilizes passage (42).

- (42) a. George Bush supports big business.
b. He's sure to veto House Bill 1711.

The analyses of concern here are those in which sentence (42b) is the nucleus and sentence (42a) is the satellite (in RST terms). Moore and Pollack argue that this utterance pair satisfies two RST relations: *Evidence*, an intentional-level relation, and *Volitional-Cause*, an information-level relation. That is, the speaker may be stating sentence (42a) as evidence for the central claim being put forth in sentence (42b) and at the same time communicating a Volitional-Cause relationship between the fact expressed by sentence (42a) and the event expressed by sentence (42b). This duality is not surprising, since one way to provide evidence for a proposition is to show that it is a consequence of another proposition that the hearer already believes. Moore and Pollack demonstrate how this connection allows the hearer to recognize a relation at one level by recognizing a relation at the other level, in either direction.

Several formalizations follow their analysis and provide mechanisms for coordinating the two levels (Moore and Paris, 1993; Asher and Lascarides, 1994). In the remainder of this thesis we will only consider the informational level of coherence resolution. It should be kept in mind, however, that our framework is compatible with the co-existence of an intentional level of coherence analysis.

2.6 Conclusions

In this section, we presented a theory of coherence relations that classifies them within three general categories originally suggested by Hume (1748): *Cause-Effect*, *Resemblance*, and *Contiguity*. The categories are distinguished by the type of arguments that are required by the relations, as well as the type of discourse inference mechanisms underlying their recognition. We also provided further motivation for the theory by suggesting ways in which the relations in each class are derivable from more primitive notions.

In the next chapter, we provide an account of VP-ellipsis that breaks the deadlock between purely syntactic and purely semantic theories of resolution. A crucial

component of this theory is our analysis and categorization of coherence relations. After providing a resolution algorithm for VP-ellipsis in Chapter 4, we discuss applications of the theory of coherence to several other language interpretation problems in Chapter 5.

Chapter 3

VP-Ellipsis and Coherence

In this chapter we address the VP-ellipsis construction, exemplified by sentence (43).

(43) Ross likes his mother, and Bill does too.

The stranded auxiliary in the second clause (the *target* clause) indicates the elision of a verb phrase. A representation for the elided VP must be recovered from the representation of another clause, in this case, the first clause (the *source* clause).

As we indicated in Chapter 1, there is an ongoing debate concerning the level of language processing at which VP-ellipsis is resolved. Inherent in syntactic accounts (e.g., Sag (1976), Williams (1977), Haik (1987), Hellan (1988), Lappin and McCord (1990), Hestvik (1993), Lappin (1993b), Fiengo and May (1994)) is the claim that VP-ellipsis is resolved at a level of syntactic structure. On the other hand, inherent in semantic accounts (e.g., Dalrymple, Shieber, and Pereira (1991), Hardt (1991a), Kehler (1993a)) is the claim that VP-ellipsis is resolved at a purely semantic level of representation. This question remains a point of contention, primarily because there are data to support both views.

In this chapter, we provide an account of VP-ellipsis that breaks the existing deadlock between the two approaches. We first show that the seemingly contradictory VP-ellipsis data exhibit a previously unnoticed systematicity, that is, a pattern exists that correlates with the type of coherence relation operative between the source and target clauses. Specifically, we show that the data support syntactic accounts when a Resemblance relation is operative between the clauses, whereas the data support semantic accounts when a Cause-Effect relation is operative. We then argue that the data is accounted for by a theory in which two distinct factors are separated out: (i) the purely semantic resolution of an anaphoric form licensed by VP-ellipsis, and (ii) the need to reconstruct missing syntactic material during Common Topic inference. The distribution of the data is shown to result from the combination of these two independently-motivated processes.

3.1 An Abstract Characterization of the Approaches

To serve as a foundation for the remaining sections of this chapter, we present an abstract characterization of the syntactic and semantic approaches to VP-ellipsis resolution. We discuss how certain instances of past work deviate from these characterizations in Section 3.6.

3.1.1 Syntactic Approaches

Syntactic accounts posit that VP-ellipsis is resolved at a level of syntactic representation. The syntactic representations for the source and target clauses of sentence

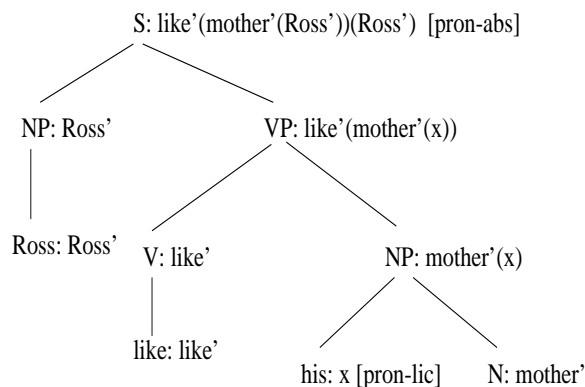


Figure 3.1: Syntactic and Semantic Representations for *Ross likes his mother.*

(43) are shown in Figures 3.1 and 3.2 respectively; the ϕ at the leaf node of the VP in Figure 3.2 indicates that the VP constituent is empty. The process of recovering




Figure 3.2: Syntactic and Semantic Representations for *Bill does.*

the elided VP can be characterized by the reconstruction of a copy of the source VP at the site of the missing target VP; the result of performing this reconstruction for sentence (43) is shown in Figure 3.3. In a syntactic reconstruction account, the

semantics of the target clause is recovered by computing it independently after the VP has been reconstructed from the source.

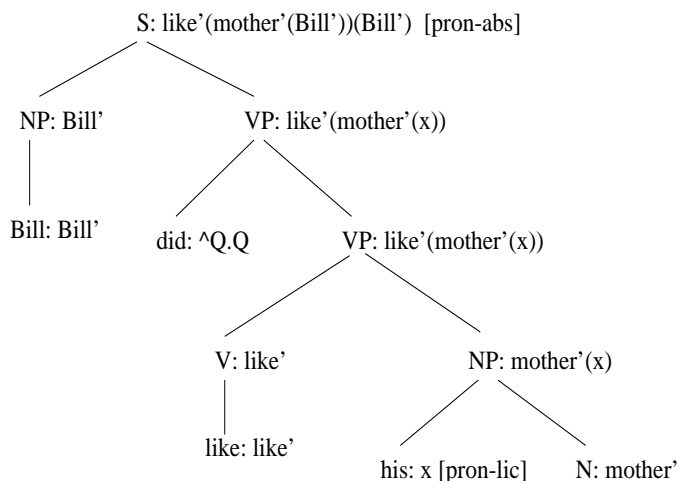


Figure 3.3: Syntactic and Semantic Representations for *Bill does [like his mother]*.

The details of actual syntactic accounts differ in various respects. While some theories posit this type of reconstruction mechanism (Williams, 1977; Kitagawa, 1991; Lappin, 1993b; Hestvik, 1993; Fiengo and May, 1994), others view VP-ellipsis as a process of deletion under suitable conditions (Sag, 1976). Furthermore, some theories operate purely at the level of surface syntax (Lappin, 1993b), whereas others operate at some level of syntactic logical form (Sag, 1976; Williams, 1977; Kitagawa, 1991; Hestvik, 1993; Fiengo and May, 1994). Common to all of these analyses, however, is the requirement that a suitable syntactic source representation be available for reconstruction in the target.

3.1.2 Semantic Approaches

Semantic accounts posit that VP-ellipsis is resolved at a purely semantic level of representation. The semantic representation for the source clause of sentence (43) is given in (44).¹

$$(44) \textit{likes}(\textit{Ross}, \textit{mother_of}(\textit{Ross}))$$

The representation of the target clause contains an uninstantiated relation that applies over the overt material in the target, as shown in (45).

¹The representation and method used here is similar to that used by Dalrymple, Shieber, and Pereira (1991), but is used here simply as an exemplar of purely semantic accounts of VP-ellipsis resolution.

$$(45) P(Bill)$$

In representation (45), P stands proxy for the missing property corresponding to the missing VP in the syntax. We solve for P by computing the relation that when applied to *Ross* results in representation (44); that is, we solve for P in equation (46).

$$(46) P(Ross) = likes(Ross, mother_of(Ross))$$

A solution to this equation is given in (47),²

$$(47) P = \lambda x.like(x, mother_of(Ross))$$

which when applied to *Bill* in the target results in meaning (48).

$$(48) likes(Bill, mother_of(Ross))$$

Some semantic accounts derive a property by abstracting over the propositional meaning of the source clause as a whole (Dalrymple, Shieber, and Pereira, 1991; Kehler, 1993a), whereas others recover a property resulting solely from the meaning of a VP (Hardt, 1992b). Common to these approaches, however, is the characteristic that they require only a suitable semantic source representation for resolution.

There are sets of data to motivate both syntactic and semantic theories that appear to be contradictory when viewed together. In the next section, we argue that there is a pattern to the data that yields a natural separation between these two types.³ We then present an account that predicts this pattern.

3.2 A Pattern in the Data

To determine whether VP-ellipsis is resolved at the level of syntax or at the level of semantics, there are two types of data that one can examine. The first type of data includes examples for which a suitable semantic representation for the source is available for resolution, but for which there is a mismatch of surface-syntactic form between the source and the target syntactic representations. In such cases, if ellipsis shows a sensitivity to syntactic form, then the syntactic approaches are evidenced, as this difference should not be manifest at the level of semantic representation. Conversely, an apparent lack of such sensitivity would provide evidence for semantic approaches. The second type of data include examples for which there is no syntactic

²Representation (47) is only one such possibility; this relation leads to the *strict* reading for the target.

³This pattern is also articulated by Kehler (1993b, 1994b).

form mismatch, but for which reconstruction of the source VP within the target would result in a syntactic constraint violation in the target. In these cases, if the source and target pair is unacceptable because of such a constraint violation, then the syntactic theories are supported; again a lack of such an effect supports a semantic analysis.

In this section, we examine six types of elliptical contexts: two in the first category (specifically, *voice alternation* and *non-VP antecedents*), and four in the second (specifically, *Condition A-C violations* and *subjacency violations*). We show that there are data in each of these contexts to support both approaches. This data exhibit a consistent pattern, however, once the nature of the coherence relationship operative between the source and target clauses is considered. Specifically, when a *Resemblance* relation holds between the two clauses (such as in sentence (43)), the data support a syntactic analysis. That is, a syntactically-parallel source VP must be available, and reconstruction of this VP in the target clause is subject to syntactic constraints. On the other hand, when a *Cause-Effect* relation holds between the clauses, the data support a semantic analysis. That is, neither a syntactic source VP nor compliance with syntactic constraints is required, instead only a suitable semantic source representation is necessary. We show this pattern in the sections that follow.⁴

3.2.1 Voice Alternation

In sentences (49)–(51), an elided target VP in the active voice receives its interpretation from a source clause in the passive voice.⁵

- (49) In March, four fireworks manufacturers asked that the decision be reversed, and on Monday the ICC did. [reverse the decision]
(from text of Rosenthal (1988), cited in Dalrymple (1991))⁶
- (50) This problem was to have been looked into, but obviously nobody did. [look into the problem] (Vincent Della Pietra, in conversation)
- (51) Of course this theory could be expressed using SDRSs, but for the sake of simplicity we have chosen not to. [express this theory using SDRSs] (from text of Lascarides and Asher (1993))

Less frequent, but still evidenced, are cases in which a target in the passive voice receives its interpretation from a source in the active voice.

⁴We discuss Hume’s third category of relation, *Contiguity*, in Section 3.7.1.

⁵The phrase shown in brackets after each example indicates the elided material under the intended interpretation.

⁶Note that the meaning of *and* corresponds with the *Result* relation in example (49), in contrast to the *Parallel* relation operative in example (43). Likewise, the meaning of *but* in examples (50) and (51) corresponds to the *Violated Expectation* relation.

- (52) In addition to inducing lethality during the first instar, it [= a chemical agent] retards embryonic development, but not to the extent that the growth cones were. [retarded] (from text of Jarecki (1992))
- (53) Actually I have implemented it [= a computer system] with a manager, but it doesn't have to be. [implemented with a manager]
(Steven Ketchpel, in conversation)

These examples are problematic for a syntactic analysis of VP-ellipsis, because the VP that needs to be reconstructed is not a surface VP constituent in the source clause. These examples, which all involve Cause-Effect relations, therefore lend support to a purely semantic account of VP-ellipsis resolution, because the predicate-argument relations in the semantic representation are the same regardless of the voice used in the syntax.

However, examples of VP-ellipsis with voice alternation are not always well-formed. For instance, the cases of Resemblance shown in (54) and (55), which are otherwise similar to examples (50) and (52), are unacceptable.

- (54) # This problem was looked into by John, and Bob did too.
- (55) # This agent retards embryonic development, and the growth cones were too.

Syntactic accounts correctly predict these cases to be unacceptable. In contrast, they are problematic for purely semantic accounts, which posit no requirement for syntactic parallelism.

Judging from this data, syntactic parallelism is apparently required in examples in which a Resemblance relation is operative, but not in examples in which a Cause-Effect relation is operative.

3.2.2 Non-VP Antecedents

In examples (56) and (57), the semantic representation for the source is evoked by a nominalization, and not a syntactic VP.⁷

- (56) This letter deserves a response, but before you do, [respond]
(Gregory Ward, personal communication)
- (57) Today there is little or no OFFICIAL harassment of lesbians and gays by the national government, although autonomous governments might. [harass lesbians and gays] (Hardt, 1993)

⁷Some speakers find these cases slightly awkward or stilted. However, most find them to be significantly better than the cases of Resemblance discussed below.

Semantic analyses successfully account for these cases presuming that nominalizations make event representations available for anaphoric reference. However, nominalizations do not appear to result in a suitable source representation in examples involving Resemblance, as shown in sentences (58) and (59).

(58) # This letter provoked a response from Bush, and Clinton did too. [respond]

(59) # There is unofficial harassment of lesbians and gays by the American government, and the Canadian government does too. [harass lesbians and gays]

A similar pattern is seen in cases in which the source representation for the ellipsis is evoked by an adjectival phrase, as illustrated in example (60).

(60) First person pronouns aren't very shiftable, although the plural ones can be. [shifted] (Barbara Partee, in conversation)

The acceptability of example (60) contrasts with the relative unacceptability of the otherwise similar case of Resemblance shown in sentence (61).

(61) # First person pronouns aren't very shiftable, and the plural ones also can't be. [shifted]

Again, the data appear to support a syntactic theory of ellipsis resolution in cases of Resemblance, whereas the data appear to support a semantic theory in cases of Cause-Effect.

3.2.3 Binding Theory Constraints

Another test for determining whether VP-ellipsis is a syntactic or semantic phenomenon is provided by examples in which reconstruction of a syntactic representation in the target would lead to a violation of one of the conditions of Binding Theory (Chomsky, 1981).

We will forgo the definitions of the conditions, referring the reader unfamiliar with them to discussions in Chomsky (1981) or Lasnik (1989). Examples that violate each of the conditions are given in (62b-d), which are to be compared with the acceptable sentence given in (62a).

(62) a. John_i likes himself_i

b. # John_i thinks Sue likes himself_i. (Condition A violation)

c. # John_i likes him_i. (Condition B violation)

d. # He_i likes John_i. (Condition C violation)

To determine whether VP-ellipsis resolution is syntactic or semantic, we can see whether target clauses violate these constraints. Sensitivity to syntactic constraints would evidence syntactic approaches, since reconstruction would presumably require adherence to them. An absence of such sensitivity would evidence semantic approaches, because such constraints do not apply at that level of representation.

It should be noted that the argument being made in this section does not require that one adopt the position that Binding Theory articulates the correct formulation of these constraints. Instead, it is only dependent on the view that there is *some* sort of constraint that is being violated in these sentences, and that such constraints are in some part syntactic in nature. We now consider the three conditions in turn.

Condition A Violations As predicted by Condition A of Binding Theory, it is generally difficult to obtain a strict reading when the source clause contains a reflexive pronoun, as shown in sentences (63) and (64).

(63) ?? John_i defended himself_i, and Bob_j did too. [defend John_i]

(64) ?? Fred_i voted for himself_i, and Gary_j did too. [vote for Fred_i]

Given appropriate semantic context, judgements improve but the examples remain somewhat stilted, as shown in sentences (65) and (66).⁸

(65) ? The alleged murderer_i defended himself_i, and his lawyer_j did too. [defended the alleged murderer_i]

(66) ? Bill Clinton_i voted for himself_i, and his campaign manager_j did too. [voted for Bill Clinton_i]

The stiltedness of reflexives under a strict reading disappears, however, in cases of Cause-Effect (from Dalrymple (1991)).⁹

(67) Bill_i defended himself_i against the accusations because his lawyer_j couldn't. [defend Bill_i]

(68) John_i voted for himself_i even though no one else_j did. [vote for John_i]

In these cases, the strict reading is readily available and perhaps preferred. There appears to be a syntactic dependency manifest in the Resemblance cases that is absent from the Cause-Effect ones.

⁸There appears to be a dialect that readily allows strict readings with reflexives. However, even for those speakers, the Cause-Effect cases given below tend to be more acceptable under the strict interpretation than the Resemblance cases.

⁹This fact was noted by Hestvik (1993) for cases of syntactic subordination.

Condition B Violations Sentences (69) and (70) are also unacceptable, as predicted by Condition B of Binding Theory.

(69) # John's_j mother introduced him_j to everyone, and he_j did too. [introduce himself_j]

(70) # John_i's lawyer defended him_i, and he_i did too. [defend himself_i]

Resolution results in a Condition B violation in the target in a syntactic reconstruction analysis, as shown by the unelided versions of sentences (69) and (70) given in sentences (71) and (72).

(71) * John's_j mother introduced him_j to everyone, and he_j introduced him_j to everyone too.

(72) * John_i's lawyer defended him_i, and he_i defended him_i too.

Semantic accounts predict sentences (69) and (70) to be acceptable. However, in examples in which a Cause-Effect relation is operative, similar sentences are in fact acceptable.

(73) John's_j mother introduced him_j to everyone before he_j had a chance to. [introduce himself_j]

(74) John_i's lawyer defended him_i because he_i couldn't. [defend himself_i]

Again, it appears that syntactic constraints apply in cases of Resemblance but not in cases of Cause-Effect.

Condition C Violations Fiengo and May (1994) and Lappin (1993b) note the unacceptability of examples such as (75) and (76).

(75) # I hit Bill_i, and he_i did too. [hit Bill_i]

(76) # Mary introduced John_j to everyone, and he_j did too. [introduced John_j to everyone]

In a syntactic reconstruction analysis, this results from a Condition C violation, as shown by the unelided versions of sentences (75) and (76) shown in sentences (77) and (78).

(77) * I hit Bill_i, and he_i hit Bill_i too.

(78) * Mary introduced John_i to everyone, and he_j introduced John_j to everyone too.

Again, semantic accounts predict sentences (75) and (76) to be acceptable. However, Dalrymple (1991) gives examples in which Condition C should apply but apparently does not.

(79) I expected Bill_i to win even when he_i didn't. [expect Bill_i to win]

(80) The lawyer defended Bill_j against the accusations because he_j couldn't. [defend Bill_j against the accusations]

Unlike the Resemblance relations operative in sentences (75) and (76), sentences (79) and (80) exhibit Cause-Effect relations. Again, it appears that syntactic constraints apply in cases of Resemblance but not in cases of Cause-Effect.

In sum, the same pattern we have seen with the syntactic form mismatches also arises with each of the Binding Theory constraints. Sensitivity to syntax appears to apply in examples of Resemblance, as if the syntactic material were being reconstructed. On the other hand, this sensitivity does not appear in examples of Cause-Effect, suggesting that in these cases only semantic material is being recovered.

3.2.4 Subjacency Violations

Haïk (1987) gives examples of apparent subjacency violations in cases of antecedent-contained deletion (ACD), exemplified by sentence (81).

(81) John read everything which Bill believes he did. [read ϕ]

A syntactic reconstruction approach correctly predicts that sentence (81) does not display a subjacency violation, since the dependency between the resulting gap and its antecedent only crosses one bounding node (i.e., S or NP).

In sentences (82) and (83), however, the reconstructed gap would be two bounding nodes away. As predicted by the subjacency constraint, these sentences are unacceptable.

(82) # John read everything which Bill believes the claim that he did. [read ϕ]

(83) # John read everything which Bill wonders why he did. [read ϕ]

This data is problematic for a purely semantic theory of ellipsis, because subjacency is a constraint on syntactic representations. However, sentence (84), which has a subordinate conjunction, does not display the expected subjacency violation (from Rooth (1981)).

(84) Which problem did you think John would solve because of the fact that Susan did?

Without ellipsis, the gap remains and a subjacency violation results, as shown in sentence (85).

(85) # Which problem did you think John would solve because of the fact that Susan solved?

It is not clear how to categorize sentences containing relatives headed by *which* with respect to our categorization of coherence relations. It appears that NP relatives and sentence-level relatives behave differently; NP relatives behave like cases of Resemblance and sentence-level ones behave like cases of Cause-Effect. As can be seen in the unfronted correlate of example (84) given in sentence (86), example (84) is derived from a sentence containing a sentence-level relative that is marked as a Cause-Effect relation.

(86) You thought John would solve which problem because of the fact that Susan did?

We will probe this issue further in Section 3.7. If we assume this dichotomy of categorization between relative pronouns, the source representations for cases of Resemblance are apparently syntactic, whereas those for cases of Cause-Effect (e.g., example (84)) are semantic.¹⁰

The data given throughout this section suggest that VP-ellipsis resolution copies a syntactic representation in examples in which a Resemblance relation is operative, and a semantic representation in examples in which a Cause-Effect relation is operative. In the following sections, we present an account that predicts this behavior. In the next section we discuss two types of information that are missing in the representations of a clause in which the VP has been elided. We then describe how these types of missing information are recovered in the context of the two types of discourse inference underlying the recognition of Resemblance and Cause-Effect relations.

¹⁰There are different approaches to handling cases of antecedent-contained deletion in the syntactic literature. In order to avoid an infinite regress of syntactic copying, accounts that apply at surface syntax (e.g., Lappin (1993b)) treat such cases as *pseudo-gapping*, positing the existence of a gap after the elided material. Accounts that apply at LF (e.g., Fiengo and May (1994)) posit that reconstruction applies after the operation of *quantifier raising*, which similarly avoids this regress. We adopt the latter of these options in our analysis. We also utilize a level of post-surface-syntactic logical form in our analysis of gapping given in Section 5.1.

3.3 Syntactic and Semantic Properties of VP-Ellipsis

In this section, we consider two questions concerning the VP-ellipsis construction: whether it results in an empty node in the syntax, and whether it licenses an independently anaphoric expression in the semantics. In Kehler (1994a), it was argued that in contrast to both gapping and event reference, VP-ellipsis has both of these properties. The results are summarized in Table 3.1. We discuss the properties of

Form	Empty Node in Syntax	Anaphoric in Semantics
Gapping	✓	
VP-Ellipsis	✓	✓
Event Reference		✓

Table 3.1: Syntactic and Semantic Properties of Ellipsis and Event Reference

VP-ellipsis here; discussions of gapping and event reference are given in Sections 5.1 and 5.2 respectively.

It is well-established (Halliday and Hasan, 1976, *inter alia*) that the form of *do* operative in VP-ellipsis is an auxiliary and not the main verb form, in contrast to various superficially similar forms of pronominal event reference (e.g. *do it*, *do that*, and *do so* anaphora). For instance, sentences (87a-f) show that VP-ellipsis is possible with other auxiliaries, but event referential forms are not.

- (87) a. Bill yelled at the press, and Hillary did too. (auxiliary *did*)
 b. Bill yelled at the press, and Hillary did it too. (main verb *did*)
 c. Bill could yell at the press, and Hillary could too.
 d. * Bill could yell at the press, and Hillary could it too.
 e. Bill will yell at the press, and Hillary will too.
 f. * Bill will yell at the press, and Hillary will it too.

Furthermore, the main verb *do* requires a non-stative direct object. This constraint applies in cases of event reference but not in cases of VP-ellipsis, as shown by sentences (88a-d).

- (88) a. Bill likes McDonald's, and Hillary does too.

- b. # Bill likes McDonald's, and Hillary does it too.
- c. Al wants to be president, and Tipper does too.
- d. # Al wants to be president, and Tipper does it too.

We therefore treat VP-ellipsis as leaving a stranded auxiliary in the syntax, commanding an empty verb phrase constituent at the site of the deletion.

It also appears that VP-ellipsis licenses an independently anaphoric form in the semantic representation. Evidence for this is provided by its tendency to pattern with other types of anaphora, such as pronouns. For instance, as described by Lakoff (1976) and Jackendoff (1972), VP-ellipsis and pronouns may be cataphoric¹¹ in similar circumstances. We first consider sentences (89a-d).

- (89) a. # Bill will ϕ , if Hillary will make a statement blasting the press.
(where $\phi = \textit{make a statement blasting the press}$)
- b. If Hillary will make a statement blasting the press, Bill will ϕ .
- c. If Hillary will ϕ , Bill will make a statement blasting the press.
- d. Bill will make a statement blasting the press if Hillary will ϕ .

Whereas sentence (89a) is unacceptable under the reading in which Bill and Hillary both make a statement blasting the press (because of a constraint on cataphoric reference in certain syntactic configurations), sentence (89c) is acceptable with such cataphoric reference. This situation is analogous to the case for object-referring pronouns, as shown by examples (90a-d).

- (90) a. # He will make a fool of himself, if Bill makes a statement blasting the press. (where $He = Bill$)
- b. If Bill makes a statement blasting the press, he will make a fool of himself.
- c. If he makes a statement blasting the press, Bill will make a fool of himself.
- d. Bill will make a fool of himself, if he makes a statement blasting the press.

Cataphora is allowable when the pronoun is embedded as in sentence (90c), as it is for VP-ellipsis in sentence (89c). Similarly, cataphora is not allowable when the ellipsis is not embedded as in sentence (90a), as is the case for VP-ellipsis in sentence (89a). This pattern causes Lakoff (1976, pg. 332) to state that “clearly these [the ellipsis

¹¹That is, they can be introduced before their antecedents are introduced.

and pronoun cases] are related phenomena, and [(89a)] should be blocked by the same constraint that blocks [(90a)]”. Jackendoff (1972, pg. 268) similarly addresses both examples using a precede-and-command rule that blocks coreference.

Another similarity between VP-ellipsis and pronouns is that each may access antecedents from clauses other than the most immediate one. Such reference is not uncommon for VP-ellipsis; Hardt (1990) reports that five percent of the examples in the Brown corpus have an antecedent that is at least two sentences back in the discourse. For instance, he cites example (91), in which the source clause is located two clauses prior to the target clause.

(91) The thought came back, the one nagging at him these past four days. He tried to *stifle it*. But the words were forming. He knew he *couldn't*.

Pronominal reference is well known to allow local but non-immediate antecedents; for instance, the subject pronoun *he* in the last sentence in passage (91) also locates its antecedent from two sentences back. Given these similarities, we conclude that VP-ellipsis licenses a pronominal form in the semantics. The idea that VP-ellipsis is a “proform” dates at least as far back as Schachter (1978), and has been adopted by various researchers since.

Figure 3.4 illustrates the syntactic and semantic properties of the elliptical clause *Hillary did*. The ϕ indicates the empty constituent in the syntax, and P indicates the anaphoric form in the semantics.

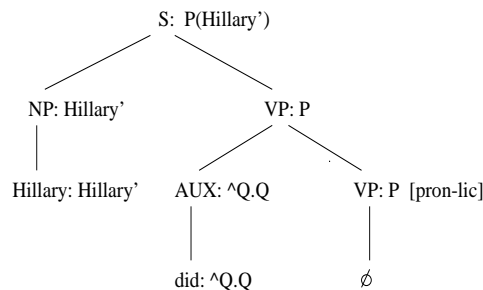


Figure 3.4: Syntactic and Semantic Representations for *Hillary did*.

As it stands, the clause-level semantics for this sentence is incomplete; one can envision two ways of recovering the semantics for the missing VP. First, the syntactic VP could be copied from the source with its corresponding semantics, from which the semantics for the newly complete structure in the target can be derived. This option is advocated by the syntactic reconstruction approach. In this case, the anaphoric expression is constrained to have the same semantics as the copied constituent. Alternatively, the anaphoric expression could be resolved purely semantically. This option

is advocated by the purely semantic approach; in this case the sentence-level semantics is recovered without copying any syntactic material. We discuss the scenarios in which either or both of these resolution processes take place in the next section.

3.4 Interaction with Discourse Inference

In this section, we argue that two distinct processes come into play during the process of resolving VP-ellipsis: (i) the resolution of the anaphoric form licensed by VP-ellipsis, and (ii) the recovery of the missing syntactic representation. Whereas the first of these is the product of general reference resolution processes, the second occurs only during the process of Common Topic inference. The interaction of these two processes results in the distribution of the data as described in Section 3.2.

Anaphora resolution is widely viewed as a process that identifies purely semantic entities as antecedents. Following this view, we posit that the form licensed by VP-ellipsis is free to seek a purely semantic antecedent without regard to any need for syntactic parallelism. Considered in isolation, this characterization follows the view of VP-ellipsis resolution advocated by the semantic approach.

Given that the anaphoric form retrieves the semantics of an elided VP, in our analysis there is no need to reconstruct the missing syntactic material to recover this semantics. However, we claim that reconstruction may be required for a different reason: to perform certain types of discourse inference.

In Chapter 2, the discourse inference process was characterized as having two general steps:

1. *Identify* and *retrieve* the arguments to the coherence relation, and
2. *Apply* the constraints of the relation to those arguments.

To allow for the identification and retrieval of the arguments to the coherence relation, the discourse inference process has access to the syntactic structure of each sentence along with its corresponding semantics. We have posited that the syntactic structure serves as a guide for the determination of (possibly subsentential) parallel elements when necessary, and provides a place from which to access their semantics. Therefore, a problem results when the discourse inference algorithms need to access nodes in the syntax that have been elided, as is the case for VP-ellipsis. We propose that the elided syntactic material is recovered when a node that the mechanisms access is missing, so that the reconstructed structure can be searched for arguments to the coherence relation.

This proposal has different ramifications for our two types of coherence relations. Recall that a distinguishing feature of Resemblance and Cause-Effect relations is the

Type of Inference	Recovers Missing Constituents	Anaphora Resolution
Common Topic	✓	✓
Coherent Situation		✓

Table 3.2: Recovery of Missing Information in the Context of Discourse Inference

type of arguments over which the relation constraints are applied. In the case of Resemblance, the arguments to the relations are sets of parallel entities and relations. Therefore, the discourse inference procedure will have to access subsentential constituents in identifying and retrieving these arguments, including the missing constituent left behind by VP-ellipsis. This will result in the reconstruction of the missing VP node, so that the Common Topic inference process can search beneath it for arguments parallel to those identified in the source. In contrast, the Coherent Situation inference process underlying the recognition of Cause-Effect relations requires only that the sentential-level propositional semantic forms be identified as arguments. These semantic forms are read off of the top-level sentence nodes of the syntactic structures. Because this node is always present (with a complete clausal semantics as a result of the resolution of the anaphoric form), no reconstruction of missing syntactic material is triggered. That is, in this case the missing VP node is never accessed by the discourse inference procedure.

To summarize, in syntactic approaches to VP-ellipsis resolution, syntactic reconstruction results from the need to recover the semantics of the target clause. In semantic approaches, the resolution of the anaphoric form is driven by the same need. The current approach shares properties of both approaches. The resolution of the anaphoric form is a purely semantic process, driven by the need to recover the missing semantics of the VP. However, syntactic reconstruction is also triggered in certain cases, *not by the need to recover the semantics of the target, but instead so that discourse inference can be carried out*. Because only inference with Resemblance relations will require access to the semantics of *sub-sentential* constituents (including the VP), the need to reconstruct will only occur in those cases. We can view the ability to elide a VP as an indication that the information required for processing is recoverable; in the case of inference with Resemblance relations, this information includes its syntactic structure as well as its semantics. The interaction between the two discourse inference mechanisms and the recovery of the two types of missing information is summarized in Table 3.2.

For the purpose of applying our analysis, we can view the requirements for anaphoric resolution and for syntactic reconstruction during Common Topic inference

as mutually-constraining. We therefore elaborate on the first step in the discourse inference process:

1. *Identify* and *retrieve* the arguments to the coherence relation, accessing the semantic arguments through the syntactic representations of the utterances. If an empty syntactic node is accessed, attempt reconstruction of the missing syntactic material. If the result is a valid syntactic structure, proceed; else fail.
2. *Apply* the constraints of the relation to those arguments.

In conclusion, we have separated out the effect of two distinct processes with respect to the distribution of VP-ellipsis data. The resolution of the anaphoric form operates on purely semantic representations. Independently, the recovery of missing syntactic constituents results from Common Topic inference. In the next section, we show that the apparent dichotomy in VP-ellipsis data arises out of the interaction between these different phenomena.

3.5 Applying the Analysis

In this section we show how the data and judgements presented in Section 3.2 follow from the analysis presented in Section 3.4. First, we consider how the discourse inference process works on the simple case of VP-ellipsis shown in sentence (92).

(92) Bill became upset, and Hillary did too.

The syntactic and semantic representations for the source and target clauses are given in Figures 3.5 and 3.6 respectively.

