My results are discussed in terms of two traditions of psycholinguistic research. One is the language-as-action tradition, in which discourse is seen as a joint activity, driven by the public and private goals of the discourse participants (H. Clark, e.g., 1996). The other is the constraint-satisfaction approach to language processing, in which the mechanisms of language processing are sensitive to an individual's past experience with regular patterns in language and the world (e.g., MacDonald et al., 1994; Trueswell and Tanenhaus, 1994).

Acknowledgments

This document represents the culmination of my formal education in linguistics, and in some way I would like to thank everyone who has taught me, inspired me, challenged me, or supported me throughout this process.

First and foremost, a big thank you to my adviser, Tom Wasow, for his endless patience and advice on doing research, writing a dissertation and pursuing a career in academia. Many thanks also to the other members of my committee: to Eve Clark, especially for her detailed comments on writing; to Herb Clark, for constantly challenging my ideas; and to Maryellen MacDonald, both for her advice on this project, and for allowing me to visit her lab during the fall 1995 semester, during which time I gained crucial experience with psycholinguistic theory and methodology. Thanks to Peter Sells for joining my orals committee, and also for his teaching and advice throughout my time at Stanford. Thanks also to Ellen Markman for serving as the university chair for the oral defense of my dissertation.

I have also benefited enormously from discussions with other scholars. I gratefully acknowledge Judith Aissen, especially for her advice on my analysis of Mapudungun, Amit Almor, David Beaver, especially for his advice on DRT, Henriëtte de Swart, Adele Goldberg, Masayo Iida, Ivan Sag, Elizabeth Traugott, Tom Veatch, and Marilyn Walker, whose advice led me to do the text analyses I present in chapter 2. Thank you to Penny Eckert and John Rickford, who got me thinking about the on-line, probabilistic interaction of multiple factors from the moment I got to Stanford (from a sociolinguistic point of view). Thanks also to my undergraduate professors for getting me started on linguistics in the first place: Donna Jo Napoli, Judy Kegl, Virginia Brennan, and Nilofar Haeri.

Thank you also to my classmates and other students who have helped me with invaluable advice and discussions: to Emily Bender, for both inspiring theoretical discussions and for her detailed comments on some chapters, to Mike Harm, for his constant willingness to answer my questions about connectionism, and to Mariela Gil, Ryan Ginstrom, Rob Malouf, Robert Thornton, Susanne Riehemann, Sarah Schuster, Scott Schwenter, and the people at SLUGS for discussions related to this dissertation.

The research for this project was funded by a Graduate Research Opportunity grant from Stanford University.

The work presented in this dissertation has also benefited from the technical assistance of several people and groups. Many thanks to Barry McLaughlin for allowing me to work in his laboratory at the University of California at Santa Cruz, and for loaning me equipment for experiments at Stanford. Thank you also to Carie Lemack for helping me to run the experiment in chapter 4. I am also very grateful to the students who run the Stanford statistics consulting service.

I am extremely thankful that as I end my career at Stanford I can say that I have enjoyed my time here. This is because of all the wonderful people I've been surrounded by. Many, many thanks to the staff, who not only are superb administrators but also have been good friends throughout my time here: Gina Wein, Michelle Murray, Tasha Newson, Kyle Wohlmut, Trudy Vizmanos, and Emma Pease. I can't imagine Stanford without you. Special thanks to Kyle for the coffee breaks and musical distractions. My sanity during the dissertation stage is largely due to the "dissertation running group": Scott Schwenter, Yukiko Morimoto, and recently Natalie Schilling-Estes and David Beaver. Thanks also to all my friends and classmates who traveled down the graduate school road with me.

My biggest thanks go to my family. My parents, Arthur Arnold and Caroline Arnold, have encouraged me, for as long as I can remember, to go through life asking questions. They have been there throughout every challenge in my life, not the least of which was graduate school. Although this dissertation is a milestone, it would mean nothing were it not against the backdrop of enormous support and love of my parents, my brother Matthew Arnold, my grandparents Catherine and Lester Scheaffer and Leona and Wiley Arnold, and every other person in my family.

Last but not least, I am grateful that my life as a graduate student has been shared with my husband Humberto Gutiérrez-Rivas. Not only is he my best friend and closest companion, but he has also helped me with discussions about Spanish, statistical analyses, and every other aspect of life. This dissertation is dedicated to him.

Table of Contents

List of Tables

List of Figures

Chapter 5: Conclusions

Chapter 1

Introduction

One of the central aspects of language use is the process of referring -- speakers[1] refer to people, places, things, and ideas, and then say things about them, and in order for communication to be successful, the listener must know what the speaker is referring to. At the same time, however, human beings have limited time, energy, and patience, and generally desire to make their communicative interchanges as economical as possible. One way to make communication more efficient is to use a shorter, less specific form of reference, such as a pronoun, when the referent is accessible to the listener. When there might be some doubt about the referent, however, the speaker needs to use a longer, more specific form of reference, such as a name or description. This tension between communication and efficiency is reflected in and embodied by Grice's maxim of quantity: "Make your contribution as informative as is required...[but not] more informative than is required" (Grice, 1975).

This dissertation addresses the following questions: What accounts for speaker choices in reference form, and how are comprehenders able to interpret referring forms as intended? These questions have concerned cognitive scientists of all types, including linguists, psychologists, philosophers, and computer scientists. This topic has attracted so

---

[1] I will use the terms "speaker" and "listener" throughout the dissertation, although in fact I usually mean for the term "speaker" to extend to writers and other producers of language, and for the term "listener" to extend to readers and other comprehenders of language. I will also use the term "discourse participants", which comprises the speaker/writer and all addressees.

much attention because it is a common yet complex aspect of human behavior. The reason that language exists at all is to communicate, which requires speakers to indicate who or what they are talking about. The existence of nominals and other referring expressions reflects the important role that reference plays in communication. Many other components of language are also used primarily to communicate information about participants and their roles: roughly speaking, verbs specify the actions, and their predicate argument structures specify the roles of the participants.

Another reason for studying reference form is that it provides a window into understanding other cognitive processes. Most researchers explain the appropriateness of particular forms of reference by appealing to the cognitive status of the referent in the mental representations of the discourse participants. For the most part, a pronoun is used when the speaker believes that the referent is already available in the listener's consciousness (i.e., when the entity is "given"). But there may be many things in the listener's consciousness, and the speaker needs to use a form that will let the listener choose the correct one. Research has suggested that at any one moment, discourse referents differ in terms of their salience in the mental representations of discourse participants, and that less-specified forms are only used when the referent is sufficiently salient.

Yet even though the concept of salience has repeatedly been invoked in the literature, it lacks a substantive definition. Things can be salient for many reasons. For example, people may tend to notice perceptually salient objects (such as a bright light), socially relevant actions (such as their boss entering the room), and the referent of linguistic pointers (such as the exclamation *Look at that!*).

In this dissertation, I will focus on linguistic information, and investigate how it can influence the salience of a referent and choices in reference form. Past accounts of reference form have generally focused on the roles of one or more factors in establishing certain referents as salient. In contrast with past work, the purpose of this dissertation is not to investigate any single factor in great detail. Instead, I will attempt to pull together many

different factors into one general framework. After reviewing past work, I will focus on five factors: Recency, Subjecthood, Focus, Parallelism, and Goal Status. The labels for these factors are for convenience only, and a full description of each term will be provided later. At the end of this chapter, I will propose a preliminary framework for understanding why all these factors influence reference form, and in subsequent chapters I will provide evidence for my proposal.

Throughout the dissertation, I will use the term "discourse" to mean any set of utterances which are meant to cohere in some way, although I use the term loosely -- a discourse may be a story, a text, a monologue or dialogue; it may consist of two clauses or it may be an hour-long discussion. I will use "anaphor" to mean a linguistic form that refers back to an entity that has been previously evoked in the discourse. The term "antecedent" refers to the linguistic form with which that entity was last mentioned, and "referent" is the entity that the anaphor refers to. For example, "He" in 1 is an anaphor.

(1) Terrence sat down. He ate lunch slowly.

The antecedent of "He" is the subject of the preceding clause, "Terrence". The referent of "He" is the person denoted by "Terrence". For simplicity, I will identify both antecedents and referents in terms of the last noun phrase used to identify them: in this case, "Terrence".

The purpose of my studies is to understand how the representations of discourse referents are influenced by properties of the linguistic expressions used to refer to them. Therefore, I will also talk about referents in terms of features of the last linguistic expression referring to them. In 1, "Terrence" is the subject and agent of the clause before "He". To express these properties, I will refer to the entity as a "subject-referent" and an "agent-referent". The anaphor and antecedent are in parallel positions, so "Terrence" is also a parallel-referent. Because the referent was mentioned recently, he is also a recent-referent.

1.1. Past work on reference

The question of how reference works linguistically and psycholinguistically is central to understanding how language works, and has been the subject of many studies in diverse fields. Researchers have suggested numerous factors and theories to account for variation in reference form. These accounts differ, but one thing they share is the concept of salience, or accessibility. Loosely speaking, all researchers have observed that pronouns are used most often when the referent is represented in a prominent way in the minds of the discourse participants, but more fully specified forms are needed when the representation of the referent is less prominent. Some of the terms that have been used to describe this phenomenon are "salience", "accessiblity", "topicality", "focus of attention",[1] or "activation". The terms are not strictly synonymous, but they are similar in that they attribute choices in reference form to the cognitive status of the referent.

But if the explanation of reference form depends on salience, then we need to know what exactly salience is. The focus of most past research has been identifying particular factors, linguistic and otherwise, that influence the salience of referents.

In the following section I will review some of the main findings about reference form and the representation of discourse referents. Multiple independent and overlapping theories have been suggested, which I will group into four categories for the purpose of presentation: 1) Linguistic theories of Salience, Accessibility, or Topicality, 2) Psycholinguistic studies of pronoun resolution and reference form, 3) the formalized model of Centering Theory, and 4) a criticism of Centering, the Information Load Hypothesis.

1.1.1. Linguistic theories of Salience/Accessibility/Topicality

Linguists have investigated two related questions: 1) in a given instance of referring, what form will be used? and 2) given a pronoun, what is its referent? Although these

---

[1] "Focus of attention" is not the same as the term "focus" which is used in descriptions of information structure, and roughly refers to the new, highlighted information in a sentence. Instead, "focus of attention" refers to the item which is currently most salient in the discourse in terms of pronoun reference, which corresponds more to the discourse topic than the focus.

questions involve somewhat different processes, researchers have discussed both in terms of the accessibility of the referent. Some authors talk about accessibility in terms of the text, some authors talk about it in terms of the cognitive status of the referents.

1.1.1.1. Topic

For the most part, people organize their speech and writing around a particular topic, they stick to that topic for awhile, and then move on to a new topic. This observation has led to the use of the category "topic" to explain various linguistic phenomena, such as intonation (Halliday, 1967), word order (Sgall, et al., 1986), and reference form (e.g., Ariel, 1990; Broadbent, 1973; Gernsbacher, 1990; Purkiss, 1978; Sanford and Garrod, 1981). However, discussing the relationship between topic and reference form is difficult to do with any precision, because of two fundamental problems with the notion of topic: 1) defining the meaning of "topic" has proven very difficult, and 2) the relationship between reference and the category "topic" (even if it can be identified) is tenuous at best. In this section, I will attempt to describe how this concept has been used, and what relationship it may have with reference form. I will conclude that the traditional concept of the category "topic" as a single entity per utterance is too problematic to make it useful. However, the idea that some entities are more topical than others has its advantages, which might be retained if topicality is considered as a gradient property, and is not limited to a single entity per utterance (Givón, 1983a).

The first problem with "topic" is that defining what it means is notoriously messy. It is often thought of as what a sentence or discourse is **about** (Reinhart, 1981), but it is also frequently associated with old information (Gundel 1974; Chafe 1976; Clark and Haviland 1977; Clark and Clark 1977). Researchers often distinguish between "sentence topic" and "discourse topic", depending on the level of discourse over which the "aboutness" is measured. In some languages (e.g., Japanese (Walker, Iida, and Cote, 1994), or Jacaltec Mayan (Aissen, 1992)), the sentence topic is in fact a grammatical position, which is partially defined on pragmatic grounds. In languages with no grammaticized topic

position, such as English, certain positions such as subject or initial position are considered highly topical (Reinhart, 1982). Although some researchers have proposed that no sentence can have more than one topic (Reinhart, 1982), more plausible proposals have viewed topicality as a graded property that is present in all discourse entities (e.g., Givón, 1983a). The "discourse topic" is what a given discourse segment is about. However, this relies on yet another hard-to-identify construct: the discourse segment. Discourse and sentence topics are related to each other, in that the topic of a sentence contributes to the interpretation of the identity of the discourse topic. However, the relation between the two is not clear, which makes it difficult to identify discourse and sentence topics consistently and categorically.

Among these difficulties, particularly troublesome is the traditional concept of "topic" as a single element in a given utterance (Reinhart, 1982).[1] In many cases, it is difficult to identify a unique topic. For example, consider the sentence in 2.

(2) In your casserole dish, place layers of hominy, meat-sauce mixture, and grated cheese.

(Rosenberg, The Impoverished Students' Book of Cookery, Drinkery, and Housekeepery:12).

The topic of this sentence could reasonably be considered to be the reader, the casserole dish, the act of placing, or the list of ingredients. At a discourse level, the text is about how to make a casserole, in which case the act of placing might be considered the most topical element. On the other hand, the implicit subject of all the instructions in the recipe is the reader. Although this problem might be solved with a variety of theoretical assumptions, the important point is that identifying a unique topic at either the sentence or discourse level can be problematic.

The problem of identifying the topic has two consequences for incorporating it into a theory of reference processing. In the first place, it poses problems for the researcher seeking to identify the construction of interest across diverse situations, which hinders the process of developing a solid theory about topics. In the second place, it is difficult to imagine how the listener might identify the topic during discourse comprehension, and if the listener cannot identify the topic, then it is doubtful that it could influence comprehension directly.

Despite the difficulties in developing a comprehensive definition of sentence or discourse topic in English, the term has been widely used in the psycholinguistic literature. Some examples of the use of "topic" are presented in 3.

---

[1] Reinhart assumes that a sentence can have 0 topics, but may not have more than 1 (1982).

(3) Uses of "topic" in the literature.

a. "...distant antecedents may nonetheless be very salient, as when they constitute the discourse topic," (Ariel, 1996:31).

b. "In natural discourse, however, topics (mainly discourse topics) constitute the most salient entities more often than not," (Ariel, 1990:23).

c. "Indeed, words and sentences that change the ongoing topic, point of view, or setting take substantially longer to comprehend than words or sentences that continue the topic, point of view, or setting," (Gernsbacher, 1990:26).

d. "Our convention will be to add TOPIC = X to the message, where X is the topic argument," (Levelt, 1989:98).

e. "Thus, since topical entities are also likely to be those referred to by pronouns, one factor contributing to the higher pronominalization propensity of animates might thus simply be that animates are more often topics than other referents," (Dahl and Fraurud, 1996:60).

f. "This finding suggests that the referent/nonreferent differences found by Gernsbacher (1989; ...) arose because the nonreferent of the pronoun was unlikely to be the topic of a following sentence, and its activation was, therefore, suppressed," (Garnham et al., 1996:531).

Most researchers who have discussed the concept of "topic" seriously have also acknowledged the difficulties in defining it. The widespread familiarity of the term and its

frequent use in explanations of linguistic phenomena reflects the intuition that some things are more important to the discourse, and that those things will be more accessible psycholinguistically.

One of the properties that has been attributed to topics is that subsequent reference to the sentence or discourse topic is more likely to be pronominal, so the topic structure may influence the interpretation of ambiguous pronouns (e.g., Ariel, 1990; Dahl and Fraurud, 1996). However, the use of "topic" to account for reference form makes the problem of defining this category even more difficult. As all researchers have acknowledged, the relationship between topic and reference is not categorical. For example, Reinhart (1982) discussed the following examples from Oehrle (1981). In these sentences, intuition suggests that if the pronoun is unstressed, it refers to the antecedent in parallel position, and if it is stressed, it refers to the nonparallel antecedent. Out of context, these interpretations are highly restricted, and appear to be "independent of the topic relations in the sentences," (Reinhart, 1982:16). The indices in 4 and 5 indicate the interpretations for unstressed pronouns.

(4) Felix$_i$ hit Max$_j$ and then he$_i$ hit Bill.

(5) Felix$_i$ hit Max$_j$ and then Bill hit him$_j$.

On the other hand, Reinhart presented Oehrle's (1981) argument that if the topic of the conjunction is specified, as in 6-8, the pronoun antecedent is identified as the topic (assuming no contrastive stress on the pronoun).

(6) As for Max, $Felix_i$ hit Max_j and then he_j hit Bill.

(7) a. Can you give me an exact description of Max's role in the fight?

b. $Felix_i$ hit Max_j and then he_j hit Bill.

(8) As for Felix, $Felix_i$ hit Max_j and then he_i hit Bill.

However, this demonstration of the role of topic in pronoun interpretation is problematic. In these examples, the topic is established either through the expression "As for X", or through a question about one character. Although these methods seem to intuitively establish one character as more central than the other, the resulting discourses in 6-8 are extremely unnatural, in particular because the second mention of that character is not pronominalized. In fact, referring to "Max" in object position in 7b seems to treat this referent as less salient. Although the question asks about "Max", the speaker then starts the sentence with "Felix". This makes both "Felix" and "Max" relatively equal in prominence, which makes the use of the pronoun "he" in 7b fairly confusing.

One of the reasons that the pronoun is confusing in 7 is because of the role of grammatical subject. In several studies, the subject has been associated with topic status. In Reinhart's influential paper on topics, she pointed out that while topics and subjects are not identical, "There is a strong preference...to place the topic in subject position," (1982:8). In addition, Ariel (1990:23-24) described a study by Broadbent (1973), which also appeared to identify subjects with topics. Although Reinhart specifically focused her discussion on sentence topics, their influence on pronoun interpretation can be seen as a discourse-level effect, as with Broadbent's stimuli, such as 9.

(9) The feedpipe lubricates the chain, and *it* should be adjusted to leave a gap half an inch between itself and the sprocket.

In Broadbent's study, "speakers understood *it* as referring to the discourse topic ('feedpipe') rather than the non-discourse topic ('chain')" (Ariel, 1990:23). Although Ariel did not define how she or Broadbent meant the term "discourse topic", this example suggests that it was defined as the grammatical subject of the first clause. A similar assumption also appeared in Ariel's discussion of Purkiss (1978). As Ariel also notes, other researchers have linked topichood to features such as repeated mention (e.g., Kameyama, 1996; Levy, 1982), or globally prominent characteristics such as being the protagonist of a narrative (e.g., Francik, 1985).

1.1.1.2. The status of topic

The notion of topic does not appear to be the best mechanism for discussing reference form. In the first place, there are few cues for identifying the topic. It is often associated with the grammatical subject, but not always. Other characteristics associated with topichood are repeated mention or identity as the protagonist, but these aren't necessary or sufficient to identify the topic either. And even if the topic can be identified, it is not a necessary condition for pronoun use.

However, despite the problems associated with the concept of topic, it is important not to throw the baby out with the bathwater. It is still true that people tend to produce discourses coherently, which usually means discussing the same referent for periods of time. This means that discourse participants will focus on some referents and actions as more important to the discourse than others. These referents and actions exhibit some of the same characteristics as what researchers have called the "topic", but one crucial difference is that there is not always a unique topic for a particular sentence or discourse segment.

1.1.1.3. Topicality as a continuum: Givón

One solution to the one-topic-per-utterance problem is to treat topicality as a continuum, as did Givón (1983a). Givón suggested that all entities are topical to greater or lesser degrees, and demonstrated that three measures of topicality correlate with the form and type of reference used. He identified the following scale of topicality, with zero anaphors as the most topical entities, and referential indefinite NPs as the least topical.[1]

(10) most continous/accessible topic

zero anaphor
unstressed/bound pronouns or grammatical agreement
stressed/independent pronouns
R-dislocated DEF-NPs
neutral-ordered DEF-NPs
L-dislocated DEF-NPs
Y-moved NPs ('contrastive topicalization')
cleft/focus constructions
referential indefinite NPs

most discontinuous/inaccessible topic (Givón, 1983:17)

In a collection of cross-linguistic text analyses, Givón and his colleagues (1983b) correlated the forms of reference in 10 with three measures of topicality: a) referential distance (how recently the entity has been mentioned), b) potential interference (how many other potential antecedents of the referring form there are), and c) persistence (how long the entity will remain in the discourse). By using both referential distance and persistence, Givón included two features in his conception of topicality: 1) the idea that topicality that reflects the status of the referent according to the discourse thus far, and 2) that notion that the way in which a speaker refers to an entity reflects the speaker's intentions about the role of that entity in the remainder of the discourse. This idea is also a component of Centering Theory, and will form part of my proposal below.

Givón's measure of "potential interference" addresses the issue of the ambiguity of the referring form, relative to the discourse situation. One problem with this measurement,

however, is defining exactly what constitutes potential interference. Givón's description is somewhat vague: "An interfering topic was counted only if it was just as **semantically compatible** (most commonly in terms of animacy, humanity, agentivity or semantic plausibility as object or subject) with the predicate of the clause of the topic under consideration," (1983:14, emphasis in original). This characterization, albeit loose, reflects the importance of constraining information from the predicate as a whole, and not just the anaphor itself. For example, if one hears "The bird saw the cat before it flew away", the interpretation of "it" is assisted by the compatibility of "the bird" and "flew". By contrast, in the sentence "The bird saw the cat before it pounced," the compatibility between "pounce" and "cat" would force a different interpretation of "it".

Givón and his colleagues argued that "topicality" can be conceived of as a continuum which can be statistically correlated with reference form. Even though Givón's measures of topicality (referential distance, potential interference, and persistence) are too rough to accurately reflect the processes of language comprehension and production (see chapter 2), the combination of these and other factors may provide a tenable characterization of topicality as a complex, multifaceted characteristic that can be applied in a gradient manner to all discourse entities. In this way, it appears that we may be able to salvage the concept of "topic".

However, by including "interference" as one of the measures of topicality, it seems that Givón has stepped beyond what was originally intended by the term "topic". The topicality of an entity, even if it is a continuous notion, seems to be a characteristic that is present in the role an entity plays in a discourse. In contrast, interference from other discourse entities is only relevant insofar as it may hinder the interpretation of referring forms. Thus, Givón's conception of "topicality" has more in common with other scales of salience or accessibility than with a traditional conception of "topic".

---

[1] Although Givón rejected this scale of topic accessibility in favor of more specific scales (such as scales of word-order, morphology, intonation, or phonological size), he claimed that the composite scale is

Givón's topicality scale differs from other accounts of reference form in that his measures of topicality concern the referring expression itself, as opposed to the cognitive status of the referent. Other scholars, such as Ariel, have been concerned with the topicality of the **referent**, while Givón's measures are meant to indicate the topicality of the referring expression. Nevertheless, Givón assumes that the text properties are associated with the cognitive status of entities, such that "What is continuing is more predictable", and "What is predictable is easier to process" (1983a:12). To a first approximation, then, the three measures of topicality used by Givón and his colleagues can be interpreted as indices of the cognitive status of the conceptual referent, and in that sense are comparable with other approaches that identify degrees of topicality with the referent and not with the referring expression.

1.1.1.4. The Accessibility Hierarchy: Ariel

Several researchers have suggested that the cognitive status of referents can be characterized in terms of a graded scale. One of the most comprehensive proposals is Ariel's Accessibility Hierarchy (e.g., 1988; 1990; 1994). Importantly, Ariel (1988, 1990) suggested that the accessibility of a referent entity is determined by multiple factors. She proposed that the four most important are those listed in 11.

---

roughly accurate. The single scale is presented here for the sake of simplicity.

(11) Factors affecting the Accessibility status of an antecedent (reproduced from Ariel, 1990:28, example 24).

a) Distance: The distance between the antecedent and the anaphor (relevant to subsequent mentions only)

b) Competition: The number of competitors on the role of antecedent.

c) Saliency: The antecedent being a salient referent, mainly whether it is a topic or a non-topic.

d) Unity: The antecedent being within vs. without the same frame/ world/point of view/segment or paragraph as the anaphor.

The first two factors are similar to Givón's measurements of "referential distance" and "potential interference", although Ariel's description of "Competition" differs from Givón's "interference" in that it does not include a discussion of semantic compatibility.

Ariel's third factor, "Saliency", addresses the difference between topical and nontopical antecedents. Her discussion reflects an assumption that "topic" is defined in terms of the grammatical subject. She also mentioned Levy's (1982) claim that topicality is influenced by the number of anaphoric references to an entity, in particular pronominal references. Although the Accessibility hierarchy is inherently a graded scale, it appears that "topic" is treated as an all-or-nothing phenomenon, implying that an entity either is the topic or not, and that a given discourse segment has one and only one topic.

The fourth area that Ariel lists is that of "Unity". This factor reflects the effect that discourse structure can have on "Working Memory", and thus reference form. Ariel suggests that choices in reference form are influenced by the discourse structure, which can be influenced by things like the passage of time within the discourse or paragraph breaks in written text. She links this factor to Fox's (1987) claim that in English, 'by using a pronoun the speaker displays an understanding that the preceding sequence has not been closed

down' (1987:18). Thus, pronouns are more natural for references to things from the same discourse segment, and fuller forms are used when the referent was last mentioned in a different segment.

Through text analysis in Hebrew and English, Ariel (1988, 1990) demonstrated different distribution patterns for different forms of reference, which she termed "accessibility markers". Her full Accessibility Marking Scale is reproduced in 12, with examples of English accessibility markers.

(12) Ariel's Accessibility Marking Scale (1990:73), with examples (1988:84)[1]

Marking Scale	**Examples**
Full name + modifier	Joan Smith, the president
Full ('namy') name	Joan Smith
Long definite description	The tall and authoritative president
Short definite description	The president
Last name	Smith
First name	Joan
Distal demonstrative + modifier	that hat we bought last year
Proximal demonstrative + modifier	this hat we bought last year
Distal demonstrative + NP	that hat
Proximate demonstrative +NP	this hat
Distal demonstrative	that
Proximate demonstrative	this
Stressed pronoun + gesture	SHE (plus gesture)
Stressed pronoun	SHE
Unstressed pronoun	she
Cliticized pronoun	(no examples in English)
Extremely High Accessibility Markers	gaps, including pro, PRO and *wh* - traces, reflexives, and Agreement

Other scholars have proposed similar scales of accessibility, such as Chafe (1994) (given > accessible > new), or Gundel et al. (1993) (in focus > activated > familiar > uniquely

[1] In a few cases the Marking Scale from Ariel, 1990 did not correspond exactly to the examples from Ariel, 1988. I filled in the missing examples or categories, following her schema.

identifiable > referential > type identifiable). These scales are similar to Ariel's Accessibility Hierarchy, but are not as comprehensive.

1.1.1.5. Prince

Another functional approach to discourse structure and referential form was provided by Prince (1992). In this paper, Prince categorized entities in terms of givenness (whether an entity is "new" or "old"), but she introduced two key concepts which are not found in Ariel (1988, 1990, 1994) or Gundel et al. (1993). First, she divided the spectrum of "givenness" into three categories: new, inferrable, and old. With the category "inferrable",[1] Prince recognized that by referring to The Bastille, for example, certain aspects of the referent, such as the door, become more accessible than completely new items would be. Second, Prince specified that the status of a discourse entity can be assessed both with respect to the discourse and with respect to the hearer's knowledge. For example, if both hearer and speaker know someone named "Samuel Davidson", the speaker may refer to this person for the first time as "Sam", without further introduction. Thus, Prince identified 6 levels of givenness, which are only partially independent of each other:

(13) Prince (1992)

HEARER-NEW	DISCOURSE-NEW
HEARER-INFERRABLE	DISCOURSE-INFERRABLE
HEARER-OLD	DISCOURSE-OLD

By distinguishing between Hearer status and Discourse status, Prince recognized the concept of "shared knowledge" discussed by Clark and Haviland (1977). In a text analysis of a letter (a relatively formal text), Prince found a tendency for grammatical subjects to be definite, and to refer to entities that were both Hearer-old and Discourse-old. However, upon analyzing her data with VARBRUL, Prince found that the correlation between Hearer-old status and subject NPs was completely dependent on the tendency for subjects to be

Discourse-old, and that Hearer information status did not contribute significantly to the outcome. It is certainly true that both Inferrability and shared knowledge between the speaker and listener are crucial for describing how reference may be made when something has not been mentioned before in a particular discourse. However, these factors do not play large roles in referring to "discourse-old" entities, which are the focus of this dissertation.