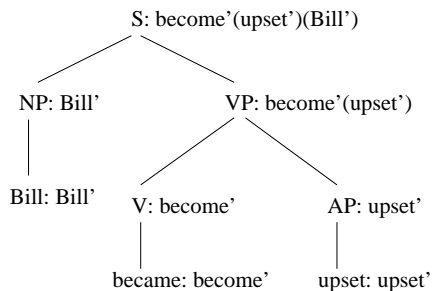


Figure 3.5: Syntactic and Semantic Representations for *Bill became upset*.

We consider the case in which a Parallel relation is inferred between the clauses assuming that *became upset* has been identified as the source clause. The definition of Parallel from Chapter 2 is repeated below.




Figure 3.6: Syntactic and Semantic Representations for *Hillary did*.

Parallel: Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where for some property vector q , $q_i(a_i)$ and $q_i(b_i)$ for all i .

The first step in establishing this relation is to identify its arguments, which are the p 's and the parallel a_i 's and b_i 's in the source and target. This identification requires a search through the syntactic structures of the utterances, and as a result the algorithm attempts to access the missing VP node of the target. Therefore reconstruction of the source syntactic material is invoked. In this case, there is a suitable syntactic source VP, so it is reconstructed from the source as shown in Figure 3.7. The parallel entities and relations can now be identified and retrieved; applying the

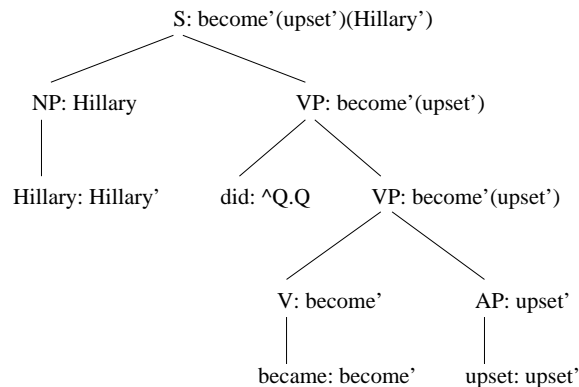


Figure 3.7: Syntactic and Semantic Representations for *Hillary did [become upset]*.

constraints to them successfully establishes the Parallel relation. Note that if the source clause had contained a pronoun that was coreferential with the subject, the target clause may have received either a strict or a sloppy interpretation as a result of semantic resolution. Even though the copied syntactic representation is identical in the source and target, it does not necessarily follow that their corresponding semantic interpretations are identical; in the case of a sloppy interpretation they would not be.

As previously stated, because constructions with Resemblance relations require that the source syntactic representation be reconstructed, the anaphora resolution

process is effectively constrained to require an antecedent corresponding to a source syntactic structure that has an appropriate form. A violation of this constraint may take one of two forms: a mismatch of syntactic form, as exemplified by the voice mismatch in sentence (93), or a syntactic constraint violation after reconstruction has occurred, as exemplified by the Condition C violation in sentence (94).

(93) # This problem was looked into by Bill, and Hillary did too. [look into the problem]

(94) # Hillary introduced Bill_j to everyone, and he_j did too. [introduce Bill to everyone]

We consider first the process of establishing a Parallel relation between the clauses of sentence (93). The syntax and semantics for the source and target clauses are given in Figures 3.8 and 3.9 respectively.

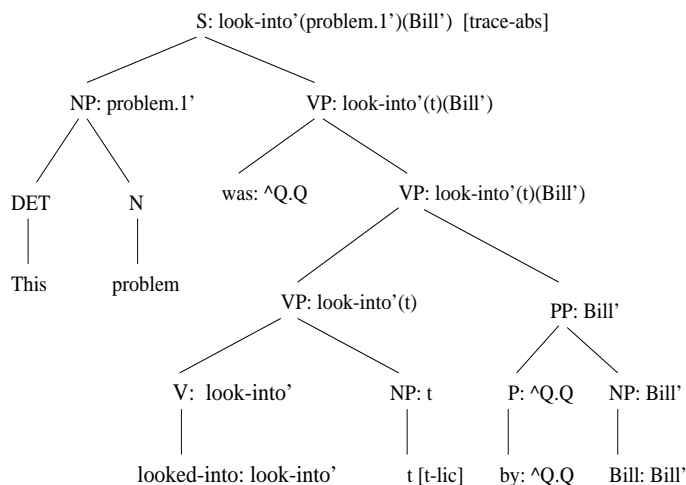


Figure 3.8: Syntactic and Semantic Representations for *This problem was looked into by Bill.*

As in our previous example, Common Topic inference will reconstruct the syntactic representation in order to identify the arguments to the relation. In this case, however, the VP in the intended source is not of a suitable form, because of the voice mismatch between the active voice auxiliary in the target and the trace-licensing passive verb phrase copied from the source. As a result, the reconstruction fails and the sentence is unacceptable.

The case is similar for sentence (94), except that in this case the potential reconstruction is unacceptable because syntactic constraints are violated in the target. The tree for the target clause is shown in Figure 3.10.¹² This tree violates Condition C;

¹²For simplicity, we will represent the term *everyone* as denoting an entity, which of course it does not.




Figure 3.9: Syntactic and Semantic Representations for *Hillary did.*

the pronominal *he* c-commands its cospecifying full NP *Bill*. Therefore, the sentence is unacceptable.

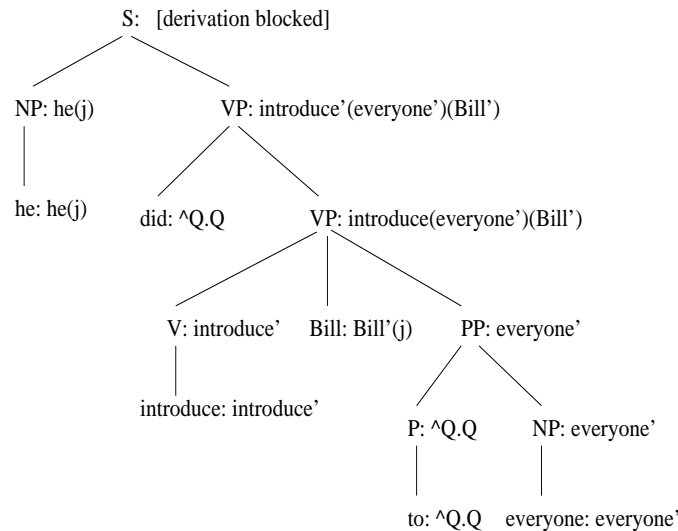


Figure 3.10: Syntactic and Semantic Representations for *he_j did [introduce Bill_j to everyone]*

In contrast with cases of Resemblance relations, Cause-Effect cases appear not to require a suitable syntactic source representation. For instance, sentence (95) is acceptable despite a passive/active voice mismatch, and sentence (96) is acceptable despite a potential Condition C violation in the target.

(95) This problem was to have been looked into, but obviously nobody did. [look into the problem]

(96) The lawyer defended Bill_j against the accusations because he_j couldn't.

The syntactic and semantic representations for the source and target clauses of sentence (95) are given in Figures 3.11 and 3.12 respectively.¹³

¹³We again simplify the analysis by representing the term *nobody* as denoting an entity.

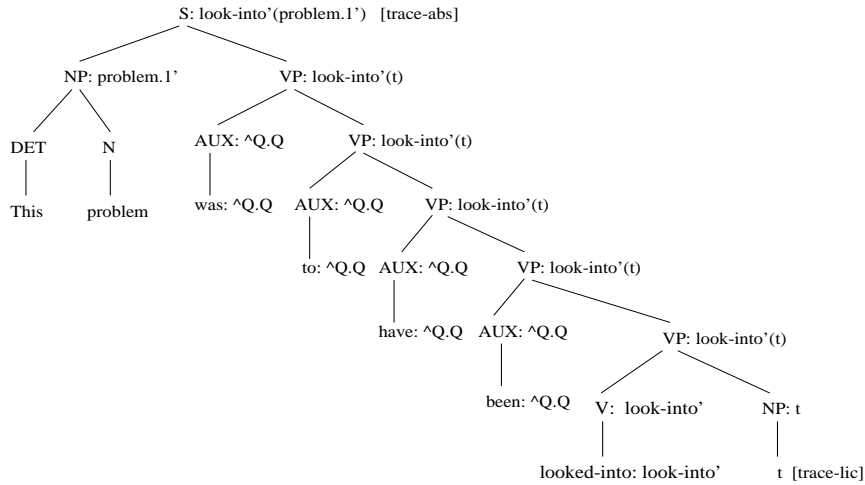


Figure 3.11: Syntactic and Semantic Representations for *This problem was to have been looked into.*

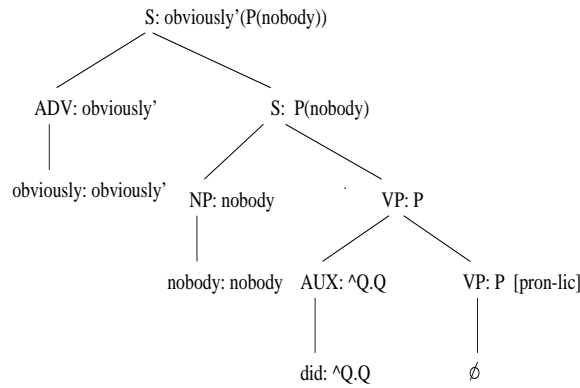


Figure 3.12: Syntactic and Semantic Representations for *obviously nobody did.*

The referential expression P in the semantics is resolved independently by locating the purely semantic antecedent $look-into(problem.1)$; the resulting syntax and semantics for the target is shown in Figure 3.13.

Coherent Situation inference attempts to establish a Cause-Effect relation between the clauses of sentence (95), specifically *Violated Expectation*. The definition of Violated Expectation presented in Chapter 2 is repeated below.

Violated Expectation: Infer P from the assertion of S_0 and Q from the assertion of S_1 , where normally $P \rightarrow \neg Q$.

Again, the first step of the process is to identify the arguments to the relation. As with the other Cause-Effect relations, these arguments are the sentential-level propositions of each sentence. These propositions are accessed from the top-level S node, which is present in the syntactic representations of both the source and target clauses.

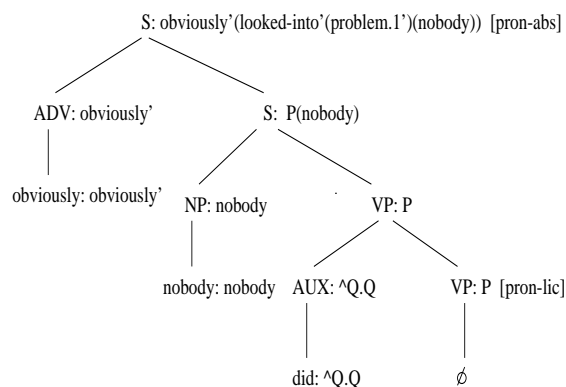


Figure 3.13: Syntactic and Semantic Representations for *obviously nobody did* after resolution.

The discourse inference mechanism reads these off and applies the constraints of the Violated Expectation relation, abducing the necessary presupposition to establish the passage as coherent. No syntactic information in the target has to be reconstructed during this process, because in contrast to Resemblance relations, the semantics of particular subsentential constituents do not have to be accessed as part of identifying the arguments to Cause-Effect relations.

The broad range of data presented in Section 3.2 is therefore predicted by the account. The need for syntactic parallelism in cases of Resemblance is a result of the mutually-constraining processes of anaphora resolution and Common Topic inference. Coherent Situation inference does not constrain the anaphora resolution process in this way, and therefore only a suitable semantic antecedent is required for VP-ellipsis.

3.6 Comparison with Past Work in VP-Ellipsis

Up to this point we have contrasted our analysis of VP-ellipsis resolution with the prototypical syntactic and semantic approaches. In this section, we compare our analysis with specific instances of past work, noting when they deviate from the prototypical accounts.

3.6.1 Syntactic Accounts

Resolution at Surface Structure

In a series of papers, Lappin (1993a; 1993b; 1994, see also Lappin and McCord (1990)), outlines a theory of VP-ellipsis resolution in which the meaning of the target VP is recovered by reconstructing a copy of the surface syntactic representation of the source VP within that of the target clause. His view is essentially the reconstruction

view outlined as the prototypical approach, assuming surface structure as the level of syntactic representation.

In arguing for this choice of level of representation, Lappin notes the existence of binding condition violations of the sort we have noted in Section 3.2.3. However, he acknowledges the existence of acceptable cases in which binding conditions would be expected to apply. Such cases include examples (97) and (98), from Dalrymple (1991).

(97) The lawyer defended Bill_i against the accusations because he_i couldn't.

(98) I expected Bill_i to win even when he_i didn't.

In addressing these cases, Lappin appeals to Evans' (1980) observation that binding condition effects may be in some part overridden by placing contrastive accent on the pronoun. In fact, it is true that the pronouns in examples such as sentences (97) and (98) do generally receive additional accent. Evans gives examples such as (99).

(99) Everyone has finally realized that Oscar is incompetent. Even HE_i has finally realized that Oscar_i is incompetent.

The claim is that in passage (99), the second sentence is acceptable if the pronoun is contrastively accented, despite the expected Condition C violation. Lappin argues that sentences (97) and (98) are acceptable for the same reason.

We remain unconvinced by this argument, however. First, while the discourse context and intended discourse effect help mediate the acceptability of sentence (99), a certain degree of stiltedness remains. In contrast, we find sentences (97) and (98) to be perfectly acceptable, with no hint of the stiltedness of passage (99) nor with any special discourse effect identified. In addition, under this argument, the unelided versions should sound as natural as the elided versions; the unelided versions of sentences (97) and (98) are given in (100) and (101), respectively.

(100) The lawyer defended Bill_i against the accusations because HE_i couldn't defend Bill_i against the accusations.

(101) I expected Bill_i to win even when HE_i didn't expect Bill_i to win.

We find the unelided versions to be much less acceptable; more accent is required for these cases than the elided versions, and the result is still stilted. Furthermore, sentences (100) and (101) display the discourse effect shared by passage (99) that is not present in sentences (97) and (98). Finally, under Lappin's argument, the Condition C violation should be suspended regardless of the construction type. Examples (102) and (103) are similar to sentences (97) and (98), but are in Common Topic constructions with contrastive accent placed on the pronouns.

(102) The lawyer defended Bill_i against the accusations, and HE_i did too.

(103) I expected Bill_i to win, and HE_i did too.

Again, in these cases the added accent may improve the acceptability of the examples somewhat, but does not make them completely felicitous. In fact, the effect here is at best like that in Evans' example (99). Again we find this effect to be notably distinct from the lack of such stiltedness in sentences (97) and (98). Nothing in Lappin's account can distinguish between these differing effects; the difference appears to correspond with our two types of coherence relationship.

Lappin also takes issue with examples having a syntactic form mismatch between source and target that support a semantic view, as exemplified in sentence (104).

(104) A lot of this material can be presented in a fairly informal and accessible fashion, and often I do. (from text of Chomsky (1982), cited in DSP)

Lappin (1993b) states that such examples are also problematic for semantic approaches such as that of DSP. He claims that abstracting over the subject in the source will result in the property shown in (105), which is not a suitable source representation because "it is the property of a deep object term rather than of a deep subject NP".

(105) $\lambda x.\exists y(\text{present}(y, x))$

The property that is needed for successful resolution is that shown in (106).

(106) $\lambda x.\text{present}(x, \text{a_lot_of_this_material})$

Lappin's argument appears to be based on a misconception about the DSP analysis. The parallelism in sentence (104) is at the semantic role level; the agent *I* in the target representation is parallel to the implicit agent in the passive source. This implicit argument is abstracted to yield the property shown in (106), which is applied to *I* in the target. Thus, the element corresponding to the syntactic subject of the source clause is never abstracted as Lappin suggests. Lappin does not address cases of syntactic mismatch other than voice alternation, such as cases with nominalized antecedents.

Resolution at Logical Form

In his extensive study of VP-ellipsis, Sag (1976) provides an account in which VP-ellipsis results from deletion under identity (more specifically, identity under "alphabetic variance") at a level of syntactic logical form. Because of this identity constraint,

the predictions of his analysis correspond to that of the prototypical syntactic analysis presented here. An exception to this arises in the case of Condition A violations; in his account reflexives can always receive either strict or sloppy readings.

Kitagawa (1991) gives an account in which reconstruction in ellipsis resolution occurs at the level of LF in Government and Binding (GB) theory. LF is a syntactic level of representation, in which certain constituents may have been moved from the positions they occupy at surface structure. His paper mainly addresses the binding of pronouns at LF; he claims that the result of copying VPs into the target is subject to binding conditions A, B, and C. His approach is essentially the prototypical syntactic approach that we outlined in Section 3.1, in which the level of syntactic representation is LF. As a result, his account also suffers from all the drawbacks of the syntactic approach to ellipsis. While he predicts the existence of binding effects, he does so in all circumstances, including the acceptable cases given earlier. Also, he does not address cases in which there is a mismatch in syntactic form between the source and the target; presumably his account would predict that such cases would be unacceptable.

Hestvik (1993) describes a proposal that is much like Kitagawa's except that he employs the operation of *reflexive raising* at LF. This operation takes place only when the reflexive is in a subordinate clause. As a result, his analysis accounts for cases in which reflexives allow for strict readings; after raising, the reflexive is c-commanded by the matrix subject and thus can be bound by it. Therefore, Condition A still applies in general, but is satisfied after this movement has taken place. Because reflexive raising only occurs in cases of subordination, the account predicts that no strict reading is available in cases of coordination.

Because cases of syntactic subordination generally co-exist with Cause-Effect relations, our account and Hestvik's make the same predictions for those cases, specifically that strict readings with reflexives should be possible. The two accounts also predict that reflexives are sloppy in cases of Resemblance, as they correspond to coordinate structures. The place in which the two accounts differ, therefore, is with cases of Cause-Effect relations arising from constructions with syntactically coordinating conjunctions. The three circumstances are shown in sentences (107a-c).

- (107) a. John voted for himself, and everyone else did too.
- b. John voted for himself, and (as a result) everyone else did.
- c. John voted for himself because everyone else did.

Both Hestvik's and our accounts predict a sloppy-only reading for sentence (107a), since the sentences are connected with a coordinating conjunction and are in a Resemblance relationship. Both accounts predict both strict and sloppy readings to be possible in sentence (107c) since the sentences are connected with a subordinating

conjunction and are in a Cause-Effect relationship. In sentence (107b), however, the sentences are coordinated syntactically, but are in a Cause-Effect relationship; therefore Hestvik predicts only a sloppy reading where we predict both are possible. While judgements get a bit delicate with such borderline examples, we find sentence (107b) to be acceptable under a strict reading, and notably more so than sentence (107a).

Hestvik does not address cases concerning a mismatch of syntactic form. His approach also patterns with the prototypical syntactic approach with respect to the other binding constraints, invoking the same argument as Lappin did with respect to Evans' work, which we have already addressed.

Fiengo and May (1994) also give an account of VP-ellipsis at the level of LF. While their account also patterns with the syntactic accounts presented here, they do address certain examples that have been used to support semantic analyses. For instance, they discuss the example of voice mismatch shown in sentence (108).

(108) This law restricting free speech should be repealed by Congress, but I can assure you that it won't. [repeal this law restricting free speech]

They address the issue by considering the question of whether a trace, such as one left behind by passivization, can serve as an antecedent of the argument of an elided VP. They conclude that it can, so that examples like sentence (108) are rendered acceptable.

The problem with this account is that it predicts that *all* cases of voice alternation should be acceptable, which we already have established is not the case in examples involving Resemblance relations. Also, it is not clear how their account could be extended to other cases of structural mismatch, such as examples with nominalized and deverbal antecedents.

As far as Binding Theory constraints are concerned, Fiengo and May posit a process called *vehicle change*, which allows pronouns, reflexives, and full NPs to be allowable reconstructions of each other. While strictly speaking binding conditions apply at the level of LF, vehicle change predicts that VP-ellipsis is allowable without regard to binding conditions, because an offending pronoun, reflexive, or full NP can in general be reconstructed as a different form that does not violate any conditions. However, they note that with respect to Condition A violations, "coordination contexts are not fully productive for strict reflexives" (pg. 212). They stipulate that this fact patterns with the type of verb used, stating that "nonextensional predicates of this [verb] class (now taken to include propositional attitude and perception verbs) ... will allow strict as well as sloppy readings" (pp. 212-213). Such verbs are taken to include "vote for", "play", "write about", and "look for"; these do not require that the entity described in their second argument actually exist. It is claimed that verbs not in this class do not readily allow strict readings for reflexives. This distinction

would therefore suggest that sentence (109a) is readily acceptable under a strict interpretation, whereas sentence (109b) is not, although we find both to strongly favor the sloppy interpretation.

- (109) a. John voted for himself and Bill did too.
b. John shaved himself and Bill did too.

This distinction between verb classes extends to Condition B cases in a similar way. Their account suggests that sentence (110a) is acceptable assuming that *him* and *John* are coreferential whereas sentence (110b) is not; again we find both cases to be unacceptable under this interpretation.

- (110) a. Mary voted for him and John did too.
b. Mary shaved him and John did too.

In sum, each of these works primarily patterns with the syntactic approach to VP-ellipsis resolution. In many cases in which an account is modified to handle data supporting a semantic view, other similar data that supports a syntactic view is no longer accounted for.

3.6.2 Semantic Accounts

Purely semantic accounts predict that VP-ellipsis resolution is not subject to any constraints on syntactic form. We discuss two different approaches, although with respect to the issues with which this chapter is concerned, each pattern strongly with the prototypical semantic approach that is described in Section 3.1.

Dalrymple, Shieber, and Pereira (1991) give an account of ellipsis resolution based on an equational analysis utilizing higher-order unification. Their account was used to illustrate the prototypical semantic approach described in Section 3.1. Because their representation is purely semantic, the account does not predict sensitivity to syntactic constraints in any cases. DSP do acknowledge the existence of both acceptable and unacceptable examples of voice alternation, concluding that the parallelism between source and target need not be syntactic. In addition, they acknowledge that there is a similar dichotomy with respect to Condition C violations, concluding that “the resolution of this puzzle remains an open question, as does its incorporation in the present analysis” (pg. 449). Given that our approach is also primarily semantic, with syntactic reconstruction being invoked by an independent process, our account could be viewed as being compatible with and incorporable in the DSP approach.

Hardt (1993) also proposes a purely semantic theory of VP-ellipsis. Unlike the DSP account (as well as the resolution algorithm we present in Chapter 4), antecedents are properties evoked directly into the discourse model by VPs. With respect to syntactic constraint violations, his account patterns with the prototypical semantic account presented earlier, in that no such constraints apply at a purely semantic level of representation.

Hardt’s account differs for the case of syntactic form mismatches, however, since the necessary antecedents in these cases are not evoked directly by VPs. In the case in which the source clause is in the passive voice and the target is in the active voice, he stipulates that what he calls the “active property”, which is the property that the source VP would have contributed if it had been in the active voice, is also placed in the discourse model as a possible referent. However, Hardt does not provide independent justification for making such a move, nor does he describe how this property is computed. Hardt represents nominalized NPs as properties, which are then made available as antecedents.

3.6.3 VP-Ellipsis and Discourse

The account presented in this thesis predicts the split between data supporting purely syntactic and purely semantic theories of VP-ellipsis resolution, making predictions that are beyond those of all previous analyses. The inability of past approaches to account for the data is a result of the fact that they operate solely within one level of language processing; we have argued that an adequate account must take discourse-level factors into consideration. In this section, we briefly compare the current work to three previous studies that explicitly tie ellipsis resolution to an account of discourse structure and coherence, namely our previous account (Kehler, 1993b) and the accounts of Prüst (1992) and Asher (1993).

A Previous Account

An analysis of VP-ellipsis was presented by Kehler (1993b) that distinguishes between two types of relationship between clauses, *parallel* and *non-parallel*. An architecture was presented in which utterances are initially parsed into syntactic propositional representations, which then disappear when they are subsequently integrated into purely semantic discourse model representations. It was posited that VP-ellipsis could access either propositional or discourse model representations: in the case of parallel constructions, the source resided in the propositional representation at the time of resolution; in the case of non-parallel constructions, the source had been integrated into the discourse model at the time of resolution. In Kehler (1994b), we showed how this architecture also accounted for facts concerning gapping constructions, which we

discuss in Section 5.1.

The current account improves upon that analysis in several respects. First, it no longer needs to be posited that syntactic representations disappear when integrated into the discourse model. In the earlier analysis, we could have dispensed with this claim in the treatment of VP-ellipsis, although at the cost of some degree of theoretical inelegance. That is, we would have to have posited that VP-ellipsis somehow searches for a semantic antecedent in the discourse model and subsequently for a syntactic antecedent as a fallback. However, the disappearance of the syntax was crucial for handling the gapping data as presented by Kehler (1994b). In that work, the infelicity of gapping in non-parallel constructions hinged on there no longer being a propositional representation available as a source in those cases (i.e., it would have already been integrated into the discourse model). In the account presented here, the gapping and VP-ellipsis data are explained in a model in which syntactic and semantic representations co-exist.

Second, the previous analysis raised various issues with regard to the interpretation of propositional representations; the main question is whether propositional representations receive a semantic interpretation. If so, it is not clear why we would need to posit a separate discourse-model-level representation. If not, we would be claiming that sentences within a parallel relationship are not interpreted until they are integrated as a unit into the discourse model, a highly questionable claim to have to make. In the current account, these questions are no longer at issue; the separation between the functions of syntax and semantics remains clear.

Third, there was a dichotomy in the previous account with respect to the level of representation from which VP-ellipsis locates and copies antecedents. That is, the same mechanism was posited to resolve at the propositional representation level in some cases, and at the discourse model representation level in others. Given this fairly fundamental distinction, it would not even be clear that we could call these processes by the same name. In the current account, we have a much more unified view of what VP-ellipsis is and how it is resolved. Two distinct factors have been separated out: the resolution of missing constituents under Common Topic inference is purely syntactic whereas the resolution of anaphoric expressions in all cases is purely semantic; the apparent dichotomy in VP-ellipsis data arises out of the interaction between these different phenomena.

Finally, it was not clear how the previous approach scaled up to more complex cases. For instance, consider the minimal pair given in sentences (111a-b), which has a non-parallel construction embedded within a parallel construction.

- (111) a. Clinton was introduced by John because Mary had refused to, and Gore was too. [introduced by John because Mary had refused to]

- b. # Clinton was introduced by John because Mary had refused to, and Fred did too. [introduced Clinton because Mary had refused to]

In each case, the second clause has the first as a source, and the third clause has the first two together as a source. Under the previous approach, resolution of the first ellipsis would have integrated both clause representations into the discourse model, and then in both cases the second ellipsis should succeed. The problem is that the previous approach was too “linear”, not scaling to cases of larger discourse structures. The current approach readily scales up to such cases.

There is one remaining difference between the earlier account and the one presented here, with respect to clauses conjoined with *but*. In the previous account these cases are all classified as *non-parallel*, resulting in the prediction that they only require semantic source representations. In our analysis, we expect cases of pure Contrast to pattern with the *parallel* class since these are Resemblance relations; this is opposed to the Violated Expectation use which is an instance of a Cause-Effect relation. The current account appears to make the correct predictions; examples (112) and (113), in which *but* has the Contrast meaning, appear to be markedly less acceptable than examples (114) and (115), in which *but* has the Violated Expectation meaning.

(112) ?? Clinton was introduced by John, but Mary didn't. [introduce Clinton]

(113) ?? This letter provoked a response from Bush, but Clinton didn't. [respond]

(114) Clinton was to have been introduced by someone, but obviously nobody did. [introduce Clinton]

(115) This letter deserves a response, but before you do, ... [respond]

In summary, the data covered in the earlier account as well as examples that conflict with that analysis are all predicted by the account given here.

Linguistic Discourse Model

Prüst (1992) also gives an account of VP-ellipsis within the context of a theory of discourse structure, namely that of the Linguistic Discourse Model (Polanyi, 1988; Scha and Polanyi, 1988) that was briefly discussed in Chapter 2. He defines a mixed representation (called *syntactic/semantic structures*) that amounts to (unapplied) semantic functions and arguments arranged in a syntactic configuration similar to the surface syntax of the sentence. Therefore, the syntactic and semantic properties of this representation are not as distinguishable as they are in our account. Prüst also gives a method for inferring parallel and contrast relationships by computing the

Most Specific Common Denominator over these structures. However, following Sag (1976), he assumes that VP-ellipsis requires syntactically-matching antecedents, and does not recognize cases that do not require a suitable syntactic source. Because of the extremely tight integration of syntactic and semantic information in his mixed representation, it is not clear how his operations could be adapted to provide a method for making the necessary abstractions in situations in which syntactic constraints do not hold.

Discourse Representation Theory

Asher (1993) provides an analysis of VP-ellipsis and event reference in the context of an account of discourse structure and coherence. Working within Discourse Representation Theory (Kamp, 1981; Kamp and Reyle, 1993), he gives explicit mechanisms for determining coherence between utterances and resolving these forms of reference. However, Asher also follows Sag in requiring that elided VPs be alphabetic variants of their referents. While Asher is working within a purely semantic framework, semantic forms for VPs are represented distinctly from the semantics of sentences; as a result he rules out cases of syntactically mismatched antecedents. In fact, his account of VP-ellipsis and that of other event referential forms (e.g., *do it* and *do that* anaphora) differ in just this regard; the latter allow for the abstraction necessary to allow syntactically mismatched antecedents. It is less clear what his account says about the syntactic constraint violation cases; presumably these constraints would not apply within the DRT framework, so the prediction would be that such constraints would not apply in any constructions. Asher, like Prüst, also defines mechanisms for determining parallelism and contrast that may be seen as algorithms for performing Common Topic inference.

3.7 Remaining Issues

In this section, we briefly address three remaining issues concerning VP-ellipsis and coherence: resolution in the context of Contiguity relations, resolution in the absence of a coherence relation between the source and target clauses, and the resolution of VP-ellipsis with situationally-evoked antecedents.

3.7.1 Contiguity Relations

Up to this point we have considered two of the three types of relation in the Hume classification, Resemblance relations and Cause-Effect relations. There is a third class of relation in his categorization, Contiguity, in which the sole relation is *Narration*. In Kehler (1994a), we grouped this relation with Cause-Effect relations; we assumed

Hobbs' definition of his *Occasion* relation which requires only that the sentence-level semantics of the clauses be checked as arguments to the relation. On the other hand, Sanders, Spooren, and Noordman (1992) classify their Temporal relations in their Additive class, suggesting a grouping with Resemblance relations. In any case, it is hard to say what our account would predict for these, because as indicated in Chapter 2, the exact constraints that need to be satisfied for the *Narration* relation remain to be specified. Examples of syntactic mismatch involving Narration do appear to be of questionable felicity, although perhaps they are slightly better than cases of Resemblance relations, as shown in passages (116a-c).

- (116) a. ?? This letter evoked a response from Bush, and then Clinton did. [respond]
- b. ?? The problem was solved by John, and then Bill did. [solve the problem]
- c. ?? Sue went to John's apartment, and then he did. [go to John's apartment]

A potential problem in judging such passages is that they involve different agents, and so these tend to be understood more as Parallel relations with an explicit "then" noting the additional temporal relations. (Note that the addition of the adverbial *too* in these sentences does not affect the meaning substantially.) In any case, the predictions for cases of Narration would follow from an understanding of the discourse inference processes underlying Narration relation recognition, which is the subject of future work.

3.7.2 Lack of Coherence Relation

The second case we consider are those in which no coherence relationship between the source and the target exists, or at least not one of the sort we have described here. For instance, in Section 3.2.4 we noted that there appears to be a distinction between uses of the relative pronoun *which*; the data presented there suggested that the NP-level use requires syntactic reconstruction, whereas the sentence-level use does not. We have no convincing explanation for why the former case indicates a Resemblance construction; perhaps the shared argument induces parallelism with respect to it. In the case of sentence-level uses, the pronoun specifies an appositive use. We saw in Section 3.2.4 that the subjacency data suggests that this patterns with Cause-Effect constructions, since the subjacency violations expected under a syntactic account were not present. This patterning is confirmed by other naturally-occurring data; in sentence (117) the expected Condition C violation is not present, and in sentence (118), ellipsis is felicitous despite a nominalized source.

- (117) We're fighting for this woman's honor, which is more than she did. [fight for this woman's honor] (Quote attributed to Groucho Marx)
- (118) Before the adjournment for their summer recess on Aug. 12, which they devoutly hope to do, leaders in both houses would like to have debated and voted on a health care bill. [adjourn] ("Health Plans Set Congress On Swift, Fateful Course", New York Times, p. 1, 7/12/94, Gregory Ward, personal communication)

Under the current analysis, it is arbitrary whether we call such uses Cause-Effect relations or not. The important matter is that they are not coherent by virtue of a relation that results in syntactic reconstruction.

3.7.3 Situationally-Evoked Antecedents

Finally, we address cases of situationally-evoked antecedents with respect to VP-ellipsis (referred to as "pragmatically controlled" antecedents by Hankamer and Sag (1976)). In these cases, the antecedent of a referring expression is not evoked by a linguistic expression, but instead by some extralinguistic yet salient event or object in the world. Hankamer and Sag (1976, henceforth H&S) used the ability to refer to situationally-evoked antecedents as a determining factor between their *surface* anaphora, which cannot refer to such antecedents, and their *deep* anaphora, which can refer to such antecedents. Surface anaphoric forms include gapping and VP-ellipsis, which are posited to refer to syntactic structures, whereas deep anaphoric forms include event referential forms such as *do it* and *do that* anaphora, which are posited to refer to purely semantic objects. They give examples (119a-b) to show the distinction; the unspoken situational context is shown in brackets.

- (119) [Hankamer points gun offstage and fires, whereupon a blood-curdling scream is heard. Sag says:]
- a. * Jorge, you shouldn't have! [surface anaphora]
 - b. Jorge, you shouldn't have done it! [deep anaphora]

However, Schachter (1977) provides a number of felicitous examples of VP-ellipsis with situationally-controlled antecedents, such as examples (120) and (121).

- (120) [John tries to kiss Mary. She says:]
John, you mustn't.
- (121) [John pours another martini for Mary. She says:]
I really shouldn't.

Schachter uses such examples to argue for a proform theory of VP-ellipsis, as does Hardt (1993). Lappin (1993b) rebuts Hardt’s argument, stating that syntactic and semantic approaches are on equal ground in handling these cases. However, neither author mentions Hankamer’s response to Schachter (Hankamer, 1978), which argues convincingly that such cases of VP-ellipsis are either formulaic or conventionalized, occurring only as “illocutionally charged expressions” and not generally as declarative statements or informational questions. For instance, the elliptical expressions examples in (122) and (123) are not felicitous, even though the contexts are the same as for Schachter’s examples (120) and (121).

(122) [John tries to kiss Mary. She says:]
* John, you’re the first man who ever has.

(123) [John pours another martini for Mary. She says:]
* John, are you aware that no one else has?

Based on these data, Hankamer argues that the ability to refer to situationally-evoked antecedents does not extend to VP-ellipsis in general, and suggests these forms should be listed in the lexicon instead of being transformationally derived.

Although we maintain that VP-ellipsis resolution is primarily a semantic process, we follow Hankamer in treating it as not being productive in its ability to refer to situationally-evoked antecedents. Based on this, we argue that contra Hankamer and Sag (1976), the questions of whether an anaphoric process requires syntactic parallelism and whether it can refer to situationally-evoked antecedents need to be separated in a general theory of anaphora. We return to this point in Section 5.2, where we claim that *do so* anaphora is another form that enforces this distinction.

3.8 Conclusion

In this chapter, we classified past approaches to VP-ellipsis resolution according to the level of representation at which VP-ellipsis is resolved. We presented data that is problematic for purely syntactic and purely semantic accounts. We showed that this data displays a previously unnoticed pattern, in which syntactic accounts are evidenced when a Resemblance relation is operative between the source and target clauses, whereas semantic accounts are evidenced when a Cause-Effect relation is operative. We then showed how these facts can be explained by two theories: the first being that VP-ellipsis leaves behind an empty constituent in the syntax as well as an anaphoric expression in the semantics, and the second being a categorization of two types of discourse inference in establishing that a passage is coherent, that is, *Common Topic* inference and *Coherent Situation* inference. We showed how the

VP-ellipsis representations and discourse inference processes interact to predict the pattern identified in the data. These predictions appear to be beyond the scope of any theory of VP-ellipsis that operates solely within a single module of language processing.

We will show how this account extends directly to data concerning *gapping* constructions in Chapter 5. In the next chapter, we describe a semantic resolution algorithm for VP-ellipsis.

Chapter 4

A Resolution Algorithm for VP-Ellipsis

In Chapter 3, we employed a prototypical semantic approach for resolving the anaphoric form licensed by VP-ellipsis. In this chapter we provide a particular algorithm for resolution, paying special attention to the derivation of strict and sloppy readings. Sentence (124) exhibits such a strict/sloppy ambiguity.

(124) Ross likes his mother, and Bill does too.

As discussed earlier, the target clause in sentence (124) may receive one of two readings, namely one in which Bill likes *John's* mother (the *strict* reading), and one in which Bill likes *his own* mother (the *sloppy* reading). This ambiguity results from the fact that the source clause contains a pronoun that co-refers with the subject.

We show that our algorithm derives the correct set of readings for a series of examples that have been used as benchmarks for previous approaches. These examples are listed below.

Cascaded Ellipsis: DSP note that sentence (125), due to Dahl (1972), has a reading in which Bill doesn't realize Bill is a fool even though his wife realizes Bill is a fool.

(125) John realizes that he is a fool, but Bill_i does not, even though his_i wife does.

This reading is predicted not to exist by what DSP call *identity-of-relations* analyses, which we discuss in Section 4.5.1. We show that our algorithm, a *non-identity* analysis like that of DSP, derives this reading in Section 4.2.1.

Stripping: In the DSP analysis, it is possible to derive the non-existent reading *Mary likes Mary’s mother* for the target of example (126), in addition to the acceptable reading *John likes Mary’s mother*.

(126) John likes his mother, and Mary’s too.

Our algorithm derives only the correct reading, as we show in Section 4.2.2.

Missing Readings: For examples in which there are N pronouns in the source clause that co-refer with the subject, one might expect that there are 2^N readings for the target. However, Dahl (1974) noticed that example (127) has only three readings, not four, where the reading in which *Harry believed that Bill loved Harry’s wife* is ‘missing’.

(127) Bill believed that he loved his wife, and Harry did too.

We show how our algorithm derives only the three existing readings in Section 4.2.3.

5-Reading Sentence: DSP use sentence (128) as a benchmark for previous analyses, which they claim has five readings.

(128) John revised his paper before the teacher did, and Bill did too.

The DSP approach derives six readings, although they appeal to an unspecified “suitable definition of generalized antecedent linking” and an additional constraint on property derivation to eliminate the sixth. We show how our algorithm derives only the correct five readings for this sentence in Section 4.2.4.

The ability of the algorithm to account for these examples is a result of two properties that distinguish it from past approaches: (i) our encoding the distinction between the semantic representations of full NPs and of the referential elements that cospecify them, and (ii) the manner in which links are established between the representations of these two types of expression.

The remainder of the chapter is organized as follows. In Section 4.1 we present our ellipsis resolution algorithm, and work through the derivations of the readings for example (124). In Section 4.2 we apply the algorithm to the foregoing examples, showing that the analysis derives the correct sets of readings without appeal to additional constraints. In Section 4.3 we describe our implementation of the algorithm, and in Section 4.4 we discuss related phenomena. In Sections 4.5 and 4.6 we compare our algorithm with past approaches.

4.1 The Account

In this section we present our algorithm for VP-ellipsis resolution.¹ We first discuss the properties of the semantic representations that we assume, and follow with a description of the algorithm itself.

4.1.1 Representation

A central claim of our analysis is that in order to account for the distribution of strict and sloppy readings in VP-ellipsis, the representation for pronouns must be distinguished from that of the full NPs to which they refer. We will represent all entities and relations, including pronouns, as unique terms and predicates in the logical representation. All unique terms have an index identifying them as such, although to improve readability these will only be displayed when necessary. A term representing a pronoun has a *link* associated with it which establishes the (one-way) relationship between the pronoun and its antecedent. In some cases, the antecedent term may itself be another pronominal term with a link. The link property is associated with the index of the term to which it refers.

For instance, the semantic representation for the sentence *Ross_i likes his_i mother* is shown in representation (129).

$$(129) \textit{likes}(Ross_{a0}, \textit{mother_of}(him_{a1}^{\rightarrow a0}))$$

The term $him_{a1}^{\rightarrow a0}$ identifies a term *him* having the index *a1*, which is linked to the term representing the antecedent of the pronoun having the index *a0*, in this case $Ross_{a0}$.²

We now consider the case in which there are several pronouns in a clause that are coreferential. Because terms can refer either to (unlinked) full NP terms as antecedents or to other linked terms (i.e., representing other pronouns), there are several possible linking relationships when there is more than one pronoun with the same antecedent in a sentence. We posit a correspondence between binding relations in syntactic representations and linking relations in semantic representations, specifically that given in rule (130).

¹A previous version of this account is provided by Kehler (1993a).

²This system bears similarities with Barwise's (1987) encoding of what he calls *restrained* variables in semantic representations; the pronominal term *him* in representation (129) being an example of such a variable. However, he chooses not to allow restrained variables to restrain other variables (although he considers doing so); the analogous ability to link terms to other linked terms in our representations is a critical facet of our analysis.

- (130) A referential element is linked to the most immediate coreferential element that c-commands it in the syntax. If there are no c-commanding elements, then it is linked to the most salient discourse referent.

That is, when a pronoun is bound by a (possibly pronominal) antecedent, the corresponding pronominal term is linked to that antecedent. The use of c-command is meant to be an approximation for the licensing of bound pronouns; in Section 4.2.3 we will see how this rule fails in just those cases when c-command fails to predict reflexivization. In the case in which a referring expression is not c-commanded by its antecedent (e.g., the antecedent is sentence external), then the most salient available referent is used, although nothing much hinges on which occurrence is selected in these cases.

To illustrate, in Section 4.2.3 we will consider sentence (131).

- (131) Bill believed that he loved his wife, and Harry did too.

Because *he* c-commands *his* in this sentence, the latter pronominal term is linked to the former, as illustrated by representation (132).

- (132) $believe(Bill_{a0}, love(he_{a1}^{-a0}, wife_of(his_{a2}^{-a1})))$

The interpretation of our logic expressions can be recovered in the obvious way. For any linked expression, the links are followed until an unlinked expression is found, which indicates the referent.

In Section 4.2.2 we show how our account correctly handles the *stripping* example; this ability results from the fact that our representation encodes the distinction between full NPs and the referential elements that refer to them. The fact that we link referential elements to their referents in the semantic representation allows us to account for the *missing readings* examples and the *5-reading* example without appeal to any additional constraints on the algorithm, as described in Sections 4.2.3 and 4.2.4 respectively. We now describe our algorithm.

4.1.2 Resolution Algorithm

Our algorithm for ellipsis resolution can be seen as embodying a “semantic reconstruction” approach. The resolution process is characterized by two steps.

1. Identify parallel elements between source and target, and
2. Construct the target form, substituting parallel elements in the source with those for the target.

As we describe the algorithm, we step through the derivation of the readings for example (124) to make the description more concrete. The representation for the source given in (129) is repeated below as representation (133); the representation of the NP *Bill* in the target is $Bill_{b_0}$.

(133) $likes(Ross_{a_0}, mother_of(his_{a_1}^{\rightarrow a_0}))$

Most of the work of the algorithm takes place in the first step; there are three distinct cases to consider. First, a parallel element in the source needs to be identified for each overt element in the target clause. Following DSP, we treat this process as a pragmatic one about which we will not have much more to say, although in all but very exceptional cases the parallel elements fill identical argument positions in the source and target logical forms.³ For sentence (124), the term $Bill_{b_0}$ in the target is determined to be parallel to the term $Ross_{a_0}$ in the source.

The second case concerns those elements in the source that do not correspond to overt elements in the target (i.e., they are part of the elided material), and which are not linked to other elements. In this case, a copy of each term is created with a new index to distinguish it in the semantic representation, although the copy has the same denotation. For sentence (124), copies are made for all of the terms in representation (133) except for $Ross_{a_0}$, since it is parallel to $Bill_{b_0}$ in the target, and for the pronominal term $his_{a_1}^{\rightarrow a_0}$, which is linked to another element.

The third case includes referential elements in the source that are linked to other elements, but which do not have overt parallel elements in the target. In this case, there is an ambiguity that we call the *refer* versus *copy* distinction. In each case we create a copy of the term with a new index; the question is how the term is linked. In the case of *referring*, the link is established between the newly created term and its parallel element in the source. For sentence (124), the reconstructed target term representing *his* is linked to the source term $his_{a_1}^{\rightarrow a_0}$; this results in the term $his_{b_1}^{\rightarrow a_1}$. In the case of *copying*, the link is established between the newly created term and the target element which is parallel to the term to which the source referring expression is linked.⁴ For sentence (124), this is the term $Bill_{b_0}$, which is parallel to $Ross_{a_0}$ (the term to which the source term $his_{a_1}^{\rightarrow a_0}$ is linked); this results in the term $his_{b_1}^{\rightarrow b_0}$. The case of *copying* can be seen as the creation of a new referring expression that mimics

³The question remains as to whether the argument positions that the overt elements in the target fill are recoverable from non-pragmatic factors. In general, an overt subject in the target will fill the first argument position if the target is in the active voice, and the second argument position if the target is in passive voice. We discuss some potentially more problematic cases in Section 4.6.

⁴This target element may also have been elided and thus copied by the algorithm; this places certain constraints on the order in which the linking in the target is established. We return to this point when we discuss our implementation in Section 4.3.

the *role* that its corresponding element plays in the source, as opposed to that element itself.

Having created the parallel elements, step two of the algorithm performs the substitutions to create the target representation. In this case there is an ambiguity depending on which of the two possible parallel elements for *his* are used, the one which *refers* to its antecedent ($his_{b_1}^{\rightarrow a_1}$) or the one which *copies* its antecedent ($his_{b_1}^{\rightarrow b_0}$). Representation (134) for the target results from performing substitution in the case that the pronominal term refers.⁵

(134) $likes(Bill_{b_0}, mother_of(his_{b_1}^{\rightarrow a_1}))$
 Reading: *Bill likes Ross's mother.*

This is the strict reading; the term his_{b_1} is linked to the term his_{a_1} , which is in turn linked to $Ross_{a_0}$.

The other choice is to perform the substitution in which the pronominal term copies, which results in representation (135) for the target.

(135) $likes(Bill_{b_0}, mother_of(his_{b_1}^{\rightarrow b_0}))$
 Reading: *Bill likes Bill's mother.*

This is the sloppy reading; the term his_{b_1} is linked to $Bill_{b_0}$.

It may not be possible to copy in some instances in which a source pronominal term is linked, specifically if the antecedent for the pronoun is intersentential. For instance, in example (136), the target clause can only be interpreted as *Bill loves Al's mother*.

(136) Al_k has a wonderful family. Ross loves his_k mother and Bill does too.

In example (136) there is no element in the target parallel to the antecedent of the source pronominal term (since the antecedent is not in the source to begin with), therefore this option does not apply, as desired. However, there are cases in which parallel elements may not be contained in the minimal clause containing the ellipsis; such cases are discussed in Section 4.4.2.

This algorithm models the resolution of the anaphoric form licensed by VP-ellipsis as discussed in Chapter 3; the resolution proceeds by performing substitution operations on semantic representations. We should note that the need to reconstruct the syntactic material during Common Topic inference does not affect the set of possible strict and sloppy readings. (Except, of course, when readings are eliminated by Binding Theory constraints, as discussed in Chapter 3.) Any choice made for linking a referential element to an antecedent in the semantics has a corresponding possibility for pronoun licensing in a derivation after syntactic reconstruction.

⁵Again, we will only display the indices of terms that are involved in link relationships.

4.2 Examples

In the previous section, we described our algorithm and showed that a basic case of strict and sloppy ambiguity in VP-ellipsis is correctly analyzed. In this section, we show how our algorithm accounts for a variety of examples that are problematic for past accounts. Section 4.2.1 demonstrates that the algorithm accounts for the *cascaded ellipsis* case, thereby retaining the advantages of the DSP approach over identity-of-relations analyses. We then show how our algorithm generates the correct readings for the *stripping*, *missing readings*, and *5-reading* examples.

4.2.1 Cascaded Ellipsis

We demonstrate that our analysis accounts for the case of cascaded ellipsis, repeated as sentence (137).

(137) John realizes that he is a fool, but Bill_i does not, even though his_i wife does.

In particular, we work through the derivation of example (137) that leads to reading (138), the reading that is problematic for identity-of-relations analyses.

(138) John realizes that *John* is a fool, but Bill does not realize that *Bill* is a fool, even though Bill's wife realizes *Bill* is a fool.

The representation for the source clause is given in (139).

(139) $realize(John_{a0}, fool(he_{a1}^{-a0}))$

The second clause is a target having the first clause as its source. For the reading we are concerned with, we want the sloppy option, so we reconstruct the source representation within the negation and *copy* with the pronominal term, resulting in representation (140).

(140) $not(realize(Bill_{b0}, fool(he_{b1}^{-b0})))$

The third clause is a target having the second clause as its source (sans the negation). In this case we want the strict reading, so we reconstruct the source representation and *refer* with the pronominal term, resulting in the representation shown in (141).

(141) $realize(wife_of(he_{c0}^{-b0}), fool(he_{c1}^{-b1}))$

This is the representation for the reading given in (138), showing that the reading not derivable by identity-of-relations analyses is derivable by our algorithm.

4.2.2 Stripping

We now show how our analysis accounts for the *stripping* example, repeated below as sentence (142).

(142) John likes his mother, and Mary's too.

In addition to the reading shown in (143a), in the DSP analysis it is possible to derive the non-existent reading shown in (143b).

(143) a. *John likes John's mother, and John likes Mary's mother.*

b. *John likes John's mother, and Mary likes Mary's mother.*

Our algorithm generates only reading (143a) for sentence (142). The representation for the source clause is given in (144).

(144) $likes(John_{a0}, mother_of(his_{a1}^{\rightarrow a0}))$

In this case, the overt element *Mary* in the target is parallel to the possessive pronoun in the source. Therefore, performing this replacement and reconstructing the form results in the representation in (145).

(145) $likes(John_{b0}, mother_of(Mary))$

Because the overt element in the target was parallel to the pronoun in the source, no pronominal elements are copied to the target and therefore no ambiguity results.

4.2.3 Missing Readings

As we previously noted, in examples in which the source VP contains N referring elements, one might expect that 2^N readings would be possible. However, Dahl (1974) notices that example (146) has only three readings, not four.

(146) Bill believed that he loved his wife, and Harry did too.

Out of the expected readings (147a)-(147d), reading (147d) is missing.

(147) a. Harry believed that *Bill* loved *Bill's* wife.

b. Harry believed that *Harry* loved *Harry's* wife.

c. Harry believed that *Harry* loved *Bill's* wife.

d. # Harry believed that *Bill* loved *Harry's* wife.

Because *he* c-commands *his* in sentence (146), the pronominal term *his* is linked to the pronominal term *he* instead of directly to *Bill*. The representation for the source clause in (146) is shown in (148).

$$(148) \textit{believe}(\textit{Bill}_{a0}, \textit{love}(\textit{he}_{a1}^{\rightarrow a0}, \textit{wife_of}(\textit{his}_{a2}^{\rightarrow a1})))$$

We show that the reading in (147d) is correctly predicted not to exist by applying the algorithm to derive all possible readings for the target clause. In each case we have two options for reconstructing each of the pronominal terms, resulting in a total of four possibilities. In the first case we *refer* with each term, yielding the *all-strict* reading given in (147a).

$$(149) \textit{believe}(\textit{Harry}_{b0}, \textit{love}(\textit{he}_{b1}^{\rightarrow a1}, \textit{wife_of}(\textit{his}_{b2}^{\rightarrow a2})))$$

Reading: *Harry believed that Bill loved Bill's wife.*

In the second case both pronominal terms *copy*, yielding the *all-sloppy* reading given in (147b).

$$(150) \textit{believe}(\textit{Harry}_{b0}, \textit{love}(\textit{he}_{b1}^{\rightarrow b0}, \textit{wife_of}(\textit{his}_{b2}^{\rightarrow b1})))$$

Reading: *Harry believed that Harry loved Harry's wife.*

In the third case we *copy* the pronominal term *he*, but *refer* with the pronominal term *his*, yielding reading (147c).

$$(151) \textit{believe}(\textit{Harry}_{b0}, \textit{love}(\textit{he}_{b1}^{\rightarrow b0}, \textit{wife_of}(\textit{his}_{b2}^{\rightarrow a2})))$$

Reading: *Harry believed that Harry loved Bill's wife.*

These three readings are the acceptable ones for the target clause in (146). The algorithm also allows for a fourth possibility, in which we *refer* with the pronominal term *he* but *copy* with the pronominal term *his*.

$$(152) \textit{believe}(\textit{Harry}_{b0}, \textit{love}(\textit{he}_{b1}^{\rightarrow a1}, \textit{wife_of}(\textit{his}_{b2}^{\rightarrow b1})))$$

Reading: *Harry believed that Bill loved Bill's wife.*

In this case the all-strict reading shown in (147a) is once again derived. The non-existent reading given in (147d) is therefore not derivable by the algorithm.

This behavior is a result of our scheme for linking pronominal terms to their referents in semantic representations. During the derivation of the representation given in (152), we chose the sloppy option for the second pronominal term. Instead of replacing the element itself with its parallel *element* in the target representation as other accounts do, we replaced it with a parallel *link*. This process was “sloppy” in that the resulting link is to the pronominal term *he* in the target; however, the

resulting effect is “strict” because that element itself refers to the pronominal term *he* in the source clause.

The use of c-command as the linking criterion explains why many speakers get all four expected readings for sentences (153) and (154), which are otherwise similar to sentence (146).

(153) Bill believed that his wife loved him, and Harry did too.

(154) Bill believed that his wife loved his brother, and Harry did too.

Since neither pronoun c-commands the other in these cases, and *Bill* c-commands both, the source clause representation would have both directly linked to *Bill* and the algorithm would derive all four readings for the target clause.

However, Sag (1976) notes that example (155) only has three readings.

(155) Edith said that finding her husband nude had upset her, and Martha did too.

Out of the four possibilities for the target, reading (156) is missing.

(156) # Martha said that finding *Martha’s* husband nude had upset *Edith*.

The algorithm derives only the three correct readings if the first instance of *her* in sentence (155) is linked to the second. Since neither pronoun c-commands the other, this violates the linking rule. However, as noted by Reinhart (1983, pp. 179-180), “experiencing” verbs such as *upset* often pose problems for linguistic analyses utilizing c-command. What is required for the linking rule is a notion of a *reflexive context* that applies to pronouns of all cases, for which c-command is an imperfect approximation. For instance, the NP *her husband* in sentence (155) is in a reflexive context, i.e., replacing ‘her husband’ with an accusative pronoun referring to Edith requires the reflexive form, as shown in clause (157).

(157) ... finding herself_i/*her_i nude had upset her_i.

Because the *her* in *her husband* in sentence (155) is in the reflexive context of the second *her*, the corresponding role link is required in the event structure.⁶ Thus, the

⁶The same predictions result if we assume that a PRO category governs *herself*, which is in turn licensed by the clause final *her*. While the link from the first occurrence of *her* will yield ambiguities in the target, the link from PRO to the second occurrence of *her* must be copied to the target during resolution. This is in accordance with the well known fact that PRO in control structures is always understood as sloppy; for instance the target clause in sentence (158) can only mean that Al wanted Al to become president, and not that Al wanted *Bill* to become president.

(158) Bill wanted (PRO) to become president, and Al did too.

algorithm works correctly for example (155); the flaw arises from using c-command in the linking rule to model reflexive contexts. In the future we would expect to revise the linking rule by adopting rules superseding c-command for predicting reflexivization.

4.2.4 5-Reading Example

DSP discuss example (159), from Gawron and Peters (1990), as a point of departure among previous analyses.

(159) John revised his paper before the teacher did, and Bill did too.

DSP claim that sentence (159) has five readings; we agree. The five readings are given in (160).

- (160) a. John revised John's paper before the teacher revised John's paper, and Bill revised John's paper before the teacher revised John's paper.
- b. John revised John's paper before the teacher revised John's paper, and Bill revised Bill's paper before the teacher revised John's paper,
- c. John revised John's paper before the teacher revised John's paper, and Bill revised Bill's paper before the teacher revised Bill's paper,
- d. John revised John's paper before the teacher revised the teacher's paper, and Bill revised John's paper before the teacher revised the teacher's paper,
- e. John revised John's paper before the teacher revised the teacher's paper, and Bill revised Bill's paper before the teacher revised the teacher's paper,

A notably absent sixth reading is given in (161).

(161) John revised John's paper before the teacher revised John's paper, and Bill revised John's paper before the teacher revised Bill's paper.

Our algorithm as stated generates only the correct five readings. We step through the derivation of these, showing that a representation for reading (161) is not derivable. We first derive the possible readings for the first ellipsis. The representation for the source clause is shown in (162).

(162) $revise(John_{a0}, paper_of(his_{a1}^{\rightarrow a0}))$

Again, the pronominal link gives us two options. First, we may use the pronominal term to *refer*, yielding the *strict* reading given in (163).

(163) $revise(teacher_{b0}, paper_of(his_{b1}^{-a1}))$

Alternatively, the pronominal term may be *copied*, yielding the *sloppy* reading shown in (164).

(164) $revise(teacher_{b0}, paper_of(his_{b1}^{-b0}))$

We now consider the readings for the second ellipsis. Unlike the case of cascaded ellipsis, in this example the second ellipsis has the entire conjoined clause as its source. We first consider the readings derived for the case in which the first ellipsis has the strict reading. The representation for the source in this case is given in (165).

(165) $before(revise(John_{a0}, paper_of(his_{a1}^{-a0})),$
 $revise(teacher_{b0}, paper_of(his_{b1}^{-a1})))$

There are two pronouns in the source, giving rise to four options. First, we can use both pronominal terms to *refer*, yielding the *all-strict* reading for the target shown in (166).

(166) $before(revise(Bill_{c0}, paper_of(his_{c1}^{-a1})),$
 $revise(teacher_{d0}, paper_of(his_{d1}^{-b1})))$

Reading: *Bill revised John's paper before the teacher revised John's paper*

In the second possibility, both pronominal terms may be *copied*, resulting in the *all-sloppy* reading shown in (167).

(167) $before(revise(Bill_{c0}, paper_of(his_{c1}^{-c0})),$
 $revise(teacher_{d0}, paper_of(his_{d1}^{-c1})))$

Reading: *Bill revised Bill's paper before the teacher revised Bill's paper*

Third, the first pronominal term may be *copied*, whereas the second pronominal term *refers*, to derive the reading shown in (168).

(168) $before(revise(Bill_{c0}, paper_of(his_{c1}^{-c0})),$
 $revise(teacher_{d0}, paper_of(his_{d1}^{-b1})))$

Reading: *Bill revised Bill's paper before the teacher revised John's paper*

Finally, the first pronominal term may be used to *refer*, whereas the second pronominal term may be *copied*, to derive the reading shown in (169).

(169) $before(revise(Bill_{c0}, paper_of(his_{c1}^{-a1})),$
 $revise(teacher_{d0}, paper_of(his_{d1}^{-c1})))$

Reading: *Bill revised John's paper before the teacher revised John's paper*

Reading (169) is the same as the *all-strict* reading shown in (166). Thus, the algorithm thus far has generated three readings, and has not generated the non-existent reading (161).

We now consider the case in which the first ellipsis receives the sloppy interpretation given in (164). The representation for the source in this case is given in (170).

$$(170) \text{ before}(\text{revise}(\text{John}_{a0}, \text{paper_of}(\text{his}_{a1}^{\rightarrow a0})), \\ \text{revise}(\text{teacher}_{b0}, \text{paper_of}(\text{his}_{b1}^{\rightarrow b0})))$$

Again, we step through the four possibilities. First, we can *refer* with both pronominal terms, yielding the *all-strict* reading for the target shown in (171).

$$(171) \text{ before}(\text{revise}(\text{Bill}_{c0}, \text{paper_of}(\text{his}_{c1}^{\rightarrow a1})), \\ \text{revise}(\text{teacher}_{d0}, \text{paper_of}(\text{his}_{d1}^{\rightarrow b1})))$$

Reading: *Bill revised John's paper before the teacher revised the teacher's paper*

Second, we can copy both pronominal terms to obtain the *all-sloppy* reading shown in (172).

$$(172) \text{ before}(\text{revise}(\text{Bill}_{c0}, \text{paper_of}(\text{his}_{c1}^{\rightarrow c0})), \\ \text{revise}(\text{teacher}_{d0}, \text{paper_of}(\text{his}_{d1}^{\rightarrow d0})))$$

Reading: *Bill revised Bill's paper before the teacher revised the teacher's paper*

Third, the first pronominal term may be *copied*, whereas the second pronominal term be used to *refer*, to derive the reading given in (173).

$$(173) \text{ before}(\text{revise}(\text{Bill}_{c0}, \text{paper_of}(\text{his}_{c1}^{\rightarrow c0})), \\ \text{revise}(\text{teacher}_{d0}, \text{paper_of}(\text{his}_{d1}^{\rightarrow b1})))$$

Reading: *Bill revised Bill's paper before the teacher revised the teacher's paper*

Finally, the first pronominal term may be used to *refer*, whereas the second pronominal term may be *copied*, to derive the reading given in (174).

$$(174) \text{ before}(\text{revise}(\text{Bill}_{c0}, \text{paper_of}(\text{his}_{c1}^{\rightarrow a1})), \\ \text{revise}(\text{teacher}_{d0}, \text{paper_of}(\text{his}_{d1}^{\rightarrow d0})))$$

Reading: *Bill revised John's paper before the teacher revised the teacher's paper*

The reading given in (174) is the same as that given in (171), and the reading given in (173) is the same as that given in (172). These two readings, along with the three readings derived earlier for the strict case, are the five readings shown in (160) for sentence (159).

As noted by DSP, the accounts of Sag (1976) and Williams (1977) can be seen to derive two readings for example (159). The analysis of Gawron and Peters (1990) derives three readings. The DSP approach derives six readings, although they appeal to an unspecified “suitable definition of generalized antecedent linking” and an additional constraint on property derivation to eliminate the sixth.

The algorithm presented here derives all and only the correct five readings for example (159). Recall that the ability of the algorithm to avoid the non-existent reading in the *missing reading* cases in Section 4.2.3 was due to our choosing to link referential elements to their referents in the semantic representation. In that case, the critical point that implicitly eliminated the missing reading occurred during a derivation in which the sloppy option was chosen, but in which the resulting behavior was strict. The ability of the algorithm to avoid generating the missing reading given in sentence (161) is due to a similar reason, one that is crucially dependent on how the strict option was implemented in the derivation of the first ellipsis (which resulted in representation (163)). When referring, we link a pronominal term to its parallel pronominal term instead of the term representing its full NP antecedent. This process resulted in a link between the terms his_{b1} and his_{a1} in representation (165), which was the source representation for the second ellipsis. When the second pronominal term was copied in deriving representation (169), the resulting behavior was strict, because the pronominal term it was linked to in the target referred to its antecedent term in the source. The representation for the non-existent reading could therefore not be derived because of this linking dependency. Accounts that do not distinguish between the representations of full NPs and pronouns cannot avoid these readings without appeal to additional constraints. In our analysis the correct readings naturally result from the distinction between referring and copying.

4.3 Implementation

We have implemented our algorithm in a Common LISP program which takes as input a (first-order) representation of the source clause, and a list of overt elements in the target and their parallel elements in the source. The program makes three passes; the first two derive the remaining parallel elements for the target (for unlinked and linked source elements respectively), and the final one performs the substitutions to generate the target readings. The program generates exactly those readings as described for each of the examples discussed above.

We step through a piece of pseudocode to familiarize the reader with the system. The two inputs to the algorithm are stored in the fields *source-rep* and *overt-elements* within the following structure; initially the fields *unlinked-elements* and *linked-elements* are empty. The fields *overt-elements* and *unlinked-elements* are to be

stored as association lists of source and target parallel elements.

STRUCT **ellipsis**

```
source-rep      ; source representation    (provided)
overt-elements  ; overt parallel elements  (provided)
unlinked-elements ; elements with no links  (not provided)
linked-elements ; elements with links     (not provided)
```

END STRUCT

The top-level function is RESOLVE-ELLIPSIS. This function calls the functions for performing the three passes. The first pass creates the parallel elements for elements in the source that are not linked, and stores the source elements that are linked in the linked-elements list. The second pass creates parallel elements for linked source elements, establishing the links for the strict and possibly sloppy readings. The third pass performs the substitutions for each of the resulting target representations.

RESOLVE-ELLIPSIS *ellipsis-struct*:

```
begin
  ellipsis-struct := CREATE-PARALLEL-ELEMENTS ellipsis-struct
  ellipsis-readings := RESOLVE-LINKED-ELEMENTS (LIST ellipsis-struct)
  target-structs := CREATE-TARGET ellipsis-readings
end
```

The first step creates parallel elements for unlinked source elements using the function CREATE-PARALLEL-ELEMENTS. The function GEN-UNIQUE generates a unique element corresponding to a given source element. Cases in which the target parallel element is overt are passed over. Elements created to correspond to unlinked elements in the source get pushed on the unlinked-elements list. Source linked elements get pushed on the linked-elements list using PUSHLINK, described below.

CREATE-PARALLEL-ELEMENTS *ellipsis-struct*:

```
begin
  foreach element in ellipsis-struct-source-rep
    case element
      in ellipsis-struct-overt-elements: continue
      unlinked in ellipsis-struct-source-rep:
        PUSH (LIST (GEN-UNIQUE element) element)
          ellipsis-struct-unlinked-elements
      linked in ellipsis-struct-source-rep:
        ellipsis-struct-linked-elements :=
          PUSHLINK element ellipsis-struct-linked-elements
    end
end
```

Our algorithm allows a target element to be linked to another target element that is itself part of the elided material. Therefore, in creating the parallel elements for linked elements, we have to make sure that we create any element that is an antecedent of a link before we create the element that links to it. PUSHLINK stores the source linked elements in an order that guarantees that they will always be popped off the list before any elements that link to them. It scans the list of linked elements, checking each member to see whether it has a link to the element it is attempting to insert. If one is found, the element is inserted before it on the list. If none are found, the element is inserted at the end of the list. Therefore, when elements are popped off of this list during RESOLVE-LINKED-ELEMENTS (described below), any element that is the antecedent of a link will be popped off the list before the element(s) that link to it. This of course will not work if there are circularities in the linking, however, in general this will not be the case.⁷

PUSHLINK *element linked-elements*

```
begin
  cycle through linked-elements and INSERT element when either
    HEAD linked-elements is linked to element or
    TAIL linked-elements is nil
  return linked-elements
end
```

The function RESOLVE-LINKED-ELEMENTS creates the strict and sloppy readings corresponding to a given element that is linked in the source. The function LINK sets the *link* property of a term to the identity of its antecedent term. No sloppy reading is created if the element that the source element is linked to does not have a parallel element in the target.⁸

⁷The (indirectly) circular linking in so-called *Bach-Peters* sentences may be problematic however, as in the example shown in (175), which has the coindexing indicated in (176).