1.1.2. Psycholinguistic Research

The functional approaches to the study of reference form, discussed above, are useful in that they identify some of the basic factors affecting reference form and provide a comprehensive account of a range of forms, from definite NPs to null anaphors and pronouns. What these approaches lack, however, is a detailed proposal of the specific factors involved in the processing of referring forms. But a number of psycholinguistic studies have demonstrated how pronoun interpretation is influenced by factors such as Implicit Causality, Subjecthood or First Mention, Parallelism, Recency, and Ambiguity. This research complements the linguistic research discussed above. Both traditions address many of the same factors, but linguists tend to concentrate on the identification of different factors, and their relation to multiple types of reference form, while psycholinguists tend to focus on how these factors are used during language production and comprehension.

1.1.2.1. Implicit Causality and Thematic Roles

Many researchers have claimed that some verbs are biased towards either the subject or the object (e.g., Au, 1986; Brown and Fish, 1983), and that this bias influences how people interpret subsequent pronouns (e.g., Caramazza et al., 1977; Caramazza and Gupta, 1979; Ehrlich, 1980; Garvey and Caramazza, 1974; McDonald and MacWhinney, 1995). The verb bias for one type of antecedent or the other is termed the "Implicit Causality" of the verb, and represents the tendency for people to attribute the cause of the event to the subject in some verbs, and the object in others. Implicit Causality is claimed to be one

---

[1] I am using this nonstandard spelling of "inferrable" following Prince (1992).

factor that can reverse the general tendency for the subject or "first mention" to be the more likely antecedent (McDonald and MacWhinney, 1995:545).

Several researchers have shown that with verbs such as "blame", which highlights the object NP as the implicit cause of the event, the processing of a subsequent pronoun is facilitated when the pronoun refers to the object, and is inhibited when the pronoun refers to the subject. The opposite effect obtains with verbs such as "please", for which the referent of the subject NP is implicitly assumed to be the cause, and therefore is considered more prominent (Caramazza et al., 1977; Caramazza and Gupta, 1979; Garvey and Caramazza, 1974). Similarly, the interpretation of ambiguous pronouns has been shown to be affected by the implicit causality of the verb (Ehrlich, 1980; Garvey and Caramazza, 1974; Grober et al., 1978; Stevenson et al., 1994). For example, in 14a the preferred interpretation of "he" is "George", but in 14b the preferred interpretation of "he" is "Walter".

(14) a. George telephoned Walter because he wanted some information.

b. George criticized Walter because he misplaced the file.

(Grober et al., 1978)

One problem with studying implicit verb causality is determining the direction of the verb's bias. Early studies (e.g., Ehrlich, 1980) defined the implicit causality as the most natural continuation to fragments of "because sentences", such as "Sue blamed Al for spilling the coffee, because she/he ...". Ehrlich (1980) asked 10 judges to complete sentences such as these, and used their completions as the baseline measure of implicit causality. This measure contains a basic circularity -- it uses pronoun assignment to determine what will be the most likely pronoun assignment.

The problem of defining Implicit Causality has been partially eliminated in more recent studies, in which researchers have defined the verb bias in terms of argument structure. For example, McDonald and MacWhinney (1995) compared "Experiencer-

Stimulus" verbs, as in 15a, with "Stimulus-Experiencer" verbs, as in 15b, where the Stimulus is interpreted as the implicit cause of the event.

(15) From McDonald and MacWhinney (1995), Experiment 1

a. Gary admired Alan time after time because he was so talented.
Experiencer **Stimulus**

b. Gary amazed Alan time after time because he was so talented.
Stimulus **Experiencer**

This independent characterization of verb bias is beneficial, particularly because it takes a step in the direction of explaining why these verbs should have this effect. Using a cross-modal priming task, McDonald and MacWhinney played sentences like 15a or b to participants, and measured the time it took them to respond to probe words like "Gary" or "Alan" at different points during the presentation. Response time was taken as an indication of the accessibility of each referent at a given point. Their biggest effect was that the first-mentioned referent was more accessible at most time points, reflecting a first-mention bias. The bias toward the stimulus-referent also had an effect, erasing the first-mention effect for stimuli like 15a, but only at the pronoun and at the end of the sentence.

Thematic roles were also used to characterize a verb's bias in Stevenson et al.'s (1994) study. They investigated the four types of sentences shown in 16.

(16) From Stevenson et al. (1994), Experiment 1

	Verb type	**Example**	**Bias**
a.	Goal/Source:	John seized the comic from Bill. John passed the comic to Bill.	goal
b.	Agent/Patient	Joseph hit Patrick. Patrick was hit by Joseph.	patient
c.	Experiencer/Stimulus:	Ken admired Geoff. Ken impressed Geoff.	stimulus
d.	Agent/goal: Agent/source:	Simon ran towards Richard. Simon ran away from Richard.	agent

Using a written sentence-completion task, they presented participants with stimuli like 17a and b, and asked them to write a continuation sentence.

(17) a. Pronoun condition: Ken admired Geoff. He .....
b. No-pronoun condition: Ken admired Geoff. .....

(Stevenson et al., 1994:527)

They analyzed sentence continuations to identify the pronoun antecedent in the pronoun condition, and to see which referent participants referred to in the no-pronoun condition. Their results showed that thematic roles did indeed influence the types of continuations that participants provided, such that participants referred more to the biased thematic role for each verb type. Their results were consistent with those of McDonald and MacWhinney, in terms of both the stimulus bias and their finding of a general first-mention bias.

In sum, thematic roles are useful for discussing the bias of individual verbs. At the same time, thematic roles are a relatively impoverished representation of participant roles. Although they provide a handle for studying verb semantics, they are mere representatives of the verb semantics, which are much more complex. This issue will be discussed in more detail in chapter 4, where I will also propose a new way of considering the effects of implicit causality and thematic roles.

1.1.2.2. Subject/First-Mention Bias

One of the most common generalizations about reference is the tendency to interpret ambiguous pronouns as co-referential with the subject, or first-mentioned entity. Stevenson et al.'s (1994) study, described above, also revealed a general tendency to interpret the pronoun as coreferential with the subject of the first clause, in the pronoun condition. In the no-pronoun condition, people tended to use pronouns to refer to the subject-referent, and to use names to refer to the object-referent.

In McDonald and MacWhinney's (1995) probe task, they found that subjects responded to a subject probe about 80 ms faster than they responded to an object probe. However, they also found that the subject advantage disappeared when the pronoun was interpreted as co-referential with the object NP. They suggested that the subject advantage is a "general accessibility advantage, and occurs whether or not anaphor resolution processes are involved.... It is pervasive and is not localized to a particular point in sentence processing," (p. 560). In contrast, they claimed that factors such as Implicit Causality and gender information are specifically involved in pronoun disambiguation, and apply only at the point of the anaphor. A similar argument was made by Garnham et al. (1996), who also found a general First-Mention advantage during a probe study of Implicit Causality.

The Subject/First Mention bias constitutes one of the core generalizations about reference form and pronoun resolution. It is the most frequently mentioned characteristic in psycholinguistic studies of reference and pronoun resolution; it reflects generalizations

about topicality which have been made in the linguistic domain; and it also plays a fundamental role in Centering Theory, discussed below.

1.1.2.3. Parallelism

Another factor, often confused with the subject bias, is the parallel function strategy. First suggested by Sheldon (1974), the parallel function strategy refers to the increased accessibility of antecedents in the same grammatical position in the preceding clause (e.g., subject vs. object). For example, in 18b, "Celia" should be more accessible as the antecedent for the subject pronoun, and "Sharon" should be more accessible as the antecedent for the object pronoun.

(18) a. Celia hugged Sharon at the train station.

b. She asked her how the trip was.

Evidence for this strategy was also found by other researchers (e.g., Corbett and Chang, 1983:289; Grober et al., 1978; Springston, 1975). For example, Springston demonstrated that pronouns with non-parallel antecedents were harder to read than those that did have parallel antecedents.

The problem with these studies is that the "parallel function strategy" was always demonstrated with pronouns in subject position, by showing that subject antecedents are easier to resolve than object antecedents. It is unclear, therefore, whether the preference for subject-subject parallelism (subject anaphor and subject antecedent) is due to parallelism, or to a subject bias. The crucial test for parallelism is the comparison of subject and object antecedents when the referring form is in object position. The potential confound between parallelism and subject biases also throws doubt on the claim of a subject bias. In chapter 2, I will present the results of a text analysis, which simultaneously supports claims about parallelism and subjecthood. These conclusions are consistent with on-line evidence from Smyth and Chambers (1996), who demonstrated parallelism effects for both subject-subject and object-object parallelism. Thus, although the original papers on parallelism didn't

adequately support their claims, it appears that parallelism is indeed a relevant factor in pronoun interpretation.

1.1.2.4. Recency/Distance

Like linguists, psycholinguists have suggested that anaphor resolution is influenced by the recency of the antecedent (also sometimes called the "distance" between the anaphor and antecedent). For psycholinguists, however, Recency is usually computed over a shorter span than for linguists, who include the preceding discourse up to as many as 20 clauses back or more (Du Bois, 1987; Givón, 1983a). For example, Clark and Sengul (1979) demonstrated that pronoun resolution was easier if the antecedent was in the previous clause than if it had occurred earlier. This finding suggests that if an entity hasn't been mentioned for awhile, the mental representation for it declines in activation, so it becomes less accessible.

There are several possible reasons why a conceptual referent becomes less available as time passes. On one hand, it may be the result of a simple decay mechanism -- if an entity hasn't been mentioned for awhile, and the representation of that entity has not been re-activated, the activation for that representation will slowly decrease and eventually disappear. On the other hand, the reason that entities become less accessible over time may be the result of interference from other referents which are mentioned in the intervening discourse. These referents may compete with the previous referents and lead to their suppression (Gernsbacher, 1990:139).

Clark and Sengul (1979) also suggested that recency effects are not linear. They found that referents mentioned in the previous clause were significantly easier to access than referents from two or three clauses back. However, they found no difference between referents from two clauses back and referents from three clauses back. They concluded that referents from the preceding clause enjoy a privileged position, in comparison with other referents.

1.1.2.5. Ambiguity

Another feature that both linguists and psycholinguists have acknowledged is the role of ambiguity. When a reader/listener is confronted with an anaphoric reference, the task of identifying the referent may involve choosing among several referents. 19 shows an example of a naturally-occurring ambiguous pronoun, which the author chose to disambiguate for purposes of humor.

(19) Unspikeable Punch is a refreshing change from booze. Contrary to expectations, people will drink a lot of it. (Punch, that is.)

(Rosenberg, The Impoverished Students' Book of Cookery, Drinkery, & Housekeepery:36)

Many researchers have also noted that the interpretation of a referring form is easier when the number of possible antecedents is lower (e.g., Ariel, 1990; Givón, 1983a; McDonald and MacWhinney, 1995)

The question is, however, what qualifies as a "possible antecedent"? One way to think of pronoun resolution is as a process whereby the comprehender needs to find a conceptual referent that matches the features of the referring expression. Under this model, there are many factors that may play a role. Most psycholinguists have discussed this issue in terms of morphological ambiguity, distinguishing between "ambiguous" and "unambiguous" pronouns in English on the basis of number and gender (e.g., MacDonald and MacWhinney, 1990; McDonald and MacWhinney, 1995). That is, English "he" is only ambiguous in a context with two or more singular, male referents. However, comprehenders clearly also need to consider other aspects of the referring expression to interpret it. This observation is reflected in Givón's definition of "Potential Interference ('ambiguity')" (see §1.1.1.3.), which stresses that an expression is only ambiguous if there is more than one referent that is semantically compatible with the entire referring

expression. Walker (1998:413) similarly proposes that selectional restrictions can influence the interpretation of pronouns, for example the antecedent of "he" in "he rode" must be animate.

Thus, I hypothesize that the processing of a referring expression is sensitive to the properties of expression as a whole, and the degree to which they match the properties of the referent and other potential referents. For example, the pronouns in 20a-c contrast with each other in terms of different types of information.

(20) Context: Lisa is in a diner.

a. Animacy:
- Lisa bought a slice of pie. She was very hungry.
- Lisa bought a slice of pie. It was freshly baked and smelled wonderful.

b. Gender:
- Lisa sat down next to Ben. She ordered a slice of pie.
- Lisa sat down next to Ben. He ordered a slice of pie.

c. Semantic fit with following predicate:
- Lisa leaned on the counter as she ate her pie. It was all wet and her arm got soaked.
- Lisa leaned on the counter as she ate her pie. It was very tasty and she finished it immediately.

As the examples in 20 illustrate, many different features of the referring expression can restrict the possible set of interpretations. Some of these properties are inherent in the pronoun itself, such as animacy and gender, while some features are available only after integrating the pronoun with its role in the predicate.

However, attempts at understanding how information from the referring expression is used during processing have yielded contradictory results. For example, in McDonald and MacWhinney's (1995) third experiment on Implicit Causality, they manipulated the ambiguity of the pronoun by controlling the number of compatible referents in the discourse situation, as in 21.

(21) Gary amazed Alan/Ellen time after time because he was so talented.

As mentioned above, subjects heard these sentences and responded to probes at several points. The general finding was that the first-mentioned (subject) NP was more accessible at almost all time points. However, when the pronoun was ambiguous, the First-Mention advantage disappeared at the pronoun. They interpreted this as evidence that both participants were briefly considered as antecedents at that point. In the unambiguous pronoun condition, however, the first-mention advantage persisted at all time points, even in the case where the pronoun did not refer to the first-mentioned NP. They interpreted this result as evidence that gender information is not used at the point where the pronoun is encountered. Their final conclusions are indeterminate, suggesting that gender information may be available at the pronoun to a certain extent, but that the use of gender information may be influenced by other factors (1995:562).

In contrast, there is clear evidence that the ambiguity of the discourse situation affects language production. For example, Francik (1985) had English-speakers tell narratives based on comic strips about two characters. In one condition, the two characters were of the same sex (e.g., a man and a boy); in another, they were of different sex (e.g., a man and a girl). An analysis of the referential forms that participants used revealed a significant tendency to use pronouns more often in the mixed-gender condition, when the pronoun was unambiguous.

1.1.3. Centering Theory

One of the most developed theories about reference and reference processing is Centering Theory, which utilizes some of the same notions of topicality and salience, but operationalizes them in a very concrete way (Grosz and Sidner, 1986; Grosz, Weinstein, and Joshi, 1995). Centering Theory grew out of artificial intelligence research, and was developed to predict the antecedent of a pronoun based on simple computations that could be implemented in a computer program. By extension, these same computations have also been considered possible mechanisms for the human language processor.

The basic idea of Centering is that all arguments in a given utterance are ranked according to their grammatical function (SUBJ > OBJ > OBJ2 > OBL). These elements are the set of "forward-looking centers", or Cfs, because they provide potential antecedents for subsequent anaphors. The Cf set contains two privileged members: 1) the "backward-looking center", or Cb, which is the entity that the utterance is "about" and which refers back to something in the previous discourse, and 2) the highest-ranked Cf, which is called the "preferred center" (Cp), and is predicted to be the Cb of the following sentence. In an ideal situation, according to Centering Theory, the Cb and the Cp are identical. Thus, Centering proposes the following four possible transitions between two utterances, where "CONTINUE" is predicted to be the most coherent and easiest to process, and "ROUGH-SHIFT" is predicted to be the most difficult.[1]

(22) Centering Transitions (Reproduced from Walker and Prince, 1996:296)

	$Cb(U_i) = Cb(U_{i-1})$	$Cb(U_i) \neq Cb(U_{i-1})$
$Cb(U_i) = Cp(U_i)$	CONTINUE	SMOOTH-SHIFT
$Cb(U_i) \neq Cp(U_i)$	RETAIN	ROUGH-SHIFT

[1] For a more complete description of Centering Theory, see Grosz et al (1995), or papers in Walker et al. (1998).

By including both backward-looking centers and forward-looking centers, Centering formalizes a concept implicit in Givón's two measurements of "referential distance" and "persistence". The idea is that speakers will place an item in a prominent position, such as subject position, if it has been important in the preceding discourse, and/or if they intend it to play a prominent role in the upcoming discourse.

The question of determining how the elements are ranked has not been thoroughly explored in Centering Theory. The basic assumption is that the hierarchy of grammatical functions determines the ranking of arguments, which means that the grammatical subject is usually the center. However, Gordon et al. (1993) claim that fronted arguments (i.e., "first mentioned" NPs) are ranked higher than subjects, and Walker et al. (1994) claim that in Japanese the highest position is the topic (marked with "wa"). Centering theorists also acknowledge that other factors, such as thematic role, may influence the ranking of forward-looking centers (e.g., Brennan, Friedman and Pollard 1987; Grosz and Sidner 1995; Grosz et al. 1995).

Some of these factors have been explored by researchers working within the framework of Centering Theory. One such factor is whether an entity has been mentioned more than once within a discourse, and especially if it has already been referred to with a pronoun (e.g., Kameyama 1996). The effect of pronominal reference on the topicality of discourse entities can be observed in 23, from Sidner (1983:298).

(23) a. I got a new hat and I decorated it with a big red bow.
I think the bow will brighten it up a lot.
If not, I think I'll still use it.

b. I got a new hat and I decorated it with a big red bow.
I think it will brighten up the hat a lot.
If not, I think I'll still use it.

Sidner suggests that the pronoun in the last line is most naturally interpreted as co-referential with the previous pronoun, which has a different referent in 23a and 23b.

In addition, the role of the center (as grammatical subject) has been experimentally investigated by Hudson-D'Zmura and Tanenhaus (1998), and Gordon et al. (1993), who both used self-paced reading tasks of short passages in English. Hudson-D'Zmura and Tanenhaus presented subjects with sentence pairs like those in 24, consisting of a context sentence and a second sentence in one of four conditions.

(24) 1. CONTEXT SENTENCE Jack apologized profusely to Josh.

2a. SUBJ. REF. (pronoun) He had been rude to Josh yesterday.
2b. SUBJ. REF. (name) Jack had been rude to Josh yesterday.
2c. OBJ. REF. (pronoun) He had been offended by Jack's comment.
2d. OBJ. REF. (name) Josh had been offended by Jack's comment.

There were two measurements of the appropriateness of each condition: 1) Participants rated each sentence pair in terms of whether the second sentence made sense with respect to the first, and 2) they answered a comprehension question. The major finding was that for subject-referent conditions, both coherence ratings and comprehension scores were higher when the referring form was a pronoun, and for object antecedents, both ratings and scores

were higher when the referring form was a name. This result was replicated in four other experiments using variations of the same methodology.

Hudson-D'Zmura and Tanenhaus theorized that processing was impaired for pronominal objects because the antecedent was mis-assigned to the "center", or subject. A problem with this analysis, however, is that both object conditions involved a non-parallel structure, while the subject conditions had a pronoun or name referring back to the character in the parallel position. This means that their findings may have come from a tendency to follow a "parallel function strategy", which would result in the same antecedent preferences. Therefore, it may not be just the prominence of the "center" that affects antecedent assignment. When the order of reference parallels that of the previous sentence, intuition suggests that two pronouns are felicitous and indeed preferred:

(25) The Lakers beat the Celtics in the final game of the season. They had lost to them at an away game earlier in the season, so this victory was important.

Thus, the results of this experiment may have been influenced by two factors: the Subjecthood/First-Mention bias, and the Parallelism bias.

Despite problems with the Hudson-D'Zmura and Tanenhaus stimuli, however, the same "repeated-name penalty" was also found by Gordon, Grosz, and Gilliom (1993), who coined the term. In contrast, Gordon et al. used 4-sentence discourses where both the subject and object anaphors appeared in the same position as their antecedents.

(26) Sample stimuli from Gordon et al. (1993), Experiment 1. Underlines are included for presentational purposes only and were not present in the experimental stimuli.

A. NAME-NAME VERSION

1. Bruno was the bully of the neighborhood.
2. Bruno chased Tommy all the way home from school one day.
3. Bruno watched Tommy hide behind a big tree and start to cry.
4. Bruno yelled at Tommy so loudly that all the neighbors came outside.

B. PRO-NAME VERSION

1. Bruno was the bully of the neighborhood.
2. He chased Tommy all the way home from school one day.
3. He watched Tommy hide behind a big tree and start to cry.
4. He yelled at Tommy so loudly that all the neighbors came outside.

C. PRO-PRO VERSION

1. Bruno was the bully of the neighborhood.
2. He chased Tommy all the way home from school one day.
3. He watched him hide behind a big tree and start to cry.
4. He yelled at him so loudly that all the neighbors came outside.

Using a self-paced reading task, they had participants read paragraphs such as those in 26. The stimuli were presented in one of three conditions: PRO-PRO (in which all subsequent references used pronouns) PRO-NAME (in which only references to the subject antecedent used pronouns), NAME-NAME (in which all subsequent references used names). The results showed that participants were equally fast in PRO-PRO and PRO-NAME, but

slower in NAME-NAME. That is, participants read stimuli faster when references to the subject-referent were pronominal, there was no difference in reading times between versions that used pronouns or names to refer to the object-referent. This suggests a repeated-name penalty just for the subject-referent, but not for the object-referent, despite the fact that intuition suggests that the PRO-NAME condition is less natural than the PRO-PRO condition.

As is evident in the studies by Gordon et al. and Hudson-D'Zmura and Tanenhaus, one of the key features of Centering Theory is that it attributes a special status to the Cb, which shares many features with what other linguists have called the "topic". Centering claims that it is more natural to use pronouns to refer to the Cb, but it doesn't make any explicit predictions about other entities in the discourse, except to say that they do not hold the same status as the Cb. Attributing special status to the top-ranked entity is a problematic aspect of Centering, for three reasons. First, the difficulty of identifying a unique Cb or Cp in all cases makes the claim difficult to test. Second, non-subjects can often be more salient than subjects. For example, Smyth and Chambers (1996) showed that in some conditions, equal repeated-name penalties can be found for both object- and subject-referents, and that reading times for a pronoun are faster when it refers to a parallel object-referent than to a nonparallel subject-referent. Third, Centering Theory stipulates that if a pronoun is used to refer to the Cb, other pronouns are acceptable (but not necessary), but it makes no predictions about if or when a pronoun will be used to refer to a non-Cb. I will discuss further criticisms of Centering Theory in relation to the Information Load Hypothesis, below.

In defense of Centering, however, this line of research has offered many contributions to the study of discourse and reference form. Centering is the only extant theory that has offered a formalized account of transitions between sentences and the use of pronouns. It is also one of the few theories that simultaneously looks at how utterances relate to the preceding discourse (via the backward-looking center) and the following

discourse (via the set of forward-looking centers). The theory is also expressed clearly and falsifiably. Centering has produced a wealth of research on local discourse structure (see Walker, Joshi, and Prince, 1998), focusing on distinctions in accessibility and prominence of entities in different discourse structures, both through text analysis (e.g., Hurewitz, 1998; Walker, Cote, and Iida, 1994; Walker and Prince, 1996), and experimentation (Brennan, 1995; Hudson-D'Zmura and Tanenhaus, 1998; Gordon et al., 1993). Centering has also provided a useful account of pronoun resolution for the artificial intelligence community, the purpose for which it was originally intended.

1.1.4. The Information Load Hypothesis

An additional criticism of Centering Theory is developed in Almor's (1995) Information Load Hypothesis. A major implication of this paper is that, contrary to the claims of Centering, reference to the most highly focused element in a discourse is not obligatorily pronominal. Derived from Grice's maxim of quantity, the Information Load Hypothesis claims that anaphors differ from each other in terms of the cost they put on the processing system. Low-cost anaphors include pronouns and NP anaphors that are more general than their antecedents, for example using "the bird" as an anaphoric expression for something previously identified as "the robin". Almor's central claim is that if a speaker uses a high-cost anaphor, there must be a justification for it. He discussed two situations in which a high-cost anaphor would be justified: 1) When the antecedent is not focused, and is therefore not accessible; 2) When the speaker wishes to express additional information about the discourse referent, for example by using the phrase "the robin" to refer to a referent previously expressed as "the bird". Repeated NPs, like the ones in Gordon et al.'s stimuli, are neither low-cost, nor do they express any additional information. Therefore, they are only functionally justified when the antecedent is low in focus.

Almor demonstrated that in fact, the processing of full NP anaphors is faster when the antecedent is salient than when the antecedent is not salient. Using a self-paced reading methodology, Almor had subjects read sentence pairs such as those in 27.

(27) Sample stimuli form Almor (1995), Experiment 1

a'. It was the robin that ate the apple. <-- FASTER READING TIMES
The bird seemed very satisfied.

a". It was the robin that ate the apple. <-- SLOWER READING TIMES
The fruit was already half rotten.

b'. What the robin ate was the apple. <-- SLOWER READING TIMES
The bird seemed very satisfied.

b". What the robin ate was the apple. <-- FASTER READING TIMES
The fruit was already half rotten.

The results showed that the NP anaphor in a' was read faster than the one in a", and the NP anaphor in b" was read faster than the one in b'. In contrast, when repeated NPs were substituted for the NP anaphors (e.g., "the robin" in a' and b', and "the apple" in a" and b"), the reverse results obtained -- a" was read faster than a', and b' was read faster than b". Almor concluded that the penalty observed by Gordon et al. was due to the infelicity of repeated names, and not a categorical prohibition against noun anaphors for focused referents.

While the Information Load Hypothesis challenged Centering Theory, it also confirmed the infelicity of using repeated names to refer to highly salient antecedents. However, this study differed from the work in Centering in the following way. The experiments by Gordon et al. (1993) and Hudson-D'Zmura and Tanenhaus (1998) demonstrated the repeated name penalty for referring to the grammatical subject, a typically topical position. In contrast, Almor demonstrated that repeated names are less natural when

the antecedent is focused by means of a cleft construction. This contrast is notable because topics and foci have traditionally been opposed in theories of information structure. However, these studies found that the referents of subjects and foci are more salient than other referents. This similarity will be explored in chapter 3.

1.1.4.1. Other uses of noun phrase anaphors

The results of Almor's study have an additional implication for the study of reference form, aside from issues that are specific to Centering Theory. That is, studies of reference form often ignore the possibility of using full NP anaphors for reasons of style, emphasis, or as Almor points out, adding information. However, as Almor's study demonstrated, even when a pronoun would be quite natural, the use of a pronoun is not obligatory.

More generally, it is easy to observe that full forms of reference are occasionally used for reasons of style, emphasis, to describe an individual in a different way, or even to manipulate the structure of the discourse. The following excerpt shows how repeated reference can be used for stylistic purposes.

(28) It was after the picnic that the town began to notice things and got mad. Tea Cake and Mrs. Mayor Starks! All the men that she could get, and fooling with somebody like Tea Cake! Another thing, Joe Starks hadn't been dead but nine months and here she goes sashaying off to a picnic in pink linen. Done quit attending church, like she used to. Gone off to Sanford in a car with Tea Cake and her all dressed in blue! It was a shame. Done took to high heel slippers and a ten dollar hat! Looking like some young girl, always in blue because Tea Cake told her to wear it. Poor Joe Starks. Bet he turns over in his grave every day. Tea Cake and Janie gone hunting. Tea Cake and Janie gone fishing. Tea Cake and Janie gone to Orlando to the movies. Tea Cake and Janie gone to a dance. Tea Cake making flower beds in Janie's yard and seeding the garden for her. Chopping down that tree she never did like by the dining room window. All those signs of possession. Tea Cake in a borrowed car teaching Janie to drive. Tea Cake and Janie playing checkers; playing coon-can; playing Florida flip on the store porch all afternoon as if nobody else was there. Day after day and week after week.

(Hurston, Their eyes were watching God, 1937:105)

In this example, it seems that the narrator wanted to emphasize not only all the things that Tea Cake and Janie did together, but the fact that they did them together, and that these activities were a repeated occurrence. This effect is achieved by repeating "Tea Cake and Janie", even though both characters are highly prominent, and are additionally referred to with pronouns and elliptical references. Another literary example of a full NP anaphor can be observed in 29, where the same character is referred to both as "Dorothy" and "the little girl".

(29) Aunt Em had been so startled by the child's laughter that she would scream and press her hand upon her heart whenever Dorothy's merry voice reached her ears; and she still looked at the little girl with wonder that she could find anything to laugh at.

(Baum, The Wonderful Wizard of Oz, 1900:12)

It has also been argued that full forms of reference can be used to manipulate the discourse structure. For example, Vonk, Hustinx and Simons (1992) suggested that non-pronominal forms of reference are one signal of a thematic shift. The evidence for this claim comes from a production study, where speakers described a cartoon which either had a thematic shift or not. Vonk et al. found that when subjects described a cartoon with a thematic shift, non-pronominal reference represented only one of several strategies that speakers used to effect the shift. When other markers were used, such as phrases signaling a change in time or place, subjects tended to use pronouns, but when no such time/place phrase was used, speakers produced a higher rate of non-pronominal reference. In addition, Vonk et al. claimed that hearers make use of information from referential form in order to deduce the structure of the discourse. They hypothesized that the appearance of a full form of reference signals a break in the thematic structure of the discourse, making preceding information less accessible than if a thematic break had not occurred. This claim was supported through two comprehension experiments, where the use of a name (in comparison with a pronoun) caused subjects to be slower in recognizing probe words related to information from the previous sentence.