(175) The man who deserves it gets the prize he wants.

(176) [The man who deserves it_i]_j gets [the prize he_j wants]_i.

⁸In the pseudocode we are modifying the list *ellipsis-structures* as we loop over it; thus the actual code is somewhat more complicated than what is given here.

RESOLVE-LINKED-ELEMENTS *ellipsis-structures*

```
begin
  for each ellipsis-structure in ellipsis-structures
    begin
      element := POP ellipsis-structure-linked-element

      ;; compute strict reading
      strict-reading := COPY ellipsis-structure
      new-element := LINK (GEN-UNIQUE element) element
      PUSH (LIST new-element element) strict-reading-unlinked-elements
      PUSH strict-reading ellipsis-structures

      ;; compute sloppy reading
      target-parallel-element := element in target parallel to source
                                term that element is linked to
      if target-parallel-element
        begin
          sloppy-reading := COPY ellipsis-structure
          new-element := LINK (GEN-UNIQUE element) target-parallel-element
          PUSH (LIST new-element element) sloppy-reading-unlinked-elements
          PUSH sloppy-reading ellipsis-structures
        end
      end
    end
end
```

Finally, the function **CREATE-TARGET** performs the substitutions to create the target representation for each reading. The association list *parallel-elements* contains all of the pairs of source and target elements; **SUBSTITUTE** replaces the source terms with their corresponding target terms in the target representation.⁹

CREATE-TARGET *ellipsis-readings*

```
begin
  for each ellipsis-reading in ellipsis-readings
    begin
      parallel-elements := APPEND ellipsis-reading-overt-elements
                                ellipsis-reading-unlinked-elements
      target := SUBSTITUTE parallel-elements ellipsis-reading
      PUSH target target-structs
    end
end
```

⁹This is the functionality of the *SUBLIS* function in Common Lisp.

```
    end
  return target-structs
end
```

In the next section, we discuss related issues and extensions to our analysis.

4.4 Extensions to the Analysis

In this section, we discuss four additional aspects of resolution that are relevant to the analysis: (i) non-pronominal referring elements, (ii) cases of extended parallelism, (iii) other types of reference that copy their antecedents, and (iv) interactions between ellipsis and quantification.

4.4.1 Non-Pronominal Referential Elements

Up to this point, we have considered examples in which the strict and sloppy ambiguities result from a pronoun in the source clause. The class of referential elements can be extended to include certain types of implicit arguments as well. For instance, consider examples (177a) and (177b), adapted from Partee (1989).

- (177) a. John went to a local bar to watch the Superbowl, and Bob did too.
b. George drove to the nearest hospital, and Fred did too.

In sentence (177a), *local* has an implicit argument that is linked to *John*, likewise for *nearest* and *George* in sentence (177b). In each case there are strict and sloppy readings for the target clause; e.g., the target in example (177a) can mean that Bob went to the bar that John went to (which is local to John, but perhaps not Bob), or a bar local to himself. Also, as is the case with pronouns, if the implicit argument is instead linked to external discourse (for instance, if “local” is interpreted as being local to the speaker instead of local to John in sentence (177a)), then there is only one reading for the target; no sloppy reading exists. Assuming that such implicit arguments are linked in semantic representations in the same way that pronouns are, the algorithm derives the correct readings for these cases.

4.4.2 Extended Parallelism

In some cases, certain parallel elements may not be contained in the minimal clause containing the ellipsis. For instance, consider example (178), from Prüst (1992).

- (178) John told a man that Mary likes him, and Bill told a boy that Susan does.

In this case, the source clause is *Mary likes him*, and the target is *Susan does*. As Prüst points out, this sentence has a sloppy reading, i.e., in which Bill told a boy that Susan likes Bill. To capture this reading, *John* and *Bill* have to be identified as parallel elements. Such extended parallelism is also required for cases of nominally-referring copying phenomena (e.g., ‘one’-anaphora), which we discuss in the next section.

4.4.3 Referring and Copying

In our algorithm, strict and sloppy readings result from the ability of a reconstructed pronominal term to either *refer* to its antecedent or to *copy* it. This distinction can be seen as a manifestation of two more general types of anaphora in natural language. To illustrate the first type, consider pronominal reference of the ordinary sort shown in example (179).

(179) Bill bought an old Ford pickup. Hillary really likes it.

Here, the pronoun ‘it’ identifies as an antecedent the entity introduced into the discourse by the noun phrase ‘an old Ford pickup’, and further co-specifies the same entity. In this case ‘it’ can be seen as *referring* to its antecedent.¹⁰

The second type of reference is illustrated by the use of ‘one’-anaphora in sentence (180).

(180) Bill bought an old Ford pickup. Now Hillary wants one.

In this case, *one* also identifies the entity introduced by the NP “an old Ford pickup” as its antecedent. However, *one* does not co-specify this entity, instead, it causes the creation of a new instantiation of an entity satisfying the same description. In this case, we can say that ‘one’ *copies* its antecedent.

The types of anaphora that may *copy* their antecedents include VP-ellipsis and event reference. For instance, the VP-ellipsis in sentence (181) refers to its antecedent, whereas in sentence (182) the VP-ellipsis copies its antecedent.

(181) Bill threw a temper tantrum, making a fool out of himself as he did.

(182) Bill threw a temper tantrum, and Hillary did too.

¹⁰To be clear, we use the word *referring* here only to establish the analogy between the process of co-specification in anaphora (in which there is no direct link manifest between the two linguistic expressions) and the process of referring in terms of our linking mechanism in semantic representations (in which there is a direct, directional link manifest between two semantic terms). We discuss this analogy below.

That is, sentence (181) describes one event, whereas sentence (182) describes two similar events. Likewise, the event referential form *do it* can display either type of anaphora, as shown in sentences (183) and (184).

(183) Bill threw a temper tantrum. Hillary can't believe he did it.

(184) Bill threw a temper tantrum. Five minutes later, Hillary did it.

The refer versus copy distinction for reconstructed pronouns in our VP-ellipsis resolution algorithm can be seen as a manifestation of these two types of reference. In the case that a pronoun *refers*, it is linked directly to its parallel antecedent term in the source, therefore co-specifying it. In the case that a pronoun *copies*, it does not co-specify its parallel source term, but instead has an identical description which establishes a link to a parallel antecedent in the target.

It would appear then that when a referential form invokes *copying* from source to target (such as VP-ellipsis), the representation of a referring expression contained in the source will (recursively) display a referring versus copying ambiguity. Given this characterization, we would predict that strict/sloppy ambiguities result from a property of the copying process in general, and not just of resolution algorithms for VP-ellipsis.

This turns out to be the case. For instance, the case of pronominal event anaphora shown in example (185) has strict and sloppy readings, as does the case of definite event anaphora given in (186).

(185) John got shot by his father. That happened to Bob too.

(186) John kissed his wife, and Bill followed his example. (Dahl, 1972)

There are also forms of nominal reference that copy their antecedents. For instance, the use of “one” anaphora can lead to strict and sloppy ambiguities, as shown in sentence (187).

(187) Although John bought a picture of his son, Bill snapped *one* himself.

Less common, but still extant, are cases in which a pronoun acts as a copying expression; these cases have been referred to as “lazy pronouns” as exemplified in sentence (188).

(188) The man who gives his paycheck to his wife is wiser than the man who spends it. (adapted from Karttunen (1969))

Finally, there are cases in which definite full NPs copy instead of refer, as shown in sentence (189).

(189) Bob actually looks back fondly on the day of his defense. Most students shut *that day* out of their mind.

Our algorithm applies directly to all of these cases also, using the same mechanisms as are employed for VP-ellipsis. In the case of object reference, the algorithm uses the noun phrase representation as a source, relying on the extended parallelism established within the sentences as a whole in the determination of parallel elements as discussed in Section 4.4.2.

In sum, the referring versus copying distinction that our algorithm relies on can be seen as a microcosm of two general types of anaphoric processes. By having this independently-motivated distinction manifest in the process of reconstructing pronominal terms in semantic representations, the correct predictions naturally result for a set of benchmark cases that are problematic for past approaches.

4.4.4 Ellipsis and Quantification

In this chapter we have concentrated on the issue of generating the correct strict and sloppy ambiguities for sentences from which a VP has been elided. None of the examples that we have considered involve quantification. DSP address a range of cases involving quantification; in fact their analysis displays a particularly elegant interaction between quantifier scoping and ellipsis resolution. Among the cases they address are those listed below.¹¹

Quantified Sources: Sentence (190) has two readings, corresponding to whether the greetings are simultaneous per each individual or the whole group.

(190) John greeted every person when Bill did.

The two readings are shown in (191a-b).

- (191) a. $when(every(x, person(x), greet(john, x)),$
 $every(x, person(x), greet(bill, x)))$
b. $every(x, person(x), when(greet(john, x), greet(bill, x)))$

In the DSP analysis, this ambiguity is a result of the two possible orderings between the processes of scoping the quantifier and resolving the ellipsis. Discharging the quantifier before resolving the ellipsis results in reading (191a), whereas resolving the ellipsis first results in reading (191b).

¹¹Shieber, Pereira, and Dalrymple (forthcoming) address additional examples of interest, as well as several of those included below.

Antecedent-Contained Deletion: Unlike sentence (190), sentence (192) only has one reading, shown in (193).

(192) John greeted every person that Bill did.

(193) $every(x, person(x) \wedge greet(bill, x), greet(john, x))$

As discussed in Chapter 3, such examples are referred to as cases of *antecedent-contained deletion* because the ellipsis site is contained in the source VP (in this case, *greeted every person that Bill did*). Reading (193) results from resolving the ellipsis before discharging the quantifier in the DSP analysis. If the quantifier is discharged first, the resulting equation has no solution, and therefore no readings result.

Quantification Parallelism: In cases of ellipsis in which more than one quantifier is elided, if the quantifiers are to be scoped separately in each clause then the resultant scoping must be the same in each. For instance, sentence (194) only has two readings, not four.

(194) John gave every student a test, and Bill did too.

This fact is predicted by the DSP account.

Quantification and Type Raising: Sentence (195) has two readings, corresponding to whether Bill revises his own paper or every student's paper.

(195) Every student revised his paper, and then Bill did.

For the first of these readings, the quantified NP *every student* is parallel to the entity-denoting NP *Bill*, which has a different type. If the meaning for *Bill* is type-raised to the type of the quantified NP, this reading is derivable.

Scoping with Indefinites: Indefinite NPs containing pronouns may give rise to several readings, as in sentence (196).

(196) John lost a book he owned, and so did Bill.

Sentence (196) has three readings: (i) Bill lost the same book as John, (ii) Bill lost a different book that John owns, and (iii) Bill lost a book that Bill owns. Such examples are discussed by Shieber, Pereira, and Dalrymple (forthcoming).

'Canadian Flag' Examples: Hirschbühler (1982) discusses what have come to be known as 'Canadian Flag' examples, such as that given in sentence (197).

(197) A Canadian flag was hanging in front of each window, and an American one was too.

In example (197), the subject of the source and target clauses do not take wide scope. This example is also addressed by Shieber, Pereira, and Dalrymple (forthcoming).

Extending our account to cover all of these examples is a subject for future work. However, a recent analysis due to Crouch (1995) gives evidence that such an extension should be possible. Crouch adopts our analysis of strict/sloppy reading determination (as put forth by Kehler (1993a)) within a substitution-based framework. As with the view manifest in the approach described here, he views ellipsis resolution as the selection of an antecedent followed by the determination of a set of substitutions to apply to it. In his account, these substitutions apply over representations in the Quasi Logical Form (QLF) used in the Core Language Engine (Alshawi et al., 1992).

Crouch claims to obtain the same coverage as DSP, without requiring order dependence between ellipsis resolution and quantifier scoping nor matching operations that are beyond second-order. His analysis covers cases of quantification and antecedent-contained deletion, with constraints on scope parallelism being maintained. Furthermore, he extends the strict/sloppy analysis to apply to quantifiers as well as pronouns, yielding the correct readings of sentences such as example (196). Finally, the analysis is also able to derive the correct readings for Canadian Flag examples such as sentence (197).

4.5 Comparison with Source-Determined Accounts

We survey past work in VP-ellipsis that has focussed on the problem of generating the correct strict and sloppy readings. Many quite varied accounts exist; in an attempt to impose some perspective we categorize these approaches as belonging to one of three classes, summarized as follows.

Source-determined, identity-of-relations: Analyses positing that the ambiguity in the target clause follows directly from an ambiguity in the source clause.

Source-determined, non-identity: Analyses positing no ambiguity in the source, but in which the target is interpreted as if it were constructed with relationships determined in some potentially ambiguous but uniform manner from those in the source.

Discourse-determined: Analyses positing that ambiguities result from copying uninstantiated representations of pronouns that are subsequently resolved in the target in accordance with discourse principles.

As noted in Chapter 3, past analyses differ with respect to the level of representation (e.g., syntactic or semantic) at which VP-ellipsis is resolved. For the purposes of our categorization, we may ignore this aspect of the analyses. In this section we discuss the first two types of analysis, which comprise the source-determined approaches. In Section 4.6, we compare source-determined accounts with discourse-determined ones.

4.5.1 Source-Determined, Identity-of-Relations Analyses

Identity-of-relations analyses¹² (Sag, 1976; Williams, 1977; Klein, 1987; Gawron and Peters, 1990; Fiengo and May, 1994) posit that source VP representations with intrasententially-referring pronouns have ambiguous derivations, although these derivations result in the same meaning for the whole source clause. The ambiguity in the target results from the ability to copy the result of each possible derivation. For example, the source VP in example (198) would have derivations that (at some level) lead to the two interpretations shown in (199) for the verb phrase (P representing the meaning of the VP).

(198) Ross likes his mother, and Bill does too.

(199) a. P = likes *Ross's* mother

b. P = likes *his own* mother

These two properties, when applied to the subject *Ross*, lead to the same reading for the source clause. However, under an identity-of-relations analysis, if the target receives a strict interpretation, then necessarily (199a) is the correct interpretation for the source VP, whereas if the target receives a sloppy reading, then (199b) is necessarily the correct interpretation.

Until recently, the identity-of-relations approach was the conventional wisdom. However, there are problems with this type of analysis, in that it fails to make the correct predictions in certain cases. Specifically, DSP discuss the *cascaded ellipsis* case repeated below as example (200).

(200) John realizes that he is a fool, but Bill_i does not, even though his_i wife does.

¹²This term is borrowed from Dalrymple, Shieber, and Pereira (1991).

Sentence (201) is an acceptable, and perhaps preferred, reading for sentence (200).

(201) John realizes that *John* is a fool, but Bill does not realize that *Bill* is a fool, even though Bill's wife realizes *Bill* is a fool.

Example (200) contains two cases of ellipsis; the reading shown in (201) results from the second clause receiving a *sloppy* interpretation using the first clause as a source, and the third clause receiving a *strict* interpretation using the second clause as a source. An identity-of-relations analysis, however, specifically predicts that the reading given in sentence (201) does not exist. Because the second clause will only have the sloppy derivation received from the first, the strict derivation that the third clause requires from the second will not be present.

Dalrymple, Shieber, and Pereira (1991) offer an analysis that does not suffer from this drawback for the same reason that our algorithm does not. We discuss their analysis in the following section.

4.5.2 Source-Determined, Non-Identity Analyses

Unlike identity-of-relations accounts, in the DSP analysis no unmotivated ambiguity is predicated to exist for VPs with pronouns, instead the ambiguity exists in the *resolution process*. We step through DSP's analysis of example (198) to acquaint the reader with their system. The representation of the source clause in this example is shown in (202).

(202) $likes(Ross, mother_of(Ross))$

The ellipsis is resolved by deriving a property P such that representation (202) results from applying it to *Ross*. This property is computed by solving equation (203).

(203) $P(Ross) = likes(Ross, mother_of(Ross))$

Strictly speaking, there are four solutions to equation (203). DSP employ a notion of *primary occurrence*, which requires that the specific occurrence in the source that is parallel to the overt occurrence in the target be abstracted over; for equation (203), this is the first occurrence of *Ross*. In accordance with this constraint, there are two solutions to equation (203), shown in (204).

(204) a. $P = \lambda x.likes(x, mother_of(Ross))$

b. $P = \lambda x.likes(x, mother_of(x))$

To generate a reading for the target clause, either one of these representations may be applied to the element in the target clause that is parallel to *Ross*, in this case *Bill*.¹³ Applying relation (204a) to *Bill* yields the strict reading in (205a); applying relation (204b) yields the sloppy reading given in (205b).

(205) a. $P(\textit{Bill}) = \textit{likes}(\textit{Bill}, \textit{mother_of}(\textit{Ross}))$

b. $P(\textit{Bill}) = \textit{likes}(\textit{Bill}, \textit{mother_of}(\textit{Bill}))$

Because strict and sloppy readings in the DSP account result from an ambiguity in the resolution process, reading (201) for sentence (200) is correctly predicted to exist. However, there are some inadequacies with the approach with respect to predicting strict and sloppy readings for other cases.

The first problem is with the *stripping* case, repeated below as sentence (206).¹⁴

(206) John_{*i*} likes his_{*i*} mother, and Mary's too.

In example (206), the possessive *Mary's* in the target is parallel to the pronominal possessive *his* in the source. Sentence (206) has only the reading given in (207).

(207) John likes John's mother, and John likes Mary's mother.

As we mentioned in Section 4.2.2, however, two readings are derived by the DSP analysis; specifically the correct one given in (207) as well as the reading given in (208).

(208) John likes John's mother, and Mary likes Mary's mother.

The representation for the source clause in DSP's system is shown in (209).

(209) $\textit{likes}(\textit{John}, \textit{mother_of}(\textit{John}))$

Here, the second occurrence of *John* is parallel to *Mary*. Extracting over this primary occurrence, two solutions result, making possible either of the two derivations shown in (210) and (211).

(210) $P = \lambda x. \textit{likes}(\textit{John}, \textit{mother_of}(x))$

$P(\textit{Mary}) = \textit{likes}(\textit{John}, \textit{mother_of}(\textit{Mary}))$

Reading: *John likes Mary's mother*

¹³The DSP analysis separates the processes of parallel element determination and ellipsis resolution. Their algorithm models the latter process, as does ours.

¹⁴DSP claim that their account handles cases of stripping as well as VP-ellipsis.

(211) $P = \lambda x.likes(x, mother_of(x))$
 $P(Mary) = likes(Mary, mother_of(Mary))$
 Reading: *Mary likes Mary’s mother*

Because only reading (210) actually exists for sentence (206), DSP’s system overgenerates in this case. This problem arises because DSP do not distinguish between full NPs and pronouns that refer to them in semantic representations. The fact that our algorithm handles this example is a result of our making this distinction.

Also, to account for *missing reading* examples and the *5-reading* sentence, DSP postulate an unspecified “suitable definition of generalized antecedent linking”, and need to impose an additional constraint on property derivation requiring that when an element is abstracted over, so must its generalized antecedent. In the *5-reading* case, such linking would have to be the result of a process that is independent of the resolution algorithm, because the elimination of the sixth reading is the result of a link established between the source and target clauses for the first ellipsis.

Our algorithm is also a non-identity approach, and therefore shares DSP’s advantages over identity-of-relations analyses. We also derive the correct readings for the stripping, missing readings, and 5-reading cases without appeal to any independent processes or additional constraints on resolution. In the next section, we contrast source-determined accounts such as ours with discourse-determined approaches.

4.6 Comparison with Discourse-Determined Accounts

In this section, we compare source-determined and discourse-determined approaches to VP-ellipsis resolution. In Section 4.6.1, we describe three accounts that fall into the discourse-determined category. We show how discourse-determined approaches make the correct predictions for a baseline set of examples in Section 4.6.2. In Section 4.6.3, we provide examples that show that discourse-determined analyses can only be maintained under certain assumptions concerning the conditions under which a VP can be elided. We argue in Section 4.6.4 that these assumptions make discourse-determined analyses reliant on the very theory that source-determined analyses attempt to provide. Finally, we discuss a series of examples that remain as problematic for source-determined approaches in Section 4.6.5.

4.6.1 Discourse-Determined Accounts

In each of the source-determined analyses we have described in the preceding section, at some level of representation (i.e., surface syntactic, syntactic logical form, or

semantic), the binding relationships for pronouns in the source clause are marked. The target is interpreted as if it were constructed with relationships determined in some uniform manner by those of the source clause at that level of representation. In a discourse-determined analysis, the binding relationships in the target are instead determined exactly as those in the source are determined, by appeal to discourse principles and pragmatics. We describe three such approaches here.

Hellan (1988) describes an account of ellipsis resolution in which syntactic material from the source is copied to the target. The binding of a copied pronoun is handled in the following way:

On the *strict* reading of the target, ... the repeated occurrence of X [the copied representation] is treated as *fully evaluated*, i.e., with a referent assigned before the ‘repetition’... On the *sloppy* reading, ... the repeated occurrence of X is treated as *non-evaluated*, i.e., as not yet having been assigned a referent. In its repeated version in the second conjunct, this amounts to Z [the copied pronoun] being free to seek a binder... (Hellan, 1988, page 228)

Hellan adds an additional constraint on binding in the sloppy case, that restricts pronouns from binding to any available entity in the discourse.

Then, after copying, [the pronoun] must find a binder; and for the copy inside S_2 [the target], this binder must necessarily be inside S_2 . (Hellan, 1988, page 233)

Pronouns in the sloppy case are free to seek binders according, presumably, to preferences dictated by a theory of discourse (in accordance with the above constraint), and therefore allows for greater freedom than is inherent in a source-determined approach.

Kitagawa (1991) provides an account in which VP-ellipsis is resolved at the level of LF in GB theory. In his account, copied pronouns are also posited to freely seek a referent, including intersentential ones. Again, presumably such reference would be guided by a set of discourse principles.¹⁵

In a series of papers, Hardt (1991a; 1991b; 1992a; 1992b) outlines an account of ellipsis that follows Kitagawa with respect to determining pronominal antecedents, although VP-resolution is taken to occur at a purely semantic level of representation. He suggests that binding relations are separately determined on the basis of a discourse model operative during the interpretation of the source and target. In the system he describes, the meaning of an elided VP is taken to be a “property”

¹⁵Although neither Hellan nor Kitagawa explicitly propose the use of discourse principles to reduce the possible overgeneration of their analyses, we presume this plausible interpretation.

in the discourse model, but one with uninstantiated variables standing proxy for the pronouns in the clause that it represents. These variables are copied into the target without the bindings they received in the source, and are resolved with respect to the discourse model active in the target clause:

Two copies of this VP, as antecedent and target in VP ellipsis, could allow the pronoun to refer to different entities, depending on the state of the current discourse model. (Hardt, 1992b, page 306)

Ambiguities in strict/sloppy readings result from having multiple possible referents for the pronoun copied to the target.

4.6.2 Predictions of Discourse-Determined Accounts

In the remainder of this section, we weigh the arguments supporting source-determined and discourse-determined approaches.¹⁶ The phenomenon that each attempts to explain concerns the readings in the target clause corresponding to the binding of the a pronoun in the source clause, the facts of which we reiterate here for the purposes of comparison. We consider example (212).

(212) Ross likes his mother, and Bill does too.

If the pronoun refers intrasententially to Ross so that the source clause is taken to mean that Ross loves his own mother (as indicated by the coindexing in sentence (213)), then the target clause is ambiguous between two readings. It might mean that Bill loves Ross's mother (the *strict* reading shown in sentence (214)) or that Bill loves his own (i.e., Bill's) mother (the *sloppy* reading shown in sentence (215)).

(213) Ross_i loves his_i mother, and Bill_j does too.

(214) Ross_i loves his_i mother, and Bill_j loves his_i mother.

(215) Ross_i loves his_i mother, and Bill_j loves his_j mother.

If *his* refers intersententially to some third person, say Al_k—that is, if the source clause is taken to mean that Ross loves Al's mother—then the target clause must mean that Bill also loves Al's mother. That is, example (216) only has the reading reflected by the indices shown in sentence (217).

(216) Ross_i loves his_k mother, and Bill_j does too.

¹⁶Some of the following material is the product of joint work with Stuart Shieber.

(217) Ross_i loves his_k mother, and Bill_j loves his_k mother.

We have already described how these facts are predicted by source-determined accounts. In the discourse-determined accounts of Hellan (1988), Kitagawa (1991), and Hardt (1991a; 1991b; 1992a; 1992b), the readings result by copying an uninstantiated pronoun from source to target and appealing to discourse-determined principles for its resolution.

Of the three analyses, only Hardt contains any explicit reference to the role of discourse processes in performing this resolution. However, much remains unsaid about the particulars of the discourse theory assumed. In recent papers (Hardt, 1992a, 1992b), no details are provided concerning the discourse principles used to predict the correct strict and sloppy readings. In earlier papers (Hardt, 1991a, 1991b), he assumes the *Centering* model of Grosz et al. (Grosz, Joshi, and Weinstein, 1983; Grosz, Joshi, and Weinstein, 1995).¹⁷

Nonetheless, an explicit theory of discourse is not necessarily required to test the predictions of a discourse-determined analysis. We can instead make use of our native intuitions concerning the felicity of discourses *not* containing ellipsis. Such intuitions, we may assume, should follow from an optimal theory of discourse processing, and we may use them as a proxy for that theory.¹⁸

For instance, consider again sentence (213), repeated in example (218), embedded in a context to favor the binding of *his* in the source clause to Al.

(218) Al has a wonderful family. Ross loves his mother and Bill does too.

Whatever factors contribute to the binding of the overt *his* to Al would affect the variable in the elided property as well. This can be seen by examining the unelided counterpart of this sentence pair, shown in passage (219).

(219) Al has a wonderful family. Ross loves his mother, and Bill loves his mother too.

¹⁷Subsequently in his thesis (Hardt, 1993), Hardt presents a different account, suggesting that the strict and sloppy distinction results from the different orders in which assumptions can be discharged for the pronoun and for the ellipsis itself. In the case that the pronoun assumption is discharged first, the strict reading results; in the case that the ellipsis assumption is discharged first, the sloppy reading results. However, it is not clear why he posits this correspondence. Unlike Hellan (1988), Hardt does not restrict the possibilities for reference under a sloppy interpretation to the target clause. Therefore, under Hardt's account, even if the ellipsis is discharged first, the strict reading is possible, because the copied pronoun can freely refer intersententially to the subject of the source clause.

¹⁸This assumes, of course, that the same model of discourse is used in determining binding relations for both elliptical and nonelliptical clauses. If this is not the case, a discourse-determined analysis of ellipsis is unfalsifiable in any case.

Assuming that the middle clause in example (219) is taken to mean that Ross loves Al's mother, it is infelicitous to interpret the final clause as meaning either that Bill loves Ross's mother or that Bill loves Bill's mother. Hardt notes this fact in discussing a similar example:

While [readings with the target pronoun switching reference] are ruled out by the alphabetic variance condition, the questionable nature of these readings can be explained without recourse to such a condition. Consider the non-elliptical counterpart of example [(218), given in (219)]. The [readings with the target pronoun switching reference] above would be no better, in the absence of deictic gestures or intonational emphasis. Therefore, these facts appear to reflect general constraints on the resolution of referential pronouns in context. (Hardt, 1991b, page 126)

Therefore, we presume a full theory of discourse would predict these facts, and whatever theory that is could then be used to predict the lack of the similar reading for the elliptical version (218).

Likewise, the correct predictions for the strict and sloppy interpretations of example (213) seemingly result from discourse-determined analyses. Assuming that *his* in the source is coreferential with *John*, the non-elided counterpart of example (213) apparently displays an identical ambiguity, as shown in sentence (220).

(220) $John_i$ loves his_i mother, and Bill loves his mother too.

Both *John* and *Bill* are salient enough that the referent of *his* in the target is ambiguous in exactly the manner required to yield both the strict and sloppy interpretations.

To summarize thus far, a purely discourse-determined analysis predicts that a sentence with ellipsis should display the same readings in a given context that the unelided form would in the same context. Examples such as those above demonstrate that a discourse-determined theory can explain at least some cases of dependencies between bindings in source and target clauses.

4.6.3 The Scope of Discourse-Determined Accounts

As it currently stands, both approaches appear to correctly predict the data. We now ask if there are examples that distinguish the two approaches. A counterexample to a discourse-determined analysis would be one in which the pronominal reference possibilities are different for its elided and unelided forms. Attempting to construct such an example makes it clear that one needs to be specific about what unelided forms are valid correspondents of elided forms. We present a set of examples that show that, at a minimum, the VP needs to be destressed within the target for ellipsis

to take place; without this assumption, the discourse-determined analyses can be counterexemplified.

For example, consider the nonelliptical example (221a). This pair of sentences is felicitous under the binding relations indicated. Its elliptical counterpart (221b), however, cannot be taken as having the meaning of (221a). Similarly, example (222b) cannot have the meaning provided by the nonelliptical version (222a).

- (221) a. Mike Tyson will always be considered one of the greats of professional boxing. After one round with Spinks_{*i*}, Tyson_{*j*} beat him_{*i*}. Now people think that no one can beat him_{*j*}.
- b. Mike Tyson will always be considered one of the greats of professional boxing. After one round with Spinks_{*i*}, Tyson_{*j*} beat him_{*i*}. Now people think that no one can.
- (222) a. The story of Lee Harvey Oswald will always be one of the most intriguing in U.S. history. When John Kennedy_{*i*} drove through Dallas, Oswald_{*j*} shot him_{*i*}. Some time afterward, Dallas nightclub owner Jack Ruby shot him_{*j*} in retaliation.
- b. The story of Lee Harvey Oswald will always be one of the most intriguing in U.S. history. When John Kennedy_{*i*} drove through Dallas, Oswald_{*j*} shot him_{*i*}. Some time afterward, Dallas nightclub owner Jack Ruby did in retaliation.

Examples (221) and (222) demonstrate that pronouns within copied VPs are not as free to seek intersentential referents as their unelided VP counterparts. Example (223) shows that this is also the case for intrasentential referents.

- (223) a. Every boy_{*i*} was hoping that Mary would ask him_{*i*} out, but the waiting is over. Last night at Bob's_{*j*} party, she asked him_{*j*} out.
- b. Every boy_{*i*} was hoping that Mary would ask him_{*i*} out, but the waiting is over. Last night at Bob's_{*j*} party, she did.

The reading in which Mary asked out Bob at Bob's party, while readily available in example (223a), is not available in its elided counterpart (223b).

The speakers we have checked with have agreed on the judgements that we have indicated here. However, there is some disagreement as to the need for added stress on the target pronoun in the unelided cases. Insofar as such stress is not required, the discourse-determined analyses are falsified by these examples. However, if any

such stress is required, even if far short of contrastive or otherwise exaggerated stress, these examples do not provide an argument against discourse-determined approaches. Instead, they only serve to make the assumptions behind a discourse-determined approach more restrictive – the assumption being that only destressed VPs can be elided. This is not an unreasonable assumption; ellipsis can be seen as the logical extension of full destressing (see Rooth (1993) for further discussion of this point).

Rooth (1993) also gives examples that potentially distinguish the two approaches, shown in (224a-b).

- (224) a. John_i's coach thinks he_i has a chance, and Bill_j does too.
 b. John_i's coach thinks he_i has a chance, and Bill_j thinks he_j has a chance too.

Rooth claims that whereas the unelided form in example (224b) gives rise to a sloppy reading, the elided form in example (224a) does not. However, like the cases discussed above, some speakers find the pronoun in example (224b) to require added stress under this interpretation. Nonetheless, these examples demonstrate that for a discourse-determined approach to be adequate, it can only apply to destressed VPs.

4.6.4 Relationship Between the Two Approaches

It turns out to be quite difficult to provide conclusive data arguing for source-determined analyses over discourse-determined ones, assuming that the latter only applies to destressed VPs. In this section, we suggest that this is not a coincidence. In fact, we argue that the type of theory that discourse-determined analyses rely on is just that which source-determined analyses attempt to provide.

Several semantic source-determined accounts, including DSP and the one presented here, cover a number of types of predicate reference, including VP-ellipsis, *do it* and *do that* anaphora, *do so anaphora*, and cases of non-pronominal definite event reference that copy their antecedents. These forms have quite different syntactic properties, but have one property in common: they indicate that their referent represents information that is *given*¹⁹ in, and therefore recoverable from, the existing discourse state. Destressing is also a well-established indicator that material is given information in the discourse (Terken, 1984; Terken and Nootboom, 1987, *inter alia*). Therefore, we might expect that a VP can be felicitously destressed only in situations in which the information it specifies is given in the discourse. As with the various

¹⁹The concept of *givenness* is a complicated one; see Prince (1981b) for discussion. Our use here corresponds most closely to her *Givenness_s*, or *discourse-old* in the sense of Prince (1992).

forms of event reference, a VP thus requires an ‘antecedent’ to license destressing that either exists in the discourse or is inferable from it.²⁰

As stated, source-determined analyses of VP-ellipsis posit a correspondence between the antecedents of elided VPs and the possible readings they license for a target clause. However, this correspondence can be generalized to one holding between a potential antecedent and the predicates that become given in the discourse as a result of its occurrence. In the case in which an antecedent is evoked by a clause containing pronouns, there are several such predicates, one corresponding to each of the strict and sloppy interpretations. It is for this reason that (semantic) source-determined analyses of VP-ellipsis resolution extend directly to the resolution of other expressions that indicate that their predicates are given in the discourse. In their generalized form, source-determined analyses posit the set of given predicates resulting from interpreting a source clause that (possibly) contains pronouns.