These examples, in addition to Almor's findings, suggest that any theory of reference must take into account the possibility of using full NPs or names for particular purposes. Because of this, it is important to realize in the following chapters that high activation of a discourse representation does not translate into absolute requirements for a particular form of reference. In all cases, the speaker/author has the prerogative to

manipulate form of reference for a particular effect. However, the effects that are possible with a particular form result from the more general distributions of these forms, through which they come to be associated with discourse functions. Furthermore, all forms of reference are not always available choices for the speaker.

1.1.5. Summary

Theories of reference form have usually focused on the status of the referent, described in terms of a property variously termed salience, accessiblity, or focus of attention. Taken together, the many studies from psycholinguistics, functional linguistics, and computational linguistics have identified a number of factors that affect reference form. These are summarized in Figure 1. Some features not discussed here are also listed, for the sake of completeness.

This figure depicts choices in reference form as falling along a continuum, such that less specified forms are to the left, and more specified forms are to the right. The factors listed have been shown or are hypothesized to affect speaker's choices of reference form. The effect of each factor pushes the speaker's preference more toward the left (i.e., less specified forms), or more to the right (i.e., more specified forms). Many different factors can influence choices in reference form on a given occasion, although the relative importance of each factor may vary from situation to situation.

Figure 1: FACTORS INFLUENCING REFERENCE FORM

less specified form		**more specified form**
<----- more	A. RECENCY OF REFERENT	less ---->
<----- more	B. REFERENT TOPICALITY Influenced by: - Subjecthood/First Mention - Protagonist Identity - Repeated Mention - Prior pronominal reference - Passivization and other marked structures (Arnold,1995).	less ---->
<----- more	C. FOCUS ON REFERENT Influenced by: - focus constructions - intonation	less ---->
<----- more	D. SEMANTIC PLAUSIBILITY OF REFERENT Influenced by: - thematic roles - perceived connection between utterances - event structure	less ---->
<----- parallel	E. PARALLELISM Influenced by: - parallel grammatical function - parallel voice (Arnold, 1996) - parallel verb or verb structure	nonparallel ---->
<----- less	F. AMBIGUITY OF REFERENT RELATIVE TO CONTEXT	more ---->
<----- no	G. OTHER USES OF FULL NP ANAPHORS - style - emphasis - to add new information	yes ---->

EXAMPLES

English Examples

(pronoun)	(name/NP)
he, she	**John, Mary, the professor**

Spanish Examples

(null/clitic)	(overt pronoun)	(name/NP)
Ø, la, lo, le	**ella, él, ellas, ellos**	**Juan, María, el profesor**

1.2. A Processing Framework

In the section above, I reviewed a wide range of factors, all of which have been linked to speaker choices of reference form or the comprehension of referring forms. On the surface, these factors seem to have little to do with each other. Why, for example, are pronouns used more when the referent has been recently mentioned, and also when the referent has played a particular thematic role? Why do both parallel structure and the topicality of the referent influence the ease of interpreting a pronoun?

In the remainder of the dissertation, I will focus on five of the factors considered above. I will confirm past findings that pronouns and other less-specified forms of reference (like null anaphors) are more natural for all five. The five factors are presented in Table 1, along with definitions of how I will be using the terms. For a general discussion of each factor, see §1.1.

Table 1. The labels and definitions for the five factors to be investigated

LABEL	**DEFINITION of the factor**
i) Recency	The property of a referent (i.e., an entity) of having been recently referred to in the discourse.
ii) Subjecthood	The property of a referent of having been last mentioned in subject position.
iii) Focus	The property of a referent of having been last mentioned as the focus of a cleft or cleft-like construction.
iv) Parallelism	The property of a referent of having been last mentioned with an NP in the same grammatical role as current referring NP.
v) Goal Status	The property of a referent of having been last mentioned as the goal argument of a verb.

The purpose of this dissertation is to address the question of what makes these five factors behave in a similar way with respect to reference form. As an answer to this question, I will provide evidence that they all share an additional property: the referents associated with these factors are also more likely to be referred to in the following discourse than other referents. Through text and corpus analyses, I will show that there is a higher probability that speakers will refer again to recent-referents, subject-referents, focus-

referents, parallel-referents, and goal-referents, in contrast with comparable referents (i.e., nonrecent-referents, object-referents, etc.). I hypothesize that the tendency toward repeated reference to certain entities is part of a more general tendency to continue to use those referents in the following interaction, whether that use is linguistic or non-linguistic. However, my findings come from purely linguistic data, and from situations with a minimum of nonlinguistic interaction (written texts and discussions from the Canadian Parliament).

Thus, I will show an association between each factor and two generalizations: a) Reference Form, such that speakers and writers tend to use pronouns and other less specified forms for subsequent reference to referents with these properties, and b) Reference Continuation, such that speakers and writers tend to refer more often overall to referents with these properties, compared to referents with contrasting properties. Throughout this dissertation, I will refer to these two generalizations as "Reference Form" and "Reference Continuation", for ease of presentation. Their relation to the five factors is depicted in Figure 2.

Figure 2: An outline of the results of the dissertation

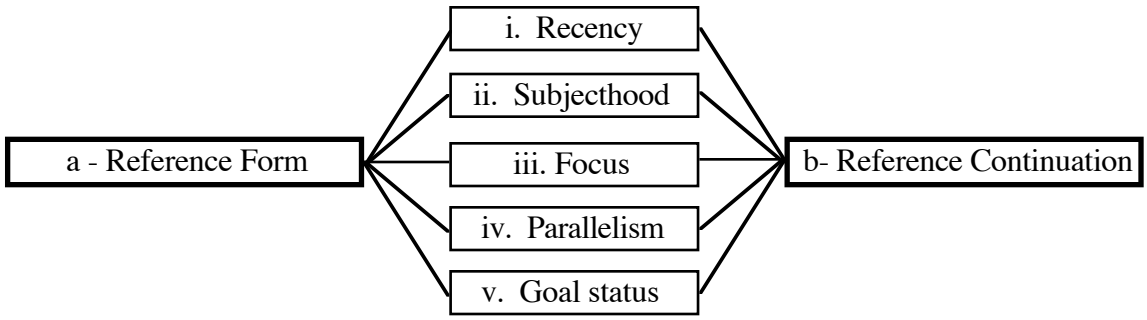

The finding that these factors are associated with Reference Form is not surprising. This essentially supports a series of past findings about reference form and pronoun interpretation. The novel result of my studies is that all five are **also** associated with an

higher overall probability that these same referents will be mentioned in the subsequent discourse. This result is important, because it offers a preliminary answer the question of why these five factors behave similarly with regard to Reference Form. That is, the five factors share at least one other property, the tendency for their referents to appear again in the following discourse.

This generalization is useful, but from the perspective of discourse processing, it is not enough. It raises two additional questions: 1) Where do these discourse patterns come from?, and 2) Why are they also associated with choices in reference form? I will consider each question in turn here, and return to them in more detail in later chapters.

1.2.1. Where do discourse patterns come from?

People use language in regular ways because language is one part of human behavior, and people behave in regular ways. As Clark (1996) has argued, discourse is just one form of joint activity. One property of joint activities is that they are goal-driven, meaning that discourse participants come to the discourse situation with a set of goals and intentions (Clark, 1996:33; Grosz and Sidner, 1986). These goals may be well defined prior to the discourse, or they may not be. The goal-oriented nature of language use offers possible explanations for numerous referential patterns.

First, speakers attempt to coordinate with their interlocutors to produce a coherent discourse (e.g., Clark, 1996; Grosz et al., 1995). People do not produce strings of unrelated sentences. Rather, they tend to talk about the same things for extended periods of time. Consequently, people refer more often to referents that have been recently mentioned than referents that have not ("i" in Figure 2, above).

The fact that speakers are pursuing larger goals also means that certain referents are more central to the speaker's intentions than others. Research has established that speakers tend to use particular constructions, such as the grammatical subject, to refer to referents that are important to the discourse (e.g., Chafe, 1976; Prince, 1992; Du Bois, 1987) or in the speaker's focus of attention (e.g., Tomlin et al., 1997). Those referents that are important to

the discourse are also those referents that the speaker will keep referring back to. These facts together account for why the referents of subjects and foci are more frequently referred to than other referents ("ii" and "iii" in Figure 2, above). These facts may also account for the fourth pattern, that speakers tend to refer more often to referents in the parallel position from the preceding clause ("iv" in Figure 2, above). It is likely that referents are put in the same semantic and pragmatic roles throughout a discourse segment. For this reason, speakers tend to use the same grammatical position to refer to a particular referent in subsequent utterances.

Why it is that speakers refer more often to goals than to sources? ("v" in Figure 2, above). This factor is different from the others, because it derives from the semantics of a proposition, rather than the forms chosen to express that proposition. That is, if John is telling a story and he says "Francine sent the final copy of her paper to the printers", he says so because this is the proposition he wanted to express. There are other ways he could have expressed it to place emphasis on one referent or another, but in all cases "the printers" would be the goal argument of the utterance. Some thematic roles reflect the perspective taken by the speaker; for example, the agent role differs for "chase" and "flee". However, the identification of a referent as a source or goal is driven purely by the semantics of the proposition (at least for verbs of transfer). Thus, the fact that speakers tend to refer more often to goals is a generalization about the things that people like to talk about. This may be because goal arguments often appear in an utterance with a source and a theme argument, where the theme moves from the source to the goal (either literally or metaphorically). If the discourse participants are concerned with what happens to the theme, then the subsequent utterance is likely to describe what the goal argument did with the theme. This pattern of reference is also sensitive to other features of the discourse structure, which I will discuss in chapter 4.

Thus, discourse patterns probably exist because of speaker's goals and intentions with regard to the discourse. But regardless of the reason, the data in the following chapters

will show that these patterns do indeed exist. That is, speakers tend to continue talking about referents that have been mentioned recently, and especially referents that occurred in subject position, as the focus of a cleft, or as the goal of a verb. Speakers also tend to refer to referents in the same syntactic position as they last referred to them.

1.2.2. Why are these factors also associated with Reference Form?

My data will also show that all five factors are associated with Reference Form. That is, pronouns and null anaphors are used more for subject-referents, focus-referents, recent-referents, parallel-referents, and goal-referent, in contrast with referents that lack these properties. Why is this, and what is the relation between Reference Form and Reference Continuation? The answer, I suggest, is that it is easier for comprehenders to access referents with these properties during the interpretation of a subsequent anaphor. And when comprehension is facilitated, speakers can use less-specified forms of reference. One reason these factors facilitate comprehension is because of their association with Reference Continuation. This association influences two related aspects of discourse processing: 1) it helps comprehenders interpret the goals and intentions of the speaker, and 2) this process and the probabilistic information associated with each factor aid comprehenders during anaphor resolution.

1.2.2.1. Interpreting the speaker's intentions

The purpose of language is to communicate. This means that comprehension does not merely involve parsing each utterance and resolving the anaphors, but also requires interpreting the contribution of that utterance relative to the discourse as a whole and the speaker's intentions (H. Clark, 1996). For example, the sentence "Samantha drank all the milk" might introduce the need for someone to buy more milk, it might begin a description of Samantha's love for milk, or it might initiate a diatribe about her inconsiderate household habits. Depending on the speaker's intentions, one referent or the other might be more important to the discourse.

Although listeners do not know the speaker's intentions, they can guess them from the context of the discourse. There are several ways a comprehender can infer the intentions of the speaker. In this example, the speaker's facial expression or intonation might indicate whether the speaker was annoyed or not, or the speaker may have prefaced this utterance with "That blasted housemate of mine!" But even in the absence of specific cues like these, listeners can still rely on information from the linguistic context, and the discourse properties associated with each referent. Following researchers like Bates and MacWhinney (1989); Du Bois (e.g., 1985, 1987), Prince (e.g., 1978, 1981, 1992), and Ward and Birner (e.g., 1995, 1996), I claim that one aspect of learning a language involves learning how different linguistic constructions are associated with the discourse status of the referent. One way of learning this association is through observing that referents in certain positions are often taken up in the following discourse: when referents of a certain type are frequently important to the following discourse, people learn an association between that type of referent and a certain level of discourse prominence.

In the example "Samantha drank all the milk," listeners may not know whether Samantha or milk was more important to the speaker's intentions on that occasion until they have observed the speaker's following actions. If the speaker followed up with "So let's buy some at the store", it would indicate that "milk" was more important on that occasion. It is also possible that the following interaction could be non-linguistic. For example, the speaker could then shake her fist at Samantha, thus indicating her displeasure. While the action is not linguistic, it is associated with the prior mention of Samantha linguistically, and thus becomes part of the listener's knowledge of the discourse properties of linguistic expressions.

Through a lifetime of experience with language, individuals learn the discourse properties of referents that are most likely to be continued in the following discourse. That is, they learn the association between each factor and Reference Continuation. In this way they learn to infer the importance of a given referent on the basis of its discourse properties.

As I will show in the following chapters, subject-referents, focus-referents, recent-referents, parallel-referents, and goal-referents are more frequently continued in a discourse than other referents. Through these frequent associations, people learn that these discourse properties are useful for discovering which referents are more important to the speaker. Learning discourse patterns like these is essentially learning generalizations about the world, specifically those generalizations that are useful for language processing.

However, these discourse properties do not provide categorical information about the speaker's intentions. When the speaker refers to Samantha using a subject or focus NP, for example, this may indicate that Samantha is important for the following discourse, but not necessarily. Instead, this information is probabilistic. People learn that when a referent has been mentioned as subject or focus, or when it appears in certain thematic roles, it has some probability of being important to the following discourse. Similarly, simply having been mentioned more recently gives a referent some probability of being continued in the discourse. These four factors (Recency, Subjecthood, Focus, Goal Status) combine with each other to yield a probabilistic interpretation of the importance of each referent to the speaker's intentions, and thus how likely it is to be mentioned or otherwise used in the following interaction.

The fifth factor, Parallelism, works similarly to the other factors, but on a much smaller scale, and its effects are transitory. Given an anaphor in object position, for example, it is more likely that the referent was last mentioned in object position of the preceding clause. Therefore, at the point where the listener knows the position of the anaphor, Parallelism provides probabilistic information that the speaker will refer to the parallel-referent from the preceding clause, such that parallel-referents are more probable than nonparallel-referents. Thus, it is different from the other factors because it does not provide information about where the discourse is going, or what the speaker's general goals and intentions are. However, Parallelism does aid comprehenders in determining the speaker's intended reference at a specific point during language comprehension.

1.2.2.2. Using probabilistic information during referent resolution

Once a person has learned the associations between linguistic constructions and discourse patterns, those linguistic associations become useful for interpreting the speaker's meaning. Among other things, they facilitate the process of interpreting anaphoric references. I further propose that one of the reasons these five factors facilitate anaphor resolution is that they are probabilistically associated with repeated mention of certain referents. This proposal rests on several assumptions about discourse processing, and is best described in terms of an activation metaphor of cognitive processes. I will describe these assumptions, and then propose how probabilistic information influences referent resolution.

First, I assume that discourse processing, and indeed all cognitive processes, involve mental representations of the discourse elements (e.g., Gernsbacher, 1990; Kintch, 1988; Morrow and Bower, 1989; Morrow et al., 1987; etc.). Psychologists have often distinguished between two levels of representation: 1) long term memory, and 2) working memory. For ease of presentation, I will continue to make the same distinction. Long-term memory is where knowledge is stored. This knowledge includes both conceptual information and linguistic knowledge (e.g., Kintch, 1988).

Discourse representations are usually assumed to be located in "Working Memory", which is a processing space that contains information about currently relevant information and processes (e.g., Baddeley, 1981, 1983; Baddeley and Hitch, 1974; Daneman and Carpenter, 1980, 1983). A crucial aspect of Baddeley's model is that working memory is a cognitive space that contains both the storage and processing of currently relevant information.

Among other things, working memory contains a model of the information shared between interlocutors, or the "common ground" (H. Clark, 1996). One part of the common ground is the discourse representation. Clark stresses the fact that a discourse includes both linguistic and non-linguistic interaction, he uses the term "discourse" only to

emphasize that language is a part of it. I continue this perspective here, using "discourse" as an event including language, but not restricted to it. A model of the discourse representation is presented in Figure 3.

Figure 3: Partial model of the Discourse Representation
(from Clark, 1996:54)

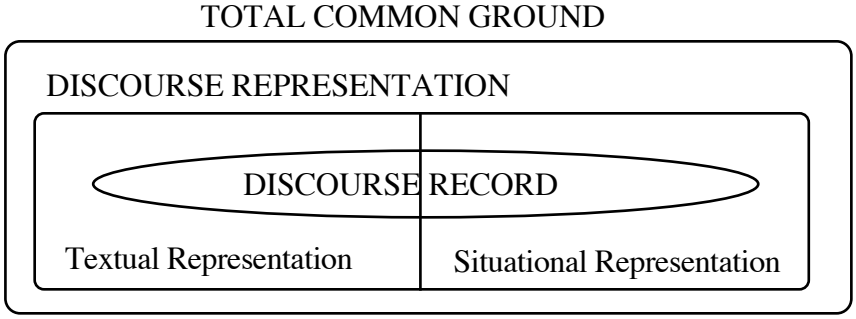

Working Memory is also where language comprehension processes take place. The linguistic input is briefly represented, long enough to parse it into syntactic structures. This information is represented in the text model, or textual representation, and is used to update the situational representation. However, the linguistic information itself does not persist for long in memory (see Fletcher, 1990, for an overview).

The situation model includes mental representations of the discourse referents and the relations between them. For example, upon hearing "The detective called the restaurant manager," a listener would minimally form a representation of the detective, the manager, and the calling event. The representation may also include inferred aspects of the situation (e.g., Kintch, 1993; Long et al., 1992; Prince, 1992).

The discourse record is a representation overlain on the text and situation models. This is a privileged section of the model, and represents the "official states and events in the current joint activity," (H. Clark, 1996:54). The information in the discourse record is that which is mutually known by all discourse participants. For example, after a transaction between a bank teller and a customer, the discourse record for each participant contains the

content of their transaction. If the customer glanced at her watch during the transaction, however, this event may not be part of the discourse record. In my discussion in the following chapters, I will focus primarily on representations within the discourse record. For simplicity, I will refer to the "text model" and "situation model" aspects of the discourse record, and will not explicitly discuss the discourse record itself, except when relevant.

Entity representations are best discussed in terms of their levels of activation in the situation model. That is, each entity representation in the mind of each discourse participant is "lit up" to varying degrees at each point of the discourse. One way of thinking of activation is in terms of a spotlight on a dramatic scene. At any given point the spotlight centers on the most important characters and/or props at that moment. Other characters and props may be at the edge of the spotlight, thus only partially lit up. Other characters are on stage, but not in the spotlight, and others are in the wings, having recently departed the stage. The spotlight/stage metaphor does not capture all the properties of discourse processing, but it provides a real-world analogue of the activation metaphor.

There are several ways that a representation can become activated. When a speaker decides to say something, it requires activating a conceptual representation of the desired message, and then activating a corresponding linguistic representation. The conceptual representation of the speaker's message may become activated for many reasons, including being triggered by something that someone else said.

Discourse comprehension also involves activating the referents. The first time a referent is mentioned, a representation is activated in the situation model. If it is mentioned a second time, this causes the listener to re-activate the representation. As time goes by, that activation decays gradually, and if the entity is not evoked again, its representation will eventually disappear from the situation model. At any given point during the decay process, the referent will be partially activated. An important characteristic of activation is that it may be partial, and therefore entity representations can vary in degree of activation at any one point.

Another important characteristic of activation is that it changes dynamically as language processing occurs over time. This means that activation rises and falls continuously as new information is encountered and integrated, and that the process of accessing a referent (for example, for the purpose of anaphor resolution) is not instantaneous.

There are a number of reasons I am adopting activation for discussing referent processing. First, I thereby continue a long tradition of using activation as a metaphor for cognitive processes. Information processing is often discussed in terms of the activation of representations at multiple levels, including both conceptual and linguistic representations. The degree to which a representation is activated reflects the ease of accessing that information, remembering it, or integrating it with other information.

For example, Marslen-Wilson (1990) discussed spoken word recognition as a process whereby lexical representations become activated dynamically through time on the basis of the phonetic input and word frequency. For example, the fragment "ele" would initially activate both "elevator" and "elephant", with the more frequent word receiving slightly more activation. As the input is processed, it eventually supports only one as the correct interpretation, resulting in full activation for that lexical item. Many researchers have also used activation to account for lexical or syntactic ambiguity resolution, where the different syntactic or semantic representations associated with an ambiguous form become activated during processing (e.g., MacDonald et al., 1994).

The activation metaphor has also been widely used in accounts of discourse processing. For example, Gernsbacher discussed discourse processing in terms of building mental representations. She claimed that the building blocks of these representations are "memory cells", which become activated by incoming stimuli (1990:1-2). Chafe (1994) suggested that discourse processing is sensitive to how referents are represented in an individual's "consciousness" (similar to an individual's situation model). He distinguished between three levels of representation: active, semi-active, and inactive. Other researchers

have also discussed referent representation in terms of degree of activation (e.g., Morrow et al., 1989).

The activation metaphor also plays a central role in connectionist models of language processing (e.g., Kawamoto, 1988; McClelland and Rumelhart, 1981; Seidenberg and McClelland, 1989). These models discuss language processing as patterns of activation over nodes. Each node represents information at one level of abstraction. For example, some nodes might represent parts of lexical meanings and pronunciations (Seidenberg and McClelland, 1989), others might represent components of the syntactic structure (see Chater and Christiansen, in press, for an overview). In these implementations, activation is a real number that varies, for example between 1 and -1.

An activation framework also has advantages over similar concepts like "salience" or "attention". The concept "salience" can be used to describe a property of the linguistic/textual structure, as opposed to a property of the mental representation of the discourse. In constrast, "activation" more clearly supports the assumption that the production and comprehension of reference form are driven by the cognitive status of discourse entities, and continues the psycholinguistic tradition of discussing mental states in terms of active and inactive information.

It is probable that activation is related to global processes like attention, in that elements in the focus of attention are also highly activated. However, I do not mean to suggest that the processes I discuss are under the conscious control of the speaker or listener. "Activation" is more appropriate for my purposes, because it is typically used for unconscious processes. In contrast, "attention" is often thought of as a conscious process, as in the request "Pay attention to what I am saying".

A further advantage to using activation for discussing reference processing is that it is conducive to considering referent status as a continuously graded phenomenon. This view, similar to those of Ariel (1990) and Givón (1983a), has several advantages over a categorical approach to salience. The first advantage is that it makes predictions about

reference to all entities, not just the most salient one. Categorical approaches claim that one entity in every utterance enjoys a privileged status as the most salient entity, and focus on the differences between this entity and the others. This perspective is evident in Centering Theory, which claims that the Cb (backward-looking center) is the most salient item in an utterance. The idea of a unique, prominent referent is also implicit in Almor's work, where he states that "...the assumption [is] that the focused discourse entity is the default antecedent because it is kept in working memory," (1995:15). Although in many cases one entity may stand out as the most highly focused one, this approach makes no predictions about pronominal reference to other entities, such as "her eyes" in 30.

(30) The sun and wind$_i$ had changed her [Aunt Em], too. They$_i$ had taken the sparkle from her eyes$_j$ and left them$_j$ a sober gray; they$_i$ had taken the red from her cheeks and lips$_k$, and they$_k$ were gray also.

(Baum, The Wonderful Wizard of Oz, ch. 1)

Activation can also be construed as a categorical concept, as Chafe does with his three levels of activation (1994). However, the concept of activation can also be easily used to describe continuously graded differences in referent status. At the same time, an activation framework can still account for the generalization that there is usually only one entity at a time with high activation, through the concept of competition. Referent representations compete for activation, so that as one representation gains activation, others lose it.

Neither does the activation metaphor require a ranking mechanism for the cognitive statuses of different referents. Centering Theory represents differences in salience in terms of a hierarchy of discourse referents (the list of forward-looking centers). This approach is computationally straightforward, but it raises the question of when the re-ranking process takes place. Currently Centering suggests it takes place at sentence boundaries, but this conflicts with the known incremental nature of language processing. In contrast, activation is a property that can vary for each individual entity, at any point during processing.

Following the preceding assumptions about activation and discourse processing, I propose that anaphor resolution works in the following way. The first time a referent is mentioned, the comprehender accesses a representation of that referent and instantiates it in a situation model. If the referent is mentioned again, the comprehender must activate the representation a second time. Re-activation can be easier or harder, depending on two things: a) the degree to which the representation is already activated, and b) the degree to which the input supports that entity as the anaphor referent. In fact, these two processes can be described as one, which I will discuss further in chapter 5.

In both cases, the activation of referent representations is influenced by the probabilistic information carried by the five factors discussed here: Recency, Subjecthood, Focus, Parallelism, and Goal Status. My data show that these factors are associated with Reference Continuation. At one level of abstraction, this means that comprehenders can use these factors as gauges of how likely it is that a referent is central to the speaker's intentions (along with other relevant sources of information).[1] I propose that the degree to which comprehenders consider a referent to be central to the speaker's intentions is reflected in the degree to which that referent representation is activated in their situation model. Thus, referents with a high probability of being central to the following discourse, like subject-referents, have more highly activated representations than those with a low probability, like object-referents.

At a different level of abstraction, these five factors can also facilitate the process of anaphor resolution itself. When comprehenders encounter an anaphor, the processing system begins trying to identify the referent. This process is influenced by many types of information, in particular the identifying information carried in the anaphoric form. For example, if the anaphor is "John", the referent must be one that is identified by this name; if the anaphor is "he", the referent must be male. But anaphor resolution is also influenced by

other factors, including the five mentioned above. Each factor supplies probabilistic information about the identity of the referent: it is more likely to be a recent-referent than not, it is more likely to be a subject-referent than an object- or oblique-referent, it is more likely to be a focus-referent than a nonfocus-referent, it is more likely to be a goal-referent than a source-referent, and it is more likely to be a parallel-referent than a nonparallel-referent. Each of these pieces of information will influence the anaphor resolution process. When the referent is one that is highly probable, resolution will be easier than when it is less probable. In addition, to the extent that these factors inform the listener about the speaker's intentions, some referents may already be more activated than others. When the referent representation is already activated, accessing it will also be facilitated.

1.2.2.3. The link between comprehension and production

The facilitation of anaphor resolution is important, because it allows speakers to use shorter, more attenuated forms to communicate with their interlocutors. Therefore, speakers can use pronouns and null anaphors more for referents with the properties listed above. For referents without these properties, comprehension will be relatively more difficult, so speakers will choose the fuller forms of reference that their addressees need.

One might ask why I am drawing a connection between the facilitation of comprehension and the speaker's choices of reference form. This is a natural connection to make. Comprehension and production are merely two sides of the communicative coin. As argued by Clark (1996), language is one type of joint activity. Speakers do not speak just to speak, they speak to perform certain actions, to achieve certain goals. This has multiple implications for how language is used, one of which is that speakers want their speech to be understood as easily as possible. Speakers design their utterances for their addressees, and choose forms of reference that will be understood. The joint nature of communication has the additional consequence that it requires the interlocutors to coordinate their actions at

---

[1] In a global sense, only four of these factors (Recency, Subjecthood, Focus, and Goal Status) are good indicators of the speaker's intentions. Parallelism is only a good indicator of the speaker's "intentions" at a

every level. One of the things they must coordinate is their mutual knowledge of what has happened during the discourse, including the cognitive status of each referent. Thus, it is plausible to assume that discourse participants will develop similar mental representations about the shared discourse events and referents. In many cases, these representations must be developed on the basis of linguistic information. The degree to which speakers keep track of their addressee's mental models is not fully known. However, evidence suggests that discourse participants usually achieve some degree of coordination (H. Clark, 1996).

The communicative nature of discourse means that speakers can take advantage of the cases when they expect that comprehension will be facilitated, and have the option of using more efficient, shorter forms. Furthermore, it has been established that speakers do make use of this option. For example, Anderson et al. (1997) showed that speakers pronounce individual words less clearly during face-to-face conversation, in comparison with cases when the discourse participants can not see each others' faces. They argue that their result arises because the speaker's face provides visual information that facilitates speech comprehension, such as place-of-articulation information about consonants (e.g., McGurk and MacDonald, 1976), making certain aspects of the acoustic input less needed. Another example of this type is Lieberman's (1963) study, which showed that speakers pronounced words less clearly when they were more predictable from the sentence context. For example, the word "nine" was pronounced more clearly in 31a than in 31b.

(31) a. The word you are about to hear is *nine*.

b. A stitch in time saves *nine*.

Speakers are sensitive to the different types of information that are available to their interlocutors, and they produce their utterances accordingly. This same idea of listener accommodation provides the basis for interpreting the results of this dissertation.

---

very specific point during processing. For further discussion of this point, see §2.4.3.