Discourse-determined analyses, on the other hand, offer no mechanisms for recovering the readings of target clauses from source clauses that contain pronouns. They instead rely on a more general, unstated theory of pronominal reference which applies after the target representation has been recovered. The examples given in the previous section show that this position can only be maintained if it is restricted to cases in which the unelided counterpart of the ellipsis is licensed to have its VP destressed. This licensing is achieved if the ‘antecedent’ of the VP is given in the discourse. Therefore, the theory that discourse-determined analyses rely on (but do not themselves provide) is one that specifies the possible binders for reconstructed pronouns in VPs that specify predicates that are given in the discourse. As stated above, this is just the theory that source-determined analyses attempt to provide.

In light of this fact, it is not surprising that it is difficult to find examples of VP-ellipsis that distinguish source-determined and discourse-determined approaches. Such an example would require an unelided counterpart containing a destressed VP that has a different set of readings than the elided case. However, both destressing and VP-ellipsis are licensed by their predicate being given in the discourse, so no such distinguishing examples are likely to exist. The difference between the two approaches is that source-determined analyses offer an account of this licensing, whereas discourse-determined analyses merely rely on such an account.

Counterexamples to source-determined analyses have been presented by proponents of discourse-determined approaches, and some of these do in fact point out problems with instances of the former type of analysis. However, it is important

²⁰ While all of these forms indicate that their antecedents are given, they differ in the extent to which this antecedent can be inferred, or *accommodated* (Lewis, 1979; Webber, 1991), from existing discourse representations. In this regard, VP-ellipsis is the most restrictive whereas destressed full VPs are the least. However, all of these forms have similar requirements for the resulting antecedent that licenses their use, whether or not this antecedent is accommodated.

to see that these examples expose a problem with our understanding of what constitutes a given predicate in the discourse in certain contexts, an understanding to which discourse-determined analyses make no contribution. We discuss these examples in the next section.

4.6.5 Remaining Problematic Examples

Hardt (1993) discusses example (225) which, on the surface, appears to be outside of the scope of theories that utilize syntactic or semantic parallelism in determining the possible strict and sloppy readings.

(225) Every boy in Mrs. Smith's class hoped she would pass him. In John's case, I think she will.

This example has a “sloppy” reading, in which the target is understood as *I think she will pass John*. However, source-determined analyses based solely on syntactic or semantic parallelism do not predict this reading, since there is no such parallelism between “every boy” (the syntactic subject, and the semantic agent or experiencer) and “John” (a possessive in an NP that is embedded in a preposed PP, with a rather vague semantic function). DSP appeal to pragmatic factors to establish parallelism (a view that we have adopted here), leaving open the possibility of an argument for why these might be considered parallel. However, without such an independently-motivated argument, it is hard to judge the plausibility of identifying such parallelism.

On the other hand, the reading is “predicted” by Hardt's discourse-determined approach, as well as Hellan's and Kitagawa's, since the unelided version of example (225) displays the “sloppy” reading without requiring additional stress on the pronoun, as shown in sentence (226).

(226) Every boy in Mrs. Smith's class hoped she would pass him. In John's_{*j*} case, I think she will pass him_{*j*}.

Under a discourse-determined analysis, whichever discourse principles license this reference possibility for sentence (226) would equally apply for the elided version in example (225).

However, it is our sense that something quite different is happening in this case. The preposed phrase *in John's case* serves a special function here, in a way that the PP *at Bob's party* in sentence (223b) does not. Specifically, the phrase *in John's case* is inherently referential; it crucially depends on context for its interpretation. It refers to a previously evoked state or event, meant to exemplify or contrast *John* with respect to some other parallel object or group of objects (in this case, every other

boy in Mrs. Smith’s class).²¹ Therefore, before the ellipsis is resolved, the meaning of *in John’s case* must be resolved. This resolution results in a representation for *John_j hoped she would pass him_j*, which is shown in (227).

(227) $hope(John_{a0}, pass(Mrs_Smith, him_{a1}^{-a0}))$

This representation serves as the source clause for the subsequent ellipsis, on analogy with the case of cascaded ellipsis discussed in Section 4.2.1. The meaning of the target of sentence (225) is then simply the strict reading, derivable by source-determined algorithms. Therefore, we don’t find this case to be a counterexample for source-determined approaches at all. We should note that while the parallel elements for the ellipsis resolution are determinable from both syntactic and semantic role parallelism, the process of identifying the parallel elements in resolving an expression like “in John’s case” is clearly a pragmatic one.

Kitagawa (1991) gives a case that is potentially more problematic. Some speakers find sentence (228), a variant of an example he gives, to have the sloppy reading shown in sentence (229).

(228) John told Mary to hand in his paper before Bill does.

(229) John_i told Mary to hand in his_i paper before Bill_j hands in his_j paper.

Example (228), assuming it has this reading, is problematic for most source-determined analyses because it necessitates that *Bill* be parallel to both *John* and *Mary*, which is a possible but unattractive option. On the other hand, in accordance with discourse-determined analyses, the unelided version shown in (229) has this reading without the need for additional stress on the pronoun.

Examples such as sentence (228) show that source-determined analyses, including the one we have presented here, do not yet make the correct predictions concerning sloppy readings. On the other hand, as we argued in the last section, discourse-determined accounts only beg the question; source-determined accounts are the only analyses that articulate the possibilities for pronominal reference in distressed verb phrases.

Finally, we note an additional problematic case that to our knowledge has gone unnoticed in the literature. The case involves coordination, in which the coordinated constituents each contain a referring expression, as in example (230).

(230) Ross_i likes his_i mother and his_i father, and Bill_j does too.

²¹This property also holds of similar constructions, such as *regarding John*, *as for John*, and *with respect to John*, for instance.

Sentence (230) has the two readings shown in sentences (231a-b).

- (231) a. Ross_i likes his_i mother and his_i father, and Bill_j likes his_i mother and his_i father too. (all-strict reading)
- b. Ross_i likes his_i mother and his_i father, and Bill_j likes his_j mother and his_j father too. (all-sloppy reading)

However, source-determined analyses predict four readings. In addition to the readings shown in (231a-b), the mixed readings shown in sentences (232a-b) are also derivable.

- (232) a. Ross_i likes his_i mother and his_i father, and Bill_j likes his_i mother and his_j father too. (mixed reading)
- b. Ross_i likes his_i mother and his_i father, and Bill_j likes his_j mother and his_i father too. (mixed reading)

That is, the readings in which the pronouns in the two coordinated constituents refer to different entities are derivable, but do not exist for example (230).

As expected under a discourse-determined approach, the readings shown in sentences (231a-b) are the only ones available for the unelided version of sentence (230), shown in example (233) (again assuming that the VP is distressed).

- (233) Ross_i likes his_i mother and his_i father, and Bill_j likes his mother and his father too.

How such examples are to be handled within source-determined analyses is a subject for future study.

4.7 Conclusions

In this chapter, we presented an analysis of VP-ellipsis resolution in terms of a straightforward algorithm that correctly predicts a wide range of phenomena. The treatment does not suffer from problems inherent in identity-of-relations analyses. Furthermore, in contrast to the approach of Dalrymple, Shieber, and Pereira (1991), the treatment directly encodes the intuitive distinction between full NPs and the referential elements that corefer with them. The correct predictions for several problematic examples of ellipsis result. The analysis also extends directly to other phenomena which copy their antecedents.

We have also characterized the distinction between source-determined and discourse-determined analyses, concluding that the determination of pronominal referents in fully destressed VPs is a problem to be distinguished from the greater problem of pronominal reference in general. In this light, source-determined analyses propose specific algorithms for resolution, whereas discourse-determined analyses merely beg the question. On the other hand, several problematic cases remain for source-determined analyses.

Chapter 5

Other Cohesive Forms

In Chapter 3, we addressed the question of whether VP-ellipsis resolution operates at the level of syntactic or semantic representation. The seemingly contradictory data was shown to exhibit a pattern which correlates with the type of coherence relation holding between the source and target clauses. We showed how a unified theory of VP-ellipsis combines with the properties of different discourse inference processes to predict data that is beyond the reach of previous theories.

In this chapter, we address several other cohesive forms, including gapping, event reference, tense interpretation, and pronoun interpretation. In Section 5.1, we show that a pattern exists with gapping data that is similar to that displayed by VP-ellipsis data, which again correlates with the nature of the coherence relation operative between clauses. We articulate the syntactic and referential properties of gapping constructions, and show that the account of discourse inference given for VP-ellipsis also predicts the gapping data without further modification.

In Section 5.2, we discuss various forms of event reference, with particular reference to the distinction between *deep* and *surface* anaphora proposed by Hankamer and Sag (1976, henceforth H&S). We discuss how resolution of these forms compares to that for VP-ellipsis and gapping, and show that the predictions that our account makes for all of these forms is more empirically adequate than those provided by the H&S dichotomy. We also discuss another form that is problematic for their categorization, specifically *do so* anaphora.

In Sections 5.3 and 5.4, we discuss the problems of tense interpretation and pronoun interpretation respectively. We again characterize past work as being of two sorts; however, instead of the dichotomy being between syntactic and semantic approaches, the distinction is between what we call *definite reference* and *coherence-based* approaches. Again, the constraints imposed by coherence relations will be shown to interact with theories of the cohesive forms to predict data that is problematic for past approaches.

5.1 Gapping

Gapping is characterized by an initial sentence (henceforth called the *source* sentence) and the elision of all but two (and in constrained circumstances, more than two) constituents in one or more subsequent *target* sentences, as exemplified in sentence (234).

(234) Bill became upset, and Hillary angry.

Gapping has been well studied within the linguistics literature (Ross, 1970; Jackendoff, 1971; Hankamer, 1971; Stillings, 1975; Kuno, 1976; Sag, 1976; Hankamer and Sag, 1976; Neijt, 1979; Neijt, 1981; Levin and Prince, 1982; Cremers, 1983; Siegel, 1984; Sag and Hankamer, 1984; Prince, 1986; Chao, 1987; Oehrle, 1987; Steedman, 1990; Gardent, 1993, *inter alia*). Most of this work has been concerned with the syntactic facts about gapping; a survey of this literature will not be attempted here, nor will a comprehensive account of all of the facts that these works address.

Instead, we address a particular phenomenon noticed by Levin and Prince (1982). They note that the pairs of conjoined sentences such as those in (235) have what they call *symmetric* and *asymmetric* readings.

- (235) a. Sue became upset and Nan became downright angry.
b. Al cleaned up the bathroom and Joe cleaned up the mess.
c. One of the students was accepted at Bryn Mawr and the high school was praised on TV.

That is, each sentence in (235) could have a *symmetric* reading in which the two events are understood as independent, or an *asymmetric* reading in which the first event is interpreted as the cause of the second event (in which *and* is paraphraseable by “and as a result”). In our terms, the symmetric readings correspond to the Resemblance relation *Parallel*, whereas the asymmetric readings correspond to the Cause-Effect relation *Result*. Levin and Prince contrast the sentences in (235) with their gapped counterparts, given in (236).

- (236) a. Sue became upset and Nan \emptyset downright angry.
b. Al cleaned up the bathroom and Joe \emptyset the mess.
c. One of the students was accepted at Bryn Mawr and the high school \emptyset praised on TV.

Unlike the sentences in (235), these sentences only have symmetric readings. For instance, whereas sentence (235a) can have a reading in which Nan became angry because of Sue's becoming upset, this reading is unavailable in (236a). This can be seen by the following contexts, again due to Levin and Prince, where gapping is acceptable in the context favoring the symmetric reading in (237), but not in the context favoring the asymmetric (causal) reading given in (238), although in both cases the non-gapped versions are acceptable.

(237) Sue and Nan had worked long and hard for Carter. When Reagan was declared the winner, Sue became upset and Nan became/∅ downright angry.

(238) Susan's histrionics in public have always gotten on Nan's nerves, but it's getting worse. Yesterday, when she couldn't have her daily Egg McMuffin because they were all out, Sue became upset and Nan became/#∅ downright angry.

The causal interpretation of the two final clauses in example (238), supported by the given context, is unavailable when gapping has applied.

The remainder of this section is organized as follows. In Section 5.1.1, we outline the syntactic and semantic properties of gapping that are relevant to the analysis. In Section 5.1.2 we show how our account of VP-ellipsis also predicts the facts that Levin and Prince note about gapping, as well as additional facts. We describe in Section 5.1.3 how the account also predicts facts concerning cases of *stripping*, and in Section 5.1.4 discuss how the account also handles more complex cases involving both gapping and VP-ellipsis. Finally, in Section 5.1.5 we compare the account to past work.

5.1.1 Syntactic and Semantic Properties of Gapping

For our analysis of gapping, we follow Sag (1976) in hypothesizing that a post-surface-structure level of syntactic representation is used as the basis for interpretation. In source clauses of gapping constructions, constituents in the source that are parallel to the overt constituents in the target are abstracted out of the clause representation. There are several reasons for positing this representation. First, it has been noted that in gapping constructions, contrastive accent is generally placed on parallel elements in both the target and the source clauses. Such accent marks the elements as focussed; abstracting these elements results in an "open proposition" that both clauses share (Sag, 1976; Wilson and Sperber, 1979; Prince, 1986; Steedman, 1990). This open proposition needs to be presupposed (or accommodated) for the gapping to be felicitous. Such presuppositional effects are well-known to occur in other constructions, for instance each of the sentences shown in (239) presuppose that *someone voted*

for Clinton; in example (239a) the focussed element is moved overtly by clefting in the surface syntax, whereas in example (239b) this is achieved by placing contrastive accent on the focussed element (and thus, under the proposal being made here, gets extraposed at a post-surface-structure level of syntax).

- (239) a. It was John who supported Clinton.
 b. JOHN supported Clinton.

Similarly, changing the accented element in turn changes the open proposition, altering the semantics for sentence pairs that have the same surface syntactic structures, as exemplified in the comparative constructions given in examples (240a-b).

- (240) a. MICKEY defended Bill more eloquently than HILLARY. [Hillary defended Bill; open proposition is $\lambda x.defend(x, Bill)$]
 b. Mickey defended BILL more eloquently than HILLARY. [Mickey defended Hillary; open proposition is $\lambda y.defend(Mickey, y)$]

The only difference between sentences (240a) and (240b) is where the accent is placed in the source clause, nevertheless the resulting semantics in the target is different in the two cases.

The same reliance on an open proposition is exemplified in uses of gapping. For instance, it would be infelicitous to open a conversation with a sentence such as (234), whereas it is perfectly felicitous in response to the question *How did the Clintons react?*. Gapping requires that the shared open proposition be recovered from the source clause representation; resolution can be characterized as the restoration of this open proposition in the gapped clause.

For simplicity, we will assume that this abstraction is achieved by fronting the constituents in the post-surface-structure, although our analysis is compatible with any possible mechanism that results in a representation with the desired properties. The syntactic and semantic representations for the source clause of example (234) after fronting are shown in Figure 5.1; the fronting leaves trace assumptions behind that are discharged when combined with their antecedents.

Target clauses in gapping constructions are represented with the overt constituents fronted out of an elided sentence node; for instance the representation of the target clause in example (234) is shown in Figure 5.2 (the empty node is indicated by ϕ). The empty constituent is reconstructed by copying the embedded sentence from the source to the target clause, along with parallel trace assumptions which are to be bound within the target. The result of this process is shown in Figure 5.3. The semantics for the embedded sentence (in this case, $\lambda x, y.become(x, y)$) is the open proposition that the two clauses share.

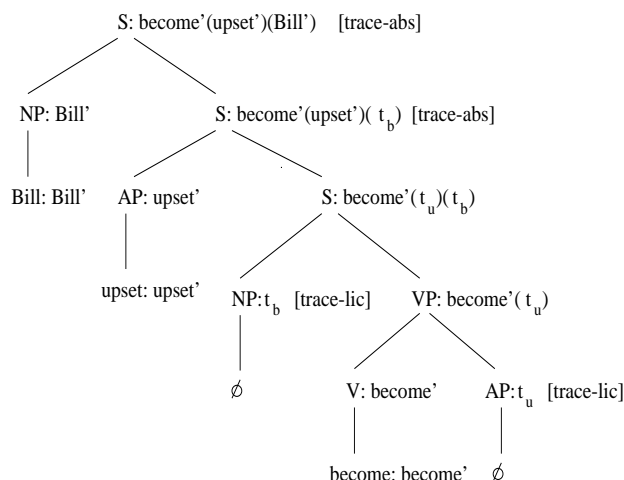


Figure 5.1: Syntactic and Semantic Representations for *Bill became upset* after fronting.

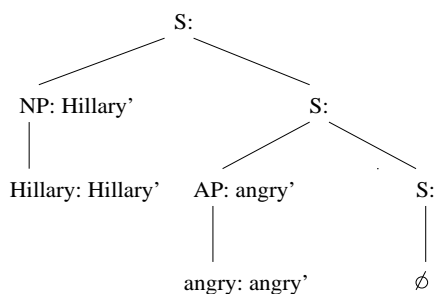


Figure 5.2: Syntactic and Semantic Representations for *Hillary angry*.

Representing gapped clauses and their reconstruction in this way has another significant advantage, in that it avoids the need to posit massive ambiguity for gapped clauses. That is, the overt elements in a gapped clause can correspond to any of a number of constituents in the source. If we posit a sentence-level structure for gapped clauses in which these elements reside in the nodes corresponding to those in which their parallel entities reside in the source representation (representing all missing material as empty nodes), then we have to posit a possible structure for the target for each of potentially many possibilities. By positing that these elements are extracted from an empty sentence node (whose semantics is obtained from the source, the structure of which is determined by its intonational properties), gapped clauses themselves are unambiguous.

Unlike VP-ellipsis, we claim that the semantics for a target clause in gapping can only be recovered by copying the missing syntax, because unlike VP-ellipsis, gapping does not result in an independently anaphoric expression in the semantics. To see this, we apply the same tests that we did for VP-ellipsis in Section 3.3 to determine

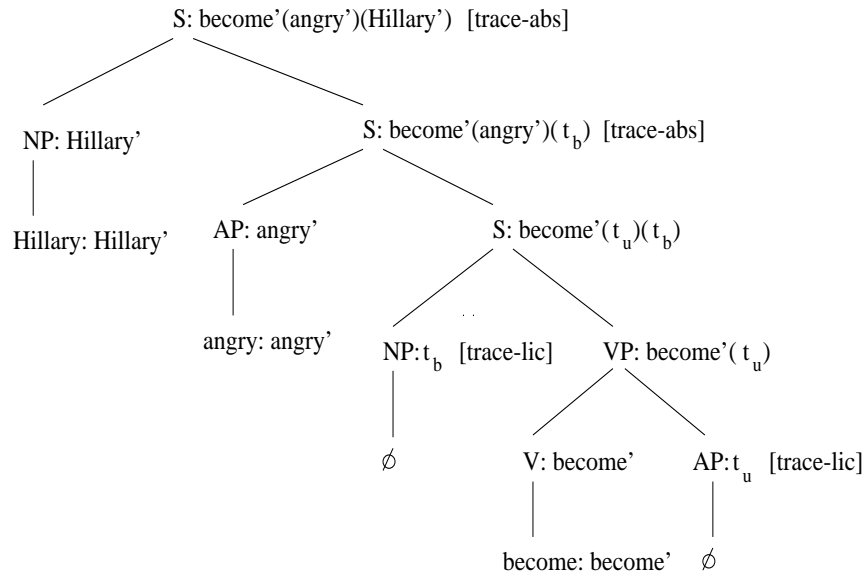


Figure 5.3: Syntactic and Semantic Representations for *Hillary angry* after reconstruction.

whether gapping behaves like a pronoun. First, gapping cannot refer cataphorically, as can pronouns and VP-ellipsis.¹

- (241) a. If he makes a statement blasting the press, Bill will make a fool of himself.
 [cataphoric reference with pronoun]
- b. If Hillary will, Bill will make a statement blasting the press. [cataphoric reference with VP-ellipsis]
- c. * If Hillary the newspaper reporters, Bill will make a statement blasting the press. [cataphoric reference with gapping]
- d. Bill will make a statement blasting the press, and Hillary the newspaper reporters. [standard gapping]

In addition, unlike pronouns and VP-ellipsis, gapping cannot locate antecedents from clauses other than the most immediate one. Therefore, under the theory developed here, the lack of an anaphoric form necessitates that the syntax needs to be reconstructed for the target to obtain a clausal semantics.

¹It is hard, if not impossible, to isolate a good test case for cataphoric reference in gapping, because of two conflicting constraints: (1) that gapping does not operate within embedded clauses, and (2) embedded clauses are necessary so as not to violate the constraints on forward reference proposed by Lakoff (1976) and Jackendoff (1972). Therefore sentences (241c) and (241d) are not a minimal pair. Nonetheless, the fact that gapping does not operate in embedded clauses (therefore requiring reference to an antecedent two clauses away) is in itself evidence against gapping behaving pronominally.

As can be seen from Figure 5.2, before copying takes place there is no sentence-level semantics for gapped clauses at all. In the next section, we describe when and how this semantics is recovered in our account.

5.1.2 The Analysis Applied to Gapping

We now show how our account handles the facts noticed by Levin and Prince. Recall that gapping constructions such as (242) are only felicitous with the symmetric (i.e., Resemblance) meaning of *and*.

(242) Bill became upset, and Hillary angry.

This fact is predicted by our account as it was presented in Chapter 3. For VP-ellipsis, there are two ways to recover the semantics of the elided clause, namely by syntactic reconstruction and by resolution of the anaphoric form. Because gapping does not license an anaphoric form, the only mechanism for recovering the semantics of the elided clause is by syntactic reconstruction. This recovery will occur in cases in which a Resemblance relation is operative, because the missing sentence node in the target will be accessed, causing the source syntactic representation to be reconstructed. Reconstruction results in a complete clausal semantics for the target as shown in Figure 5.3; the arguments to the Parallel relation can then be identified and reasoning carried out. In the case in which a Cause-Effect relation is operative, however, no such reconstruction of syntactic material takes place. Because reconstruction is the only mechanism for recovering a clause-level semantics of the gapped clause, the gapping is unacceptable. Therefore, the facts are predicted without modification to the theory as developed to account for VP-ellipsis.

Another test of our theory is provided by the behavior of reflexives that are part of the deleted material in the target.² Consider sentences (243a-b), assuming that *Al* in the target is parallel to *Bill* in each.

- (243) a. Bill bought his wife a book on health care, and Al a book on the environment.
b. Bill bought himself a book on health care, and Al a book on the environment.

Sentence (243a) exhibits a strict/sloppy ambiguity; that is, Al may have bought a book for Bill's wife or his own wife. However, there is a strong preference in sentence (243b) for the sloppy interpretation (if it allows a strict reading at all), in which Al

²I thank Arild Hestvik for bringing this issue to my attention.

buys *Al* a book. The effect in these cases is essentially the same as that noted for the behavior of reflexives in VP-ellipsis discussed in Chapter 3. This fact about gapping is also predicted by the account, because syntactic reconstruction requires that syntactic constraints be met in the target, whereas this is not so for purely semantic resolution.

While Levin and Prince limit their discussion to sentences conjoined with *and*, we might consider cases involving other coordinations as well. We start by turning our attention the coordinating conjunction *or*. Like *and*, *or* also has Resemblance and Cause-Effect uses. Consider example (244).³

(244) John will go to New York, or Bill will go to Boston.

This sentence has two readings: a symmetric (disjunctive) reading, and an asymmetric causal reading (e.g., to express a threat of the form *If A doesn't happen then B will!*). Like the case with *and*, gapping in clauses conjoined by *or* is allowable in the symmetric case, but not in the asymmetric case. This is shown by the (gapped) examples given in (245) and (246); the former is embedded in a context favoring the parallel reading, and the latter is embedded in a context favoring the causal reading.

(245) (John's and Bill's boss speaking): A meeting should not be scheduled on Thursday, since one of our people will be heading out of town. Either John will go to New York, or Bill (will go)/ \emptyset to Boston.

(246) (John's and Bill's boss speaking): Listen, John, you might not want to be transferred to New York, and I can't force you to go. But I can transfer Bill, and you can imagine what life would be like around here for you if Bill's not here. Now either you will go to New York, or Bill will go/ $\# \emptyset$ to Boston!

Therefore, it appears that our account holds up for these cases also.

However, Sag (1976) notes that gapping is highly stilted with the coordinating conjunction *but*, as in sentence (247).

(247) $\#$ John voted for Clinton but Tom \emptyset Bush.

Indeed, while it is often cited that gapping only applies to coordinate structures, the examples discussed in the literature are almost invariably conjoined by either *and*, *or*, or *nor*, and are rarely conjoined by *but*.

Our account predicts that gapping is acceptable with the Contrast meaning of *but*, but not the Violated Expectation meaning. However, to most (but not all) speakers, sentence (247) is still stilted under a purely contrastive reading. There appear to be other constraints at play here, however. Sentence (248), from Halliday and Hasan (1976, pg. 203), appears to be much more acceptable.

³I thank Gregory Ward for an extended discussion that helped sharpen my intuitions regarding these sentences.

(248) Sybil takes coffee very strong but Joan rather weak.

This example differs from the sentence (247) in that one of the overt constituents in an adverbial. This pattern appears to be somewhat general; sentence (249) is felicitous in contrast to the infelicity of the otherwise similar sentence (247).

(249) John voted for Clinton confidently, but Tom somewhat reluctantly.

Therefore, gapping appears to be acceptable with the purely contrastive meaning of *but* when the full VP is shared by the antecedent and gapped clauses.⁴ We have no complete account of this fact to give here, but sentences such as (248) and (249) show that gapping is acceptable with the contrast meaning of *but* in at least some instances.

Finally, further evidence for this proposal is the fact that gapping is unacceptable with other Cause-Effect relations, as in examples (251a-d).

- (251) a. # John voted for Clinton because Tom \emptyset Bush.
b. # John voted for Clinton even though Tom \emptyset Bush.
c. # John voted for Clinton despite the fact that Tom \emptyset Bush.
d. # John voted for Clinton although Tom \emptyset Bush.

The account predicts an interesting dichotomy between gapping and VP-ellipsis. That is, while the pattern between the two forms is quite similar with respect to the two types of coherence relation, there is a notable difference. In the cases in which VP-ellipsis is *most* restricted (that is, it requires a matching syntactic antecedent), gapping is *least* restricted (that is, it is felicitous). Conversely, in the cases in which VP-ellipsis is *least* restricted (that is, it only requires a suitable semantic source representation without regard to syntactic parallelism), gapping is *most* restricted (that is, it is not felicitous at all). In our account, this contrast naturally results solely from the fact that VP-ellipsis licenses an anaphoric form in the semantics, and gapping does not. The remainder of the relevant properties of these two forms, as well as the account of discourse inference during interpretation, remain constant.

To conclude this section, the gapping facts noticed by Levin and Prince, as well as those for other conjunctions, are predicted by our account as developed for VP-ellipsis without further modification.

⁴This may also be true when only the verb is shared; compare sentence (250) with sentence (247).

(250) John voted for Clinton, but Tom \emptyset for Bush.

5.1.3 Another Form of Local Ellipsis: Stripping

In this section, we briefly discuss the related *stripping* construction, exemplified by sentence (252).

(252) Bill became upset, and Hillary too.

The stripping construction is similar to gapping except that there is only one bare constituent in the target; unlike VP-ellipsis there is no stranded auxiliary in the target. As with the bare constituents in gapping, the bare constituent in stripping generally receives contrastive accent. As noted earlier, the focus marking brought about by such accent is used in determining the correct open proposition shared by the source and target clauses (recall the discussion of examples (240a-b) above). We thus propose the same kind of fronting of the overt target constituent in the post-surface syntax of stripping as we did for gapping.

Like gapping, stripping does not license an anaphoric form in the semantics; it also does not behave like a pronoun in the same way that VP-ellipsis does. We therefore might predict that stripping is also acceptable in cases of Resemblance but not in cases of Cause-Effect, which appears to be the case.

(253) Bill became upset, $\left\{ \begin{array}{l} \text{and also} \\ \text{but not} \\ \# \text{ and (as a result)} \\ \# \text{ because} \\ \# \text{ even though} \\ \# \text{ despite the fact that} \end{array} \right\}$ Hillary.

As a final note, stripping is also possible in comparative deletion constructions, as in examples (240a-b). A comprehensive analysis of gapping, stripping, and pseudo-gapping in such cases requires an articulation of a syntax and semantics for these constructions, which is beyond the scope of this study. Therefore, we will not delve any further into such constructions here.

In summary, gapping and related constructions are predicted to be infelicitous in those cases in which Cause-Effect relations are operative, as there is no mechanism for recovering the sentential semantics of the elided clause.

5.1.4 Interaction of Gapping and VP-Ellipsis

In this section we address cases in which gapping and VP-ellipsis interact, as exemplified by the pair of examples shown in sentences (254) and (255), adapted from Sag (1976, page 291).

(254) Barbara likes George, and Hillary \emptyset Bill, although she doesn't know why she does.

(255) ?? Barbara likes George, and Hillary \emptyset Bill, and Tipper does too.

The account of gapping that he gives correctly predicts that sentence (255) is infelicitous, but incorrectly predicts that sentence (254) is also. The reason is that in each case, the account predicts that the source and target clauses for the VP-ellipsis in the final clause would have to be alphabetic variants for his deletion operation to be allowable, which in these cases, they are not. The representations for sentence (255) are shown in (256), using Sag's representation scheme.

(256)

$$\begin{aligned} \{Barbara, George\} \subset \hat{x}\hat{y}[x, \lambda r(r \text{ like } y)] \& \\ \{Hillary, Bill\} \subset \hat{w}\hat{z}[w, \lambda s(s \text{ like } z)] \\ \dots[Tipper, \lambda t(t \text{ like } him_j)] \end{aligned}$$

The two lambda expressions in the first two clauses are alphabetic variants and therefore the acceptability of gapping is predicted, but those in the second and third are not, so it is predicted that the VP-ellipsis in each case is infelicitous.

Sag then suggests a weakening of his alphabetic variance condition, with the result that both of the above examples would be predicted to be acceptable. He does not consider the case in which the judgements are correct as stated, and therefore considers no solutions that would predict this dichotomy.

The felicity of sentence (254) and the infelicity of sentence (255) are exactly what the analysis given here predicts. The representation for the second clause after resolution is given in Figure 5.4. In example (255), the third clause is in a Resemblance relationship with the second (and the first, for that matter), so the algorithm attempts to retrieve the antecedent syntactic structure. The VP representation in its current state is not suitable for copying for VP-ellipsis resolution, because two constituents have been extracted out of the sentence (and therefore one has been extracted out of the VP). Therefore, the ellipsis in the third clause fails since the trace will not find an antecedent in the target. On the other hand, the third clause in example (254) is related to the second by a Cause-Effect relation. Therefore, there is no requirement to reconstruct the syntax of the VP; semantically-based VP-ellipsis resolution only makes reference to the sentence-level propositional semantics and therefore succeeds. Thus, the apparent paradox between examples (254) and (255) is just what is predicted by our account.

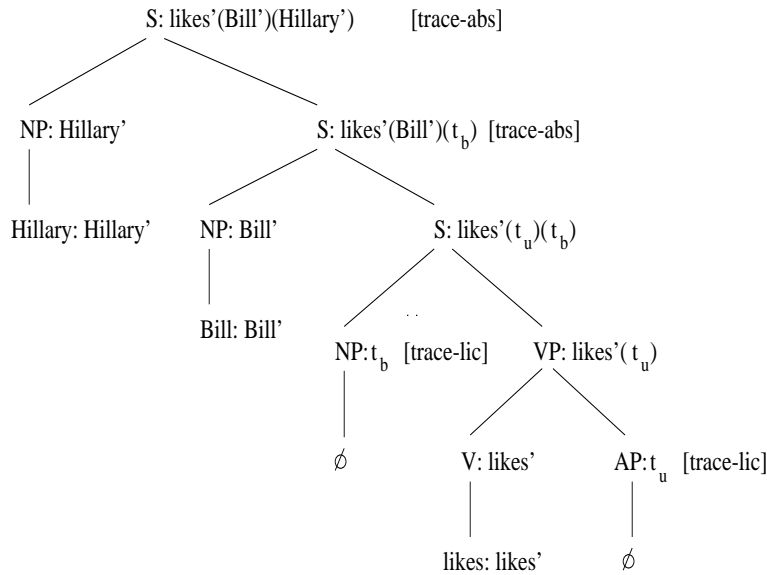


Figure 5.4: Syntactic and Semantic Representations for *Hillary Bill* after reconstruction.

5.1.5 Comparison to Past Work

Levin and Prince give an analysis of the facts they note about gapping by utilizing the *ordered entailment* framework of Wilson and Sperber (1979). In that framework, processing a sentence results in a computation of foreground and background entailments. The background entailments are those “presupposed” propositions resulting from applying constituent-to-variable replacement rules on focussed constituents. Specifically, the *First Background Entailment* (FBE) is the open proposition resulting from replacing a minimal tonically stressed (or clefted) constituent with a variable. For instance, sentence (257a), with the indicated stress on *Bill*, has as its First Background Entailment form (257b), along with other background entailments (257c) and (257d).