1.3. A look ahead

In the following chapters, I will report on five factors: Recency, Subjecthood, Focus, Parallelism, and Goal Status. For each of these, I will provide evidence for the double association in Figure 2. First, I will confirm that they are associated with an increased use of less specified forms of reference. Although reference can be made using a wide variety of forms, the distinctions I am principally concerned with are between full NP anaphors, overt pronouns, and null anaphors. Second, I will demonstrate that for all five factors there is an increased probability that certain referents will continue in the following discourse. This evidence comes from analyses of written texts and spoken corpora, from arating questionnaire, and from a discourse-completion experiment.

In chapter 2, I will present text analyses in three languages of the distribution of forms of reference, with respect to Recency, Subjecthood, and Parallelism. In chapter 3, I will focus on the contrast between subjects and the foci of clefts in English. With a rating questionnaire and corpus analysis, I will demonstrate that subjects and foci have similar functions, and suggest some implications for theories of information structure. In chapter 4, I will discuss source and goal arguments, and their relation to both Reference Form and overall Reference Continuation. The data in this chapter come from a discourse-completion experiment and a corpus analysis.

In chapter 5, I will return to the question of the underlying motivation for the double association between reference form and overall subsequent reference. I will propose a more detailed account of how probabilistic information from these five factors influences anaphor resolution. This proposal will be largely speculative, but will build on constraint-based models of language processing (e.g., MacDonald et al., 1994, Trueswell and Tanenhaus, 1994; Tanenhaus and Trueswell, 1995). I will suggest that probabilistic information about discourse patterns is a characteristic that gets stored with abstract linguistic representations, and thus becomes available for use during language processing.

Throughout this dissertation, I assume that choices in reference form are influenced by multiple constraints. I am only investigating five factors here, but there are clearly other factors that are relevant to reference form. Some other factors are linguistic, such as the ambiguity of potential referring forms. Others may be nonlinguistic, such as the prominence of visual information available to all discourse participants. My approach shares features with Accessibility Theory (Ariel, 1990), which also suggests that multiple factors influence the choice of reference form. The primary difference between my work and other studies is that I will show that these five linguistic features share the characteristic that they are frequently associated with subsequent reference.

Chapter 2

Cross-linguistic evidence for multiple constraints on reference form

The framework for this dissertation, presented in chapter 1, assumes that choices in reference form depend on multiple factors. What are these factors? Are all of them necessary? It is costly to investigate many different factors with psycholinguistic experiments, and any attempt to do so risks an uninterpretable outcome. Text analysis is useful as an exploratory tool, since it affords the simultaneous investigation of multiple structures.

A second advantage to text analysis is that it allows the comparison of several languages, as I will do here by investigating English, Spanish, and Mapudungun. This methodology is particularly useful for looking at Spanish and Mapudungun, which would be more difficult to study experimentally. Neither language has a long tradition of experimental studies to provide the foundation for experimentation. Furthermore, in the

United States it would be virtually impossible to find native speakers of Mapudungun, which is spoken by the Mapuche people of southern Chile and Argentina.

Throughout this chapter, I will demonstrate in detail the correlation of multiple factors with reference form. The constraints I will discuss stem from very different sources of information, but the one thing that they have in common is that they can be interpreted in terms of their influence on the reader's mental representation of referents in the text.

The main findings of these studies will show that reference form in all three languages is correlated with three of the factors described in chapter 1: a) Recency, b) Subjecthood, and c) Parallelism. I will also show that subsequent reference is more likely to occur a) when the referent was mentioned recently, b) when the referent was last mentioned in subject position, and c) when the anaphor and antecedent are in parallel positions. These findings together will support the claim I set out in chapter 1, that these three factors are associated with both Reference Form and Reference Continuation.

2.1 About the text analysis methodology

These text analyses follow in a tradition of similar analyses on reference form and discourse structure. Such analyses are particularly popular among functional linguists, and have been conducted on texts in a number of different languages with diverse structures. For example, Du Bois (1985, 1987) and others working within the theory of Preferred Argument Structure have shown for many languages that grammatical functions are specialized in terms of pragmatic function. Givón and his colleagues (1983b) demonstrated that reference form correlates cross-linguistically with three features of "Topic Continuity": a) the recency with which an entity has been mentioned in the text, b) the level of competition (i.e., how many entities could conceivably be the referent for a particular form), and c) the referent's level of "persistence" - that is, how long the entity persists throughout the following discourse, as a measure of how important the speaker considers the entity to

be. The constraints of competition and recency have also been demonstrated for Hebrew and English by Ariel (1988, 1995).[1]

The role of text analysis in the study of language has grown over the last two decades, in part due to improved technologies that can mechanize some aspects of the analysis; this has made possible the investigation of large, multi-million-word corpora (Gibson et al., 1994). The terms "text analysis" and "corpus analysis" are both used to refer to distributional studies of texts, using spoken or written data. I will use the term "text analysis" to refer to the analysis of a narrative text as a whole, where nearly every reference is coded and analyzed. In chapters 3 and 4, I will also present what I call "corpus analysis", namely, the study of a selected sample of utterances from a much larger corpus.

Both text analysis and corpus analysis have been used for multiple purposes. For example, the correlation between structure and function, even if it is not one-to-one, is typically taken as an indication that the structure is used for that function. Since pronouns frequently refer to the subject-referent of the previous clause, Subjecthood is assumed to influence the use of pronominal reference. As a result, distributional frequencies are often assumed to be a reflection of the psycholinguistic processes that underlie language use. This type of evidence has been brought to bear in many studies of reference and discourse, such as Du Bois (1987), Givón (1983b), Ariel (1988), or other studies such as Hawkins (1994), Siewierska (1993), and Wasow (1997).

A second function of text or corpus analysis is as a measure of the frequencies of particular structures or meanings. Estimations of "real" frequency information (where "real" means the frequency of use that an individual might encounter) are important because recent models of language processing have claimed that listeners make use of frequency during language processing (e.g., Bates and MacWhinney, 1989; MacDonald et al., 1994; Trueswell and Tanenhaus, 1994). For example, some ambiguous words, such as "ball", have one meaning that is more frequent than other meanings (in this case, the "ball as round

---

[1] For more information about these constraints, see the discussion in chapter 1.

object" meaning is more frequent than the "ball as a formal dance party" meaning). The frequency of each meaning affects lexical processing, in that it is easier to access the more frequent meaning than the less frequent ones (e.g., Rayner and Frazier 1989). In addition, some syntactically ambiguous structures are resolved more frequently with one meaning than another. In sentences such as 1a or b, the "with-" phrase is temporarily ambiguous, and could modify either the verb or the object NP.

(1) a. The woman expected the bus with air-conditioning but was at the wrong stop.

b. The woman expected the bus with anticipation but was at the wrong stop.

(Spivey-Knowlton and Sedivy, 1995)

Since prepositional phrases more frequently modify the verb phrase than the object NP, verb-modifying PPs are usually easier to understand than object-modifying PPs (Spivey-Knowlton and Sedivy, 1995). In order to test claims about how frequency of word meaning or syntactic structure affects processing, however, researchers need independent measures of frequency. In this regard corpus analysis plays an important role, since it provides one estimate of the structural and lexical frequencies that people experience. This purpose for corpus analysis will surface again in chapters 3 and 4.

The text analyses in this chapter fulfill both the functions to some degree. The primary goal was to investigate the correlations between certain variables and reference form. I hypothesized that any correlations between reference form and different factors would indicate factors relevant to the production and comprehension of reference form. I investigated texts in several languages to make possible a cross-linguistic comparison of these factors. Second, I used these text analyses to provide supporting evidence for the primary theme of this dissertation. That is, I looked for general patterns of reference. Which types of referents are mentioned more often? Are certain factors also related to the

probability that certain referents would be continued in the following discourse? That is, are there factors that correlate with Reference Continuation?

After describing the particular methods used in this study, I will present the results for all three analyses.

2.2 Methods

In order to investigate reference form in English, Spanish, and Mapudungun, I performed three text analyses, one in each language. The corpora used for these analyses were collections of written narrative texts that contained mostly third-person referents. The few references to first or second person were excluded from the analyses. The reason for focusing on third-person referents was that I am primarily interested in how linguistic, textual characteristics affect reference form. Discourse structure and reference form can also be affected by nonlinguistic features such as eye gaze, physical location or context, and the identity of the discourse participants. By looking only at written texts and only third-person referents, however, I have limited the analysis to cases where the reference form is primarily affected by linguistic information.

The texts chosen for this study were all narratives, primarily traditional tales or children's stories. These narratives are ideal for this kind of study because the clause structure is typically simpler than in other written genres, and most instances of reference concern concrete people and objects, rather than complex abstract ideas. These properties simplified the analysis. Another reason for using this genre was that there is not a wide variety of published texts in Mapudungun, but there is a sizable collection of traditional tales. In all cases, the texts were published in written form. Some of the Mapuche tales were originally told orally, and published in transcription. However, these texts appear with fully formed sentences with no disfluencies, indicating that the oral narratives were edited before publication, and for this reason fall into the category of "written text". The narratives used for each language are listed in Table 1.

Table 1. Texts used

STORY	# clauses[1]	TYPE
English (1378 clauses total)		
Peter Pan, chapter 2 (pp. 25-29)	174	- classic children's story
The Wonderful Wizard of Oz, chapter 1	161	- classic children's story
Georgie the Gentle Ghost	160	- children's narrative tale
Sylvester and the Magic Pebble	228	- children's narrative tale
The Terrible Hodag[2]	270	- traditional tale (Wisconsin)
Spanish: (524 clauses total)		
Caipora, el Padremonte[3]	121	- traditional tale (Mexico)
El Irupe	62	- traditional tale (Argentina)
El Palto del Anima *(The Widow's Avocado Tree)*	53	- traditional tale (Chile)
Las Lágrimas de Potira *(Potira's tears)*	60	- traditional tale (Brazil)
María Tolete	117	- traditional tale (Venezuela)
Las Lágrimas del Sombrerón *(The tears of the Sombrerón)*	111	- traditional tale (Guatemala)
Mapudungun: (731 clauses total)[3]		
Füta peñi ka kude pangküll *(Big Brother and Female Lion)*	72	- traditional Mapuche tale
Kuse pangküll ngürü engu *(The old female puma and the fox)*	177	- traditional Mapuche tale
Kiñe wentru l̲an̲turkei *(A widowed man)*	144	- traditional Mapuche tale
Missionary	60	- traditional Mapuche tale
An Old Man	98	- traditional Mapuche tale
Kiñe Wentru Nierkey Kiñe Kon̲a *(A man had a servant)*	180	- traditional Mapuche tale

[1] The number of clauses in each database includes quotes and other clauses which were excluded from the final analysis. Quotes count as a single clause, regardless of the length of the quote.

[2] It is interesting to note that "The Terrible Hodag" and "Caipora, el padremonte" have similar themes, even though they are in different languages and from distant cultures. Both stories are about lumbermen who go into the forest and encounter mythical beasts of the forest who help only the good men, the ones who treat the forest well.

[3] All examples in Mapudungun will be written in the Unified Alphabet, one of the standard alphabets used in Chile for Mapudungun. Most of the characters are the same as the IPA, except for the following.

Unified Alphabet	IPA
ü	ø
ll	ʎ
ng	N
d	Q
q	χ

Underlined characters (n̲, l̲, t̲) denote dentals (e.g., n̲ denotes a dental nasal).

The analyses in each language followed similar procedures, although the analysis for English, which had the largest database, was also more detailed than the other two. The Spanish analysis was also slightly more detailed than that for Mapudungun. I will first describe the methods for analyzing English, and then briefly describe the differences in the Spanish and Mapudungun analyses.

2.2.1. English text analysis methodology

The purpose of this study was to analyze how different forms of reference were used, according to different properties of the referent. Therefore, the first thing I looked at was the form of each referring NP. I was particularly interested in the choice between pronouns and full names or descriptions. Other forms were sometimes used (e.g., reflexive, elliptical, null, or possessive), but these were excluded from the final analysis. The full set of codings for NP form is listed in Table 2.

Table 2. Codings for NP form

CODING	DESCRIPTION	EXAMPLES
full NP	names or descriptions	Sylvester, the woman, an amazing pebble
pronouns	personal pronouns and demonstratives	it, she, he, they, him, her, them, this, that
reflexive	reflexive pronouns	himself, herself, itself, themselves
elliptical	a deleted argument in a conjoined phrase	Aunt Em dropped her work and Ø came to the door.
Ø	a null argument that's controlled by another argument.	• He ran about the room, now here, now there, Ø barking loudly. • Each day, he and the other lumberjacks went into the forest Ø to cut down trees.
possessive (full NP or name)		Dorothy's eyes
possessive (pronominal)		her eyes

Each instance of reference was also coded for a number of additional properties, including a) the grammatical function of the anaphor (i.e., the NP itself), b) the grammatical function of the antecedent (i.e., the grammatical function of the NP with which the conceptual referent was last mentioned) if it was in the previous clause, and c) the recency with which the referent was last mentioned. A full list of these codings is presented in Table 3.

Table 3. Codings for additional properties

A. Grammatical Function of Anaphor

CODING	EXAMPLES
Subject	They took the rockers from the Whittakers' porch . . .
Object	They took the rockers from the Whittakers' porch . . .
Oblique	They took the rockers from the Whittakers' porch . . .

B. For references to something last mentioned in the previous clause (or last main clause): What was the Grammatical Function of the antecedent?

CODING	EXAMPLES
Subject-referent	But **Georgie** never scared anybody. He was much too shy for that.
Object-referent	They took **the rockers** from the Whittakers' porch because they were old and antique.
Oblique referent	They belonged to **a huge creature**. It had the head of an ox, the feet of a bear, the back of a dinosaur, and the tail of an alligator.

C. Recency of last reference to the same entity

CODING	*DESCRIPTION*/ EXAMPLES
same clause	She took off her hat.
1 clause back	Through the branches, he saw a giant red eye. Then he saw another red eye.
last main clause[1]	*The referent was not in the previous clause, but only subordinate clauses intervene. Example:* "Timber!" they shouted as each tree crashed to the ground. The men worked hard from sunup to sundown.
1 clause plus quote	Ole Swenson told them. "Who is he?" asked the lumberjacks.
2 clauses back	As he got close to the Hodag's den he felt the ground shake. Suddenly he saw the giant beast...
2 clauses plus quote	
3 clauses back	
3 clauses plus quote	
4 clauses back	
4 clauses plus quote	
5 clauses back	
5 clauses plus quote	
6 or more clauses back	
new	*hasn't been mentioned before*
generics	*Note: Generics were combined with "new" for the analysis, because they do not have a true antecedent in the discourse. Examples:* To his great surprise the rain stopped. It didn't stop gradually, as rains usually do. They sniffed the rock on Strawberry Hill, but it smelled like a rock.

[1] The inclusion of the category "last main clause" meant that some references had two possible codings for Recency. For example, the referent of "he" in d was last referred to as the null subject of the preceding clause, and also as the subject of the last main clause, in a.

a. He was thinking so hard
b. he forgot
c. Ø to watch out for the Hodag.
d. Suddenly he looked up.

In cases like these, main clauses always took precedence over subordinate clauses, so this one was coded as "last main clause".

EXCLUDED: The analysis of Recency did not include cases where an anaphor referred to the entire previous clause, a superset of previously mentioned entities, or a subset of a previously mentioned referent.

the entire previous clause	All day they cut the branches off the trees and put logs into piles. It was hard work.
superset	But Georgie and Miss Oliver went to be with the cow in the meadow. When they got to the barn, they all knew...
subset	. . .to cut down trees. . . . . as each tree crashed to the ground.

For the purposes of the analysis, a clause was defined as the maximal phrase headed by a VP (under some analyses of syntactic structure). A clause could be finite or nonfinite, matrix or subordinate. The grammatical functions of noun phrases were determined according to whether the NP was an argument of the verb or not, and if so, which position it filled. Each clause was also identified in terms of the type of clause, according to the distinctions in Table 4.

Table 4. Codings for type of clause

CODING	EXAMPLES
main	But when he said the same thing, holding the pebble in his hoof, the sky turned black.
subordinate-finite	But when he said the same thing, holding the pebble in his hoof, the sky turned black.
subordinate-nonfinite	But when he said the same thing, holding the pebble in his hoof, the sky turned black.

Some NPs were excluded from analysis. Predicate nominatives, as in 2, below, and appositives, as in 3, were excluded, since these NPs were not used referentially.

(2) Sylvester was a rock.

(3) And there was Sylvester, a rock on Strawberry Hill, with the magic pebble lying right beside him on the ground, . . .

Some NPs were excluded from the analysis because they did not form a part of a clause, as in 4b.

(4) a. while Mrs. Duncan set out the picnic food on the rock --

b. alfalfa sandwiches, pickled oats, sassafras salad, timothy compote.

The dummy subjects of atmospheric verbs, such as "it rained" or "it was cold" were also excluded, as were noun phrases that referred to time, such as "yesterday", or "three hours later". In addition, reference to first or second person (I, you, we) was excluded, for the reasons mentioned above.

When deciding how recently an entity was mentioned, I included cases where an entity was previously referred to with a different noun phrase (in contrast with Brown, 1983:319, who excluded these cases, which he called "synonyms"). 5 illustrates an example of anaphoric reference with a different NP.

(5) As he got close to the Hodag's den
he felt the ground shake.
Suddenly he saw the giant beast.

The decision to include these cases follows logically from the view that reference is a cognitive event. Although my focus is on linguistic reference, the linguistic form is interpreted relative to the listener's representation of the discourse. The interesting thing about reference is that it can take place with a wide variety of forms, for example using "Clinton", "Bill", "he", or "the president" to refer to the same person. Since the purpose of this study is to understand which factors influence the choice among referring forms, it follows that NP anaphors, as well as pronouns, should be included as one of the many choices for reference.

2.2.2. Spanish text analysis methodology

The text analysis of Spanish was similar to that of English, except in the following ways. The system of reference in Spanish differs from that in English, in that there are at least three possible forms of reference: full noun phrase, pronoun, and null ("pro-dropped") subjects or cliticized objects. These three forms of reference offer increasingly less information -- the full noun phrase or name provides specific information as to the identity of the referent, and pronouns provide person, number and gender information (e.g., third singular female). In the case of null subjects, the morphological marking on the verb indicates the person and number of the verb,[1] such as first singular (1sg), or third plural (3pl). Cliticized objects are categorized with null subjects because both forms represent the least marked form of reference that is commonly used. In the case of indirect objects, clitics also provide only person and number information, although direct object clitics additionally provide gender information.

[1] In some tenses and moods, the morphology is ambiguous for person. For example, in the present subjunctive, the same forms are used for 1st and 3rd person singular, and in all tenses and moods the same forms are used for 3rd person singular and the formal 2nd person singular (usted).

(6) Full Noun Phrase (denotes a specific entity)

a. Subject- El gato miró al ratón.
The cat looked-at(3sg) PREP-the rat.

b. Object- El gato miró al ratón.
The cat looked-at(3sg) PREP-the rat.

(7) Pronouns (person, number, and gender)

a. Subject- El miró al ratón.
He looked-at(3sg) PREP-the rat.

b. Object- El gato miró a él.
The cat looked-at(3sg) PREP him.

(8) Null/ Clitic (person, number and occasionally gender)

a. Subject- Ø miró al ratón.
Ø looked-at(3sg) PREP-the rat.

b. Object- El gato lo miró.
The cat him looked-at(3sg).

Because I am interested in the choice among these three forms of reference, only subjects and objects (both direct and indirect) were included in the study. Although oblique arguments can vary between full NPs and pronouns, null or clitic reference types are not

possible with oblique arguments[1].

Apart from the coding of the referring form, the Spanish text analysis was similar to the English analysis, but less detailed. One difference was that recency was only coded according to the four categories in 9.

(9) Categories of "Recency" for Spanish

a. previous clause

b. 2-5 clauses back

c. 6 or more clauses back

d. new

In addition, quotations were discounted when ascertaining recency, because the size of the Spanish corpus was too small to investigate subtle differences like the effect of quotations on recency. This simplification undoubtedly leaves some variation unaccounted for, since the investigation of quotes in English demonstrates that quotes do contribute to recency effects on reference form (see §2.3.1.4 below).

A second difference between the English and Spanish methodologies was that in Spanish, there was no category "last main clause". This difference is only relevant in cases like 10c, where a subordinate clause intervenes between an anaphor and its antecedent.

---

[1] Note that double-object verbs in Spanish have two objects: a direct object (theme) and an indirect object (goal). Both objects can be expressed as full NPs, pronouns, or clitics. An example is:

Abrió	el	morral,	sacó	un	atado	de tabaco
Opened(3s)	*the*	*pack,*	*took-out(3sg)*	*a*	*bundle of tobacco*	

y	se	lo	alcanzó.
and	*3sg(goal)*	*3sg(theme)*	*reached(3s).*

He opened the pack, took out a bundle of tobacco, and handed it to him.

(10) a. Cada mañana, muy temprano, dos compadres iban juntos al monte

Every morning, very early, two companions went together to the mountain peak

b. a cortar leña.

to cut wood.

c. El monte era una belleza.

The mountain peak was a beauty.

In the Spanish analysis, the last reference to the concept "el monte" in 10c is coded as occurring 2 clause back. In contrast, in the English analysis it would have been coded as "oblique of last main clause".

2.2.3. Mapudungun text analysis methodology

The Mapudungun analysis was similar to the Spanish analysis, except that it only included main clauses. There are several types of subordinate clauses in Mapudungun, some of which allow specification of all arguments, and other types of clauses where the subject argument is controlled by a higher clause. Since this study investigated the variation between null referents and full NPs, I did not want to include cases where a full NP argument was disallowed. It was beyond the scope of this study to identify which subordinate clauses allowed explicit subject arguments and which did not, so only references that occurred in matrix clauses were analyzed.

However, references in subordinate clauses were still included for the purpose of analyzing the last mention of a referent, as in 11.

(11) a. dew nie-lu mapu
Once have-verbal.noun land
After he got his land,

b. dewma-y ñi ruka
make-3s poss(3s) house
He built his house.

In this example, the clause in 11a is nonfinite, so the references corresponding to "he" and "land" are not included in the analysis. However, the clause in 11b contains a reference to "he", the person last mentioned in the previous clause. The information in the subordinate clause is relevant for this purpose, and the reference to "he" in 11b is coded as having the last reference as subject of the previous clause.

Mapudungun is a language with a rich set of verbal suffixes that mark, among other things, person and number for both subject (obligatorily) and object (optionally). This system of person/number marking is one of the features that allows Mapudungun to have both subjects and objects that are null, just like subjects in Spanish. Pronouns also exist in Mapudungun, but are used only for subjects. As in the Spanish analysis, I only investigated references in subject and object position, since oblique arguments occur only as full NPs. Thus, in Mapudungun I looked at the form of reference, which could be a null referent, a pronoun (for subject only), or a full NP.

The most salient feature of Mapudungun is that it displays an inverse system, which means that there are two sets of person/number morphemes for transitive verbs, with the choice between these two determined by the relative prominence of the arguments. The generalization is that the more prominent argument must be the grammatical subject at all times, where prominence is determined by means of the hierarchy in 12.

(12) Prominence Hierarchy for Mapudungun:

1st person, 2nd person > 3rd person proximate > 3rd person obviative

The term "proximate" simply refers to the argument that is more salient in the discourse, and "obviative" refers to the argument that is less salient. This distinction is not morphologically marked, nor does it depend on animacy or other features. Instead, the choice between proximate and obviative is gleaned from the discourse structure. A complete definition of how the more salient argument is chosen is beyond the scope of this study, but roughly speaking, it can be said that choice between direct and inverse form is sensitive to the same factors that influence reference form: the more salient character must have appeared previously in the discourse, usually recently. In contrast, the less salient character is often new or has not appeared recently in the discourse. In addition, global discourse factors affect relative ranking, so the main character of a story is usually ranked higher than other characters. However, this generalization can be modified if a secondary character becomes locally prominent during certain discourse segments (for an example, see Rivano-Fischer 1991:126-127). Although higher-ranked characters tend to be animate, it is also possible for inanimate entities to be ranked higher than animates (see Arnold, 1996).

The relative ranking of two arguments in a transitive clause determines whether the verb will appear in the **direct** voice, or the **inverse** voice. In most cases the verbal form is pre-determined, since 1st and 2nd person are **always** considered “more salient” than 3rd persons, and thus must always be the grammatical subject. This means that whenever one chooses to express the proposition “She saw me” (for example), the inverse construction must be used. To say “I saw her,” on the other hand, the direct form of the verb is obligatory. With two 3rd-person arguments, the verbal form depends on their relative salience. The distribution of the direct and inverse forms in Mapudungun is shown in 13.

(13)

Direct:		Inverse:[1]	
Actor	Undergoer	Actor	Undergoer
1	3	3	1
2	3	3	2
3 prox	3 obv	3 obv	3 prox
		1	2
		2	1

For example, if I wished to say "I called my mother", I would use the direct form:

(14)	Mutrüm-fi-n	ñi	ñuke
	Call-OBJ(3)-SUBJ(1s)	POSS	mother
	I called my mother.		

On the other hand, to say "My mother called me," I would need to use the inverse form:

(15)	Mutrüm-e-n-ew	ñi	ñuke
	Call-INV-SUBJ(1s)-OBJ(3)	POSS	mother
	My mother called me.		

In both cases "I" is the subject, and "my mother" is the object.

The inverse system in Mapudungun affects the current text analysis in three ways. First, it is doubly important to exclude 1st and 2nd person references from the analysis (or analyze them separately, which was beyond the scope of this study), since these references are not only limited to null or pronominal forms, but are also required to be the subject of a transitive event. These grammatical restrictions are likely to interact with the pragmatic

features of interest to this study and obscure their effect. Second, since the relative prominence of two arguments affects the whole clause, it was necessary to exclude not just 1st or 2nd person references, but the entire clause they appeared in. Third, the grammaticization of pragmatic factors for grammatical functions means that generalizations about the role of Subjecthood in choosing reference forms are expected to be heightened in Mapudungun, since the grammatical function itself is defined by many of the same pragmatic considerations that underlie variation in reference form.

2.3. Results and Discussion

Three factors emerged from my analyses. The first was the recency with which an entity has been mentioned. While some researchers have discussed recency in terms of the distance an anaphor must travel to find a textual antecedent, a more plausible interpretation of recency is as a measure of how activated or accessible the referent is in the minds of the discourse participants. In §2.3.1, I will show how recency of mention correlates with reference form in all three languages, and develop a detailed picture of various aspects of recency in English.

The second and third factors reflect structural properties of the referent and referring form. In §2.3.2, I will demonstrate how the grammatical function of the anaphor and the relative positions of the anaphor and antecedent are correlated with choices in reference form. Although these factors are often confounded in experimental studies, this analysis considers the effect of each one.

2.3.1. Recency

As mentioned in chapter 1, scholars from many traditions have studied reference form and described it in terms of the cognitive status of the referent. It is acknowledged that pronouns are interpretable only when their referent is retrievable based on context, which usually means that the referent is present in the listener's focus of attention. But what

---

[1] Interactions between 1st and 2nd person also fall into the category "Inverse", although their morphology

causes something to be in the focus of attention? One factor is the recency with which the referent was last mentioned.

The results from the analyses of English, Spanish, and Mapudungun are shown in Table 5. In this table, as in the rest of the results section, the results about reference form are presented in terms of the percentage of pronouns or null anaphors. The reason for comparing pronouns (in English) and null anaphors (in Spanish or Mapudungun) is that these are the least-specified forms of reference for each language. In Spanish and Mapudungun, as in other languages with null reference (e.g., Sacapultec, Du Bois (1987:814)), overt pronouns are rarely used. The percentages for each case are calculated over the entire set of references for each category. In English, the entire set of references includes pronouns and full names or descriptions. In Spanish and Mapudungun, the entire set of references includes null anaphors (i.e., only verbal morphology), pronouns, and full names or descriptions.

Table 5. Correlations between reference form and recency of mention.

A. ENGLISH

χ^2=616, DF=3, p<.001

	n	%pronoun	
previous clause/ last main clause	336	lllllllllllllllllllllllllllllllllllllllllllllllllllllllllllll lllllllllllllllllllllllllll	86%
2-5 clauses back	174	lllllllllllllllllllllllllllllllll	32%
6+ clauses back	197	l	1%
new	257		0%

---

is slightly different from other forms. For a fuller discussion, see Arnold (1994).

B. SPANISH

χ^2=394, DF=3, p<.001)

	n	%null	
previous clause	228	IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII IIIIIIIIIIIIIIIIIIIIII	79%
2-5 clauses back	111	IIIIIIIIIIIIIIIIIIIIIIIIIIII	28%
6+ clauses back	107	III	3%
new	246		0%

C. MAPUDUNGUN

χ^2=278, DF=3, p<.001)

	n	%null	
previous clause	393	IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII IIIIIIIIIIIIIIIIIIIII	78%
2-5 clauses back	93	IIIIIIIIIIIIIIIIIIIIII	22%
6+ clauses back	22	IIIII	5%
new	111		0%

These results support previous claims (e.g., Du Bois, 1987; Ariel, 1988; Givón, 1983a) that reference to entities which have recently appeared in the discourse is more likely to be null or pronominal. This correlation between recency and form of reference supports hypotheses that entities that have occurred recently in the discourse are more accessible, and more in the focus of attention. From a processing perspective, their accessibility can be considered in terms of having more highly activated mental representations in the minds of the discourse participants.

2.3.1.1. The choice of categories in text analysis: a methodological issue

One striking aspect of the analyses for English and Spanish is that most of the variation in reference form occurs when the referent has been mentioned within the last five clauses. In English, there is a only one instance of a pronoun referring to something which has not been mentioned for more than 5 clauses, shown in 16.

(16) Dorothy felt as if she were going up in a balloon. The north and south winds met where the house stood, and made it the exact center of the cyclone. In the middle of a cyclone the air is generally still, but the great pressure of the wind on every side of the house raised it up higher and higher, until it was at the very top of the cyclone; and there it remained and was carried miles and miles away as easily as you could carry a feather. It was very dark, and the wind howled horribly around her, but Dorothy found she was riding quite easily.

(Baum, The Wonderful Wizard of Oz, 1900:15)

However, even this example may be explained away by claiming that this usage is in fact an example of cataphora, referring forward to "Dorothy". There are also several properties of this example that allow a post-hoc explanation of why such long-distance reference is possible: Dorothy is the main character of the narrative, and in that sense may be said to be more accessible (Morrow et al., 1989). Also, all intervening referents are inanimate, so the pronoun is not ambiguous in this context. Furthermore, this entire section is about the same event, and the intervening material provides a description which is implicitly from Dorothy's perspective.

In any case, as Table 5 shows, the majority of pronominal references occur when the referent has been mentioned within the previous five clauses. This suggests that researchers should look here for the relevant correlations between reference form and recency of mention, and therefore several recent studies of reference have clearly looked in the wrong

places. For example, researchers in Centering Theory have claimed that pronominal reference depends on how an entity was last mentioned, but they consider only the previous clause, both theoretically (Grosz et al., 1995) and experimentally (Hudson D'Zmura and Tanenhaus, 1998; Gordon et al., 1993). But, as Table 5 shows, this limitation misses the difference between referents mentioned in the previous clause and referents mentioned before that.

While it is possible to focus too narrowly, as Centering does, it is also possible to focus too broadly, and miss much of the interesting detail. This is the case in Du Bois (1987), who followed Chafe's tripartite system of coding reference as either "new" (= it had not appeared in the text), "accessible" (= it had appeared but not within the last 20 clauses) or "given" (= it had appeared within the last 20 clauses) (see Chafe, 1994, chapter 6, for an updated description). By distinguishing only between entities that had been mentioned within the last 20 clauses and those which had not been mentioned for 20 clauses, Du Bois missed the fact that there is a large, graded effect of recency for referents mentioned within a space of 5 clauses or so, but not for referents that haven't been mentioned for more than 5 clauses.

2.3.1.2. A closer look at English

The data in Table 5 showed that pronouns were almost never used in English when the referent had not been mentioned for more than 5 clauses. But what patterns of pronoun use occur for referents mentioned 2-5 clauses back? A closer look at the English data reveals that the percentage of pronominal references falls off steadily as the distance to the last mention increases, as shown in Figure 1.[1]

---

[1] The data in Figure 1 exclude all cases of reference where there is a quote intervening. This issue will be discussed in detail in §2.3.1.3. This analysis contrasts with the earlier analysis of Recency in all languages, which simply ignored quotes for purposes of measuring Recency.

As in the previous analysis, the category "one clause since last mention of referent" includes the category "last main clause", in which the referent was mentioned in the last main clause, with one or more subordinate clauses intervening.

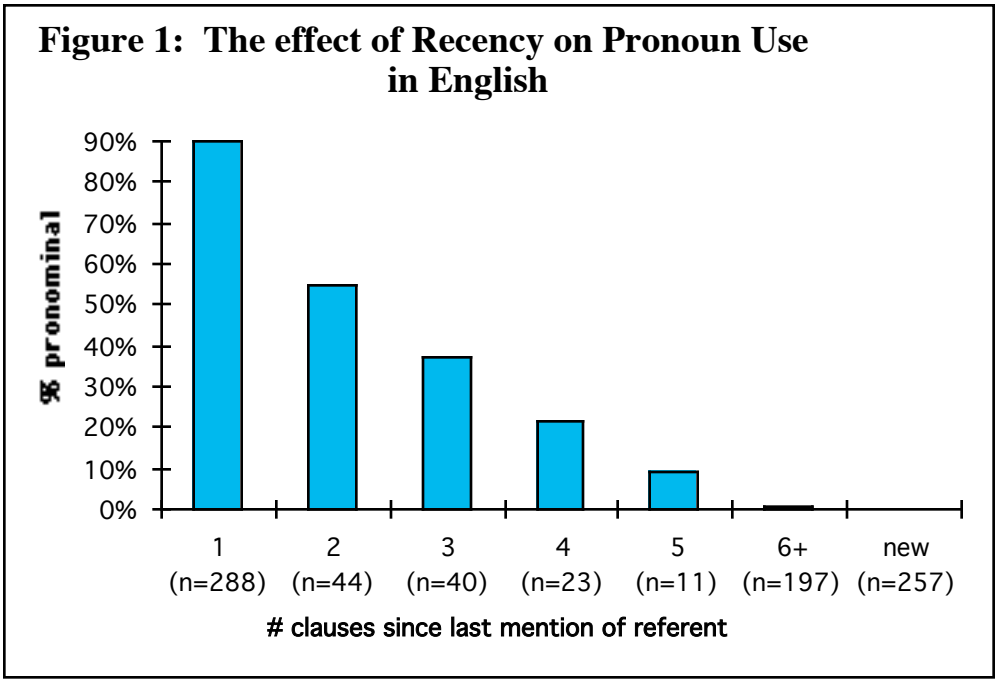

As time passes, pronouns are used less and less, and full names or descriptions are used more and more. Since recency of mention can be assumed to affect the activation or accessibility of the representation of an entity, these data provide a detailed picture of how reference is made to entities of differing activation.

This picture is consistent with the data on written English presented by Brown (1983) in Givón's volume on Topic Continuity. However, the methodology of my study differed from that in Brown's. Brown, along with the other authors in Givón (1983b), used recency as a dependent, rather than an independent variable, and investigated the distribution of recency for each type of referring form. In contrast, I counted the referential form as the dependent variable, investigating the variation in reference form under different conditions, like recency.

There are three reasons why my methodology is preferable. First, it is more natural than that of Givón and his colleagues (1983b), since it reflects the theoretical direction of causality: production of a referring form occurs in the context of information such as how

recently an entity had been mentioned. Therefore it is reasonable to believe that recency of mention affects the form of reference, and not that the form of reference affects recency of mention.

Second, the choice of methodology affects the information it reveals. For example, Brown reported that in his corpus, the average "referential distance" was 1.72 for unstressed pronouns and 2.27 clauses for demonstrative pronouns. From this we as readers know that the antecedent of most pronouns was within 1 or 2 clauses, but we don't know what the maximum referential distance is, nor what the preferred form of reference is in a particular situation. In contrast, the data in Figure 1 tell us that reference to something last mentioned in the previous clause is overwhelmingly pronominal, and reference to something that hasn't been mentioned for 6 clauses is almost always a full name or description.

The third advantage of my methodology is that it demonstrates the graded nature of recency. In many text analyses that investigate recency, most authors have distinguished between just a few categories of recency. The most basic contrast is that of "given" (i.e., it has been mentioned before) versus "new" (i.e., it has not been mentioned in this discourse). Most authors contrast at least three levels of recency, for example Chafe's system of "given" vs. "accessible" vs. "new" (1994). This system might suggest that there are three discrete cognitive categories for discourse referents. However, my analysis suggests that the effect of recency is graded.

2.3.1.3. The privileged status of the previous clause

Even though recency displays graded properties, the results in Figure 1 show that the relationship between recency and reference form is not completely linear. When recency is measured in terms of number of clauses, the percentage of pronominal reference decreases at a slower and slower rate as the recency of the last mention decreases. The greatest difference occurs between one clause and two clauses back, suggesting that there is something special about the cognitive status of information in the previous clause. This observation is consistent with Clark and Sengul's (1979) argument that entities mentioned in

the previous clause have a privileged position. They found in their first experiment that people read passages faster when the referents of pronouns or definite NPs were mentioned in the previous clause, but found no reliable difference between reading times for referents mentioned 2 or 3 clauses back.

Clark and Sengul also found evidence that the privileged position was the previous clause, rather than the previous sentence. Their second experiment showed that when the referent had been mentioned two clauses prior to the anaphor, there was no difference whether that clause was part of the previous sentence or not. Furthermore, their third experiment showed that entities in the previous clause were more accessible than other entities, regardless of whether they were in a main or subordinate clause. This supports Givón's claim that the clause is the basic unit of processing (1983a:7).

The results of my text analysis provide further support for the claim that clause type has little or no effect on the privileged position of the previous clause. I looked at reference to entities in the previous clause, and compared cases where the previous clause was a main clause with cases where it was a subordinate clause. Table 6 shows that when the referent was last mentioned in the previous clause, pronouns were used about 90% of the time, with no significant difference between main and subordinate clauses ($\chi^2(1)=2.8$, $p=.1$).

Table 6. Referents in main vs. subordinate clauses

	total	% pro	
previous clause (main)	194	90%	] n.s.
previous clause (subordinate)	60	97%	

While there is no difference between these two categories, together they have a higher rate of pronoun use than the next category, "last main clause", described below ($\chi^2(1)=4.7$, $p<.02$).

At the same time, there is evidence that main and subordinate clauses have different influences on the accessibility of referents that were mentioned **before** the previous clause. To investigate this difference, I compared two categories of anaphoric reference. First, I looked at cases where the referent was not mentioned in the previous clause, but was mentioned in the last main clause. 17 illustrates the "last main clause" type of reference, where the only intervening clause is a subordinate one.

(17) "Timber!" they shouted [as each tree crashed to the ground]. The men worked hard from sunup to sundown.

In some cases there was more than one clause intervening, but all were subordinate. I compared this category with cases where the referent had not been mentioned for two clauses, and the intervening clause was a main clause. The data in Table 7 show that when a main clause intervenes between an anaphor and its antecedent ("two clauses back"), pronoun use is lower than when one or more subordinate clauses intervene ("last main clause") ($\chi^2(1)=5.24$, $p<.02$).

Table 7. Intervening clauses: main vs. subordinate clauses.

	total	% pro	
last main clause	34	79%	] p<.02
two clauses back	44	55%	

It is clear from these data that recency does not affect reference form in a linear fashion: main clauses produce more interference with the accessibility of referents than subordinate clauses.

One explanation for this finding is that information in subordinate clauses is interpreted as supporting the propositions in the main clause. Thus, a new main clause forces listeners to shift their focus of attention to a new topic, whereas subordinate clauses

maintain the focus of attention of the previous clause. When the focus of attention remains on the information in the main clause, it is easier for the listener to access that information. An alternative explanation for the data in Table 7 could be that by mentioning something in a main clause, the speaker or writer indicates that it will play an important role in the upcoming discourse. This could result in higher activation for the referents in the readers' mental representation, which would interfere with information from preceding clauses. In contrast, information from subordinate clauses may be less activated and interfere less. However, if this were the case, one would also expect more pronoun use for reference to entities a preceding main clause than entities in a preceding subordinate clause. But this was not the case in the English text analysis, as Table 6 shows.

2.3.1.4. The effect of quoted utterances

With respect to recency, an additional question whether quotations interfere with the accessibility of referents as much as nonquoted utterances. According to some theories, discourse is structured in a hierarchical fashion, with some segments embedded inside others (e.g., Grosz and Sidner, 1986; Polanyi, 1994; and Fox, 1987). This view is formalized in Grosz and Sidner's (1986) description of Centering Theory, which consists of a) the discourse structure, b) the intentional structure, and c) the attentional structure. The discourse structure is composed of hierarchically organized discourse segments that correspond to discourse segment purposes in the intentional structure, and focus spaces in the attentional structure.

The attentional structure in Centering Theory is a model of the speaker and listener's focus of attention at each stage of the discourse. It is linked to the intentional structure: at any point in the discourse, the focus of attention is on referents that are relevant to the current intention. The intentional and attentional structures are defined by means of a focus stack, which consists of focus spaces for different segments of the discourse. The focus space is defined as that which constrains referring expressions like definite NPs, and limits the listener's search for referents for anaphoric elements (such as pronouns). The focus

space that is on the top of the stack at any given moment constrains reference in the current discourse. When the discourse enters an embedded segment, a new focus space is layered on top of the one for superordinate discourse segments. When an embedded segment ends, the focus space is "popped" off the stack, leaving the superordinate focus space at the top. The claim is that returning to the higher focus space allows pronominal reference to something that is hierarchically recent but which has not been mentioned for some time. An example of how pronominal reference can "ignore" an embedded structure is shown in 18.

(18) An excerpt from a corpus of financial advice dialogues (Pollack, Hirschberg, and Webber, 1982, as cited in Walker, 1998).

a. C: Okay Harry, I have a problem that uh my - with today's economy *my daughter is working,*
b. H: I missed your name.
c. C: Hank.
d. H: Go ahead Hank
e. C: *as well as her uh husband.*
They have a child
and they bring the child to us every day for babysitting.

In this example, C says "...my daughter is working," and then H interrupts with a side-sequence, which runs for three turns. After the side-sequence ends in 18d, C picks up the utterance from 18a, and uses the pronoun "her" as if there had been no interruption.

In contrast, there is also evidence to support the idea that as additional material is added to the discourse, accessing a distant referent becomes more difficult, regardless of the discourse structure. Walker (1998) argued that although pronominal reference can "ignore" an embedded structure as in 18, the same dialogue feels less natural if the interruption is longer, as in 19.

(19) (Walker, 1998).

a. C: Okay Harry, I have a problem that uh my - with today's economy *my daughter is working,*

b. H: I missed your name.

c. C: Hank.

d. H: Is that H A N K?

e. C: Yes

f. H: Go ahead Hank

g. C: *as well as her uh husband.* They have a child and they bring the child to us every day for babysitting.

Through an analysis of several naturally-occurring examples, Walker showed that both pronouns and full NPs can occur under four different conditions of hierarchical and linear recency. She argued that Grosz and Sidner's stack model of hierarchical discourse structure does not predict choices between pronouns and fuller forms, and argues instead for her "cache model". The cache model, essentially an implementation of Baddeley's working memory, is a limited capacity storage area for currently relevant information. Referent representations in the cache model are extremely accessible. As new information comes into the cache model, old information gets pushed out into "main memory".

Walker argued that processes of Centering (see §1.1.3) take place with respect to information in the cache. This proposal contrasts with the claim of hierarchical models that they take place with respect to information in the focus space at the top of the focus stack (Grosz and Sidner, 1986). The primary evidence for the focus stack model comes from cases like 18, called "return pops", where a pronoun is used to refer to something hierarchically but not linearly recent. Walker argues the cache model can account for these

cases, as well as reference to linearly recent referents. In the case of a return pop, the information from the hierarchically recent segment is retrieved into the cache, making it more accessible. At the same time, the information from the linearly recent segment remains in the cache until it is displaced by new information. Thus, Walker claimed that the accessibility of discourse referents can be influenced by both linear and hierarchical recency.

Walker (1998) was primarily concerned with the **possible** forms of reference under different conditions of discourse structure, and so she concluded that the stack model does not accurately predict where pronouns are an option. At the same time, she pointed out the need for future work to investigate how discourse structure relates to the frequency of different forms of reference. The present corpus analysis is ideal for investigating this question.

Although Walker (1998) and Grosz and Sidner (1986) discussed focus spaces in terms of the discourse intentions, identifying the intentional structure can be difficult. For this reason, I chose instead to examine quoted material as clearly marked examples of embedding. Quotes constitute a separate plane of discourse from the baseline story; they contain different referents, represent different speakers, and often take place at a different time than the main narrative event. In addition, writing conventions have developed punctuation specifically for setting quotes apart from the main discourse. For these reasons, one can argue that quotes are distinct discourse segments, embedded within the main narrative line.

There are three different predictions about how embedded material should influence pronoun use. Hierarchical theories of discourse structure would predict that quotes should be discounted from an analysis of the main flow of discourse, since they represent a separate plane of discourse. That is, a hierarchical theory would predict that there should be no difference whether a quote intervenes between an anaphor and its antecedent or not. This could be called the strong version of the Invisible Embeddings hypothesis. However, as

Walker observed, linear recency can also play a role, in combination with hierarchical recency. This could be called the weak version of the Invisible Embeddings hypothesis. Finally, a theory of recency as a linear phenomenon would predict that quotes should behave like nonquoted passages in all respects.

I investigated these hypotheses in the English corpus by coding both a) the number of nonquoted clauses since the referent was last mentioned, and b) whether there were one or more quotes intervening as well. The results reported in Figure 1 only include nonquoted clauses. The percentage of pronominal references in each category of recency, with and without quotes, are presented in Table 8 and Figure 2. These data show that when a quote increases the linear distance between an anaphor and the last time it was mentioned, the percentage of pronominalization falls. The comparison between "no quote" and "plus quote(s)" was significant for "one clause", "two clauses", and "three clauses" ($\chi^2(1)$'s > 4, p's < .05), although not for the categories "four clauses" and "five clauses" ($\chi^2(1)$'s < 1, p's > 1).

Table 8. The effect of quotes on reference form (English)

	no quote		plus quote(s)	
one clause	90%	(n=288)	60%	(n=48)
two clauses	55%	(n=44)	26%	(n=31)
three clauses	38%	(n=40)	7%	(n=14)
four clauses	22%	(n=23)	11%	(n=9)
five clauses	9%	(n=11)	0	(n=2)

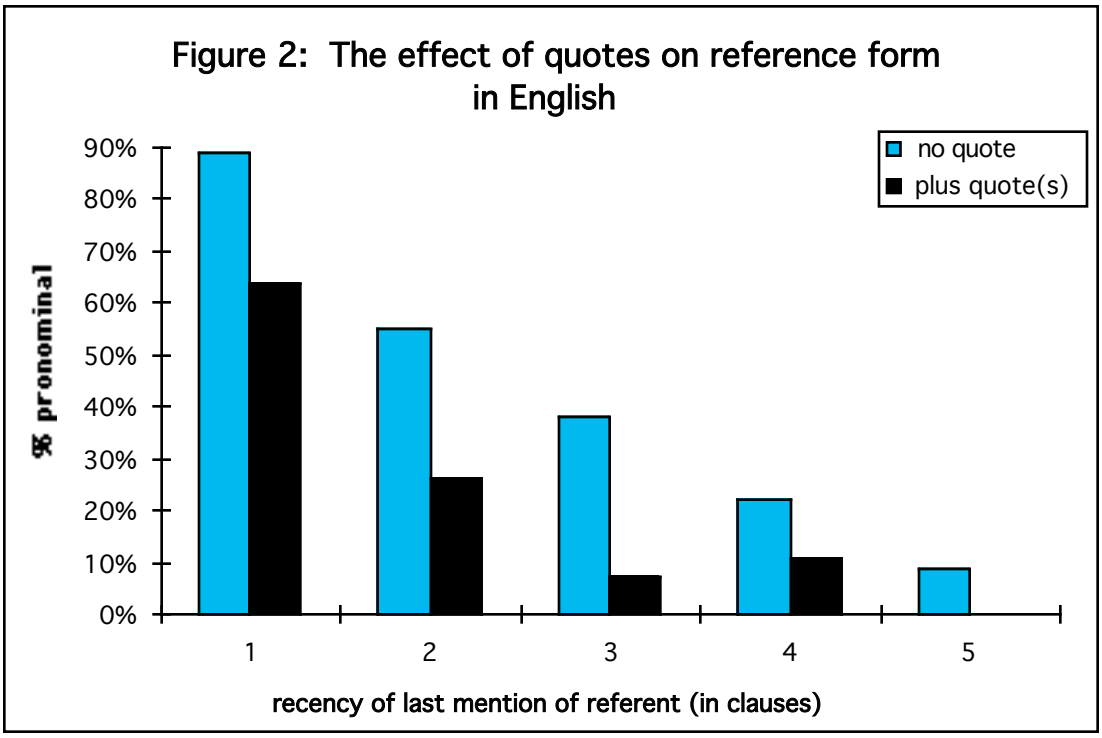

The strong version of the Invisible Embeddings hypothesis would predict that pronouns would be used equally as often in cases where a quote intervened and cases where none did. However, pronouns were used less often in each category of recency when quoted material intervened. The results suggest that it is incorrect to make strong about the ability of anaphoric reference to "skip over" embedded segments, thus providing evidence against the strong version of the Invisible Embeddings hypothesis.

Neither do these results strongly support the weak version of the Invisible Embeddings hypothesis. This version would predict that the presence of quoted material between an anaphor and the antecedent would reduce the rate of pronominalization, but not as much as the presence of another clause. The crucial comparisons are depicted in Figure 3. The weak Invisible Embeddings hypothesis would predict that the differences between each pair should be negative.

Figure 3: The differences between adding quoted and nonquoted material

	% pronominalization			
one clause plus quote(s)	60%	(n=48)	⌋	-5%
two clauses	55%	(n=44)		
two clauses plus quote(s)	26%	(n=31)	⌋	+8%
three clauses	38%	(n=40)		
three clauses plus quote(s)	7%	(n=14)	⌋	+15%
four clauses	22%	(n=23)		
four clauses plus quote(s)	11%	(n=9)	⌋	-2%
five clauses	9%	(n=11)		

However, only two of these comparisons showed negative differences, and more importantly, none of these comparisons was significant (all $\chi 2$'s < 1). It could be that the low cell totals contributed to this nonsignificance. But as they stand, these data suggest that quoted material may interfere with the accessibility of previously-mentioned referents as much as non-quoted material.

In sum, these data supported neither the strong nor the weak version of the Invisible Embeddings Hypothesis for quoted material. It is possible that quotes do not function the same way as other types of embeddings. This question would have to be investigated through an explicit comparison of different types of embedded structures. If they do function similarly, however, these results would indicate that the accessibility of discourse referents is influenced more by linear recency than hierarchical recency.

2.3.1.5. Recency: Conclusion

The results of a text analysis in three languages support the idea that reference form depends on how recently the referent was last mentioned in the discourse. Underspecified forms such as pronouns or null anaphors are more natural when the referent was mentioned in the previous clause, but fuller forms are needed when the referent hasn't been mentioned for more than 5 clauses. A closer look at English reveals a detailed picture of recency,

including the difference between main and subordinate clauses, and the dominance of linear recency over hierarchical recency. Figure 4 brings all of these features together, and presents an overall picture of how recency correlates with the rate of pronominalization in English.

Figure 4: The effect of Recency on reference form in English

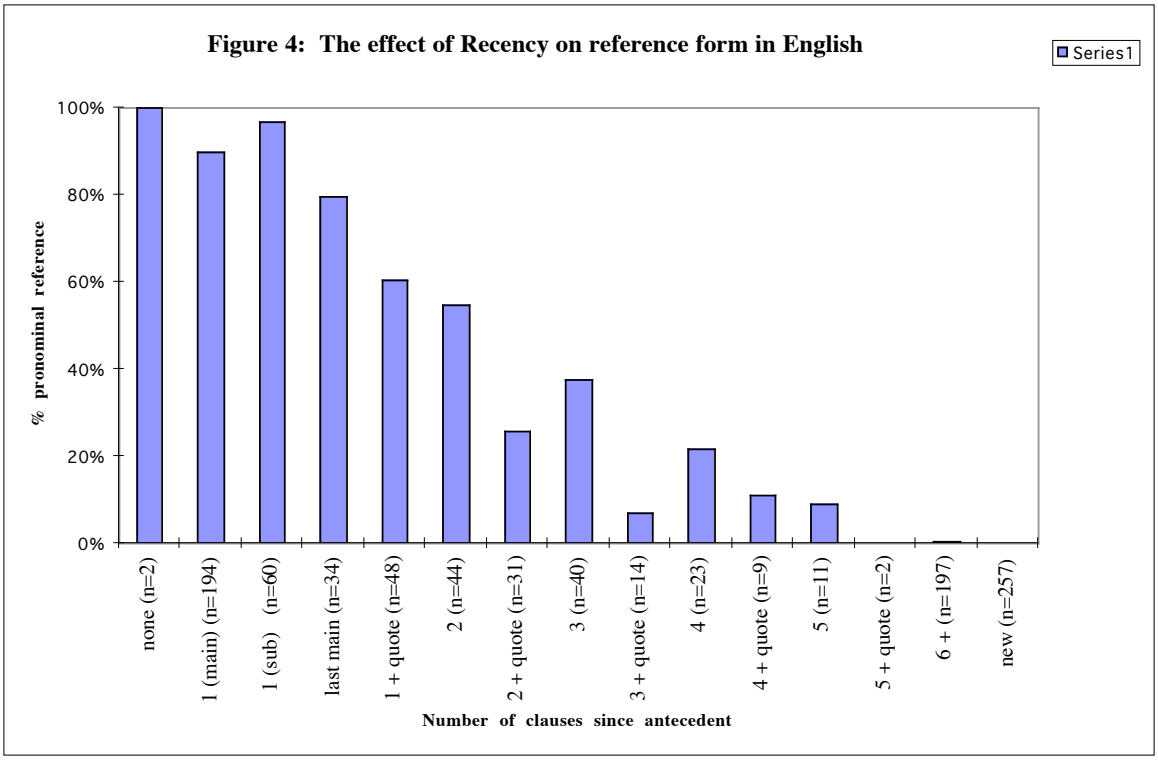

2.3.2. Subjecthood and Parallelism

I now turn to the second and third factors observed in this text analysis: Subjecthood and Parallelism. The following data show that structural properties of the anaphor and antecedent correlate with choices in reference form.

2.3.2.1. Subjecthood

Past research on reference form has claimed that entities mentioned in subject position are more salient and more accessible than other entities, and the comprehension of reference to these entities is easier when the referring expression is a pronoun than when it is a full NP (e.g., Gordon et al., 1993; Hudson-D'Zmura and Tanenhaus, 1998). A related claim grants privileged position to First-Mentioned referents (e.g., Stevenson et al., 1994; Gernsbacher and Hargreaves, 1988, 1992), reflecting claims that sentence-initial referents are highly topical (Reinhart, 1982). The higher accessibility of grammatical subjects may be due to either their grammatical position or initial sentence position, but investigation of this issue is beyond the scope of this study. Teasing these two factors apart is difficult, for different reasons in each language. In English, Subjecthood and First Mention are highly correlated, making a distributional comparison of the two difficult. Word order is more flexible in Spanish and Mapudungun, but the relative position of NPs can only be measured when arguments are overt. Since both languages have a high rate of null reference, the influence of sentence position on reference form is difficult to ascertain. Therefore, I only looked at Subjecthood in this analysis.