- (257) a. **BILL'S** father writes books.
 b. Someone's father writes books.
 c. Someone writes books.
 d. Someone does something.

As Levin and Prince put it, an utterer of a sentence such as (257a) is taking (257b) to be in the hearer's consciousness at speech time (i.e., “given”).

Levin and Prince utilize this account in postulating the following rule applying to gapping:

(258) Discourse Function of Gapping:

Upon hearing a gapped sentence, a Hearer infers that the Speaker intends that both (all) the conjuncts of the Gapped sentence share a single open proposition as their First Background Entailment, i.e., as that which is appropriately in the hearer's consciousness at that point in the discourse. The open proposition consists of (the representation of) the material deleted in the second (through nth) conjunct, with variables replacing (the representation of) the constituents remaining in the second (through nth) conjunct. The foreground is, of course, the new information.

To see how this principle is used in predicting the facts concerning gapping and causal implicature, consider again sentence (259).

(259) Sue became upset and Nan became downright angry.

Because the corresponding elements in both the source and the target clauses are contrastively stressed in gapping structures, under the symmetric reading in the Levin and Prince account, the two clauses share the First Background Entailment given in (260).

(260) Someone became something. [open proposition is: $\lambda x, y. become(x, y)$]

Under the causal reading, they claim that there are a number of possibilities for what the First Background Entailments are, for example, one set of possibilities for the two clauses in sentence (259) may be as given in (261a-b) respectively.

(261) a. Something happened.

b. Nan did something.

Because two First Background Entailments are required for causal implicature, the Discourse Function of Gapping rule accounts for why gapped sentences don't yield causal implicatures.⁵

Our account improves upon the Levin and Prince account in several respects. First, where Levin and Prince have to stipulate the Discourse Function of Gapping rule, in our account this behavior is the result of more fundamental and independently-motivated processes. Second, we have accounted for the data with an analysis that is unified with an account of VP-ellipsis, which itself does not have the same pragmatic

⁵For causation, it seems that both clauses could still share the same FBE, specifically one of the form *something happened*. But even in this case, the Discourse Function of Gapping rule is not satisfied, since this open proposition does not contain two variables standing proxy for the non-elided constituents.

effects that gapping does. Finally, we have extended the argument to show why gapping fails in constructions involving a variety of other conjunctions. However, both accounts share the fact that these data are handled by appeal to discourse reasoning, and not only surface-syntactic issues.

We will forgo a comparison between our analysis and the remainder of the approaches cited at the beginning of this section, and instead make a more general point. In many of these works, it is stipulated that gapping does not apply in subordinated clauses (such as in sentences (251a-d)) as it does in coordinated ones. Of course, such a statement is not in itself a theory, but instead simply a descriptive fact. Furthermore, the fact that gapping is unacceptable with coordinating conjunctions indicating Cause-Effect relations suggests that the purely syntactic split between coordinating and subordinating conjunctions is not the correct one to make. That is, this stipulation does not address the fact that gapping is infelicitous in Cause-Effect uses with coordinating conjunctions. The analysis presented here accounts for these cases as well as those involving syntactic subordination, and does so in a theoretically-motivated and non-stipulatory way.

Finally, in Chapter 3 we discussed two theories that tie ellipsis resolution to broader theories of discourse structure and coherence, those of Prüst (1992) and Asher (1993). Prüst addresses gapping, but does not acknowledge the infelicity of gapping with Cause-Effect relations, and therefore provides no account for it. Furthermore, it appears that neither Asher nor Prüst can account for the infelicity of mixed gapping/VP-ellipsis cases as discussed in Section 5.1.4.

5.1.6 Conclusions

In this section, we have provided an analysis of gapping that accounts for the facts noted by Levin and Prince (1982), as well as additional facts in constructions with conjunctions other than *and*. The correct predictions result from assuming the account of VP-ellipsis presented in Chapter 3 without modification; the differences between gapping and VP-ellipsis result only from the independently-motivated fact that gapping is not referential whereas VP-ellipsis is. The analysis was also shown to extend to the stripping construction, as well as the cases of mixed gapping/VP-ellipsis constructions noted by Sag (1976). The work is compatible with and extends previous syntactic accounts of gapping, although it provides a more adequate division between the data than the purely syntactic distinction between coordination and subordination often stipulated in that work.

5.2 Event Reference

In this section we discuss various forms of event reference, with particular reference to the distinction between *deep* and *surface* anaphora proposed by Hankamer and Sag (1976, henceforth H&S). After discussing the H&S dichotomy, we show how the syntax and semantics of event referential forms differs from that of elliptical forms like gapping and VP-ellipsis, and show how they interact with our architecture. We show that our account makes a better separation between different anaphoric behaviors than theirs; furthermore this separation does not arise by stipulating a categorization of types of anaphora but instead results from more basic facts concerning the different constructions. We then discuss another form that is problematic for their architecture, *do so* anaphora.

5.2.1 The Sag and Hankamer Dichotomy

In their classic study of anaphora, Hankamer and Sag (1976) argue for a categorical distinction between so-called *deep* and *surface* anaphora. Surface anaphors are ‘syntactically controlled’, requiring a linguistic antecedent of an appropriate syntactic form. Examples of surface-anaphoric forms include VP-ellipsis, gapping, and stripping. Deep anaphors, on the other hand, may be ‘pragmatically controlled’ in that they do not require a linguistic antecedent at all, but only an appropriately constructed referent in the discourse model. Examples of deep-anaphoric forms include pronominals and event referential forms like *do it* and *do that* anaphora.

In Sag and Hankamer (1984), this dichotomy is revised to distinguish between two types of anaphoric process, *ellipsis* (their earlier surface anaphora) and *model-interpretive anaphora* (their earlier deep anaphora). The former process obtains antecedents from *propositional representations*, which maintain the surface syntactic constituent structure of a sentence. On the other hand, model-interpretive anaphora (MIA) locates antecedents in a *discourse model*, where the representations are purely semantic in nature.⁶

The distinction between the two types of anaphora is illustrated in sentences (262a-c). As we did for VP-ellipsis and gapping, we will refer to the clause containing the anaphor as the *target* clause, and the clause giving rise to the antecedent as the *source* clause.

- (262) a. The decision was reversed by the FBI, and the appeal was too. [VP-ellipsis
(surface)]

⁶For historical reasons, we will continue to use the original ‘surface’ and ‘deep’ terminology, as introduced in Hankamer and Sag (1976).

- b. # The decision was reversed by the FBI, and the ICC did too. [VP-ellipsis (surface)]
- c. The decision was reversed by the FBI, and the ICC did it too. [event anaphora (deep)]

In Hankamer and Sag’s theory, example (262a) is acceptable because the source representation is a surface VP and is therefore retrievable from its propositional representation. Example (262b) is unacceptable because the putative source *reverse the decision* is not a surface VP in its propositional representation. Sentence (262c), on the other hand, is acceptable because *do it* is a deep anaphor, and therefore is interpreted with respect to a discourse model, in which a purely semantic representation for *reverse the decision* is assumed to exist.

In the H&S dichotomy of anaphora, the requirement that there be a syntactic antecedent for surface anaphora implies that the antecedent must be linguistic, i.e., that surface anaphora cannot be what they call ‘pragmatically-controlled’.⁷ The unacceptability of pragmatically-controlled VP-ellipsis is shown in example (263), from Hankamer and Sag (1976).

(263) [Hankamer points gun offstage and fires, whereupon a blood-curdling scream is heard. Sag says:]

- a. * Jorge, you shouldn’t have! [VP-ellipsis (surface)]
- b. Jorge, you shouldn’t have done it! [*do it* anaphora (deep)]

In sum, two hallmarks of surface anaphora are that (1) there must be a linguistically-evoked antecedent, and (2) that antecedent must be of an appropriate syntactic form. On the other hand, deep anaphora only requires a suitable semantic antecedent in the discourse model, and allow for such antecedents to be pragmatically-controlled.

5.2.2 Syntax and Semantics of Event Reference

We now consider the syntactic and semantic properties of event referential forms such as *do it* and *do that*. As we noted in Chapter 3, it is the main verb *do* that is operative in these forms of anaphora, in contrast to the auxiliary *do* operative in VP-ellipsis. For instance, other auxiliaries can appear in elided forms as in sentences (264a-c), but cannot be followed by *it* or *that* as in sentences (265a-c).

(264) a. Bill signed the legislation, and Al did too. [auxiliary *do*]

⁷We referred to such antecedents as “situationally-evoked” in Section 3.7.3.

- b. Bill has signed the legislation, and Al has too.
- c. Bill will sign the legislation, and Al will too.

- (265) a. Bill signed the legislation, and Al did (it / that) too. [main verb *do*]
 b. * Bill has signed the legislation, and Al has (it / that) too.
 c. * Bill will sign the legislation, and Al will (it / that) too.

Furthermore, *do* requires that its object be non-stative; a pronominal object cannot refer to a state as VP-ellipsis can, as shown in sentences (266).

- (266) a. Bill likes McDonald's, and Hillary does too.
 b. # Bill likes McDonald's, and Hillary does (it / that) too.
 c. Al wants to be president, and Tipper does too.
 d. # Al wants to be president, and Tipper does (it / that) too.

In these event referential forms it is the pronoun that is anaphoric; the fact that the pronouns refer to events results from the type constraints imposed by the main verb *do*. These forms are thus analyzed as the transitive main verb *do* followed by a referential direct object; they constitute full verb phrases in the syntax. These forms are therefore anaphoric in the semantics, but do not leave behind an empty constituent in the syntax. The syntax and semantics of the sentence *Hillary did it* is shown in Figure 5.5.




Figure 5.5: Syntax and Semantics of *Hillary did it*

We can now apply our account to cases of event reference. Again as H&S note, although elliptical sentences such as (262b) are unacceptable due to a voice mismatch, similar examples with event referential forms such as (262c) are much more acceptable. Given the syntactic and semantic properties we have outlined, this fact is predicted

by our analysis. In our system, the need for syntactic parallelism results from a missing constituent being reconstructed in the syntax. Since these forms have no missing constituents, the anaphoric form in the semantics is free to obtain a referent without regard to syntactic parallelism or the type of discourse inference employed. In the next section, we describe the dichotomy of anaphoric forms that results from our analyses.

5.2.3 A New Dichotomy

Sag and Hankamer categorize the different types of anaphora based on their referential properties, including whether they require a syntactic antecedent and whether they allow pragmatically-controlled antecedents. Among the *surface* anaphors are gapping, stripping, VP-ellipsis, and *do so* anaphora; these all require syntactic antecedents in their account, and therefore do not allow for pragmatic control. Among the *deep* anaphors are *do it* and *do that* anaphora; these only require purely semantic discourse model antecedents and allow for pragmatic control.

Several types of data that we have presented in this thesis are problematic for this dichotomy. First, we have seen cases of VP-ellipsis that require only suitable semantic antecedents without the need for syntactic parallelism; these were the cases that were related by Cause-Effect relations. On the other hand, cases of Resemblance showed the need for syntactic parallelism, so a move to categorize VP-ellipsis among the deep anaphors is not sufficient either. Second, while we agree with S&H that gapping requires a syntactic antecedent, there is nothing in their theory to explain why gapping is infelicitous in cases of Cause-Effect relations, where nonetheless a suitable syntactic antecedent is available.

Our architecture provides a different way of dividing the types of anaphora. This division does not arise by stipulating a categorization of types of anaphora as primitive notions as Sag and Hankamer do, but instead results from more basic facts concerning the syntactic and semantic properties of the constructions. Specifically, the division results from two properties of the analysis: (1) the syntax of empty constituents is reconstructed during Common Topic inference, which imposes the need for a suitable syntactic source, and (2) anaphoric forms refer independently to semantic forms in the discourse model. These facts result in the predictions that (i) gapping and VP-ellipsis behave as surface anaphora in cases in which a Resemblance relation is operative, (ii) VP-ellipsis behaves like deep anaphora in cases in which a Cause-Effect relation is operative, (iii) gapping is not felicitous at all in cases of Cause-Effect, and (iv) event referential forms such as *do it* and *do that* always operate as deep anaphora, since there is no missing constituent to be reconstructed.

The aforementioned discussion makes no mention of the issue of pragmatically-controlled antecedents. In the next section, we argue that the requirement for lin-

guistically-evoked antecedents needs to be distinguished from the requirement for a syntactically parallel antecedent. The basis for this argument is another form of event reference, *do so* anaphora.

5.2.4 Another Problem for Sag and Hankamer: ‘Do So’ Anaphora

The *do so* construction is exemplified in sentence (267).

(267) Bill signed the legislation, and Al did so too.

The *do so* construction is superficially similar to the *do it* and *do that* constructions, in that it is the main verb *do* that is operative. For instance, the use of *so* is not felicitous with other auxiliaries as is VP-ellipsis, as shown in sentences (268a-c).

- (268) a. Bill signed the legislation, and Al did so too. [main verb *do*]
b. * Bill has signed the legislation, and Al has so too.
c. * Bill will sign the legislation, and Al will so too.

Also, like *do it* and *do that*, *do so* requires an activity as antecedent, as shown in examples (269).

- (269) a. Bill likes McDonald’s, and Hillary does too.
b. # Bill likes McDonald’s, and Hillary does so too.
c. Al wants to be president, and Tipper does too.
d. # Al wants to be president, and Tipper does so too.

Despite these similarities with other deep-anaphoric forms, H&S treat the anaphor *so*, and consequently the form *do so*, as a surface anaphor. The basis for this classification is the fact that *do so* lacks the ability to be pragmatically controlled, as illustrated by example (119), repeated below as (270c).

- (270) [Hankamer points gun offstage and fires, whereupon a blood-curdling scream is heard. Sag says:]
a. * Jorge, you shouldn’t have! [surface]
b. Jorge, you shouldn’t have done it! [deep]

c. * Jorge, you shouldn't have done so! [surface]

Because *do so* fails to be felicitous with a pragmatically-controlled antecedent, it patterns with surface anaphora in this respect.

However, *do so* does not satisfy the other defining characteristic of surface anaphora, which is the requirement for a syntactically-parallel antecedent. An examination of naturally-occurring data has turned up many instances of *do so* where no acceptable surface-syntactic VP is available, which we describe below.

Voice Alternation: In each of the following examples, there is a voice mismatch between the source and target clauses: the main verb of the source clause occurs in the passive voice and the main verb of the target clause in the active voice.

(271) ... since regardless of which bit is initially assigned, it will be flipped if more information is gained by doing so. [flipping it] (text of Magerman (1994, page 29))

(272) Section 1 provides the examples to be derived by Gapping, and a formulation of Gapping capable of doing so. [deriving the examples] (text of Neijt (1981))

(273) As an imperial statute the British North America Act could be amended only by the British Parliament, which did so on several occasions. [amend an imperial statute] (Groliers Encyclopedia)

(274) That design, federal officials say, has to be reevaluated. To do so, federal officials have ordered more borings to obtain a more accurate picture of the soil conditions, and, they hope... [reevaluate that design] (newspaper article)

In these examples, the source clause is not of a suitable form at the surface-syntactic level to license deletion in the target. On the other hand, we might expect that a suitable semantic antecedent is available in the discourse model (where presumably the distinction between passive and active is lost), and therefore *do so* patterns like deep anaphora in these cases.

Nominalized antecedents: In each of the following examples, the antecedent of *do so* is evoked by a nominalization contained in the source clause.

(275) The defection of the seven moderates, who knew they were incurring the wrath of many colleagues in doing so, signaled that it may be harder to sell the GOP message on the crime bill than it was on the stimulus package. [defecting] (Washington Post article)

- (276) For example, in the dialogue of Figure 2, the purpose of the subdialogue marked (3) is to support the agents' successful completion of the act of removing the pump of the air compressor; the corresponding Shared-Plan, marked (P3) in Figure 3, specifies the beliefs and intentions that the agents must hold to do so. [successfully complete...] (from text of Lochbaum (1994))
- (277) Even though an Israeli response is justified, I don't think it was in their best interests to do so right now. [respond] (token provided by Dan Hardt)

In each case, *do so* felicitously refers to an event in the discourse model, even though the event was evoked by a non-VP constituent at the level of syntax.

Split antecedents: The following examples are cases of 'split antecedents', in which the event being referred to depends on which entity serves as the agent of the referent.

- (278) The survey results, released by county officials this week, also showed that most of the teenagers who drank alcohol, smoked marijuana or had sex started doing so between the ages of 13 and 16. [drinking alcohol / smoking marijuana / having sex] (Washington Post article)
- (279) Fortunately, the first person to die in 1990 and the first couple to file for divorce in 1990 were allowed to do so anonymously. [die / file for divorce] (Roeper, cited in DSP)

We assume that no syntactic machinery exists for copying VPs from multiple source clauses to the syntax of the target clause, nor for mapping over the semantic entities specified by the group-denoting subject. The resolution of these cases is therefore presumed to be performed at the semantic level.

Binding conditions: A consequence of a deletion-based account of anaphora is that, prior to deletion, the target clause must be a valid syntactic structure with regard to various syntactic constraints. However, under such an account, the target clause in the following example would violate Condition C of Binding Theory in that a pronominal *c*-commands a non-pronominal NP.

- (280) Liszt was one of the great altruists in the history of music: he performed the large piano works of Robert Schumann and Frederic Chopin when they were physically unable to do so; [perform the large piano works of Robert Schumann/Frederic Chopin] (Groliers Encyclopedia)

Polarity alternation: In the following examples, there is switch in polarity from source to target, suggesting that the negation is scoped out of the VP at the time of resolution; furthermore, in some cases an article, quantifier, or conjunction must alternate with the switch.

- (281) The United States and the USSR recognize no territorial claims in Antarctica and make none, although reserving the right to do so. [make territorial claims in Antarctica] (Groliers Encyclopedia)
- (282) Although this causes no problem in the first pregnancy, it can do so in the second pregnancy. [cause *a* problem] (Groliers Encyclopedia)
- (283) Although he never built any rockets, he encouraged a group of young engineers who were doing so; one of the young men was Sergei Korolev, who became the chief designer of Soviet spacecraft in the 1950s and 1960s. [building *some* rockets] (Groliers Encyclopedia)
- (284) According to Docetism, the eternal Son of God did not really become human or suffer on the cross; he only appeared to do so. [become human *and* suffer on the cross] (Groliers Encyclopedia)

Other form mismatches: The following examples also display a syntactic form mismatch between the source and target clauses.

- (285) There was a lot more negativity to dwell on, if anyone wished to do so. [dwell on more negativity] (newspaper article)
- (286) With or without the celebration, Belcourt is well worth seeing, and you can do so year round. [see Belcourt] (newspaper article)

Cornish (1992) cites sentence (287) as a case where the antecedent is embedded within a deverbal adjectival phrase, noting that it is problematic for the Hankamer and Sag dichotomy.

- (287) ... He went on to claim that the allegedly high-spending Labour authorities had, by so doing, damaged industry and lost jobs. [spending highly]

Again, the necessary event representations in these cases are presumably available in the discourse model, but the necessary syntactic VP sources required under the hypothesis that *do so* is licensed by syntactic parallelism are not present.

In sum, *do so* does not require that there be a suitably-matching syntactic antecedent in the source clause, thus, it correlates with deep anaphora rather than surface anaphora in this respect. On the other hand, it is the case that *do so* is infelicitous with pragmatically-controlled antecedents. Therefore, we conclude that the requirement for syntactic parallelism and the ability to take non-linguistic antecedents need to be distinguished in a general theory of anaphora. We now discuss an analysis of *do so* that accounts for its anaphoric properties.⁸

Kehler and Ward (1995) provide an account where *do so* is analyzed as an intransitive *do* that is modified by an adverbial *so* that indicates that the modified event is discourse-old (Prince, 1992).⁹ Examples are provided that show that *do so* is just one instance of this use of *so*; this *so* may modify other verbs that denote a previously mentioned event or a generalization of one. In this sense the intransitive act of *doing* can be seen as the most general event which subsumes all specific events. This progression from specific to general can be seen in examples (288)–(290).

(288) “...And with complete premeditation resolved that His Imperial Majesty Haile Selassie should be strangled because he was head of the feudal system.” He was *so strangled* on Aug. 26, 1975, in his bed most cruelly. [Chicago Tribune 12/15/94]

(289) We dropped bombs on a few strategic targets one night, and in *so attacking* them, essentially started a war.

(290) Ethel offered to grade my papers for me and, by *so doing*, endeared herself to me for life. (Bouton, 1970)

In passage (288), the use of *so* indicates that the strangling has already been evoked into the discourse model, i.e., it is discourse-old. In passage (289), the *so* indicates that the attacking is discourse-old; here attacking is a superclass of the act of dropping a bomb. Finally, passage (290) contains the use of *so doing* (which is the preposed correlate of *do so*); here, the *so* indicates that the *doing* is discourse-old.

This analysis accounts for the referential properties of *do so* that we noted earlier. In this analysis, event referents are taken to be purely semantic entities in the discourse model, so we would expect no requirement for syntactic parallelism. On the other hand, we would also expect that the form cannot take pragmatically-controlled antecedents, as such referents are not discourse-old. That is, while such referents may be situationally salient, they have not been evoked into the discourse model; the

⁸This analysis is the result of joint work with Gregory Ward.

⁹A loose analogy can be drawn with the role of a definite determiner in indicating that the entity specified by the NP it modifies is discourse-old.

creation of discourse model referents from a non-linguistic sources requires *accommodation* (Lewis, 1979; Webber, 1991). Therefore, these properties are just what we would expect given the analysis presented here.

In sum, the facts concerning *do so* show that the requirement for syntactic parallelism and the ability to take pragmatically-controlled antecedents need to be distinguished in a general theory of anaphora. The *so* in *do so* signals that the modified event is discourse-old; therefore, it does not require syntactic parallelism, but cannot be used to refer to pragmatically-controlled antecedents.

5.2.5 Conclusions

To summarize this section, we have shown how our architecture improves on the dichotomies of anaphora proposed by Hankamer and Sag (1976) and Sag and Hankamer (1984). The division between the types of anaphora does not arise by stipulating a dichotomy as Sag and Hankamer do, but instead results from more basic facts concerning the syntactic and semantic properties of the constructions along with the account of discourse inference presented in Chapter 2. The resulting division makes the correct predictions for gapping and VP-ellipsis data that are problematic for their account. The analysis is also compatible with the anaphoric properties of *do so* anaphora. These properties show that that the requirement for syntactic parallelism and ability to take pragmatically-controlled antecedents need to be distinguished in a general theory of anaphora.

5.3 Tense Interpretation

In previous sections, we addressed outstanding problems in VP-ellipsis, gapping, and event reference interpretation with respect to theories of anaphora that require either syntactic or semantic parallelism for resolution. Data that are problematic for past approaches were accounted for by a theory in which both types of resolution come into play.

In this section as well as Section 5.4, we discuss two other linguistic phenomena that affect the cohesiveness of a discourse, tense interpretation and pronoun interpretation. We again characterize past work as being of two sorts. In these cases, however, instead of the dichotomy being between syntactic and semantic approaches, the distinction is between what we call *definite reference* and *coherence-based* approaches.

In this section we claim that to adequately account for the determination of temporal relations that hold between events in a discourse, a theory of tense interpretation should include aspects of each type of analysis. Once again, constraints imposed by coherence relations will be shown to interact with those imposed by a theory of tense interpretation to account for data that is problematic for past theories. In this case, however, the constraints do not correlate with Hume's three-way categorization of relations in any fundamental way; they instead derive from basic properties of the world, for instance the fact that causes precede effects.

5.3.1 Background and Past Work

Tense interpretation has received much attention in linguistics (Partee, 1984; Hinrichs, 1986; Nerbonne, 1986, inter alia) and natural language processing (Webber, 1988; Kameyama, Passonneau, and Poesio, 1993; Lascarides and Asher, 1993, inter alia). Several researchers (Partee, 1984; Hinrichs, 1986; Nerbonne, 1986; Webber, 1988) have sought to explain the temporal relations induced by tense by treating it as anaphoric, drawing on Reichenbach's separation between event, speech, and reference times (Reichenbach, 1947). To account for the forward progression of time induced by successive simple past tenses in a narrative, they treat the simple past as referring to a time evoked by a previous past tense. We will refer to these works as *definite reference* approaches. On the other hand, Lascarides and Asher (1993) take the view that temporal relations are resolved purely as a by-product of reasoning about coherence relations holding between utterances, and in doing so, argue that treating simple and complex tenses as anaphoric is unnecessary. This approach parallels the treatment of pronoun resolution espoused by Hobbs (1979), in which pronouns are modeled as free variables that are bound as a by-product of coherence resolution. We will refer to such works as *coherence-based* approaches. The Temporal Centering framework (Kameyama, Passonneau, and Poesio, 1993) integrates aspects of both

approaches, but patterns with the first in treating tense as anaphoric.

To limit the scope of this discussion and to make it more concrete, we focus on two analyses: the definite reference approach of Hinrichs (1986), and the coherence-based approach of Lascarides and Asher (1993). We discuss them in terms of examples (291a-b), taken from Lascarides and Asher (1993).

- (291) a. Max slipped. He spilt a bucket of water.
b. Max slipped. He had spilt a bucket of water.

Passage (291a) is understood as a *Narration*, with the events displaying forward movement in time. On the other hand, the clauses in passage (291b) are related by the *Explanation* relation, with the events displaying backward movement in time. How the two works address these examples is discussed below.

A Definite Reference Approach Definite reference approaches generally utilize Reichenbach's (1947) analysis of tense. We describe his system first, and follow with a discussion of a related proposal due to Hinrichs (1986).

Reichenbach makes the distinction between three times: the *point of speech*, which is the time that the utterance was made, the *point of the event*, which is the time the event is taken to have occurred, and the *point of reference*, which provides a temporal perspective point from which to interpret the point of the event. The relations between these times for the simple past and the past perfect are shown in (292).

- (292) a. Simple Past: $point\ of\ event = point\ of\ reference < point\ of\ speech$
b. Past Perfect: $point\ of\ event < point\ of\ reference < point\ of\ speech$

The Reichenbach analysis can be utilized in treating tense as anaphoric by taking the *point of reference* of the event specified by the second clause as referring to the *point of the event* of the event specified by the first clause. Since the events introduced with the past perfect have their point of the event ordered before their point of reference as indicated in (292b), the second event in passage (291b) is understood as occurring before the first. That is, with the times as marked in (293a),¹⁰ the temporal relations inferred are as indicated in (293b).

- (293) a. Max slipped ($t1_{POE}$). He had ($t2_{POR}$) spilt ($t2_{POE}$) a bucket of water.
b. $t2_{POE} < t2_{POR} = t1_{POE} < t2_{POS}$

¹⁰Here, *POE* stands for *point of event*, *POR* for *point of reference*, and *POS* for *point of speech*.

With this interpretation, the Reichenbachian account makes the correct predictions with regard to the past perfect in sentence (291b). However, as it stands, this account does not predict the forward progression of time for sequences in the simple past such as in example (291a), since it will predict that an event time specified by the second clause occurs at the same time as the event time of the first clause.¹¹

Hinrichs (1986) proposes an account that predicts that the events described by sequences of utterances in the simple past occur in temporal succession, that is, in the order in which they are introduced into the discourse. He says:

It turns out that if two sentences in the past tense *both* contain events that can be identified as either an accomplishment or an achievement,¹² then the events are understood as happening in succession. (Hinrichs, 1986, page 68)

We argue later in this section that this claim is an overstatement, but it does hold true for passage (291a).

As with most other analyses, events in Hinrichs' system introduce an event time into the discourse. However, to account for the forward movement of time in narratives, he also proposes that accomplishments and achievements introduce another reference time that is temporally ordered after the time of the event itself. A subsequent past tense then refers to this time "ensuring that two consecutive accomplishments or achievements in a discourse are always ordered in a temporal sequence." This process for example (291a) is depicted in Figure 5.6. Here, the first sentence introduces both the event time tI as well as the new reference point tI' . The event time evoked by the subsequent clause then refers to this reference point. Since the reference point is ordered after the event time for the first clause, the forward movement of time is captured.

Hinrichs does not discuss the use of the past perfect as is present in example (291b), but if we assume that he follows the Reichenbachian account, his analysis makes the correct predictions for sentences (291a-b). However, there are other ways of accounting for this ordering of event times; in the next section we describe an approach that handles it through coherence resolution.

A Coherence-Based Approach Lascarides and Asher (1993) give an account in which tenses are not treated as forms of definite reference; instead, temporal relations are recovered purely as a by-product of coherence resolution. The simple past and

¹¹Or, in an analysis in which the reference time is not taken to be referential for the simple past, the resulting prediction is only that both events occurred before the speech time.

¹²We limit the scope of this section by restricting the discussion to accomplishments and achievements (Vendler, 1967).

Max slipped (t1). He split (t2) a bucket of water.




Figure 5.6: Event times in Hinrichs' system

past perfect tenses therefore do not differ in their referential properties, but instead only differ with respect to their “discourse roles”.

Their account of passage (291a) follows directly from inferring a Narration relation between the two clauses. They provide the two rules for Narrations given below; the first being a default rule (which is therefore defeasible), and the second an indefeasible axiom.

Narration If the clause β currently being processed is to be attached by a discourse relation to the clause α that is part of the text processed so far, then normally, $Narration(\alpha, \beta)$ holds.

Axiom on Narration If $Narration(\alpha, \beta)$ holds, and α and β describe the eventualities e_1 and e_2 respectively, then e_1 occurs before e_2

In understanding passage (291a), a Narration relation is inferred between the first and second sentences without any information to the contrary. As a result, the events are understood as occurring in the order in which they were introduced into the discourse.

Things get more complicated in their account of the past perfect. Because the past perfect is treated as sententially equivalent to the simple past, the difference between the two forms will have to be accounted for elsewhere. They do this by first postulating the *Connections When Changing Tense* Law given in (294).

(294) **Connections When Changing Tense (CCT)**

$$\Box((\langle \tau, \alpha, \beta \rangle \wedge sp(\alpha) \wedge pp(\beta) \rightarrow C_{pp}(\alpha, \beta))$$

Here, $\langle \tau, \alpha, \beta \rangle$ means that some discourse relation holds between propositions α and β , $sp(\alpha)$ means that α is in the simple past, $pp(\beta)$ means that β is in the past perfect, and $C_{pp}(\alpha, \beta)$ means “that α and β are connected by the kind of discourse relation allowed between simple pasts and pluperfects”; in other words, those relations compatible with the backward movement of time. This last part is summarized in the rule given in (295).

$$(295) C_{pp}(\alpha, \beta) \rightarrow \text{Elaboration}(\alpha, \beta) \vee \text{Explanation}(\alpha, \beta) \vee \text{Parallel}(\alpha, \beta) \\ \vee \text{Contrast}(\alpha, \beta)$$

For the particular example given in passage (291b), Lascarides and Asher also introduce a *Slipping Law*, as given in (296),¹³

(296) **Slipping Law**

$$\langle \tau, \alpha, \beta \rangle \wedge C_{pp}(\alpha, \beta) \wedge \text{Info}(\alpha, \beta) > \text{Explanation}(\beta, \alpha)$$

which they gloss as follows:

... if spilling the water and slipping are connected so that either the spilling explains, elaborates, parallels or contrasts the slipping, then normally the spilling explains the slipping. (Lascarides and Asher, 1993, pg. 472).

Note that this law makes crucial reference to the C_{pp} predicate, so it only applies in cases in which the first clause is in the simple past, and the second clause is in the past perfect. The predicate $\text{Info}(\alpha, \beta)$ is glossed by “ α describes Max slipping and β describes Max spilling a bucket of water.”

Given these rules, interpreting passage (291b) causes the CCT Law to be satisfied, which then yields $C_{pp}(\alpha, \beta)$. This in turn causes the Slipping Law to be satisfied, which yields $\text{Explanation}(\beta, \alpha)$. Inherent in the definition of *Explanation* is the fact that the first argument (in this case β) precedes the second (i.e., α) since causes precede effects; therefore the correct temporal inferences result.

In the next section, we give examples that are problematic for both the Hinrichs and the Lascarides and Asher approaches. We give an account of tense interpretation that combines aspects of the definite reference approaches and the coherence-based approaches with a clean interface between them. The result is a theory that accounts for the problematic data without recourse to special, unmotivated rules or principles.¹⁴

5.3.2 Problematic Data

We argue that aspects of both definite reference and coherence-based analyses are necessary to account for the recovery of temporal relations. To demonstrate our approach we address examples (297a-d); passages (297a-b) are repeated from passages (291a-b).

(297) a. Max slipped. He spilt a bucket of water.

b. Max slipped. He had spilt a bucket of water.

¹³The $>$ operator is used for default implication.

¹⁴This account is also described by Kehler (1994c).

- c. Max slipped because he spilt a bucket of water.
- d. Max slipped because he had spilt a bucket of water.

As previously noted, passage (297a) is understood as a *narrative*, indicating that the spilling was subsequent to the slipping. Passages (297b-d) are instead understood as the second clause *explaining* the first, indicating that the reverse temporal ordering holds. We present two problems with the analyses we are considering. The first problem arises from treating the simple past as anaphoric. Specifically, if a treatment such as Hinrichs’s is used to explain the forward progression of time in example (297a), then it must be explained why sentence (297c) is as felicitous as sentence (297d). That is, one would predict a clash of temporal relations for sentence (297c): the simple pasts will induce the forward progression of time, but the conjunction indicates an Explanation relation, which implies the reverse temporal ordering. Thus, we would expect that only sentence (297d) would be felicitous, since in this sentence the temporal relations induced by the tenses and by the conjunction agree. But instead, sentence (297c) and sentence (297d) are both felicitous, having the same reading.

The second problem arises from assuming that all temporal relations are recovered solely from reasoning with coherence relations. Specifically, because the use of the simple past in passage (297c) is as felicitous as the past perfect in passage (297d) under the *explanation* interpretation (in these cases indicated explicitly by *because*), then it must be explained why passage (297a) is not understood as an *explanation* as is passage (297b), where in each case the relationship needs to be inferred. That is, in a purely coherence-based framework such as that of Lascarides and Asher, discourse relations are proven from domain information whether or not an explicit conjunction constrains the possible relation. Therefore, there is no explanation for why passages (297a-b) are not understood analogously to passages (297c-d). We present our analysis in the next section, and account for these facts in Section 5.3.4.

5.3.3 The Account

In our account, we first postulate rules characterizing the referential nature of tense, which may constrain the temporal relations holding between the events specified in the discourse. We then follow by specifying the role of discourse relations in further constraining the temporal relations between clauses.

The rules we propose for governing tense are given below.

1. *Main verb* tenses are indefinitely referential, creating a new temporal entity under constraints imposed by its type (i.e., past, present, or future) in relation to

a *discourse reference time*¹⁵ t_R . For instance, a main verb past tense introduces a new temporal entity t under the constraint *prior-to*(t, t_R). For simple tenses t_R is the speech time, and therefore simple tenses are not anaphoric.

2. Tensed *auxiliaries* in complex tenses are anaphoric, identifying t_R as a previously existing temporal entity. The indefinite main verb tense is then ordered with respect to this t_R .

The tenses used may not completely specify the temporal relations between the described events. We claim that these relations may be further refined by constraints imposed by the coherence relation operative between clauses, as long as the coherence relation does not require a temporal ordering that contradicts any ordering induced by the tenses (that is, the temporal ordering induced by the tenses constrains the set of coherence relations that can be inferred, cf. Kameyama, Passonneau, and Poesio (1993)). We describe coherence relations that are relevant to the examples in this section and give temporal constraints for them.¹⁶

Narration: The *Narration* relation is characterized by a series of events displaying forward movement of time, such as in passage (297a). As did Lascarides and Asher (1993), we capture this ordering as a constraint imposed by the *Narration* coherence relation itself:

(298) If *Narration*(P, Q) then $t_P < t_Q$

Furthermore, again following Lascarides and Asher (1993), we treat *Narration* as a “default” relation, i.e., the relation assumed to hold without an explicit cue to the contrary.

Parallel: The *Parallel* relation relates utterances that share a common topic. This relation does not impose constraints on the temporal relations between the events beyond those provided by the tenses themselves. For instance, if passage (297a) was uttered in response to the question *What bad things happened to Max today?* (inducing a *Parallel* relation instead of *Narration*), a temporal ordering among the sentences is no longer implied.

Elaboration: Utterances standing in the *Elaboration* relation each describe the same event, therefore imposing the constraint that the times be the same.

¹⁵This term is borrowed from Kameyama et al. (1993).

¹⁶We assume here that the two clauses in question are related directly by a coherence relation. This may not be the case; for instance the use of a past perfect may signal the start of an embedded discourse segment, as in Webber’s flower shop example (Webber, 1988; Kameyama, Passonneau, and Poesio, 1993). How this account is to be extended to address coherence at the discourse segment level is the subject of future work.

(299) If *Elaboration*(P, Q) then $t_P = t_Q$

Result: The *Result* relation imposes the same temporal constraints as *Narration*, since causes precede effects.

(300) If *Result*(P, Q) then $t_P < t_Q$

Explanation: The *Explanation* relation denotes a *Result* relationship with reversed clause ordering, as in sentences (297b-d). Therefore, the second event is constrained to preceding the first.

(301) If *Explanation*(P, Q) then $t_Q < t_P$

To summarize the analysis, we claim that tense operates as indefinite reference with respect to a possibly anaphorically-resolved discourse reference time. The temporal relations specified may be further refined as a by-product of establishing the coherence relationship operative between clauses, *Narration* being but one such relation.

5.3.4 Examples

We now analyze the examples (297a-d) as presented in Section 5.3.2, using this approach.

In passage (297a), the two clauses are in the simple past tense. In this analysis, each evoke a new time into the discourse, ordered in the past with respect to the speech time. No ordering is inferred between the times as a result of the tenses used. Instead, the implicit ordering on the times results solely from understanding it as a *Narration*, which is the default interpretation given that no other relation was cued.

Passage (297b) is similar to passage (297a), except that the second clause is in the past perfect. In this case, the auxiliary *had* refers to the event time of the slipping, and thus the past tense on *spill* creates a temporal entity constrained to precede that time. This necessitates a coherence relation that is consistent with this temporal order, in this case, *Explanation*.

Passage (297c) is similar to passage (297a), except that a relation other than *Narration* is cued, in this case *Explanation* (cued by *because*). As with passage (297a), the times as evoked by the simple pasts are not ordered with respect to each other. They are then ordered by the *Explanation* relation itself, resulting in the backward progression of time.

Finally, in passage (297d), the tense orders the times in backward progression (as in passage (297b)), and the *Explanation* relation inferred (cued by ‘because’, as in passage (297c)) is consistent with that ordering.

Restating the first problem noted in Section 5.3.2, if treating the simple past as anaphoric is used to account for the forward progression of time in passage (297a), then one would expect the existence of the Explanation relation in passage (297c) to cause a temporal clash, where in fact passage (297c) is perfectly felicitous. No clash of temporal relations is predicted by our account, because the use of the simple pasts do *not* in themselves imply a specific ordering between them. The Narration relation orders the times in forward progression in passage (297a) and the Explanation relation orders them in backward progression in passage (297c). The Parallel relation would specify no ordering. (See the potential context for passage (297a) given in Section 5.3.3.)

Restating the second problem noted in Section 5.3.2, if temporal relations can be recovered solely from reasoning with coherence relations, and the use of the simple past in passage (297c) is as felicitous as the past perfect in passage (297d) under the Explanation interpretation, then it remains to be explained why passage (297a) is not understood as an Explanation as is passage (297b), where in each case the relationship needs to be inferred. In our account, this is explained by our treating Narration as the default relation in absence of a specific cue to the contrary. The use of the past perfect (as in passage (297b)) is one such cue since it implies reversed temporal ordering; the use of an explicit conjunction indicating a coherence relation other than Narration (as in passages (297c-d)) is another such cue. While passage (297a) could be understood as an Explanation on semantic grounds, the hearer assumes Narration since no other relation is cued.

5.3.5 Comparison to Definite Reference Approaches

Recall that in Hinrichs’s approach, events are taken to evoke reference times that are ordered after the time of the event itself. This allows for the account to predict the forward movement of time in narratives. There are two problems with this approach; the first being simply that such a stipulation is not well-motivated. That is, there is no independently-motivated justification for such a proposal except to account for the phenomenon at hand. The second problem is more concrete: it makes the wrong predictions in certain cases; we have already seen this in its predicting a temporal clash in sentence (297c).

There are also other scenarios which violate the Hinrichs account. Not surprisingly, these counterexamples also involve coherence relations other than Narration. Several such examples are given in Dowty (1986), although these are presented in reference to his own account. He gives a rule for “interpreting successive sentences in a discourse temporally”, called the *temporal discourse interpretation principle* (TDIP), which essentially states that in the absence of definite time adverbials indicating otherwise, that the reference time for a sentence is interpreted as a time which immediately

follows the reference time of the previous sentence. This rule is essentially stating the temporal constraints that are inherent in the definition of the Narration coherence relation. Dowty goes on to cite three types of apparent counterexamples to this principle.¹⁷ The first is given in passage (302), due to Dry.

(302) John knelt at the edge of the stream and washed his face and hands. He washed slowly, feeling the welcome sensation of the icy water on his parched skin.

Example (302) violates Dowty's principle as originally stated because the events are understood to have occurred at the same time. The operative coherence relation in example (302) is *Elaboration*, since the second sentence merely elaborates the information conveyed by the first. As stated earlier, because the same event is being described in both clauses, Elaboration constrains the times to be the same.

The second example he gives is another case in which the clauses are not understood as a Narration, given in passage (303).

(303) At the signal, everyone went to work at once. Mary searched the room for any of the items on the list that might be there. John went next door to do the same in Bill's apartment. Susan organized the rest of the players to canvass the block.

These sentences are related by the Parallel relation, which does not impose temporal constraints between the event times. The fact that they occurred at the same time in this case is due to the indication to this effect provided by the first sentence.

Lastly, he considers a case that he attributes to Kamp, given in passage (304).

(304) Pedro dined at Madame Gilbert's. First there was an hors d'oeuvre. Then the fish. After that the butler brought a glazed chicken. The repast ended with a flaming dessert.

¹⁷These examples cause him to modify the circumstances in which the TDIP applies:

It seems, therefore, that the TDIP must be modified to allow that if the discourse itself conveys some implication as to how events are to be specifically ordered, this should take priority over the third part of the rule that orders reference times successively. In other words, this third part is the "default case" to be followed when neither time adverbials nor entailments and implicatures of the discourse itself give clues to the ordering of events. (Dowty, 1986, page 58)

In our account, the "implication as to how events are to be specifically ordered" simply results from the temporal constraints applied by the coherence relation that is recognized, which in each of the following examples is a relation other than Narration.

This is again a case of *Elaboration*, here holding between the first sentence and a *discourse segment* containing the remaining sentences. While the remaining sentences are related by Narration (and thus a forward movement of time is inferred), these sentences are not related to the first by Narration. This example is one that shows that all of these analyses, including ours, need to be scaled up to include temporal *intervals*; at the segment level, the Elaboration relation constrains the dining and the subsequent series of events to occurring over the same temporal interval.

In any case, taking Dowty’s modification to the TDIP as restricting its application to Narrations, these examples are no longer counterexamples. They are all counterexamples for Hinrichs, however, since in each case we have a series of accomplishments or achievements that does not display the forward movement of time that his account of tense will induce. In the account presented here, these examples are unproblematic, since the simple past tenses themselves do not impose any ordering; this work is done as a result of constraints on the coherence relation inferred. The correct predictions result without stipulating the introduction of additional reference times into the discourse that are not otherwise motivated.

Webber (1988) also presents a definite-reference account, although one that handles a greater range of examples than does the Hinrichs approach. In her framework, a simple past tense is anaphoric, but can refer to one of three times associated with a previously evoked event: the time of the event itself, the *preparatory* phase, and the *consequent* phase (these latter terms come from an ontology of events used by Moens and Steedman (1988) and Passonneau (1988)). Example (305) illustrates the first of these three possibilities.

- (305) a. John played the piano.
b. Mary played the kazoo.

Webber claims that this example is understood with the two events happening at the same time, as would be predicted by treating the past tense in the second sentence as referring to the time evoked by the first sentence (this is the first option for reference described above). She elaborates this by stating “whether this is further interpreted as two simultaneous events or a single event of their playing a duet depends on context and, perhaps, world knowledge as well”. However, we do not find example (305) to necessarily imply that the two playings are contemporaneous. For instance, if these sentences are spoken in response to the question *What instruments did John and Mary each play today?*, passage (305) only implies that the playings each happened some time that day, and not necessarily at the same time. (Therefore, we would also include this aspect of interpretation as one which is determined by context and world knowledge.) In our framework, the past tenses in passage (305) both evoke new event

times into the discourse model. The sentences are related by the *Parallel* relation, which imposes no further constraints on the times evoked.

Example (306) illustrates the case in which a past tense refers to the *consequent phase* of a previously evoked event.

- (306) a. John went into the florist shop.
b. He picked out three red roses, two white ones and one pale pink.

Because the consequent phase of an event is ordered after the event itself, such reference induces the forward movement of time. In our framework, this fact results from understanding this text as a *Narration*.

Finally, example (307) illustrates the case in which a past tense refers to the *preparatory phase* of a previously evoked event.

- (307) a. John bought Mary some flowers.
b. He picked out three red roses, two white ones and one pale pink.

In this case, the time of the picking event is understood as an initial step in the buying event, which is captured by the link to the preparatory phase. However, in this case the second sentence is understood as beginning an *Elaboration* of the event described in the first sentence. For instance, it might be completed as in passage (308).

- (308) a. John bought Mary some flowers.
b. He picked out three red roses, two white ones and one pale pink.
c. He walked up to the register and paid for them.

In this case, the second and third sentences form a *Narration* which as a segment serve as an *Elaboration* of the first sentence. In our account, if we extend the constraint that *Elaboration* places on the events to account for temporal intervals, the fact that the picking event occurred at the beginning of a more complex buying event results from the coherence constraints.

Given these three possibilities, it is not clear how Webber's analysis would handle example (297c). In that example, the second event is understood to occur before the first, even though it uses the simple past. This interpretation is not compatible with any of the three alternatives that Webber provides. It is closest to the possibility of referring to the preparatory phase of the first event, but this analysis cannot be maintained; the analysis would then fail to distinguish between the case in which the event happened strictly earlier than the evoked event (as in example (297c)), and the

case in which the event was part of (and thus temporally overlapping with) a more complex event that had been evoked (as in example (307)).

As a final point, the choice among the three possible referents for a particular tense in Webber's analysis cannot be determined without first comprehending the passage. Comprehending the passage requires that one identify the coherence relations operative between the utterances that comprise it. However, the constraints imposed by the coherence relations result in the same predictions regarding temporal relations as the tense-as-reference framework. Therefore, the tense-as-reference account provides no new information about temporal relations: the temporal relations have to be identified independently before the correct choice among the three options can be determined.

5.3.6 Comparison to Coherence-Based Approaches

We see several advantages of this approach over the purely coherence-based approach of Lascarides and Asher (1993, henceforth L&A). First, L&A note the incoherence of example (309)

(309) ? Max poured a cup of coffee. He had entered the room.

in arguing that the past perfect should not be treated as anaphoric:

Theories that analyse the distinction between the simple past and pluperfect purely in terms of different relations between reference times and event times, rather than in terms of event-connections, fail to explain why [(297b)] is acceptable but [(309)] is awkward. (Lascarides and Asher, 1993, pg. 470)

Example (309) indeed shows that coherence relations need to be utilized to account for temporal relations, but it does not bear on the issue of whether the past perfect is anaphoric. The incoherence of example (309) is predicted by both their and our accounts by virtue of the fact that there is no coherence relation that corresponds to Narration with reverse temporal ordering.¹⁸ Recall that in addressing this example, L&A specify a special rule (the *Connections When Changing Tense (CCT)* Law) that stipulates that a sentence containing the simple past followed by a sentence containing the past perfect can be related only by a subset of the otherwise possible coherence relations. However, this subset contains just those relations that are predicted to be possible in our account by treating the past perfect as anaphoric; they are the ones that do not constrain the temporal order of the events against displaying backward

¹⁸For instance, in the same way that Explanation corresponds to Cause-Effect with reverse temporal ordering.

progression of time. Therefore, we see no advantages to adopting their rule; furthermore, they do not comment on what other laws have to be stipulated to account for the facts concerning other possible tense combinations.

Second, to explain why the Explanation relation can be inferred for passage (297b) but not for passage (297a), L&A stipulate that their causal *Slipping Law* (stating that spilling can cause slipping) requires that the CCT Law be satisfied. This constraint is imposed only to require that the second clause contain the past perfect instead of the simple past. However, this does not explain why the use of the simple past is perfectly coherent when the Explanation relationship is cued overtly as it is in sentence (297c). (As previously noted, in their framework discourse relations should be inferrable whether or not they are cued by an overt conjunction.) Furthermore, they do not adequately explain why CCT must be satisfied for this causal law and not for those supporting similar examples for which they successfully infer an unsigned Explanation relation (see discussion of example (2), pg. 463).

Third, the L&A account does not explain why the past perfect cannot stand alone nor discourses generally be opened with it; consider stating sentence (310) in isolation.

(310) Max had spilt a bucket of water.

Intuitively, such usage is infelicitous because of a dependency on a contextually-salient time which has not been previously introduced.¹⁹ This fact is not captured by the L&A account because sentences containing the past perfect are treated as sententially equivalent to those containing the simple past. On the other hand, sentences in the simple past are perfectly felicitous in standing alone or opening a discourse, introducing an asymmetry in accounts treating the simple past as anaphoric to a previously evoked time. All of these facts are explained by the account given here.

5.3.7 Conclusion

In this section, we considered two types of approaches to tense interpretation, *definite-reference* approaches and *coherence-based* approaches. We discussed problems with the definite reference account of Hinrichs (1986) and coherence-based account of Lascarides and Asher (1993). Each analysis has evidence in support of it, but were forced to make additional stipulations to cover the data: for Hinrichs, it was the postulation of a second reference time evoked by the simple past, and for Lascarides and Asher, it was special rules to allow for an account in which the past perfect is not anaphoric.

¹⁹One does find the past perfect used in this manner as a scene-setting device in certain genres. However, it is this contextual dependency that is being flouted to achieve the desired rhetorical effect. The effect is analogous to the use of a pronoun in a story before its referent has been introduced.

In addition to the inelegancy of these stipulations, we have provided examples that are problematic for them.

We have given an account of temporal relations that combines aspects of both types of approach. Main verb tenses are resolved indefinitely with respect to a possibly anaphorically-resolved discourse reference time, and the resultant temporal relations may be further refined by constraints that coherence relations impose. The account covers the problematic data while also avoiding the need to make unmotivated stipulations. Such stipulations result from trying to cover complicated and perhaps contradictory data within the confines of a single module of language processing; by articulating a simpler theory of the linguistic phenomena along with the role of coherence resolution, the desired predictions result from independently motivated principles.

This study leaves several topics for future work. One topic is the interaction between tense and other types of temporal modifiers and connectives. In previous definite-reference approaches, the simple past tense has been taken to *refer* to times evoked by such expressions; here, the natural analysis is only that the indefinitely evoked time be *consistent* with any temporal information provided elsewhere (that is, while the two cannot contradict, no reference is involved). Other topics have already been noted in the preceding sections, including the extension of the analysis to handle intervals, an account of statives and additional types of events, and how tense interpretation interacts with the determination of discourse structure.

5.4 Pronoun Interpretation

In the previous section, we discussed two types of approach to tense interpretation that exist in the literature, namely *definite reference* approaches and *coherence-based* approaches. This same dichotomy of approaches exists in the pronoun resolution literature. Therefore, it seems natural to examine the pronoun interpretation problem in light of the themes put forth in this thesis.

The purpose of this section is to make a single point: that pronoun interpretation research, like the other phenomena addressed in this thesis, may also benefit from theories which combine a unified account of the linguistic phenomenon with a characterization of the role of discourse inference mechanisms that establish coherence relations. We show that the data that has been cited as support for competing analyses exhibits a pattern; once again this pattern correlates with the nature of the coherence relationship operative between clauses. Given the scope of the problem, however, we do not provide an in-depth account of pronoun interpretation here. Instead, we only provide a brief sketch of some properties that such an account in our framework might have, and compare these properties with related work.

5.4.1 Past Work

There are dichotomous views in the literature concerning the process by which pronoun resolution is performed; we again make the distinction between *definite reference* and *coherence-based* approaches. Definite reference theories treat pronominal reference as an independent process, often incorporating a notion of *focus* or *salience* to rank potential referents (Kameyama, 1988; Passonneau, 1991; Grosz, Joshi, and Weinstein, 1995, *inter alia*). Alternatively, purely coherence-based theories have suggested that pronominal reference is determined as a by-product of the determination of the coherence relations holding between sentences (Hobbs, 1979; Wilensky, 1978). We discuss each type of approach, and the type of data that supports each, below.

Definite Reference Approaches Definite reference approaches characterize pronominal reference as a phenomena with its own underlying resolution processes. Various factors guide these processes; among these is usually a notion of *focus* or *salience* that is used in heuristics for determining likely referents of pronouns. Such salience ranking is often determined in part by surface syntactic information. For example, the *Centering* framework of Grosz, Joshi, and Weinstein (1995) utilizes grammatical role information in specifying the accessibility of potential referents on an ordered forward-looking center list. Kameyama's (1988) work contains rules for property sharing utilizing grammatical roles. Passonneau (1991), in looking at the distribution of 'it' and 'that' for NP reference in naturally-occurring texts, concludes that both

syntactic form of the antecedent and its grammatical role are needed to adequately account for the data.

Evidence for a focus-based approach is provided by discourses such as (311), from Grosz, Joshi, and Weinstein (1995).

- (311) a. Terry really goofs sometimes.
- b. Yesterday was a beautiful day and he was excited about trying out his new sailboat.
- c. He wanted Tony to join him on a sailing expedition.
- d. He called him at 6AM.
- e. He was sick and furious at being woken up so early.

Sentence (311e) causes the hearer to be misled: while context indicates that the intended referent for *He* is Tony, readers tend to assign its referent to be Terry initially. Thus, such examples display a garden path effect; the hearer is apparently employing an interpretation mechanism that favors Terry, not waiting for semantic information to disambiguate the referent. Such examples suggest that there is something more than just reasoning involved in pronoun interpretation. While various definite-reference accounts differ in the way in which they implement this idea, common to them is that they favor the subject of the previous sentence (i.e., Terry) as the referent for the subsequent pronoun.

Here, we discuss the Centering theory of Grosz, Joshi, and Weinstein (1995, henceforth GJW) as an exemplar of a definite-reference approach. In Centering, each utterance in a discourse has exactly one backward-looking center (denoted C_b) and a partially-ordered set of forward-looking centers (C_f^1, \dots, C_f^n). The notation $C_b(U_n)$ is used to refer to the C_b of sentence n , and $C_f(U_n)$ to refer to the C_f list of sentence n . Following Brennan, Friedman, and Pollard (1987), we refer to the highest-ranked forward-looking center C_f^1 of sentence n as $C_p(U_n)$. Roughly speaking, $C_f(U_n)$ contains all entities that are referred to in sentence n ; amongst this list is $C_b(U_n)$. $C_b(U_{n+1})$ is the most highly ranked element in $C_f(U_n)$ that is realized in U_{n+1} . The rules specifying how entities are ranked on the C_f list are not fully developed, but factors that have been suggested to affect ranking include surface order, grammatical role, and pronominalization. Three intersentential relationships between a pair of sentences n and $n+1$ are defined:

1. Center Continuation: $C_b(U_{n+1}) = C_b(U_n) = C_p(U_{n+1})$.
2. Center Retaining: $C_b(U_{n+1}) = C_b(U_n)$, but $C_b(U_{n+1}) \neq C_p(U_{n+1})$.

3. Center Shifting: $C_b(U_{n+1}) \neq C_b(U_n)$.

Rules 1 and 2 impose constraints on center realization and movement respectively.

Rule 1: If any element of $C_f(U_n)$ is realized by a pronoun in U_{n+1} then the $C_b(U_{n+1})$ must be realized by a pronoun also.

Rule 2: *Sequences* of continuations are preferred over *sequences* of retaining; and sequences of retaining are to be preferred over sequences of shifting.

In particular, a pair of continuations across U_n and across U_{n+1} , represented as $\text{Cont}(U_n, U_{n+1})$ and $\text{Cont}(U_{n+1}, U_{n+2})$ respectively, is preferred over a pair of retentions, $\text{Ret}(U_n, U_{n+1})$ and $\text{Ret}(U_{n+1}, U_{n+2})$. The case is analogous for a pair of retentions and a pair of shifts.

Rule 1 itself does not predict the problem with the pronoun in sentence (311e) being assigned Terry as its referent, because it does not make any prediction when there is only one pronominalized entity. The algorithm of Brennan, Friedman, and Pollard (1987, henceforth BFP) utilizes Rule 2 in making predictions for pronominal reference.²⁰ The assignment of pronominal referents is that which yields the most preferred relation in Rule 2 (assuming Rule 1 is followed). This correctly accounts for the oddness of sentence (311e), since *he* coreferring with Tony constitutes a *Shift* relation whereas *he* coreferring with Terry constitutes a *Continue* relation.²¹

Sidner (1983) notes that there are counterexamples to focus-based accounts (including hers) which are induced by parallelism effects as in example (312).

(312) a. The green *Whitierleaf* is most commonly found near the wild rose.

b. The wild violet is found near it too.

According to focus-based theories, the entity specified by *the green Whitierleaf* is the most salient item in the discourse. Furthermore, there is no semantic reason that it cannot be the referent of *it* in sentence (312b). However, the referent of *it* in example (312) is understood to be the entity specified by *the wild rose*, which occupies a much less salient position in the previous sentence. It appears that the parallelism extant between the sentences in passage (312) is affecting the way in which the pronoun is interpreted.

Kameyama (1986) discusses example (312), as well as example (313).

²⁰BFP augment the hierarchy with a fourth relation, termed *Shifting-1*, to distinguish whether or not the backward-looking center is the highest-ranked member of the C_f list in a Shift. This distinction does not come into play in interpreting example (311).

²¹In actuality, however, the BFP algorithm does not always chose the most salient element in the previous discourse, and therefore diverges in some circumstances from other focus-based accounts. For more details, see the discussion in Suri and McCoy (1994) of their example (6), and how it contrasts with their example (1).

- (313) a. Carl is talking to Tom in the Lab.
b. Terry wants to talk to him too.

In light of these examples, she suggests that Centering be augmented with a *property-sharing* constraint which “requires that two pronominal elements realizing the same *Cb* in adjacent utterances share a certain common grammatical property.” In these cases, the constraint proposes that the two non-subject phrases cospecify (i.e., *him* and *Tom*).

Common to all of these definite reference accounts is the idea that pronoun interpretation is an independent process governed by heuristics relying on notions of salience and focus, which generally appeal to surface string position and grammatical role as factors affecting salience.

Coherence-Based Approaches In coherence-based approaches (Hobbs, 1979; Wilensky, 1978), pronoun interpretation is seen as a by-product of more general reasoning about the most likely interpretation of an utterance. In Hobbs’ account, for example, pronouns are modeled as free variables and are assigned to objects during the process of establishing coherence relations. Therefore, potential referents of pronouns are those which result in valid proofs of coherence. For instance, consider passages (314a-b), adapted from Winograd (1972).

- (314) The city council denied the demonstrators a permit because
a. they *feared* violence.
b. they *advocated* violence.

Each of these continuations is readily interpretable, despite the fact that in clause (314a) *they* refers to the city council and in clause (314b) *they* refers to the demonstrators; in each case an Explanation relation can be proven. The only difference between the two completions is the verb used; the difference in preference of referent appears to be solely attributable to their differing semantic implications.²²

However, Hobbs himself acknowledges the power of grammatical role-based heuristics, noting that upon hearing example (315),

²²This difference does not necessarily entail that there is *no* garden-path effect for sentence (314b); it is difficult to make claims about such effects without experimental evidence. The more pertinent question for us is whether the semantics of the verbs and the coherence cue provided by *because* affects the heuristics for resolution that would otherwise apply. This question is the subject of studies of the so-called “implicit causality” of verbs in the psycholinguistics literature, which suggest that such effects may exist (Garvey and Caramazza, 1974; Garvey, Caramazza, and Yates, 1976; Ehrlich, 1980, *inter alia*).

(315) John can open Bill's safe. He ...

one is likely to assume that *John* is the referent of *He*. The existence of a garden-path effect in example (316), in which *He* refers to *Bill* instead of *John*, suggests that pronominal reference resolution is not guided by coherence considerations alone.

(316) John can open Bill's safe. He's going to have to get the combination changed soon.

As focus-based theories would predict, the reader assigns *John* as the referent of *He*, and double-takes when semantic information later suggests otherwise. Hobbs attributes this effect to the fact that most coherence relations involve the same agent in each clause being related, and since agents are generally expressed as subjects, the hearer is inclined to make an initial guess that the previous subject is the referent of a succeeding pronoun. Hobbs acknowledges, however, that this point does not fully explain the heuristic.

To summarize, we have discussed two types of approach to pronoun interpretation, specifically *definite reference* approaches and *coherence-based* approaches. Each type of approach has a corresponding set of examples that provides support for it.

5.4.2 Accounting for the Role of Coherence Relations

The examples that we have discussed exhibit a familiar pattern with respect to the nature of the coherence relation operative between clauses. The garden path effects in cases of *Contiguity* like passage (311) are used as the basis of arguments for focus-based approaches. Coherence-based accounts do not adequately explain such effects. On the other hand, such effects may be mediated by a cue indicating a *Cause-Effect* relation as is present in sentences (314a-b); the integral role that semantic information plays in coherence-based accounts may help explain such effects. Cases of *Resemblance* such as those in passages (312) and (313) pose a problem for both types of approach, and have led researchers like Kameyama to augment a focus-based theory. As with the other phenomena discussed in this thesis, this pattern suggests that the way in which pronouns are resolved depends on the discourse inference mechanisms operative in establishing coherence, but in a way that goes beyond simply preferring an assignment that makes the most sense semantically.

It is already well-established that pronominal reference and coherence are inter-related phenomena. Pronouns have the interesting property that while they are a more ambiguous form than their counterpart non-pronominal definite descriptions, their use is generally preferred over the less-ambiguous form when their referent is salient in the discourse. In fact, in a set of reading-time experiments, Gordon, Grosz, and Gilliom (1993) found that pronouns were interpreted more quickly than proper

names, despite the fact that the pronouns were ambiguous and the proper names were not. This result suggests that pronouns serve to ease the processing load in interpreting interrelated sentences, apparently offsetting the greater load resulting from the need to disambiguate between potential referents. It is conceivable that the way in which pronouns ease the processing load may differ with respect to different types of discourse inference.

An account of pronoun resolution with these features is beyond the scope of this study; here we sketch some properties of our framework that such an account might utilize. We informally characterize the *default referent* for an (unstressed) pronoun as the one requiring the least amount of computational resources to access. The amount of computational resources might differ with respect to the type of discourse inference employed. In the case of Resemblance relations, Common Topic inference proceeds by identifying parallel elements as arguments to the relation; this identification requires reference to the syntactic structures of the utterances being related. In this case, the default referent might be the one which maximizes parallelism, that is, the entity denoted by the pronoun's parallel element in the clause containing the antecedent. For instance, in example (313), *Tom* and *him* are parallel elements, and so if *Tom* is identified as argument a_i to the Parallel relation (for some i), *him* will be identified as argument b_i . Parallelism is then maximized if these two phrases corefer, whether or not the resulting interpretation makes the most sense semantically. This process essentially results in the same predictions as Kameyama's property-sharing heuristic. Note that the reliance on grammatical notions inherent in the Centering approach falls out from the fact that the syntax is used to guide the identification of parallel elements in Common Topic inference.²³

In cases of Cause-Effect relations, Coherent Situation inference proceeds using the proposition-level semantics of the clause. In this case, we might think of the default referent of a pronoun as being the referent that allows for the most plausible Cause-Effect relation to be established. This process would lead to assigning *the city council* as the referent of *they* in example (314a), and *the demonstrators* as the referent of *they* in example (314b).

This is not to say, however, that *no* constraints on salience come into play in cases of Cause-Effect relations; for instance, in cases in which the semantic information

²³This analysis also suggests why Kameyama's heuristic does not apply as readily to other cases, such as example (317) in which *him* is taken to refer *Carl*.

- (317) a. Carl is talking to Tom in the Lab.
b. Terry will speak to him tomorrow about interrupting Tom's work.

While there is syntactic parallelism between the two clauses of example (317), they are not related by a *Resemblance* relation, and therefore Common Topic inference is not employed.

needed for resolution is delayed significantly, we still see what looks to be a focus-based effect as in sentence (318).

- (318) ? John can open Bill's safe, because he often leaves the combination where other people can see it. (where *he* = Bill)

On the other hand, there still appears to be a difference between cases in which a Cause-Effect relation is marked, as it is in examples (314a-b), and similar cases in which the relation is left implicit, as in sentences (319a-b).

- (319) The city council denied the demonstrators a permit.
- a. They *feared* violence.
 - b. ? They *advocated* violence. (where *They* = the demonstrators)

Sentence (319b) is more jarring than sentence (314b), even though the same coherence relation is eventually established for each.²⁴ We can explain this effect by appealing to a hypothesis we made in our account of tense interpretation given in Section 5.3, in which we suggested that hearers assume a Narration relation without an explicit cue to the contrary. This hypothesis predicts garden-path effects in cases in which another relation holds but is not cued, such as in example (319b), because in this case the hearer is presumably using the resolution strategy for Narrations when the pronoun is encountered. This case contrasts with sentence (314b), in which the Cause-Effect relation is cued by *because* before the pronoun is encountered, perhaps altering the resolution strategy. Further work is required to identify the appropriate division of labor between the pronoun resolution and discourse inference processes for these cases.

As we indicated in Chapter 2, the constraints imposed by the Contiguity relation Narration are less well understood, as are the discourse inference mechanisms underlying its recognition. We have characterized Narration as a mechanism for describing a complex sequence of events centered around some system of entities; the salient entities provide the points of connection between utterances. From this perspective, pronominal reference can be seen as establishing connections to salient entities as necessitated by the discourse inference process. We may stand to learn more about this process from the focus-based literature on pronoun resolution; for instance, the determination of salient entities presumably utilizes syntactic information such as

²⁴See McKoon, Greene, and Ratcliff (1994) for a description of psycholinguistic experiments that established this difference for pairs of examples similar to sentences (314b) and (319b).

grammatical role. While this syntactic information is available to the discourse inference process in our model, we have no more to say about the mechanics of this process.