I also limited this analysis to cases where the referent was last mentioned in the previous clause or last main clause, because I expected that the effect of Subjecthood would be biggest in the immediately following clause. My analysis compared cases where the referent had last been mentioned as a subject with cases where it had last been mentioned as an object or oblique. I will call these "subject-referents", "object-referents", and "oblique referents", respectively. Because the Mapudungun database only included cases where the anaphor was in the main clause, the data in this figure from English and Spanish are also

restricted to main clauses. However, analyses of all clauses in English and Spanish yield similar results.[1]

Figure 5: The effect of Subjecthood on reference form in three languages

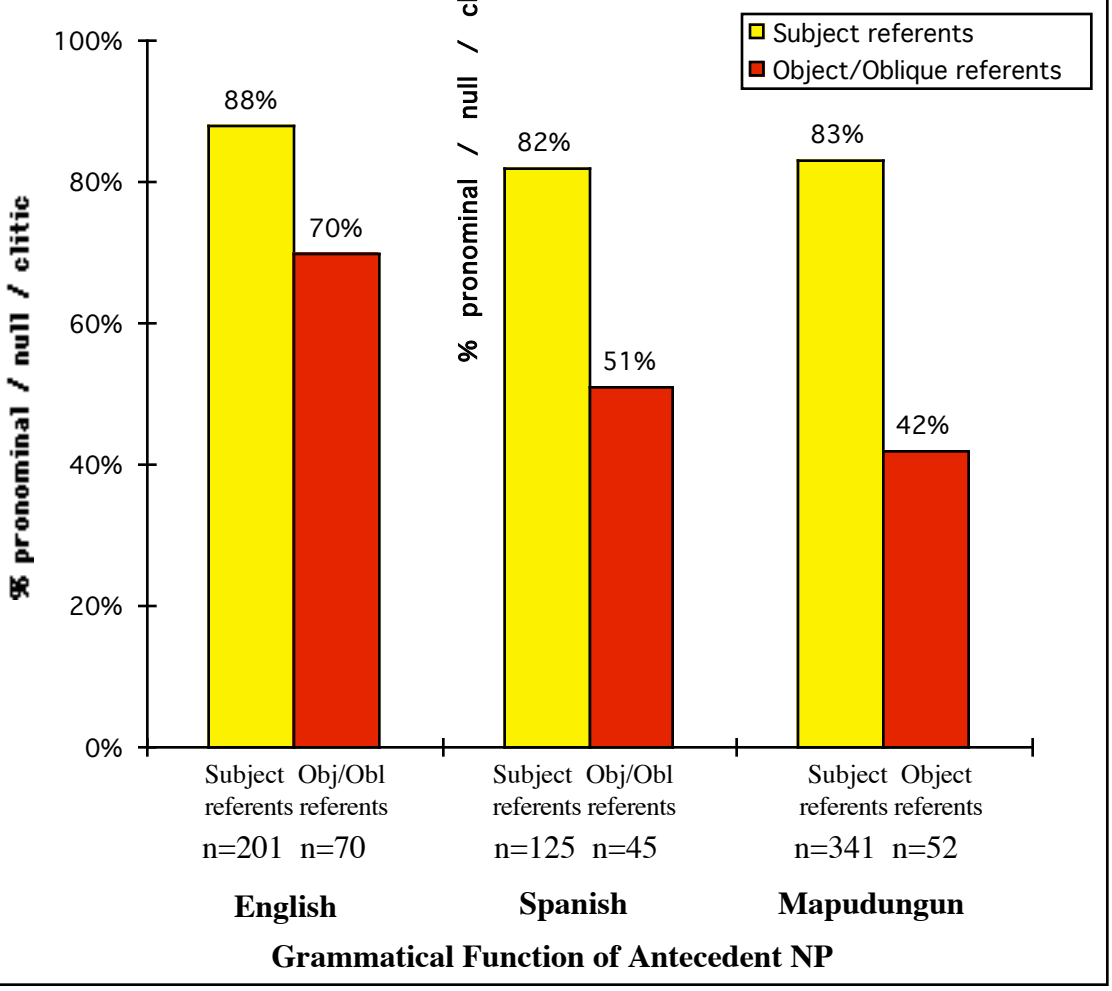

[1] An analysis of all clauses in English and Spanish yields a similar result to the data in Figure 4. Both comparisons are significant (English: χ^2=4.87, DF=1, p<.05; Spanish: χ^2=11.3, DF=1, p<.001).

	GF of referent	% pronominal/null	n
English	Subject	89%	246
	Object/Oblique	77%	90
Spanish	Subject	84%	168
	Object/Oblique	65%	63

Figure 5 shows that the percentage of pronominal or null references is higher for reference to subject-referents than object- or oblique-referents in English (χ^2=4.87, DF=1, p<.05) , Spanish (χ^2=11.3, DF=1, p<.001), and Mapudungun (χ^2=43, DF=1, p<.001). Assuming that less-specified forms of reference are used more often when the antecedent is more accessible, these data support claims that the subject is a prominent position.

2.3.2.2. Parallelism

A second claim that has appeared in the literature on reference form and pronoun resolution is that parallel structure between an anaphor and its antecedent facilitates pronominal reference. For example, in 20a, the pronoun is claimed to be easier to resolve because it refers to something in parallel position, whereas in 20b, it is more difficult because the reference is not parallel.

(20) a. In line at the post office, Jane tapped Emily on the shoulder. She just wanted to point out that the line was moving forward.

b. In line at the post office, Jane tapped Emily on the shoulder. She turned around and glared at her before moving forward.

One problem with this claim, however, is that in 20a (the parallel version) the antecedent is the subject, and in 20b (the nonparallel version) the antecedent is the object. This raises the question of whether claims of parallelism are merely a re-characterization of the subject bias, and conversely, whether the subject bias is in fact the result of parallelism.

I addressed this question with the text analysis by investigating reference form in terms of the grammatical function of both the anaphor and the antecedent, including all grammatical functions. Examples of parallel reference between subjects, objects, and obliques are shown in 21.

(21)

a. Subject Parallelism

As he was studying this remarkable pebble,

he began to shiver, probably from excitement, . . .

(Steig, Sylvester and the Magic Pebble, 1969)

b. Object Parallelism

Then Ole Swenson tied all the carts to the Hodag's tail.

The Hodag pulled them through the forest to the sawmill.

(C. Arnold, The Terrible Hodag, 1989)

c. Oblique Parallelism

At last she crawled over the swaying floor to her bed,

and lay down upon it; . . .

(Baum, The Wonderful Wizard of Oz, 1900:15)

I conducted this analysis in all three languages. Again, since the Mapudungun database only included cases where the anaphor was in a main clause, I restricted the English and Spanish analyses to main clauses as well, for comparison (but see footnote 11, below).

If both Parallelism and Subjecthood influence reference form, then one would expect to see a combined effect of the two factors, such that reference form is pronominal or null more often both when the antecedent is parallel and when the antecedent is the subject. The results, displayed in Table 9, support this prediction. There is a significant effect of parallelism in all three languages, but for English and Spanish the parallelism effect only shows up for subject-to-subject reference (where "subject-to-subject" means that the anaphoric referring form occurs in subject position, and refers to something last mentioned in subject position as well). In Mapudungun there is evidence for both subject and object parallelism, but the effect is stronger for subject-to-subject reference. These data suggest that there is a general subject bias, but that parallelism between an object or an oblique and

its referent can neutralize the subject bias by increasing the rate of null/pronominal reference to parallel objects and obliques. In Mapudungun the effect of parallelism is even stronger, and even overpowers the subject bias in the case of object-to-object reference. Yet even for Mapudungun, the rate of null reference for object-to-object reference is not as high as for subject-to-subject reference, suggesting that both Subjecthood and parallelism are influencing the outcome.

Table 9. Parallelism effects in English, Spanish, and Mapudungun

English		% pronoun	N
Subject ($\chi^2(1)=21$, $p<.001$)	Subject-to-Subject	88%	180
	Subject-to-Object/Obl.	59%	44
Object ($\chi^2(1)=.22$, n.s.)	Object-to-Object	92%	12
	Object-to-Subject/Obl.	86%	14
Oblique ($\chi^2(1)=1.1$, n.s.)	Oblique-to-Oblique	100%	5
	Oblique-to-Subj./Obj.	81%	16

Spanish		%null/clitic	N
Subject ($\chi^2(1)=32$, $p<.001$)	Subject-to-Subject	83%	102
	Subject-to-Object/Obl.	21%	19
Object ($\chi^2(1)=0.7$, n.s.)	Object-to-Object	83%	12
	Object-to-Subject/Obl.	92%	26

Mapudungun		% null	N
Subject ($\chi^2(1)=82$, $p<.001$)	Subject-to-Subject	86%	328
	Subject-to-Object	14%	28
Object ($\chi^2(1)=12$, $p<.001$)	Object-to-Object	75%	24
	Object-to-Subject	15%	13

The only problem with investigating each type of parallel reference was that it yielded small totals in some cells. Therefore these data are only suggestive. This problem was

compounded by limiting the analysis to cases where the anaphor was in a main clause. However, an analysis of all clauses in English and Spanish yielded similar results.[1]

2.4. The discourse patterns of Recency, Subjecthood, and Parallelism

The preceding results show that Reference Form in three languages is correlated with three of the factors mentioned in chapter 1: a) Recency of mention, b) Subjecthood, and c) Parallelism. These results support other findings of Recency (e.g., Givón, 1983a; Clark and Sengul, 1979), Subjecthood (e.g., Gordon et al., 1993), and Parallelism (e.g., Sheldon, 1974). This type of finding has typically taken as evidence that entities that have been mentioned recently, in subject position, or in parallel position are somehow more prominent than other entities, so that less-specified forms of reference are preferred.

The question remains, however, why are these three factors associated with prominence? Recency, Subjecthood, and Parallelism have very little in common, other than the fact that they are all associated with choices in reference form. Why do they all have similar effects on language use?

The data in the following sections will show that all three factors share an additional property: they also correlate with the most frequent types of reference. All discourse referents are not equal. Some are mentioned once, and disappear. Others are more important to the discourse. At any given point in a text, there are many referents that have

---

[1] The pattern shown in Table 9 is not limited to only main clauses. For English and Spanish it is possible to include all clauses in the analysis, in which case a similar pattern emerges:

		English		**Spanish**	
		% pronoun	N	%null/clitic	N
Subject (χ^2's (1)>14, p's <.001)	Subject-to-Subject	90%	213	85%	136
	Subject-to-Object/Obl.	69%	59	42%	24
Object (χ^2's (1) <1.0, n.s.)	Object-to-Object	92%	13	90%	21
	Object-to-Subject/Obl.	88%	26	92%	25
Oblique (χ^2(1)=.71, n.s.)	Oblique-to-Oblique	100%	6		
	Oblique-to-Subj./Obj.	79%	19		

already been mentioned, but some of them are more likely to be mentioned again than others. I will show that in all three languages, writers use attenuated forms for precisely those referents that they also tend to refer to more often. These data will provide the first pieces of support for my central claim, that these factors display a double association with Reference Form and Reference Continuation.

2.4.1. Recency

The results in §2.3.1 showed that writers tend to use less-specified forms of reference for entities that have been mentioned recently. Why does recency have this effect? At a general level, and as many researchers have argued, it is because the more recently mentioned referents are more accessible than those that have not been mentioned for a while. More recently mentioned referents are more in the focus of attention of the discourse participants, and are thus more accessible.

One way of thinking of recency is as the result of a decay mechanism. That is, when something is mentioned, the discourse participants are forced to access that entity and integrate it into their mental models of the discourse. As time goes by, if that entity is not mentioned again, the mental representation of it will begin to fade and ultimately disappear, especially if new information is competing for limited attentional resources. When the representation of a discourse entity is freshly activated, less-specified forms of reference are more natural, but as time goes by and the representation fades, more fully specified forms of reference are needed to re-access the referent. Such a view is consistent with connectionist models that have been proposed to account for other aspects of language processing (e.g., McClelland and Rumelhart, 1981; Seidenberg and McClelland, 1989), and is implicit in the use of the term "decay" for this phenomenon (Givón, 1983a).

While this view is likely to contain some truth, it may not be the entire story of why recently mentioned things are salient. Instead, the influence of recency may stem from the fact that people tend to talk about the things that they were just talking about. As I mentioned in §1.2, discourses tend to be organized around the goals and intentions of the

participants. For this reason, people tend to talk about ideas for extended periods of time, rather than stringing unrelated sentences together. This results in a pattern whereby people refer to recently mentioned things more than they refer to other things.

This generalization is supported by observing the patterns of reference in the corpus analyses in English, Spanish, and Mapudungun. In all three languages, reference to entities last mentioned in the previous clause is the most frequent type of reference. That is, of all the references in each database, a large percentage refer to entities in the previous clause. Table 10 shows the percentage of all references to entities, according to how recently they were mentioned. Here the number of references includes **all** references. References to a subset or superset of recently mentioned entities are excluded, as are references to entire clauses.

Table 10. The distribution of references according to the recency of the referent.

A. ENGLISH			
Recency of mention	# refs.	% of all references	
previous clause/ last main clause	336	|||||||||||||||||||||||||||||||||||	**35%**
2-5 clauses back	174	||||||||||||||||||	18%
6+ clauses back	197	||||||||||||||||||||	20%
new	257	|||||||||||||||||||||||||||	27%

B. SPANISH			
Recency of mention	# refs.	% of all references	
previous clause	231	|||||||||||||||||||||||||||||||||	**33%**
2-5 clauses back	108	||||||||||||||||	16%
6+ clauses back	107	|||||||||||||||	15%
new	246	||||||||||||||||||||||||||||||||||||	36%

C. MAPUDUNGUN			
Recency of mention	# refs.	% of all references (main clauses only)	
previous clause	434	||||||||||||||||||||||||||||||||||||||||||||||||||||||||| ||||||	**63%**
2-5 clauses back	83	|||||||||||||||	15%
6+ clauses back	22	||||	4%
new	111	||||||||||||||||||	18%

As the data in Table 10 show, the largest percentage of references to something which has appeared previously in the discourse refers to something in the previous clause. This

tendency is strongest in Mapudungun, where this category represents 63% of all references. In English and Spanish the effect is smaller, but in these languages approximately one third of all references have referents that were mentioned in the previous clause. This is the largest portion of all the references to given information.

The correlation between Recency and Reference Continuation may be difficult to appreciate from the figures in Table 10, which presents recency in terms of 4 categories. It is highly implausible to suppose that the language processing system categorizes recency of mention in terms of "one clause back", "2-5 clauses", or "6 or more clauses". The actual system is probably continuous, but the results of the discourse analysis can only be presented categorically. A finer-grained analysis of recency in English supports the claim that frequency of subsequent reference drops off gradually over time since the last mention, as shown in Table 11.

Table 11. Patterns of reference by recency in English

<table>
<tr><th># of clauses since last reference</th><th># of references</th><th>% of all references</th><th></th></tr>
<tr><td>1</td><td>336</td><td>||||||||||||||||||||||||||||
|||||||</td><td>35%</td></tr>
<tr><td>2</td><td>75</td><td>||||||||</td><td>8%</td></tr>
<tr><td>3</td><td>54</td><td>||||||</td><td>6%</td></tr>
<tr><td>4</td><td>32</td><td>|||</td><td>3%</td></tr>
<tr><td>5</td><td>13</td><td>|</td><td>1%</td></tr>
<tr><td>6+</td><td>197</td><td>||||||||||||||||||||</td><td>20%</td></tr>
<tr><td>new</td><td>257</td><td>|||||||||||||||||||||||||||</td><td>27%</td></tr>
</table>

I did not conduct a clause-by-clause analysis for the category "6+", but is probable that if I did, the percentage of references would continue to drop off asymptotically.

Tables 10 and 11 show that the probability of subsequent reference is correlated with recency. Table 5 in §2.3.1 showed the same correlation between pronominalization and recency. This parallelism provides the first piece of evidence for the main claim of this dissertation: Recency is associated with 2 things: a) Reference Form, and b) Reference

Continuation. That is, writers use pronouns more for the type of referents that they tend to mention again: recent referents.

Why can writers do this? It is because they know that recently mentioned referents are more activated in the mind of the reader, which facilitates the comprehension of the anaphoric forms. The more activated the referent, the easier it is to interpret pronouns and null anaphors. Since speakers and writers desire efficient communication, they choose attenuated forms only when referents are activated for the listener or reader.

The data presented here also provide another way of thinking about why referents become activated for the reader. During text comprehension, readers are trying to retrieve the message that the writer intended. At one level, this involves constructing a mental representation of the discourse referents. From the writer's perspective, some referents are more central to their message and therefore are more activated than others in their minds. Because the activated referents are the ones that are central to the writer's message, they are the ones that the writer intends to mention in the following sentences. The tendency for writers to continue to refer to referents that have been recently mentioned stems from their attempt to present a coherent text.

During text comprehension, the reader tries to construct a representation of these referents that is similar to what the writer intended. But without knowing what was in the writer's head, the reader can only go by information from the text, such as the recency of referents. When a referent has been mentioned recently at a given point in the discourse, the reader knows it has a high probability of being central to the writer's message, and that it has a high probability of being mentioned again. Therefore, it is to the reader's advantage to consider that referent highly activated.

This view is not mutually exclusive with the decay mechanism described above, because it also assumes that activation automatically decays if the representation is left "unattended", as it were. But this account suggests why the language processing system has evolved a sensitivity to recency. That is, discourse participants focus on certain

characters and items as likely elements in the upcoming discourse. When a referent has just been mentioned, it is highly activated, because there is a high probability that it will play a role in the immediately following discourse. However, as time goes by and the writer does not mention this character, it becomes increasingly less probable that it will be mentioned again. This gets translated into lower activation for the reader's representation of that referent.

Further support for this position comes from evidence that the effects of recency are not linear (Clark and Sengul, 1979). A simple recency-as-decay account would predict that the rate of pronominalization would fall steadily as recency decreases, without consideration of other factors. However, the analysis in English showed that the effects of Recency are modulated by other aspects of the discourse structure (see Figure 4). First, pronouns were used in English more often when one or more subordinate clauses intervened between the anaphor and antecedent than when one main cause intervened. Second, data on the effects of Subjecthood and Parallelism show that when the referent appeared in the previous clause, these factors overpowered any effect of Recency.

A nonlinear view of Recency is also supported by research on the hierarchical nature of discourse (e.g., Walker, 1998; Grosz and Sidner, 1986; Polanyi, 1994; Fox, 1997). Although my results did not support a hierarchical analysis of quotes, examples like 18 (§2.3.1.4) suggest that some embedded segments can pose less of a threat to the main flow of the discourse than non-embedded sections. This suggests that the activation of discourse entities does not slip away at a constant rate, but rather is regulated by other factors.

Further support comes from Greene et al.'s (1994) study on text comprehension. They studied texts in which two characters become separated, and later reunited. They argued that readers keep track of information that is shared by two characters in a text, and that when the characters reunite, the information in their common ground becomes more accessible. They found that in these situations, readers were able to use the character's common ground to interpret pronouns that had no recently mentioned referent.

2.4.2. Subjecthood

The results of the text analyses in all three languages showed that Reference Form was also associated with the grammatical function of the last reference to the conceptual referent. Less specified forms of reference were used most often when the last reference to that entity was in subject position. This raises the question of why it is that subject-referents should be more accessible than other referents.

Again, the data show that writers tend to refer more often to the referents of subjects, in comparison with the referents of objects or obliques. Table 12 shows that the number of subsequent references to the subject-referents far exceeds the number of references to the object- or oblique-referents. This pattern is especially pronounced for Mapudungun, but this may stem from the fact that the Mapudungun analysis did not include oblique arguments or any referring forms from subordinate clauses.

Table 12. Patterns of reference by grammatical function

<table>
<tr><th></th><th>Grammatical Function of Referent</th><th># of refs</th><th>% total references</th><th></th></tr>
<tr><td><u>English</u></td><td>Subject</td><td>246</td><td>||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
|||||||||||||||</td><td>73%</td></tr>
<tr><td></td><td>Object/Oblique</td><td>90</td><td>|||||||||||||||||||||||||||</td><td>27%</td></tr>
<tr><td><u>Spanish</u></td><td>Subject</td><td>168</td><td>||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
|||||||||||||||</td><td>73%</td></tr>
<tr><td></td><td>Object/Oblique</td><td>63</td><td>|||||||||||||||||||||||||||</td><td>27%</td></tr>
<tr><td><u>Mapudungun</u></td><td>Subject</td><td>341</td><td>||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
|||||||||||||||||||||||||||||</td><td>88%</td></tr>
<tr><td></td><td>Object</td><td>52</td><td>|||||||||||||</td><td>13%</td></tr>
</table>

These results show that the referents of subjects are far more likely to receive subsequent reference than other referents. These data do not tell us why this pattern exists, but some clues can be found in past research. As many researchers have noted, there are certain characters and objects which are more central to a discourse at any given point (e.g., Ariel, 1990; Chafe, 1994; Grosz et al., 1995). It has also been well-established that speakers

choose certain constructions, depending on the discourse status of the referents (e.g., Givón, 1983a; Ward and Birner, 1996). In particular, speakers and writers tend put central characters in subject position (e.g., Du Bois, 1987; Prince, 1994; Tomlin et al., 1997), a position which has also been associated with the sentence-level topic in English (Reinhart, 1982). It has also been established that speakers and writers refer more often to characters that are central to the discourse (Francik, 1985). Thus, it seems natural that subject-referents are subsequently mentioned more often than other characters.

Considering the results from Recency and Subjecthood together, a general pattern of reference emerges. Writers are more likely to refer to something that was mentioned in the previous clause than other referents. In particular, they are more likely to refer to something that was last mentioned in subject position of the preceding clause. These results are compatible with the corpus analysis results in chapter 3, where the Subjecthood factor will be discussed further.

2.4.3. Parallelism

The text analyses, described above, also showed that when the anaphor and its antecedent are in the same grammatical position, the use of pronouns and null reference increases. This effect combined with the effect of Subjecthood, such that pronominal and null reference was especially frequent when both anaphor and antecedent were in subject position. When Subjecthood and Parallelism led to different preferred forms of reference, the effect of Subjecthood was weakened in English and Spanish, and nullified in Mapudungun. But why does Parallelism play the role it does?

As with Recency and Subjecthood, the text analyses reveal a double association with Reference Form and Reference Continuation. Figures 6-8 show that for each anaphor position, the highest proportion of references have parallel antecedents. This pattern holds for English ($\chi^2(4)=25$, $p<.001$), Spanish ($\chi^2(2)=33$, $p<.001$), and Mapudungun ($\chi^2(1)=27$, $p<.001$).

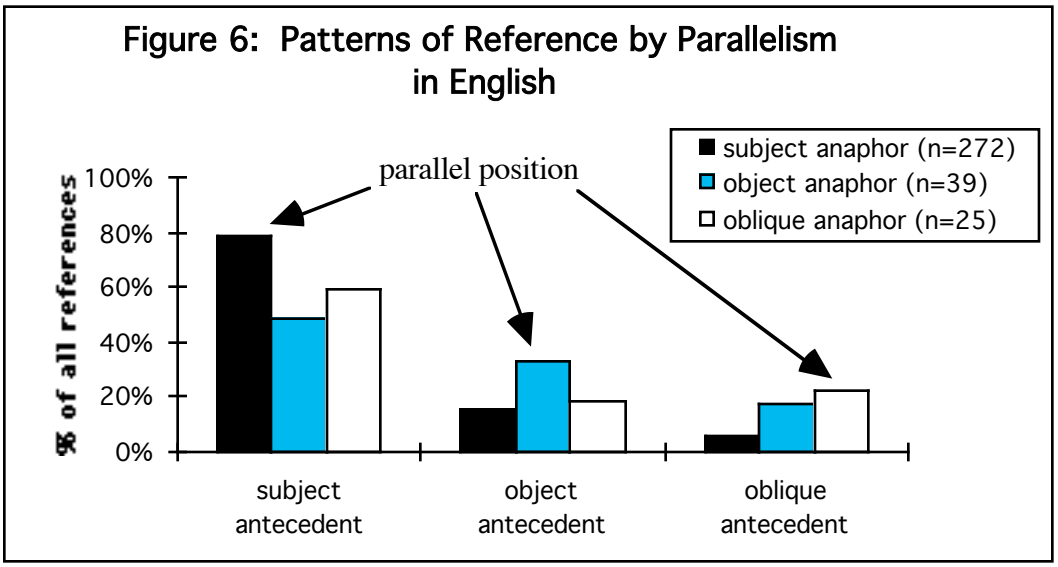
Figure 6: Patterns of Reference by Parallelism in English
parallel position
subject anaphor (n=272)
object anaphor (n=39)
oblique anaphor (n=25)
% of all references
100%
80%
60%
40%
20%
0%
subject antecedent
object antecedent
oblique antecedent

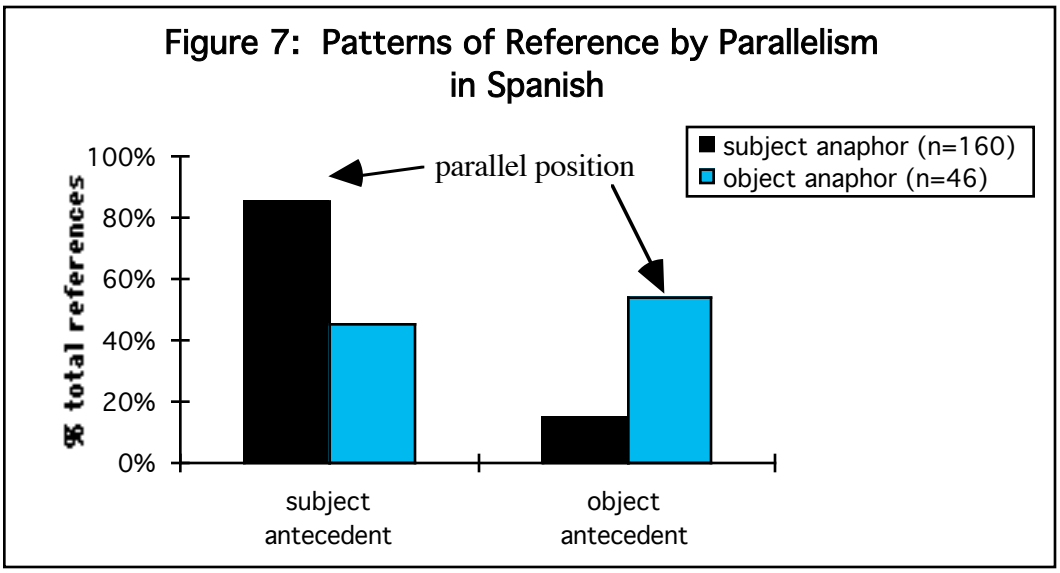
Figure 7: Patterns of Reference by Parallelism in Spanish
parallel position
subject anaphor (n=160)
object anaphor (n=46)
% total references
100%
80%
60%
40%
20%
0%
subject antecedent
object antecedent

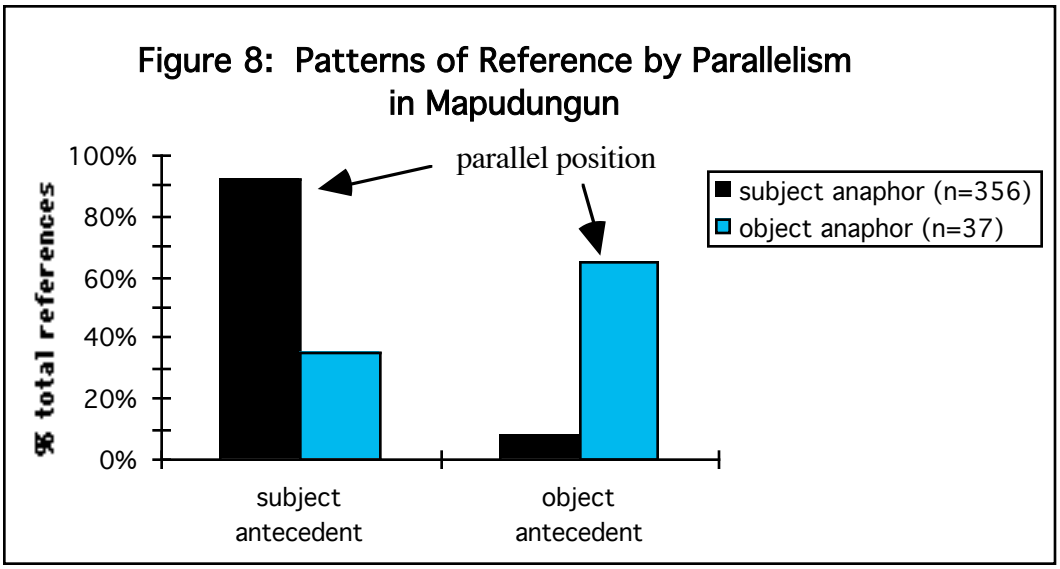

Figures 6-8 reveal an overall effect of the Subject Bias. In all three graphs, the subject antecedent cluster is the tallest. This shows that writers referred to the referents of subjects more often than other referents. However, this effect is the strongest in the case where the reference is parallel, when both anaphor and antecedent are in subject position. In the cases of the object and oblique anaphors, the bias to refer to the subject is reduced, and there is an increased rate of anaphoric references to antecedents in parallel position.

These data show that parallel reference is statistically more frequent than nonparallel reference. Why does this pattern exist? It could be that people tend to place things in similar roles in successive sentences. That is, if something was mentioned as an object or an oblique, it is likely to be playing a role in the discourse that will make it likely to appear again in the same grammatical role. For example, in 22 "the puppy" gets passed from one person to another, playing the same thematic (and grammatical) role in each utterance.

(22) Carol picked up the puppy and passed it to John, who held it for awhile. Then Susie took it out for a walk. By the end of the evening everyone had handled the animal at some point or another.

At the same time, by using the same grammatical role for repeated references to a particular referent, the speaker also maintains the discourse status of that referent, relative to other entities. In both cases, maintaining the same semantic and discourse perspectives on a referent promotes coherence within a discourse segment.

The increased tendency for anaphors to have parallel antecedents mirrors the finding in §2.3.2.2 that pronouns are used more for parallel antecedents than non-parallel ones. That is, parallelism is doubly associated with both Reference Form, and Reference Continuation. What is the relationship between these findings?

As with Recency and Subjecthood, I argue that writers use pronouns more for parallel reference because it is easier for readers to process parallel reference than nonparallel reference. The reason it's easier is because parallel reference is a more common discourse pattern than nonparallel reference. Readers know that writers tend to refer to referents using the same grammatical positions in subsequent utterances. Therefore, with respect to any given anaphor, parallel-referents are more probable than nonparallel-referents.