We summarize by making the same meta-theoretical point that we have made throughout this thesis. That is, it may be possible to develop a unified account of a particular linguistic phenomenon (in this case, pronoun resolution), but only if the nature of the coherence relation operative between clauses is factored into the analysis.

5.4.3 Mixed Approaches

In this section, we describe previous work that is motivated by similar types of data as the discussion given here, but that attempt to cover all of the data with a single theory that does not make the distinctions that we advocate.

Kameyama

Kameyama (1994) gives an account of pronoun resolution that combines the effects of focusing with those of common-sense preferences. She describes four preferences in pronoun interpretation, which are paraphrased below.

1. Subject Antecedent Preference: prefer a subject referent over a non-subject one.
2. Pronominal Chain Preference: prefer a referent that was also pronominalized.
3. Grammatical Parallelism Preference: prefer a referent that occupies a similar grammatical role.
4. Commonsense Preference: prefer a referent that leads to a plausible interpretation.

She outlines examples supporting each of these preferences; in the examples below the entities indicated were the preferred referents in a survey that Kameyama performed. Examples (320) and (321) provide support for the *subject antecedent* preference.

(320) John hit Bill. Mary told him to go home [him = John].

(321) Bill was hit by John. Mary told him to go home [him = Bill].

Note that the semantics in both of these cases is the same; they only differ with respect to the voice of the first clause. The preferred referent for the pronoun is

different, however, corresponding to which person is made more salient (here, the one placed in subject position).

Examples (322) and (323) are cases in which the *pronominal chain* preference applies.

(322) Babar went to a bakery. He greeted the baker. He pointed to a blueberry pie.
[He = Babar]

(323) Babar went to a bakery. The baker greeted him. He pointed to a blueberry pie. [He = the baker]

In passage (322), both the subject antecedent preference and the pronominal chain preference predict the correct referent. In passage (323), the two preferences make different predictions; in this case the subject antecedent preference wins out (although three out of thirteen respondents chose *Babar* as the referent).

Example (324) is of the now familiar sort that supports the *grammatical parallelism* preference.

(324) John hit Bill. Mary hit him too. [him = Bill]

Finally, example (325) supports the *commonsense* preference; in this case semantic information is able to override the grammatically-based preferences.

(325) John hit Bill. He was severely injured. [him = Bill]

There are clear parallels to be drawn between these four preferences and the three types of relation that we have outlined. Kameyama's *subject antecedent* and *pronominal chain* preferences correspond to two preferences in Centering theory: a preference to refer to the C_p of the prior utterance, and a preference to refer to the C_b of the prior utterance. These preferences are indicative of the inherent ambiguity between Continuing and Shifting after a Retain in Centering. (This ambiguity is manifest in example (323), in which not all of the respondents agree on the preferred referent.) The clauses in these examples are related by the Narration relation, thereby falling into the *Contiguity* category. The *grammatical parallelism* preference is needed in examples in which the clauses are related by a *Resemblance* relation, for instance the *Parallel* relation operative in example (324). Lastly, the *commonsense preference* is needed in cases in which a *Cause-Effect* relationship holds, as in example (325).

As Kameyama notes, these preferences often contradict. She provides a list of *preference interactions*, specifying which preferences have precedence over others. However, it appears to us that the overriding preference depends on the nature of the coherence relationship between the clauses; the examples Kameyama has chosen bear

a striking similarity to the examples previously discussed in terms of the correlation between resolution strategy and coherence relationship. Therefore, Kameyama’s well-motivated preferences appear not to all apply to all examples; for instance, we might not want to apply the grammatical parallelism preference in cases of Cause-Effect. A better tack might be to tie the preferences to aspects of the coherence resolution process instead of having the preferences compete in interpreting all examples.

Psycholinguistic Studies

There is also an ongoing debate in the psycholinguistics literature concerning various pronoun resolution strategies that have been proposed. For instance, Crawley, Stevenson, and Kleinman (1990) argue for the primacy of the *subject assignment* strategy²⁵ over the *parallel function* strategy²⁶ in pronoun interpretation. Smyth (1994) argues for a “feature matching” strategy that embodies the opposite claim. Stevenson, Nelson, and Stenning (1993) posit that both strategies operate together, with each producing a candidate antecedent. They suggest that a ‘competition mechanism’ selects the final choice from these, citing the similar Competition Model proposed by McWhinney, Bates, and Kliegl (1984). Each study draws their conclusions from experiments using different sets of data.

There appears to be two options being considered in these works: (i) that either a single strategy is employed for pronoun resolution that will explain all of the data, or (ii) that several strategies are employed that ‘compete’ in some unwieldy and unpredictable fashion. As with other literatures discussed in this thesis, we see a continuing tradeoff between theories that sacrifice coverage in pursuit of unity and vice versa. What needs to be considered is a theory that is unified, but in which several other factors interact in *well-defined* ways to predict the varying distribution of the data.

We are suggesting in particular that the type of the coherence relation needs to be factored into the analysis. To our knowledge, past studies have not adequately controlled for these relations (although some have controlled for the *conjunction* between clauses, e.g., Ehrlich (1980)). Smyth (1994) does note that “the semantic structure of a conjunction ... imposes constraints on the discourse model that the listener constructs to interpret the sentence.” He considers the distinction between sentences (326a-b), noting that the causal reading indicated in sentence (326a) makes nonsubject assignment less felicitous than subject assignment, whereas in sentence (326b) the nonsubject assignment is preferred under the parallel interpretation.

(326) a. Phil tickled Stanley, and (so) Liz poked him. (*him* = Phil)

²⁵That is, the strategy that appears to be used in cases of Contiguity.

²⁶That is, the strategy that appears to be used in cases of Resemblance.

- b. Phil tickled Stanley, and Liz poked him. (*him* = Stanley)

Again, this fact correlates with our distinction between Coherent Situation inference, which is operative in sentence (326a), and Common Topic inference, which is operative in sentence (326b). Further psycholinguistic experimentation is required to test the correlation between pronoun interpretation strategy and coherence relation type.

5.4.4 Conclusion

In this section, we have shown that different sets of data used to support competing analyses of pronoun resolution correlate with an independently-motivated distinction between Hume's three broad classes of coherence relations. Therefore, the same distinction that has been used in accounts of other phenomena in this thesis may also have applications here. This pattern provides initial evidence that a unified account of pronoun resolution may require that the nature of the coherence relation operative between clauses be factored into the analysis. We have sketched some properties that such an account might have, but a more fully-articulated analysis is a subject for future work.

Chapter 6

Conclusion

6.1 Summary

In this thesis, we presented analyses and algorithms for resolving a variety of cohesive phenomena in natural language. Past work has attempted to explain the complicated behavior of these expressions with theories that operate within a single module of language processing. We have argued that such approaches cannot be maintained; in particular, the data we present strongly suggest that the nature of the coherence relation operative between clauses needs to be taken into account. We have provided analyses that combine unified accounts of each cohesive form with a unified characterization of the role of discourse inference processes underlying the establishment of such relations. In each case, we showed how the interaction of these two aspects of language processing results in a theory that is more empirically adequate than previous accounts.

As a foundation for these analyses, we presented a categorization of coherence relations and the discourse inference mechanisms used to apply the constraints they impose. Unlike past categorizations of relations, our categorization is based on particular properties of the relation definitions and of their corresponding discourse inference procedures. We also suggested ways in which the specific relations in each category may be derivable from more basic principles.

We utilized this theory to break the deadlock between *syntactic* and *semantic* approaches to resolving VP-ellipsis. We showed that the data exhibits a pattern with respect to our categorization of coherence relations, and presented an account which predicts this pattern. We demonstrated that a similar pattern exists for gapping constructions, although the distribution of gapping was shown to be most restricted in situations where VP-ellipsis is least restricted, and vice-versa. Our account also predicted this pattern without further modification; the differences between the distributions of the two forms result solely from an independently-supportable distinction

between their referential properties. Our analysis also makes the correct predictions with respect to forms of event reference. We demonstrated that the accounts of VP-ellipsis, gapping, and event reference result in a more independently-motivated and empirically-adequate distinction among types of anaphoric processes than do the classic works of Hankamer and Sag (1976) and Sag and Hankamer (1984).

We also utilized discourse inference processes in analyses that break the deadlock between *definite-reference* and *coherence-based* approaches to tense and pronoun interpretation. We claimed that definite-reference analyses do not adequately account for the full role of coherence resolution, whereas coherence-based accounts fail to acknowledge the full role the properties of the linguistic forms in determining their interaction with context. In the case of tense interpretation, we presented a theory that interacts with discourse inference processes to predict data that is problematic for past approaches in both categories. In the case of pronoun interpretation, we demonstrated that a pattern is exhibited by data commonly cited in the literature, one which again correlates with the nature of the type of coherence relation that is operative between clauses. This pattern suggests that a unified account of pronoun resolution may also require that the type of coherence relation be factored into the analysis. We made some preliminary proposals for how such a pattern might result from the properties of the different types of discourse inference.

In addition, we presented an account of VP-ellipsis resolution that predicts the correct set of strict and sloppy readings for a number of benchmark examples that are problematic for past approaches. The correct readings can be seen to result from a general distinction between *referring* and *copying* in anaphoric processes. The account extends to other types of reference, such as event reference and ‘one’-anaphora. We also analyzed the difference between source-determined and discourse-determined analyses to resolving these forms, concluding that discourse-determined analyses assume the very theory that source-determined analyses attempt to provide.

We believe this work has achieved two main objectives. First, our analyses of each of the phenomena we address contribute to their respective literatures, adding to our theoretical understanding of each. These advances in turn provide the basis for algorithms with better coverage than past approaches. Second, we have drawn attention to the more general impact of coherence processes in interpreting natural language. In the past, these processes have generally been studied in isolation; our work suggests that their role is considerably more far-reaching, and that future research in a variety of areas should take notice of it.

6.2 Future Directions

Finally, we identify three lines of future research suggested by the work presented here. Each of these constitutes a part of a broader investigation into the relationship between coherence and the interpretation of cohesive forms.

The first line of research would expand the study to additional types of expressions that affect the cohesion of a discourse. This category not only includes other types of ellipsis and anaphora, but also the use of certain syntactic constructions that have been linked to pragmatic functions, for instance, topicalization (Prince, 1981a), VP-preposing (Ward, 1990), and VP-inversion (Birner, 1994). The fact that a language affords its speakers several syntactic means to express the same idea suggests that the choices among them may have different ramifications for discourse inference processes. Using the studies we have presented as a basis for this work, one might investigate how the use of these other forms might also be seen to interact with the process of establishing the coherence of a discourse.

A second line of research would re-examine our models of discourse processing in light of the role of cohesive forms in language comprehension. A striking feature of cohesive forms is that they appear to reduce processing load during interpretation despite the fact that they are generally more ambiguous than their more explicit counterparts. For instance, passage (327) contains uses of pronominal object reference, event reference, and ellipsis; it also does not mark the temporal progression of time between the described events.

(327) Clinton said he couldn't give the speech because of another engagement. He asked Gore if he would do it. He agreed to.

Some of these forms have been replaced by their less ambiguous counterparts in passage (328). Despite greater explicitness, the change has reduced the coherence of the passage, making it more difficult to understand.

(328) Clinton said he couldn't give the speech because of another engagement. Then, Clinton asked Gore if he would give the speech. Then, Gore agreed to give the speech.

The difference between passages (327) and (328) suggest that cohesive forms assist the understanding process to an extent that offsets the added complexity resulting from the need to disambiguate them. However, few approaches to language interpretation are poised to take advantage of this effect; instead the use of such forms only complicates the understanding process. It may prove fruitful to investigate language processing architectures in which the positive effects of using cohesive forms might be more directly manifest.

A third, related line of research would reconsider approaches to computing coherence relations in light of cohesive forms and related cues to coherence. The abductive model due to Hobbs et al. (1993) is general and powerful; the cost of this power is that it relies on purely semantic, “AI-complete” reasoning processes. However, overt cues to coherence are often provided in naturally-occurring discourses, indicated by choice of linguistic form, syntactic form, connectives and cue words, and intonation. A theory which takes greater advantage of these cues may be more appropriate for the application of less powerful but more easily computable language processing techniques.

References

- Allen, J.F. and C.R. Perrault. 1980. Analyzing intention in utterances. *Artificial Intelligence*, 15(3):143–178. Reprinted in Grosz et al. 1986: 441-458.
- Alshawi, Hiyan, David Carter, Richard Crouch, Stephen Pulman, Manny Rayner, and Arnold Smith. 1992. Clare: A contextual reasoning and cooperative response framework for the core language engine. Technical Report CRC-028, SRI International, Cambridge Computer Science Research Centre.
- Asher, Nicholas. 1993. *Reference to Abstract Objects in Discourse*. SLAP 50, Dordrecht, Kluwer.
- Asher, Nicholas and Alex Lascarides. 1994. Intentions and information in discourse. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL-94)*, pages 34–41, Las Cruces, June.
- Barwise, Jon. 1987. Noun phrases, generalized quantifiers, and anaphora. In Pater Gärdenfors, editor, *Generalized Quantifiers*. D. Reidel, pages 1–29.
- Birner, Betty J. 1994. Information status and word order: An analysis of English inversion. *Language*, 70(2):233–259.
- Bouton, Lawrence P. 1970. Do so: Do + adverb. In J. M. Sadock, editor, *Studies Presented to R. B. Lees by his Students*. Linguistics Research, Edmonton, pages 17–38.
- Brennan, Susan E., Marilyn W. Friedman, and Carl J. Pollard. 1987. A centering approach to pronouns. In *Proceedings of the 25th Meeting of the Association for Computational Linguistics*, pages 155–162.
- Chao, Wynn. 1987. *On Ellipsis*. Ph.D. thesis, University of Massachusetts at Amherst.
- Chomsky, Noam. 1981. *Lectures in Government and Binding*. Foris, Dordrecht.
- Cohen, Philip R. and C. Raymond Perrault. 1979. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212. Reprinted in Grosz et al. 1986: 423-440.
- Cornish, Francis. 1992. *So Be It*: The discourse-semantic roles of *so* and *it*. *Journal of Semantics*, 9:163–178.
- Crawley, R. A., R. J. Stevenson, and D. Kleinman. 1990. The use of heuristic strategies in the interpretation of pronouns. *Journal of Psycholinguistic Research*, 19:245–264.

- Cremers, Crit. 1983. On the form and interpretation of ellipsis. In Alice G. B. ter Meulen, editor, *Studies in Modeltheoretic Semantics*. Foris, Dordrecht, pages 145–160.
- Crouch, Richard. 1995. Ellipsis and quantification: A substitutional approach. In *Proceedings of the Seventh Conference of the European Chapter of the Association for Computational Linguistics (EACL-95)*, pages 229–236, Dublin, Ireland, March.
- Dahl, Östen. 1972. On so-called “sloppy” identity. *Gothenburg Papers in Theoretical Linguistics*, 11. University of Göteborg.
- Dahl, Östen. 1974. How to open a sentence: Abstraction in natural language. In *Logical Grammar Reports, No. 12*. University of Göteborg.
- Dalrymple, Mary. 1991. Against reconstruction in ellipsis. Technical Report SSL-91-114, Xerox.
- Dalrymple, Mary, Stuart M. Shieber, and Fernando Pereira. 1991. Ellipsis and higher-order unification. *Linguistics and Philosophy*, 14:399–452.
- Dowty, David. 1986. The effects of aspectual class on the temporal structure of discourse: Semantics or pragmatics? *Linguistics and Philosophy*, 9:37–61.
- Ehrlich, Kate. 1980. Comprehension of pronouns. *Quarterly Journal of Experimental Psychology*, 32:247–255.
- Evans, Gareth. 1980. Pronouns. *Linguistic Inquiry*, 11:337–362.
- Fiengo, Robert and Robert May. 1994. *Indices and Identity*. MIT Press, Cambridge, MA.
- Gardent, Claire. 1993. A multi-level approach to gapping. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340*. University of Stuttgart, University of Tuebingen, and IBM Germany.
- Garvey, C. and A. Caramazza. 1974. Implicit causality in verbs. *Linguistic Inquiry*, 5:549–564.
- Garvey, C., A. Caramazza, and J. Yates. 1976. Factors underlying assignment of pronoun antecedents. *Cognition*, 3:227–243.
- Gawron, Mark and Stanley Peters. 1990. *Anaphora and Quantification in Situation Semantics*. CSLI/University of Chicago Press, Stanford University. CSLI Lecture Notes, Number 19.
- Gordon, Peter C., Barbara J. Grosz, and Laura A. Gilliom. 1993. Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17(3):311–347.

- Grosz, Barbara. 1977. The representation and use of focus in dialogue understanding. Technical Report Technical Note 151, Stanford Research Institute.
- Grosz, Barbara and Candace Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Grosz, Barbara J., Karen Sparck Jones, and Bonnie Lynn Webber, editors. 1986. *Readings in Natural Language Processing*. Morgan Kaufmann Publishers, Inc.
- Grosz, Barbara J., Aravind Joshi, and Scott Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics*. To appear.
- Grosz, Barbara J., Aravind K. Joshi, and Scott Weinstein. 1983. Providing a unified account of definite noun phrases in English. In *Proceedings of the 21st Conference of the Association for Computational Linguistics (ACL-83)*, Cambridge, MA.
- Haïk, Isabelle. 1987. Bound variables that need to be. *Linguistics and Philosophy*, 10:503–530.
- Halliday, M.A.K. and Ruqaiya Hasan. 1976. *Cohesion in English*. Longman's, London. English Language Series, Title No. 9.
- Hankamer, Jorge. 1971. *Constraints on deletion in syntax*. Ph.D. thesis, Yale.
- Hankamer, Jorge. 1978. On the nontransformational derivation of some null VP anaphors. *Linguistic Inquiry*, 9:66–74.
- Hankamer, Jorge and Ivan Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry*, 7:391–426.
- Hardt, Daniel. 1990. A corpus-based survey of VP ellipsis. Unpublished Manuscript.
- Hardt, Daniel. 1991a. A discourse model approach to VP ellipsis. In *Proceedings of the AAAI Symposium on Discourse Structure in Natural Language Understanding and Generation*, Asilomar, CA, November.
- Hardt, Daniel. 1991b. Towards a discourse level account of VP ellipsis. In *Proceedings of the Eastern States Conference on Linguistics (ESCOL 1991)*, pages 121–129.
- Hardt, Daniel. 1992a. Verb phrase ellipsis and semantic identity. In *Semantics and Linguistic theory II*, Ohio State University, May.
- Hardt, Daniel. 1992b. VP ellipsis and contextual interpretation. In *Proceedings of the International Conference on Computational Linguistics (COLING-92)*, Nantes, July.
- Hardt, Daniel. 1993. *Verb Phrase Ellipsis: Form, Meaning, and Processing*. Ph.D. thesis, University of Pennsylvania.

- Hellan, Lars. 1988. *Anaphora in Norwegian and the Theory of Grammar*. Studies in Generative Grammar 32. Foris, Dordrecht.
- Hestvik, Arild. 1993. Strict reflexives and the subordination effect. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340*. University of Stuttgart, University of Tuebingen, and IBM Germany.
- Hinrichs, Erhard. 1986. Temporal anaphora in discourses of english. *Linguistics and Philosophy*, 9:63–82.
- Hirchbühler, Paul. 1982. VP deletion and across-the-board quantifier scope. In James Pustejovsky and Peter Sells, editors, *Proceedings of NELS 12*. GLSA, University of Massachusetts, Amherst.
- Hobbs, Jerry. 1978. Why is discourse coherent? Technical Report 176, SRI International.
- Hobbs, Jerry. 1979. Coherence and coreference. *Cognitive Science*, 3:67–90.
- Hobbs, Jerry. 1990. *Literature and Cognition*. CSLI Lecture Notes 21.
- Hobbs, Jerry R. 1993. Intention, information, and structure in discourse: A first draft. In *Burning Issues in Discourse, NATO Advanced Research Workshop*, pages 41–66, Maratea, Italy, April 13-15.
- Hobbs, Jerry R., Mark E. Stickel, Douglas E. Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- Hovy, Eduard H. 1990. Parsimonious and profligate approaches to the question of discourse structure relations. In *Proceedings of the Fifth International Workshop on Natural Language Generation*, pages 128–136, Dawson, PA, June.
- Hovy, Eduard H. 1993. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63:341–384.
- Hume, David. 1748. *An Inquiry Concerning Human Understanding*. The Liberal Arts Press, New York, 1955 edition.
- Jackendoff, Ray S. 1971. Gapping and related rules. *Linguistic Inquiry*, 2:21–35.
- Jackendoff, Ray S. 1972. *Semantic Interpretation in Generative Grammar*. MIT Press.
- Kameyama, Megumi. 1986. A property-sharing constraint in centering. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pages 200–206, New York.

- Kameyama, Megumi. 1988. Japanese zero pronominal binding: Where syntax and discourse meet. In William J. Poser, editor, *Papers from the Second International Workshop on Japanese Syntax*, pages 47–74. CLSI.
- Kameyama, Megumi. 1994. Indefeasible semantics and defeasible pragmatics. Technical Report 544, SRI International.
- Kameyama, Megumi, Rebecca Passonneau, and Massimo Poesio. 1993. Temporal centering. In *Proceedings of the 31st Conference of the Association for Computational Linguistics (ACL-93)*, pages 70–77, Columbus, Ohio, June.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam, pages 277–321.
- Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic: An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers, Dordrecht.
- Karttunen, Lauri. 1969. Pronouns and variables. In *Papers from the Fifth Regional Meeting of the Chicago Linguistics Society*.
- Kehler, Andrew. 1993a. A discourse copying algorithm for ellipsis and anaphora resolution. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics (EACL-93)*, pages 203–212, Utrecht, the Netherlands, April.
- Kehler, Andrew. 1993b. The effect of establishing coherence in ellipsis and anaphora resolution. In *Proceedings of the 31st Conference of the Association for Computational Linguistics (ACL-93)*, pages 62–69, Columbus, Ohio, June.
- Kehler, Andrew. 1994a. Common topics and coherent situations: Interpreting ellipsis in the context of discourse inference. In *Proceedings of the 32nd Conference of the Association for Computational Linguistics (ACL-94)*, Las Cruces, New Mexico, June.
- Kehler, Andrew. 1994b. A discourse processing account of gapping and causal implicature. Manuscript presented at the Annual Meeting of the Linguistic Society of America, Boston, January.
- Kehler, Andrew. 1994c. Temporal relations: Reference or discourse coherence? In *Proceedings of the 32nd Conference of the Association for Computational Linguistics (Student Session)*, Las Cruces, New Mexico, June.
- Kehler, Andrew and Gregory Ward. 1995. On the anaphoric status of *do so*. Paper presented at the Annual Meeting of the Linguistic Society of America, New Orleans, January.

- Kitagawa, Yoshihisa. 1991. Copying identity. *Natural Language and Linguistic Theory*, 9:497–536.
- Klein, Ewan. 1987. VP-Ellipsis in DR theory. In Groenindijk and Stokhof, editors, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*. Foris.
- Kuno, Susumu. 1976. Gapping: A functional analysis. *Linguistic Inquiry*, 7:300–318.
- Lakoff, George. 1976. Pronouns and reference. In J. McCawley, editor, *Syntax and Semantics, Volume 7: Notes from the Linguistic Underground*. Academic Press, New York. Originally distributed by Indiana University Linguistics Club, 1968.
- Lappin, Shalom. 1993a. Ellipsis resolution at S-structure. In *Proceedings of NELS 23*, Amherst, MA. GLSA, University of Massachusetts.
- Lappin, Shalom. 1993b. The syntactic basis of ellipsis resolution. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340*. University of Stuttgart, University of Tuebingen, and IBM Germany.
- Lappin, Shalom. 1994. The interpretation of ellipsis. In Shalom Lappin, editor, *Handbook of Contemporary Syntactic Theory*. Blackwell.
- Lappin, Shalom and Michael McCord. 1990. Anaphora resolution in slot grammar. *Computational Linguistics*, 16:197–212.
- Lascarides, Alex and Nicolas Asher. 1993. Temporal interpretation, discourse relations, and common sense entailment. *Linguistics and Philosophy*, 16(5):437–493.
- Lasnik, Howard. 1989. *Essays on Anaphora*. Kluwer.
- Levin, Nancy and Ellen Prince. 1982. Gapping and causal implicature. Presented at the Annual Meeting of the Linguistic Society of America.
- Lewis, David. 1979. Scorekeeping in a language game. In R. Bauerle, editor, *Semantics from a Different Point of View*. Berlin: Springer Verlag.
- Longacre, Robert E. 1983. *The Grammar of Discourse*. Plenum Press.
- Mann, William C. and Sandra A. Thompson. 1987. Rhetorical structure theory: A theory of text organization. Technical Report RS-87-190, Information Sciences Institute.
- McKoon, Gail, Steven B. Greene, and Roger Ratcliff. 1994. Discourse models, pronoun resolution, and the implicit causality of verbs. Draft, Northwestern University and Princeton University.

- McWhinney, B., E. Bates, and R. Kliegl. 1984. Cue validity and sentence interpretation in english, german, and italian. *Journal of Verbal Learning and Verbal Behavior*, 23:127–150.
- Moens, Marc and Mark Steedman. 1988. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28.
- Moore, Johanna D. and Cecile L. Paris. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4):651–694.
- Moore, Johanna D. and Martha E. Pollack. 1992. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics*, 18(4):537–544.
- Neijt, Anneke. 1979. *Gapping: A Contribution to Sentence Grammar*. Foris, Dordrecht.
- Neijt, Anneke. 1981. Gaps and remnants – sentence grammar aspects of gapping. *Linguistic Analysis*, 8(1):69–93.
- Nerbonne, John. 1986. Reference time and time in narration. *Linguistics and Philosophy*, 9:83–95.
- Oehrle, Richard T. 1987. Boolean properties in the analysis of gapping. In Geoffrey J. Huck and Almerindo E. Ojeda, editors, *Syntax and Semantics XX: Discontinuous Constituency*. Academic Press, New York, pages 201–240.
- Partee, Barbara. 1984. Nominal and temporal anaphora. *Linguistics and Philosophy*, 7:243–286.
- Partee, Barbara. 1989. Deictic and anaphoric pieces of meaning. LSA Conference on Logic and Linguistics, Santa Cruz. July 19-21.
- Passonneau, Rebecca. 1988. A computational model of the semantics of tense and aspect. *Computational Linguistics*, 14(2):44–60.
- Passonneau, Rebecca. 1991. Persistence of linguistic form in discourse processing. In *Proceedings of the Fall Symposium on Discourse Structure in Natural Language Understanding and Generation*, Monterey, CA, November.
- Pereira, Fernando. 1990. Categorical semantics and scoping. *Computational Linguistics*, 16(1):1–10.
- Polanyi, Livia. 1988. A formal model of the structure of discourse. *Journal of Pragmatics*, 12.
- Prince, Ellen. 1981a. Topicalization, focus-movement, and yiddish-movement: A pragmatic differentiation. In *Berkeley Linguistics Society 7*, pages 249–264, Berkeley, CA.

- Prince, Ellen. 1981b. Toward a taxonomy of given-new information. In Peter Cole, editor, *Radical Pragmatics*. Academic Press, New York, New York, pages 223–255.
- Prince, Ellen. 1986. On the syntactic marking of presupposed open propositions. In *Papers from the Parasession on Pragmatics and Grammatical Theory at the 22nd Regional Meeting of the Chicago Linguistics Society*, pages 208–222, Chicago, IL.
- Prince, Ellen. 1992. The ZPG letter: Subjects, definiteness, and information-status. In Sandra Thompson and William Mann, editors, *Discourse Description: Diverse Analyses of a Fundraising Text*. John Benjamins B.V., Philadelphia/Amsterdam, pages 295–325.
- Prüst, Hub. 1992. *On Discourse Structuring, VP Anaphora, and Gapping*. Ph.D. thesis, University of Amsterdam.
- Rambow, Owen. 1993. Rhetoric as knowledge. In *Proceedings of the Workshop on Intentionality and Structure in Discourse Relations*, pages 102–105, Ohio State University, June.
- Reichenbach, Hans. 1947. *Elements of Symbolic Logic*. Macmillan, New York.
- Reinhart, Tanya. 1983. *Anaphora and Semantic Interpretation*. University of Chicago Press, Chicago.
- Rooth, Mats. 1981. A comparison of three theories of verb phrase ellipsis. University of Massachusetts Occasional Papers in Linguistics, Volume 7.
- Rooth, Mats. 1993. Ellipsis redundancy and reduction redundancy. In S. Berman and A. Hestvik, editors, *Proceedings of the Stuttgart Workshop on Ellipsis: Arbeitspapiere des Sonderforschungsbereich 340, Bericht Nr. 29-1992, SFB 340*. University of Stuttgart, University of Tuebingen, and IBM Germany.
- Ross, John R. 1970. Gapping and the order of constituents. In Manfred Bierwisch and Karl Erich Heidolph, editors, *Progress in Linguistics*. Mouton, The Hague, pages 249–259.
- Sag, Ivan. 1976. *Deletion and Logical Form*. Ph.D. thesis, MIT.
- Sag, Ivan and Jorge Hankamer. 1984. Toward a theory of anaphoric processing. *Linguistics and Philosophy*, 7:325–345.
- Sanders, Ted J. M., Wilbert P. M. Spooren, and Leo G. M. Noordman. 1992. Toward a taxonomy of coherence relations. *Discourse Processes*, 15:1–35.
- Scha, Remko and Livia Polanyi. 1988. An augmented context free grammar for discourse. In *Proceedings of the International Conference on Computational Linguistics (COLING-88)*, pages 573–577, Budapest, August.

- Schachter, Paul. 1977. Does she or doesn't she? *Linguistic Inquiry*, 8(4):762–767.
- Schachter, Paul. 1978. English propredicates. *Linguistic Analysis*, 4(3):187–224.
- Shieber, Stuart M., Fernando C. N. Pereira, and Mary Dalrymple. forthcoming. Interactions of scope and ellipsis. *Linguistics and Philosophy*.
- Sidner, Candace. 1983. Focusing in the comprehension of definite anaphora. In M. Brady and R. Berwick, editors, *Computational Models of Discourse*. MIT Press, pages 267–330.
- Siegel, Muffy E. A. 1984. Gapping and interpretation. *Linguistic Inquiry*, 15:523–530.
- Smyth, R. 1994. Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, 23:197–229.
- Steedman, Mark. 1990. Gapping as constituent coordination. *Linguistics and Philosophy*, 13(2):207–263.
- Stevenson, R. J., A. W. R. Nelson, and K. Stenning. 1993. Strategies in pronoun comprehension. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society*, pages 976–981.
- Stillings, Justine T. 1975. The formulation of gapping in English as evidence for variable types in syntactic transformations. *Linguistic Analysis*, 1(2):247–273.
- Suri, Linda Z. and Kathleen F. McCoy. 1994. RAFT/RAPR and centering: A comparison and discussion of problems related to processing complex sentences. *Computational Linguistics (Squibs and Discussions)*, 20(2):301–317.
- Terken, J. and S. G. Nootboom. 1987. Opposite effects of accentuation and deaccentuation on verification latencies for given and new information. *Language and Cognitive Processes*, 2(3/4):145–163.
- Terken, Jacques. 1984. The distribution of pitch accents in instructions as a function of discourse structure. *Language and Speech*, 27(3):53–73.
- Vendler, Zeno. 1967. *Linguistics in Philosophy*. Cornell University Press.
- Ward, Gregory. 1990. The discourse functions of VP preposing. *Language*, 66:742–763.
- Webber, Bonnie Lynn. 1988. Tense as discourse anaphor. *Computational Linguistics*, 14(2):61–73.
- Webber, Bonnie Lynn. 1991. Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes*, 6(2):107–135.

- Wilensky, Robert. 1978. *Understanding Goal-Based Stories*. Ph.D. thesis, Yale. Research Report #140.
- Williams, Edwin. 1977. Discourse and logical form. *Linguistic Inquiry*, 8(1).
- Wilson, Deirdre and Dan Sperber. 1979. Ordered entailments: An alternative to presuppositional theories. In C. K. Oh and D. Dinneen, editors, *Syntax and Semantics XI: Presupposition*. Academic Press, pages 299–323.
- Winograd, Terry. 1972. *Understanding Natural Language*. Academic press, New York.

Sources of Examples

- Chomsky, Noam. 1982. *Noam Chomsky on the Generative Enterprise*. Foris, Dordrecht.
- Jarecki, Jill. 1992. The role of synaptic activity during embryonic neuromuscular junction development in *drosophila melanogaster*. Unpublished Yale University Prospectus, August.
- Lascarides, Alex and Nicolas Asher. 1993. Temporal interpretation, discourse relations, and common sense entailment. *Linguistics and Philosophy*, 16(5):437–493.
- Lochbaum, K. E. 1994. A model of plans to support inter-agent communication. In *Proceedings of the AAAI Workshop on Planning for Inter-agent Communication*, pages 26–32, Seattle, Washington.
- Magerman, David. 1994. *Natural Language Parsing as Statistical Pattern Recognition*. Ph.D. thesis, Stanford University. Stanford University CS Technical Report CS-TR-94-1502.
- Neijt, Anneke. 1981. Gaps and remnants – sentence grammar aspects of gapping. *Linguistic Analysis*, 8(1):69–93.
- Rosenthal, Harry F. 1988. Agency orders UPS to resume fireworks deliveries—after July 4th. Associated Press Newswire, June.