For example, when readers come to the end of 23a, they may know that the subject of the following sentence is imminent. Even if they don't, once they encounter the anaphor, they identify it as an NP in subject position. This knowledge, combined with the knowledge that writers tend to maintain parallel reference, will cause the subject of 23a to become partially activated.

(23) a. Fluffy clouds of smoke chugged out of the engine's smokestack.

b. They reminded him of the Hodag.

(C. Arnold, The Terrible Hodag, 1989:20)

As the reader continues to process 23b, inferential processes will support the interpretation of the referent of "they" as the subject-referent. Because the parallelism information supports the same interpretation as subsequent inferential processes, the referent is relatively easy to access.

A similar phenomenon happens in 24; as readers access the verb "showed" in 24b, they access the knowledge that it is almost always used as a transitive verb. Therefore they know that the following word is probably a grammatical object, so their experience with parallelism may begin to activate the object-referent from 24a. As the anaphor itself is encountered, the reader becomes certain that it is in object position, so parallelism effects become even stronger.

(24) a. They found the boss man

b. and showed him

c. that they had finished their job.

(C. Arnold, The Terrible Hodag, 1989:28)

25b shows an example of parallel oblique reference. In this case, the word "upon" signals the imminence of an oblique reference, and may begin the activation of the parallel oblique referent from 25a.

(25) a. At last she crawled over the swaying floor to her bed,

b. and lay down upon it;

(Baum, The Wonderful Wizard of Oz, 1900:15)

In all of these examples, the parallelism effects are both weak and transitory. They are not sufficient by themselves to identify the correct referent, but they can facilitate or inhibit the process of referent resolution.

Note that Parallelism is a different kind of factor than Recency or Subjecthood. For the two latter factors, it is reasonable to think of them as indicators of the salience of their referents. That is, readers can infer that subject-referents or recently mentioned referents are likely to have been actively represented in the writer's mind. In contrast, it is harder to defend the claim that parallel referents are more important to the writer's goals than nonparallel referents. A writer knows the importance of any given entity from the moment it is mentioned (or before), but parallelism information does not become relevant until the position of the anaphor is known. In example 24, above, the referent of "the boss man" in 24a would not have been less salient if the anaphor in 24b were put in subject position.

Therefore, a connection between Parallelism and Reference Form cannot be made via a global concept like salience. However, the account I presented above shows how Parallelism can influence how readers represent parallel referents during anaphor resolution: Since parallel referents are probabilistically more likely to be the referents, their representations become partially activated when the anaphor position is identified. This momentary and partial activation facilitates the comprehension of the referring form.

Writers know that parallelism facilitates comprehension, and therefore they are able to use less-specified forms of reference for parallel referents. In this way, parallelism effects on reference form can be explained in terms of general patterns in discourse structure, and not an idiosyncratic function that says "Use a pronoun to refer to something in the same position, otherwise use a fuller form".

2.4.4. Overall picture of reference form

The data in this chapter provide the first pieces of evidence for my hypothesis. For three factors, and in three different languages, I found striking parallels. Each factor was associated with two things: a) an increased use of less-specified forms (Reference Form), and b) an increased probability of subsequent reference to the same referents (Reference Continuation). This pattern of results is depicted in Figure 9, following the schema from chapter 1.

Figure 9: The double association for Recency, Subjecthood, and Parallellism

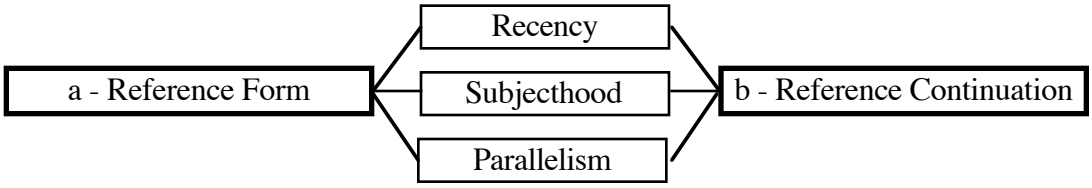

For each of these three factors, I have speculated on different motivations for why there is an increased probability of subsequent reference. For Subjecthood, there is evidence that the pragmatic function of subject position makes it specialized for highly accessible referents. For Recency, this pattern falls out of the more general tendency to talk about things for extended periods of time. For Parallelism, discourse referents may tend to play the same semantic and discourse roles across utterances, which results in subsequent references in the same grammatical positions. In all cases, however, the fundamental generalization is that speakers aim to be coherent. This general goal underlies all observed discourse patterns.

The fact that all these patterns exist has strong implications for language comprehension. As H. Clark (1996) has argued, one of the tasks that discourse participants need to accomplish is the coordination of the shared information in the discourse. At one

level, this means that discourse participants need to coordinate their representations of common information in the discourse, including the relative "salience", or activation of each representation. How can they do this?

The double association depicted in Figure 9 suggests a mechanism for this coordination. During language comprehension the reader or listener is trying to construct a mental model that approximates that of the speaker or writer. Thus, comprehenders need to find out what elements are activated in the mind of the writer, or in other words, what elements the writer is intending to talk about. Clues to this activation are present in the text, as measured for example by Recency or Subjecthood. Referents with these properties are the ones likely to be mentioned again, and therefore should be represented with high activation. Parallelism effects are smaller, and only become relevant when the reader encounters the anaphor. But Parallelism also helps the reader interpret the anaphor in the way the writer intended, by partially activating parallel referents, the types of referents writers tend to refer to in that situation.

The higher the activation of a referent representation, the easier it will be for the reader to access it as the referent of an anaphor. Thus, comprehension will be facilitated. When comprehension is facilitated, as I argued in chapter 1, the speaker/ writer has the option of using less specified forms of reference, which have the added benefit of making communication more efficient for all discourse participants.

The proposal, at this point, is necessarily incomplete. There are multiple questions about the representations involved. In particular, how do the speaker's representations relate to the listener's representations? How do referents become activated? These questions will be addressed in more detail in chapter 5.

The main conclusion so far is this: text analyses have shown a similarity between the referents of subjects, recently mentioned NPs, and parallel NPs. In all cases, there is a double association with Reference Form and Reference Continuation. Furthermore, this pattern was replicated across three different languages, with three different systems of

reference. In the following chapters, I will present further evidence for Subjecthood, as well as evidence that this double association extends to the factors Focus and Goal Status.

Chapter 3

Topic, Focus, and reference form

In this chapter I will investigate the ways in which the categories of "topic" and "focus" affect choices in reference form. These categories of information structure deserve special attention because they are invoked to explain many phenomena, yet are notoriously difficult to define. Furthermore, an inspection of past research reveals a discrepancy between linguistic analyses of information structure and psycholinguistic studies on reference: while linguists have treated "topic" and "focus" as opposing categories, psycholinguistic experiments suggest that comprehenders represent the referents of both topic and focus in similar ways, such that both are more activated than other referents in the discourse. The studies in this chapter will address this discrepancy, and propose an explanation for why the categories of topic and focus have similar effects on language processing.

First I will describe a rating experiment that directly compared grammatical subjects with the focus of clefts, two constructions that have been associated with "topic" and "focus", respectively. I will demonstrate that both categories make their referents more salient, such that it is more natural to use a pronoun for subsequent reference. I will then

describe a corpus analysis that investigated the similarity between these categories, and show that the referents of both subjects and the focus of clefts are more likely than other entities to be mentioned in the following discourse. This finding will corroborate the results from chapter 2, which showed that grammatical subjects are doubly associated with Reference Form and Reference Continuation. The findings from this chapter will also show that the same characteristics apply to the foci of clefts, thus providing evidence of this double association for two of the factors presented in chapter 1: Subjecthood and Focus.

3.1. Topic and Focus: similar or different?

To investigate the categories of "topic" and "focus", I will look at grammatical subjects and the foci of clefts, as instances of these categories. "Topic" and "focus" are particularly worth comparing, because they have traditionally been contrasted in the linguistic literature. Although the category "topic" lacks a clear-cut definition, it is often thought of as the "given" portion of an utterance (Gundel, 1974; Chafe, 1976; Tomlin et al., 1997), and is often contrasted with the focus (e.g., Sgall et al., 1986; Vallduví, 1993; see de Swart and de Hoop, 1995, for an overview)

In contrast to the topic, the focus traditionally refers to the new, informative part of the utterance. One prototypical example of focus occurs with cleft constructions. Clefts are generally assumed to exhibit a clear differentiation between what is new, the "focus", and what is given or presupposed (e.g., Chomsky, 1970). Examples are given in 1.

(1) a. What I found on the desk was a three-volume series on starfish.

|____| |__|

given/ presupposition focus

b. It was the book on mollusks that I found in the drawer.

|___| |___|

focus given/ presupposition

In addition, the focus of clefts is often contrasted with more topical constructions, for example in Givón's (1983a) proposed scale of "topicality", presented in chapter 1. In this scale, the focus of clefts is in the second lowest category on the scale, reflecting the fact that the element in this position is usually new or contrastive, and usually expressed with a full description or name. For a review of focus see Tomlin et al. (1997) or Vallduví (1993); see also the discussion in chapter 1.

At the same time, research has suggested that subjects and foci may have similar effects during language processing, such that the referents of both are more salient than other referents. First, there is a wealth of evidence that subject-referents are more salient than other referents in non-clefted utterances. This generalization is supported by the text analysis in chapter 2, which showed an effect of Subjecthood in English, Spanish, and Mapudungun. This result is also consistent with the findings of past research (e.g., Gordon et al., 1993; see chapter 1). Second, there is growing evidence that the focus-referent is more salient than other referents in a clefted utterance. As discussed in chapter 1, Almor's (1995) experiments showed in sentences like "What the bird ate was the fruit", the referent of the focus ("the fruit") was more accessible than the nonfocus ("the bird"). Almor's results are consistent with Just and Carpenter's (1977) findings and Sidner's (1983) discussion of the focus of clefts.

Thus, there is a discrepancy between linguistic work on information structure and psycholinguistic experiments. The former treats topics and foci as different, the latter suggests they may be similar . One reason for this discrepancy may be that linguistic studies of clefts have generally been concerned with the discourse conditions necessary for the cleft construction (e.g., Prince, 1978), or the discourse properties of the referents of subjects (e.g., Prince, 1992). In contrast, psycholinguistic studies have investigated the effect of grammatical function and cleft construction on the cognitive status of the referents and how this affects subsequent references (e.g., Almor, 1995; Gordon et al., 1993).

The studies in this chapter investigated the hypothesis that the grammatical subject and the focus of clefts are similar, in the sense that they both increase the accessibility of their referents. This comparison is important to make, for two reasons. First, psycholinguistic research has looked at either subjects (e.g., Gordon et al., 1993) or the focus of clefts (e.g., Almor, 1995), but not both. Second, while Almor's findings suggest that the foci of clefts are more accessible than nonfoci, his study did not explicitly investigate reference form. In the first part of this chapter, I will report on a discourse rating study that demonstrates a parallel between the factors Subjecthood and Focus.

3.2. Experiment: Topic and Focus

The study of categories like "topic" and "focus" is only beneficial to the extent that these categories are well defined. Although sentential topics are not always the same as grammatical subjects, they have often been associated with this role, especially when they are in sentence-initial position (Lambrecht, 1994:131; Reinhart, 1982). For this reason, I will use the grammatical subject in this study as an operationalization of sentential topic. To avoid confusion, I will use the term "grammatical subject" during the discussion of this experiment. Similarly, I will operationalize the category "focus" as the focus of cleft-like constructions such as "the one she saw was **Ann**". Although this is not strictly a wh-cleft, it has been called a "cleft with a lexical head" (Prince, 1978), and the term "pseudo-cleft" has

been used for both this construction (Chafe, 1976) and wh-clefts (Akmajian 1970). For the purposes of this study, I will categorize this construction with wh-clefts.

3.2.1. Methods

3.2.1.1. Procedure

The method for this experiment was a rating questionnaire that elicited off-line judgments about naturalness. Participants were asked to read short paragraphs and rate them for naturalness on a 7-point scale.[1] Each item was a three-sentence story. The first sentence set the context, but didn't mention any character by name. The second sentence introduced two characters by first name, either with or without a cleft construction. The third sentence started with an anaphoric reference to one of the characters from the second sentence, either as a pronoun or a name. A sample stimulus set is shown in Figure 1.

Figure 1: Sample stimuli for the Topic/ Focus experiment

A. Without cleft

i. The guests were nervously standing around in the living room, trying to decide which person to talk to.
ii. Ann decided to say hi to Emily first.
iii. Emily/She looked like the friendliest person in the group.
iii'. Ann/She hated to be in a room full of people where no-one was talking.

B. With cleft

i. The guests were nervously standing around in the living room, trying to decide which person to talk to.
ii. The one Ann decided to say hi to first was Emily.
iii. Emily/She looked like the friendliest person in the group.
iii'. Ann/She hated to be in a room full of people where no-one was talking.

[1] The instructions to the participants were these:

> Please read the following paragraphs. Some of them sound better than others. Rate each paragraph on how natural it sounds. If it sounds like something you'd say easily, circle 7. If it sounds like something really unnatural, something that you'd never say and you wouldn't expect anyone else to say it that way, rate it a 1. Be sure to rate the paragraphs on how natural they sound, not how likely it is that the events would occur.

There were 12 items, which were combined with 10 fillers. The fillers also had three sentences, but followed a different structure from the stimulus items. The stimuli and fillers were combined in three different orders for each list, to avoid ordering effects.

Thus, the experiment followed a (2x2x2) design, where the factors were 1) focus construction (yes vs. no); 2) 3rd-sentence continuation (1st mention vs. 2nd mention), and 3) type of reference for 3rd-sentence subject (noun vs. pronoun). There were 8 versions of each item, which were presented in a factorial design. That is, all items appeared an equal number of times in each condition, but only once for a given participant. Half the participants saw stimuli with "the one..."[1] constructions, and half saw stimuli with no focus construction.

Based on the findings of Gordon et al. (1993) and others, I expected that for the non-focused stimuli, the pronominal versions would be rated higher when the referent was the 1st mentioned character of the second sentence (the subject-referent), but not when the referent was the 2nd mentioned character. Given the findings of Almor (1996), I expected the opposite pattern for the focused stimuli, where I expected the pronominal versions to be rated higher for the 2nd mentioned entity (the focus-referent).

3.2.1.2. Participants

144 participants rated 12 stimuli and 10 fillers in exchange for one candy bar. Time of participation was approximately 10 minutes.

3.2.2. Results

Some participants were excluded from the analysis. One participant was excluded because he was not a native speaker of English, and one was excluded because he was a native speaker of South African English, not American English. "Native speaker" was defined as having begun to learn English by age 5. Two participants were excluded because they rated all stimuli items equally, and four random participants were excluded in order to make the cells even. The total number of participants included in the analysis was 136.

The ratings for each participant were first normalized to eliminate any variation in the baseline ratings for each subject. That is, some participants tended to rate the sentences higher overall than others. The noise resulting from this variation was reduced by taking advantage of the fact that all participants saw the exact same set of filler stimuli. Each individual's average score for all the filler stimuli was subtracted from the average filler score for all the participants, and this normalization score was then added to each stimulus score for that individual. The normalized scores were then submitted to a three-way ANOVA analysis. When a participant skipped an item for some reason, the average score for that person's stimuli items was substituted.

The results of the ANOVA showed a main effect for focus construction ($F1(1,134) = 5.469$, $p < .03$; $F2(1,11) = 38.638$, $p < .001$) and a main effect for 3rd-sentence continuation ($F1(1,134) = 8.996$, $p < 0.005$; $F2(1,11) = 4.736$; $p < .06$). The crucial result, however, was the three-way interaction among focus construction, pronominalization, and 3rd-sentence continuation, which was significant by subjects ($F1(1,134) = 6.690$, $p < .02$). The interaction did not quite reach significance by items ($F2(1,11) = 4.047$, $p < .07$), but the small number of items (12) suggests this is due to low power. No other factors were significant (p's >.1). The average ratings are shown in Table 1.

Table 1. Normalized average scores for each category in the Topic/Focus experiment. The higher rating is underlined.

	WITHOUT-CLEFT			WITH-CLEFT		
	pronoun		name	pronoun		name
NP1 referent	<u>5.14</u>	>	5.00	4.40	<	<u>4.68</u>
NP2 referent	5.19	<	<u>5.35</u>	<u>4.80</u>	>	4.71

[1] For variation, two items used the construction "the person..." instead of "the one...".

The average ratings for each category showed that when the second sentence contained an explicit focus construction, pronouns were rated higher for referring to the 2nd mentioned character, which was the focused entity, while names were preferred for referring to the nonfocused entity. Examples are shown in Figure 2.

Figure 2: Examples of preferences in Topic/Focus experiment in the With-Cleft condition

A. NP1 referent (nonfocus-referent)

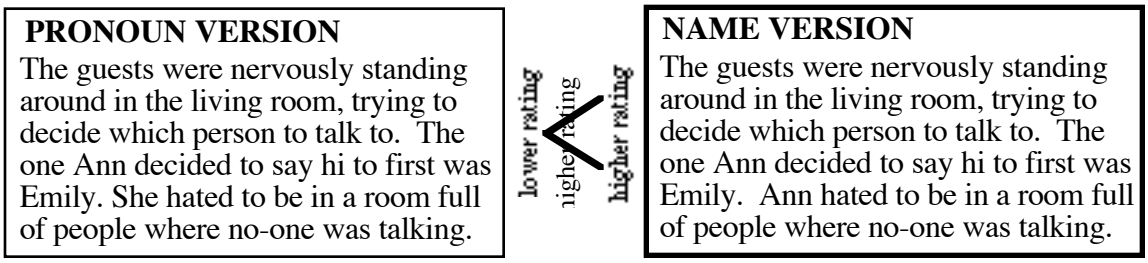

B. NP2 referent (focus-referent)

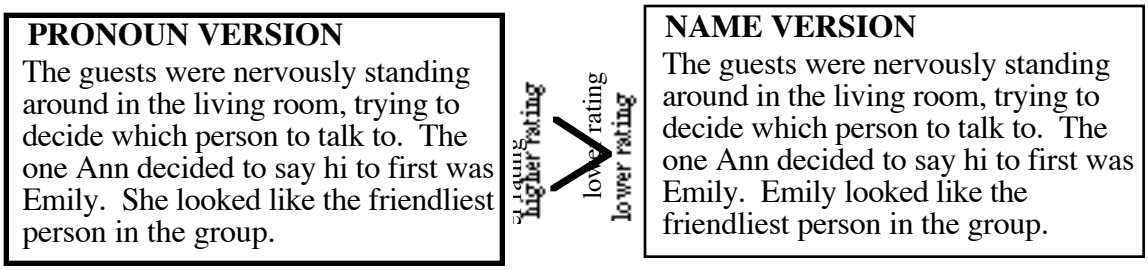

In contrast, when the second sentence did not contain the focus construction, the preferences were reversed: pronominal stimuli were rated higher for the 1st mentioned entity (the topic), and the stimuli with names were rated higher for the 2nd mentioned entity. Examples are shown in Figure 3.

Figure 3: Examples of preferences in Topic/Focus experiment in the Without-Cleft condition

A. NP1 referent (subject-referent)

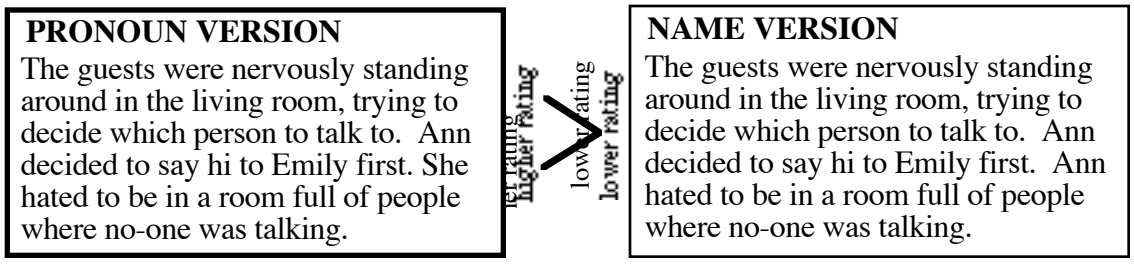

B. NP2 referent (nonsubject-referent)

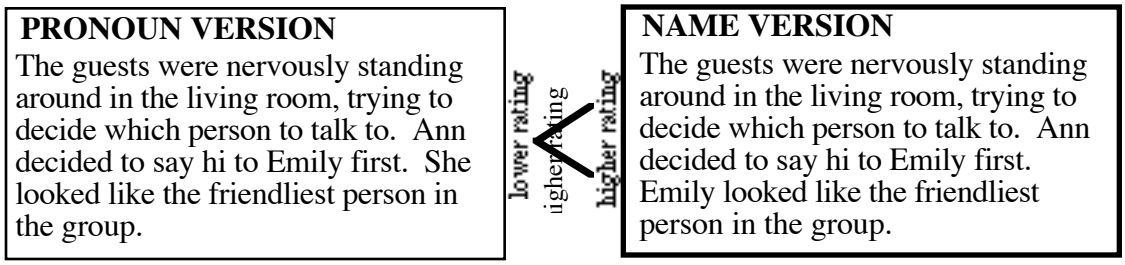

The results for this experiment show that readers preferred stimuli that used pronouns for both topical referents (subject-referents) and focused referents (the focus of "the one" constructions), but they preferred stimuli that used names for references to other entities. These results corroborated the findings of Almor (1996) and Gordon et al. (1993). If these preferences for reference form are assumed to reflect the mental representations of the referents, they suggest that the referents of both subjects and the focus of clefts are more activated than other arguments. Thus, despite the traditional opposition between topic and focus, the two categories may be more similar to each other than previously thought.

At the same time, the results of this study raise the question of why it should be that both constructions increase the accessibility of their referents. It is this question that the following corpus analysis aimed to answer.

3.3. Corpus Analysis: Subject vs. Focus

The results of the Subject/Focus questionnaire suggested that both subjects and the focus of clefts lead their referents to be more prominently represented than other referents in an utterance. What is it about these two constructions that does this? A common explanation for these two facts cannot depend on either properties of topics or properties of foci. An explanation in terms of "aboutness" or "givenness" would only account for the prominence of subjects, and an explanation in terms of the prominence of new information would only account for the foci of clefts.

The hypothesis that I investigated with this corpus analysis was that both subjects and the focus of clefts signal that there is a high likelihood that their referents will be mentioned again in the subsequent discourse. That is, both constructions may be pointers to the topic of the **following** utterance. In a "normal", nonclefted utterance, the best bet for the topic of the following utterance is the topic of the current one, since speakers usually talk about the same thing for extended periods of time. On the other hand, a clefted utterance is a marked construction that the speaker may employ to indicate that the topic will shift to the referent of the focus. This hypothesis is consistent with Sgall et al.'s proposal (1986:58) that the focus of one utterance is related to the topic of the next.[1] If this is the case, I expected that after nonclefted utterances, speakers would refer to the referent of the subject more often than to other elements in the utterance, but that after clefted utterances, speakers would refer more often to the focus of the cleft.

To investigate this hypothesis I conducted a corpus analysis, using the Aligned-Hansard Corpus from 1986. The Aligned-Hansard corpus is a collection of transcripts from the Canadian Parliament, so the discourse it represents is natural and communicative, albeit formal. The discourse is spoken, although one might imagine that the speakers may have had prepared notes at their disposal. However, the transcripts are doubtlessly edited, as

---

[1] This proposal in undoubtedly influenced by a similar comment by Weil, an earlier scholar in the Prague school tradition (1844, 1887, as cited by Tomlin et al., 1997).

they contain no disfluencies or partial sentences, and some segments may be translated from French.

Using the Aligned-Hansard corpus, I compared nonclefted utterances, as in 2a, with object-clefted wh-clefts, as in 2b.

(2) a. We must have legislation and charters which will allow people to come before the courts with their concerns and grievances with some effect and justice.

b. What we called for in the committee were five different principles to be enshrined in this particular piece of legislation.

My hypothesis was that speakers would refer to the subject-referents and focus-referents more frequently than to other referents from the preceding utterance.

3.3.1. Method

I analyzed the 1986 Aligned-Hansard corpus by extracting two types of utterances: wh-clefts (also sometimes called "pseudo-clefts") and nonclefted sentences. To find wh-clefts I searched for "What" (with a capital W) that was not followed by a question-mark. I only considered those utterances that clefted the object or object-of-PP, because I was interested in the comparison between the embedded subject and the focused NP. I also discarded cases where there was no following utterance, or where the following utterance belonged to another speaker. This method yielded 146 total examples of wh-clefts. The comparison set of data was a random sample of non-clefted sentences, which I assembled by opening each file from 1986, scrolling down three pages, and taking the utterance where my cursor landed. In this set I excluded questions, sentences with dummy subjects or nonreferential subjects (like the generic "one"), and cases where there was no following sentence by the same speaker. This method resulted in 263 examples.

For each sentence in my sample, S1, I identified the next "independent" utterance, S2. I defined "independent utterance" as a clause which is finite, not a sentential

complement of the matrix clause, and not a relative clause. I then coded S2 for the first reference to something that had been referred to in S1 (if any). The categories I found in this analysis are presented in Table 2.

Table 2. Identity (in S1) of the referent of the first referring NP in S2

Nonclefted Sentences (S1)	
1. Subject	matrix clause Subject
2. Part of Subject	part of matrix clause subject (e.g., "I" when previous subject was "my Party")
3. Related to Subject	something related to the matrix clause subject, such as another member of a set
4. Object/Obj of PP	matrix clause object or object of PP
5. Part of Object	part of matrix clause object
6. Subordinate Subject	subject of a subordinate clause
7. Subordinate Other	another part of the subordinate clause
8. Whole	the whole sentence, or an entire non-NP constituent within the sentence (such as an entire PP, VP, or subordinate S).
9. None	there is no referent in common with S2
Clefted Sentences (S1)	
1. Embedded Subject	subject of embedded clause (e.g., What we need is...)
2. Related to Embedded Subj.	Related to subject of embedded clause
3. Focused NP	NP in the focus of the Wh-cleft
4. Part of Focused NP	part of the NP in the focus of the Wh-cleft
5. Related to Focused NP	related to the NP in the focus of the Wh-cleft
6. Embedded Object	object of embedded clause
7. Unclear	
8. Whole	the whole sentence, or an entire non-NP constituent within the sentence (such as an entire PP, VP, or subordinate S).
9. None	there is no referent in common with S2

As mentioned above, I looked at S2 to find the first NP that had the same referent as any of the NPs in S1, and if so, which one. The hypothesis behind this methodology was that

certain categories (namely, subjects and focused NPs) may be indicators that their referents are likely to be mentioned again in the following discourse. If more than one referent from S1 was mentioned in S2, only the first one was analyzed. Once a referent was identified, I noted how it had been referred to in the **first** utterance (S1) -- that is, what the grammatical function of the referring expression was in S1. The following examples show how particular cases were coded.

Table 3. Examples from Corpus Analysis

first referent	**example**	
subj	S1:	The charter is giving individuals a new opportunity to seek redress when their rights are infringed upon.
	S2:	It also imposes on all an obligation to be tolerant of the rights of others.
subj	S1:	Actually, the countries of South-East Asia have recently penetrated the fur market.
	S2:	Are we going to let them have a free rein?
obj of PP	S1:	I have talked to other whites who have worked in the South African police force
	S2:	and through them I have learned something about the enormously tragic environment of the prison system in South Africa and the harshness of the treatment to which many . . .
focus	S1:	What the Government finally got was a deck hand whose name is Mr. Lander.
	S2:	He has been busying himself rearranging the deck chairs on the deck of the Titanic.
focus	S1:	What we have now is a breath of fresh air,
	S2:	and the people of Canada welcome it.

The codings for each utterance were tabulated in terms of four categories. For nonclefted utterances, these categories were: 1) Subject, 2) Object, 3) Other (references to other referents or the whole utterance in S1), and 4) No referent from S1. For clefted utterances, the categories were: 1) Embedded Subject, 2) Focus, 3) Other, and 4) No referent from S1 or Unclear. The categories Subject, Embedded Subject, Object, and Focus

included cases termed "part of Subject/ Focus/ etc." or "related to Subject/ Focus/ etc.". Examples of these categories are provided in Table (4).

Table 4. Examples of "Part of" and "Related to" codings.

(1) Part of Subject

a. The unfortunate aspect of the dilemma facing our producers is that the problem is not of their own doing. Our farmers are the innocent victims of the costly madness of subsidies which has gripped the European Economic Community and,

b. From time to time my Party speaks about the need to democratize these centres of power and put them in the hands of ordinary Canadians. I suggest that those are some constructive and positive means by which we could seek to do that.

(2) Part of Object

I would like to quote the words of the Hon. Member, the Liberal health critic. He said: "Higher prices might be justified if it led to more drug research in Canada. ...."

(3) Part of Focus

What the Hon. Member has not answered with respect to Challenge '86 '86, has not refuted, is the allegation, the observation, that there has been a significant shifting of funds from the non-profit service and municipal sector to private enterprise. I have no problem with funding training opportunities in private enterprise.

(4) Related to Embedded Subject

What the Government is entering into now is a bilateral negotiation by appointing trade envoys. The U.S. administration said to Canada that . . .

(5) Related to Focus

What we are really talking about is a farm crisis. Just the other day the Committee on Agriculture met with Dr. Lillian Walker and Dr. James Walker.

One might wonder about the rationale for categorizing cases where the speaker makes reference to a part of an argument or something related to an argument. The idea behind this choice is that this corpus analysis is measuring discourse patterns -- that is, how often speakers continue talking about the referents of subjects, and how often speakers continue talking about the referents of the focus of clefts. Whatever the speaker's choice for continuing a discourse, the listener will interpret the following utterances in terms of what

came before, looking for coherence in the discourse. In some cases coherence is established by means of reference to something previously mentioned, in some cases comprehenders engage in what Clark and Haviland (1977) call bridging inferences. If the speaker mentions something that represents only a subpart of the referent of an NP in the previous sentence, or something that relates to a referent in the first sentence, the listener will need to access the representation of that referent in order to establish a connection (Loewenstein et al., 1998). This idea is also reflected in claims by Almor (1995:33) and Prince (1981:251) that suggest that reference to one member of a set causes the listener to access the other members of the set as well. Therefore, to understand how discourses are frequently continued, it is necessary to include cases where an argument is the topic of the discourse, even if it is not explicitly or fully mentioned.

3.3.2. Results

I tabulated the number of continuations in each of the categories mentioned above. For nonclefted utterances, I calculated the percentage of continuations with subject-referents, object-referents, other-referents, or no referents from S1. For clefted utterances, I calculated the percentage of continuations with embedded-subject-referents, focus-referents, other-referents, or no referents from S1. I compared clefted and nonclefted utterances with each other, comparing the categories subject-referent with embedded-subject referent, and the categories object-referent with focus-referent. A chi-squared analysis of this distribution showed that these four categories behaved differently with respect to nonclefts and clefts ($\chi^2(3)=82$, $p<.001$).

The Reference Continuation patterns for Clefts and Nonclefts are presented in Figure 4, illustrating that the most frequent type of continuation for nonclefted utterances is with reference to the subject-referent, but for clefted utterances it is with reference to the focus-referent. The percentages of each bar do not add up to 100%, because they were calculated out of all references in each category. However, here I am only showing

references to the subject or object in nonclefted utterances or the embedded subject or focus in clefted utterances.

Figure 4: Reference Continuation for clefted and nonclefted utterances

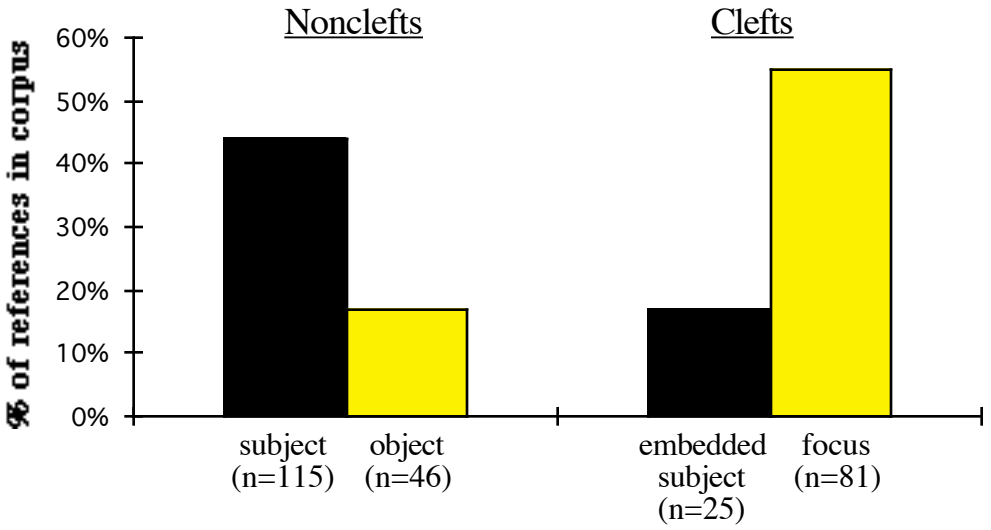

This figure shows that wh-clefted and nonclefted utterances are associated with quite different patterns in discourse. For nonclefted utterances, the subject-referent (usually the first-mentioned referent) has a much higher probability of being mentioned in the following clause than any other referent, whereas for clefted utterances, it is the focus-referent (usually not the first-mentioned referent) that has the highest probability of being mentioned in the following clause. The corpus analysis thus parallels the experiment in showing a similarity between the categories subject and focus.

3.4. Discussion

There were two main findings in this chapter. First, the rating experiment showed that readers preferred passages that used pronouns to refer to the subject-referents of

nonclefted sentences and the focus-referents of cleft or cleft-like constructions. Second, the corpus analysis showed that the referents of both subjects and foci were the entities most likely to be referred to again in the following sentence. These results indicate that the factors Subjecthood and Focus are both associated with two things: a) Reference Form, and b) Reference Continuation. This association, depicted in Figure 5, follows the same pattern that emerged in the studies in chapter 2.

Figure 5: The double association for Subjecthood and Focus

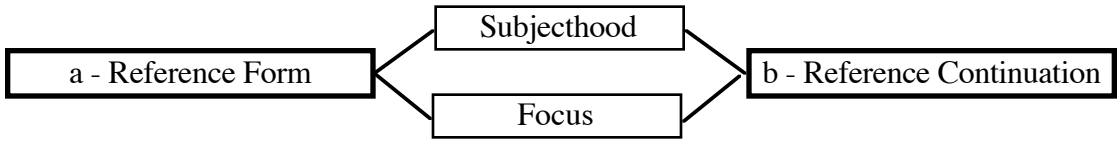

These results raise a number of questions. Why are both factors associated with Reference Form? I conducted the corpus analysis to try to understand why pronouns were considered more natural for reference to both subject-referents and focus-referents. The corpus analysis showed that both were also associated with Reference Continuation. But why do speakers refer more often to the referents of subjects and foci? And why does Reference Continuation consistently pattern with Reference Form? I will consider each of these questions separately.

3.4.1. The relationship of Subjecthood and Focus with Reference Continuation

First I will address the question of why speakers refer more often to subject-referents and focus-referents. That is, why do the patterns found in the corpus analysis exist? I argue that they arise because language use is a joint activity (H. Clark, 1996:23). There are two properties of joint activities that drive patterns of reference: 1) joint activities are goal-driven, and 2) joint activities require the coordination of common ground. These properties lead to the following four generalizations:

1. For any given discourse participant, some referent representations are highly activated.
2. The highly activated referents are the ones that participant will refer to more.
3. The highly activated referents are the ones that participant will refer to using subjects and foci.
4. Other discourse participants can use the speaker's choice of words as an indication of the speaker's representation of the discourse referents, and construct their own representations to match the speaker's.

I will address each of these claims in turn, and discuss why they result from the generalization that language use is a joint activity.

3.4.1.1. For any given discourse participant, some referent representations are highly activated.

Joint activities are goal-driven (H. Clark, 1996). Because of this, they are organized around certain themes, and certain referents will play more important roles than others. For example, in a discussion about how to make Turkey Tetrazzini, "turkey" is more important than "pot". In a discussion about painting, "paint" is more important than "water". In any given joint activity, there may be many goals, some private and some public. Crucially, however, there is one principal goal, which H. Clark calls the "domain goal". This is the publicly established goal that all participants in the activity are nominally working towards (H. Clark, 1996:34).

What does this mean in terms of the speaker's mental representation? Speakers do not merely interact with the world; they also maintain mental representations of the objects and events in the current situation, both those that are physically present and those that are described linguistically. This level of representation is sometimes called a "mental model", or "situation model" (H. Clark, 1996:53). The cognitive representations of discourse entities may be understood via the mechanism of activation, introduced in chapter 1. Thus, speakers come to a discourse situation with a conceptual representation of their goals and

intentions, in which some referents are more highly activated than others. As the discourse proceeds, the activation of referents will change, as the speaker's goals and intentions change.

3.4.1.2. The highly activated referents are the ones that are frequently referred to.

The goal-driven nature of language use has a further consequence: the conversation will revolve around the referents that are important to the goals of the speaker. Since those referents are also the ones that are highly activated for the speaker, this means that speakers will refer more to the referents that are highly activated for them. Thus, in a discussion about how to make Turkey Tetrazzini, "turkey" will be mentioned more often than "pot", and in a discussion about painting, "paint" will be mentioned more often than "water".

Consider the following illustration of how activation and frequency of reference relate to each other. Imagine a couple, Jane and Mark, who are planning a dinner party. Jane says to Mark, "Marsha called. She can't come tonight, because she has the flu." As Jane initiates the conversation, she has a set of goals and intentions. These goals and the referents associated with them are present in her mind. That is, they are activated. But she may have many things in her mental representation, and not all are equally activated. She may also be thinking about vacuuming the house, or the implications of Marsha not coming (for example, perhaps Marsha was supposed to give another guest a ride). The representations of the vacuum and the other guest are also active, but less so than her representation of Marsha. As Jane begins the conversation, she intends to talk about Marsha, so her representation of Marsha is highly activated. This is also the referent she will most likely repeatedly refer to during the conversation that ensues.

Thus, from Jane's perspective we have the following scenario. Certain referents are more activated in her mind, because they are important to her goals and intentions. Because these referents are important to the situation, she will tend to refer to them more frequently than to less important referents.

3.4.1.3. People use subjects and foci for referents that are highly activated in their models

The fact that language is a joint activity leads to a third generalization about language use: the highly-activated referents in the speaker's mind are the ones the speaker will refer to in subject or focus position. The reason for this lies in the most basic property of joint activities: they necessarily take place between 2 or more people, who need to coordinate with each other at multiple levels. In linguistic exchanges, participants need to coordinate to establish reference (Clark and Marshall, 1981; Clark and Wilkes-Gibbs, 1986), negotiate the meaning of their utterances (H. Clark, 1994, Part IV), and establish other aspects of the common ground (H. Clark, 1996). One of the things they need to coordinate is the discourse status of each referent.

In the preceding discussion, I focused on the speaker, Jane, and the representations she has as a result of her goals and intentions as she starts a conversation. But this picture ignores what is going on in the mind of Mark, her interlocutor. Mark does not have the same representations as Jane, but in order for communication to be successful, the two of them must believe that their mental representations are fairly similar (see H. Clark, 1996:49). For their representations to be similar, they also must coordinate the cognitive status of each referent. How can they do this?

Linguistic expressions offer one tool for Jane and Mark to coordinate the cognitive status of discourse referents. As many researchers have noted, linguistic expressions are pragmatically specialized (see, among others, Du Bois, 1985, 1987; Prince, 1992; Ward and Birner 1995, 1996). Speakers choose particular forms and expressions according to the discourse role of the information they are expressing. Two types of such pragmatically specialized expressions are grammatical functions and cleft constructions.

Grammatical functions have received much attention, both in terms of their pragmatic specialization and their psychological effect on discourse comprehension. In particular, subject position is associated with referents that are topical (e.g., Reinhart, 1982),

given (e.g., Du Bois, 1987; Prince, 1992), and psychologically salient (e.g., Gordon et al., 1993). The focus of cleft-like constructions has received less attention, but has also been noted as an informationally specialized position (e.g., Prince, 1978), that is also psychologically salient (e.g., Almor, 1995; Sidner, 1983).

Although subjects and foci play different pragmatic roles, my results demonstrate that their roles have some properties in common. The results of the elicitation experiment are consistent with previous claims that both subjects and foci make their referents more salient, since pronoun usage has generally been linked with the salience of referents. The results of the corpus analysis further suggest that the referents of subjects and foci are highly likely to be referred to again. Because subjects and foci share this property, they are both useful mechanisms for the discourse participants to coordinate their mental representations.

To illustrate this coordination, let's return to the example above. Jane knows which referents are central to her goals and intentions, and therefore which referents are highly activated for her. Mark has his own representations, but he doesn't know what Jane is thinking until she speaks. When Jane says "Marsha called," Mark knows that "Marsha" is the subject of the sentence, and therefore is highly activated for Jane. In order to cooperatively participate in the conversation, Mark also instantiates a highly active representation of Marsha in his mind. In this way, Jane and Mark have achieved a certain level of coordination in their conversation.

Just as there are two participants in this fictitious interaction, there are also two sides to their coordination. First, let's consider the processes for the speaker. The results of the corpus analysis revealed that speakers tend to refer again to the referents of subjects and clefts. As I said earlier, speakers tend to refer to some referents more often because they are important to their goals and intentions. This means that speakers use subjects and clefts more often to refer to referents that are important, and therefore highly activated.

So why do speakers choose subjects and foci to refer to expressions that are activated? This issue has two underlying questions. The first question is a diachronic one: why did the grammatical function and the cleft construction evolve to have the functions they do? I can merely speculate on the answer to this question. The salience of subjects may derive from the fact that they are often used to indicate the agent role of a proposition. If agents are salient, salience could have become grammaticized into the grammatical subject.

The second question is one of on-line processes of language production. Why do speakers choose subject and focus position on any given occasion? Many factors are likely to influence this choice, but two are particularly notable. First, grammatical constraints in a language restrict the range of options available for expressing a particular idea. For example, there is a limited number of ways to express the idea "Marsha called (me)," which might include "I was called by Marsha," or "That was Marsha on the phone." Second, the speaker is influenced by the pragmatic function of the available expressions. One reason for this may be that their pragmatic specialization may result from the cognitive processes of language production. Constraints on planning and production may lead speakers to produce highly activated referents early in an utterance (Arnold et al., 1998; Tomlin et al., 1997), or in a prominent place like focus position. A second reason that the speaker places important referents in subject or focus position may be as a signal to the listener that those referents should have highly activated representations (Clark and Haviland, 1977). In either case, the result is that speakers use subject and focus positions more often for referents that are highly activated in their mental representations.

3.4.1.4. Listeners are sensitive to the pragmatic role of subjects and foci

The other side to coordination during discourse concerns the listener. How does the listener discover which referents are more activated for the speaker? One way is by attending to the speaker's use of expressions like subjects and foci: once a referent has

been mentioned in these positions, the listener knows that it should be represented with high activation. But why do subjects and foci play this role?

This question can be answered at one level in terms of topicality. Topical referents are those that the discourse is about (Reinhart, 1982). From the perspective of one discourse participant, these are the referents most central to his or her goals and intentions. Subjects are often used for highly topical referents, those that the discourse is about. While foci are not topical like subjects are, they might signal a topic shift, such that the focus will be the topic of the following utterance (Sgall et al., 1986:58). But this story raises another question: how do these constructions come to be associated with discourse patterns like this?

The answer, I propose, is through an individual's experience with the language. For example, Jane's experience with English leads her to realize that at any given point in a discourse, some referents are more central than others. She realizes this because some referents get referred to more often than others. The reference may be indirect, as with the categories "part of" and "related to" in the corpus analysis.[1] Thus, her understanding of "topical" is directly related to her understanding of how often an entity is referred to during a discourse segment. This conception of "topical", similar to Givón's topicality continuum (Givón, 1983a), allows referents to be topical to varying degrees. It also treats topicality as a dynamic property: the relative importance of discourse referents varies over time. Furthermore, referent activation varies as discourse participants negotiate the topic of the discourse, the perspective they take on it, and the direction the discourse is going. This dynamic nature of language use adds several degrees of complexity to the discussion of these issues, so I have simplified the example above by considering only a single point in time at the beginning of a conversation.

---

[1] The "reference" may also be non-linguistic, such as pointing at something or otherwise bringing it to the attention of the other discourse participants. See §5.3.2.3 for a discussion of this point.

With this conception of topicality, I propose that people learn an association between subjects, foci, and the way their referents are used in discourse. One thing they have in common is that both subject- and focus-referents are very likely to be mentioned in the following discourse. Each time a person hears a subject or focus used, this constitutes an exemplar of that category. Through a lifetime of collecting exemplars of subjects and foci, people develop categories that include the information about how they are used in discourse.[1] In other words, people learn that subject-referents and focus-referents are associated with a high degree of topicality in the following discourse.

An important feature of this association is that it is probabilistic. That is, subjects and foci are not **always** used for referents that will be topical in the following discourse. Thus, it would not be efficient for people to draw a categorical association between these categories and topichood. Instead, the relationship is partial and probabilistic, and the categories for subjects and foci should include this information. Furthermore, my conception of topicality is not the traditional one. Topicality here is an index of the degree to which referents are central to the goals and intentions of the speaker.

I further propose that the regular association of Subjecthood and Focus with Reference Continuation makes the choice of these expressions informative for the listener. This interpretation draws on recent proposals for constraint-based models of language processing. These models suggest that as people gain experience with language, they also gain experience with how words and structures are associated with particular meanings and functions. In discourse processing, this idea can be extended to consider that people also learn through experience to associate structures with particular discourse patterns. This information is stored as a part of the subject and focus categories, and is available for use in novel discourse situations.

---

[1] I am leaving it open whether the process of category construction continues throughout the life of an individual, or whether it slows down or ceases in adulthood. It is possible that once a critical number of examples has been encountered, the category is so robust that further exemplars do little to change it.

This proposal applies to the problem of coordination as follows: From the listener's perspective, once a referent has been mentioned in subject or focus position, there is a relatively high probability that it will be mentioned again. This information is co-extensive with knowing that the referent is important to the speaker: if the referent is important to the speaker, it is likely to be mentioned repeatedly during the conversation, and vice versa. Thus, once a referent is mentioned as subject or focus, the listener can infer that it is important to the speaker. With this knowledge, the listener can instantiate a representation of the referent with high activation. In this way, speakers and listeners coordinate the activation of discourse referents.

3.4.1.5. Summary

In sum, I have argued for several generalizations on the basis of language use as a joint activity. Speakers refer more often to things that are important to their goals and intentions, which are the things that are activated in the speaker's representation. They also tend to use pragmatically specialized expressions, like subject and focus, to refer to those things. The regularity of using subjects and foci for highly activated referents is stored in the categories subject and focus. This makes these expressions good tools for the coordination of the cognitive status of referents in the representations of speakers and listeners. In the following section I will address the question of how this influences the choice of reference form for subject-referents and focus-referents.

3.4.2. The relationship of Subjecthood and Focus with Reference Form

The results of the rating experiment raise the question of why readers prefer pronouns for referring to both subjects and foci. This question is relevant to the issue of why speakers use pronouns on some occasions, but fuller forms of reference on others.

In chapter 1, I showed how this question has been addressed by appealing to concepts such as salience or accessibility. In this dissertation I have appealed instead to activation as the mechanism for representing the cognitive status of referents. While there

are many reasons for a referent to become activated in the mind of the speaker, one way is when it is central to the speaker's goals and intentions.

I have also presented the problem of reference as a problem of coordination, following H. Clark and his colleagues (e.g., H. Clark, 1996; Clark and Marshall 1981; Clark and Wilkes-Gibbs, 1986). Subject, focus, and other expressions all provide tools for coordinating the activation of mental representations for each discourse referent. One reason that subjects and foci can function as tools in this way is because the referents associated with these positions have a higher probability of being referred to again than other referents. This probability allows the listener to infer which referents are activated for the speaker, and thereby allows speakers and listeners to coordinate the level of activation for each discourse referent.

An important part of this proposal is that the association of subject- and focus-referents with high activation develops through the association of their categories with frequent subsequent reference, and thus the perception that these referents are topical for the speaker. Therefore, once a referent has been mentioned in subject or focus position, there is a relatively high probability that it will be mentioned again in the following utterance. This property facilitates the processing of a subsequent reference, if it does occur.

There are two reasons that reference processing is facilitated for subject- and focus-referents, both based on the assumption that reference processing involves activating the referent of the anaphor. First, when the anaphor itself is encountered, there is a higher probability that the subject- or focus-referent is the referent of that anaphor. The probabilistic information associated with these categories is relevant at this point in processing, such that subject- and focus-referents will become partially activated, based on the probability that they are the referent of that anaphor. Thus, if the anaphor does indeed refer to the subject-referent or focus-referent, processing will be facilitated. If it doesn't, anaphor resolution will be inhibited.

Second, the process of activating the referent is easier if the referent is **already** activated. As I argued above, discourse participants work to coordinate their representations of the discourse entities. The referents of subjects and foci are often the ones that are activated for the speaker, because they are often the ones that are important to the speaker's goals and intentions for the following discourse. The listener expects the speaker to produce a coherent discourse, such that subsequent utterances should relate to each other in some way. Thus, it is advantageous for the listener to construct similar representations of the discourse referents. Although the listener does not know for sure which ones are activated for the speaker, probabilistic information associated with subject and focus categories allows the listener to make good guesses about the speaker's representation. Listeners therefore partially activate subject- and focus-referents even before a subsequent anaphoric reference is encountered.

The ease of accessing subject- and focus-referents has important implications for the speaker's choice of reference form. When comprehension is facilitated, the listener needs less specific information to correctly access the speaker's intended referent. Therefore, less specific forms of reference, like pronouns, will work. Previous research on speech processing has shown that speakers do in fact take advantage of the ability to use less specified forms when comprehension is facilitated. The referents of subjects and foci are easier to access, therefore these positions are associated with an increased use of pronouns for subsequent reference.

3.5. Conclusions

3.5.1. Summary

Speakers talk about things because they are trying to communicate a message, and certain referents are more central to that message than others. This results in a consistent pattern of discourse in which speakers refer to certain referents more often than to others. Speakers also use certain structures, like subjects and foci, for those same referents that they tend to refer to again. Thus, the referents of subjects and foci are associated with a high

probability of subsequent mention. From the listener's perspective, this high probability makes subsequent references easier to understand. And, where comprehension is facilitated, the speaker can use less specified forms of reference to communicate with the listener.

To illustrate this proposal, let's return to the conversation between Jane and Mark. When Jane initiates the conversation with Mark, she has certain goals and intentions. The referents that are central to Jane's goals will be highly activated in her mind, and the referents that are less central to her goals will be less activated, but still present in her mental representation of the situation. As before, let's assume that Jane has two goals: one is to communicate the contents of her conversation with Marsha, and one is to vacuum the house. At this point, however, Jane knows that these referents are only active for her, and are not part of her common ground with Mark. A partial representation of their situation models is shown in Figure 6. For simplicity I have listed referents using names and phrases, but in actuality these are meant to be non-linguistic, conceptual referents. Referents with relatively high activation are represented with capital letters.[1]

Figure 6: Some objects in Jane and Mark's situation model

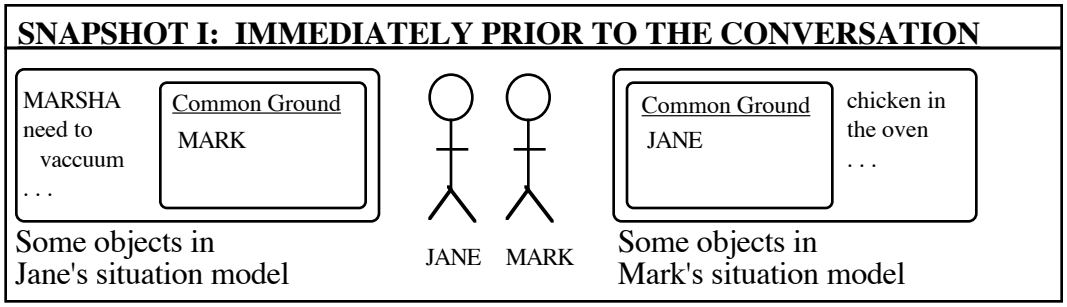

[1] The simplified representation of activation in this figure should not be taken to imply any commitment to the relative activation of the discourse referents. That is, activation is likely to be graded, but the binary distinction between capital and lowercase fonts only allows for two levels of activation, a simplification sufficient for the present purposes.

At the beginning of the conversation, Mark also has a certain set of referents activated in his situation model. For example, he may be concerned about a chicken that he has in the oven. When Mark and Jane come in contact with each other, they also begin to keep track of information that is shared between the two of them. That is, they begin to construct a common ground (H. Clark, 1996). At this point, descriptions of Jane's and Mark's situation models become increasingly complex. However, some features of their interaction are supported by my findings. At the start of their conversation, Jane says to Mark "Marsha called." As she says this, this statement becomes a part of their common ground, and draws attention to their common acquaintance, Marsha. At the same time, they both know that Jane has referred to Marsha in subject position. This is one of the cues from which Mark can infer that Marsha is likely to play a part in the following discourse, so this referent becomes highly activated for Mark.

Figure 7: Some objects in Jane and Mark's situation model

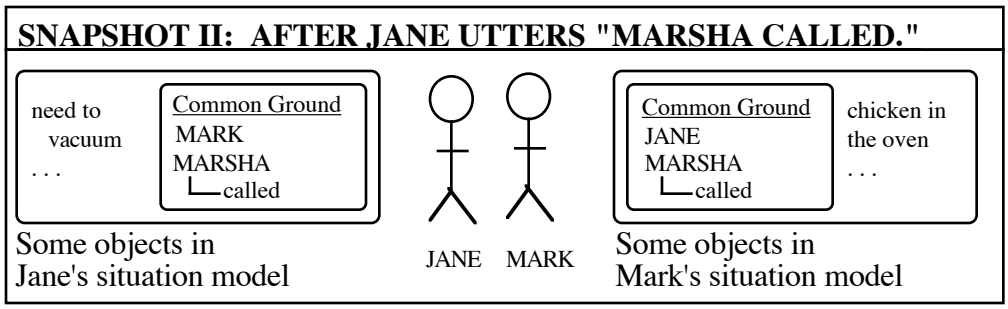

After Jane has established Marsha as a mutually activated referent, either she or Mark can refer to Marsha with only a pronoun, and expect that the reduced form of reference will be adequate for communication. In this way, the association of Marsha with subject position mediates the status of Marsha in Jane and Mark's common ground.

3.5.2. Theoretical implications of this approach

This account explains how the association of Subjecthood and Focus with Reference Continuation also results in the association of these factors with Reference Form. This account is summarized in Figure 8.

Figure 8: How subject and focus coordinate the cognitive status of a discourse referent for discourse participants

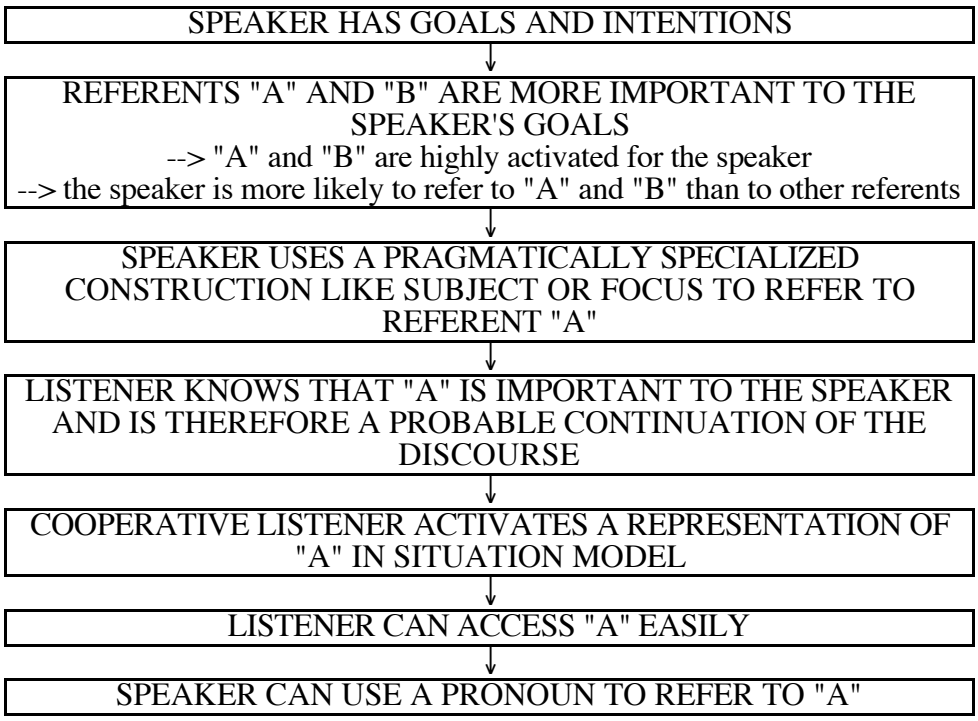

I have used activation as a measure of a referent's cognitive status, and therefore the thing that needs to be coordinated. This approach is consistent with past accounts of reference form that relied on the concepts salience or accessibility. In my proposal, highly activated referents are those deemed salient by the discourse participants, and hence more accessible than less activated referents.

At the same time, my proposal contrasts with previous accounts, in the following ways.

First, the activation account does not depend on "salience" as a primitive. The goal of previous accounts has generally been to identify factors that lead a particular referent to

be salient or accessible (e.g., Ariel, 1990; Gordon et al., 1993; McDonald and MacWhinney, 1995; Stevenson et al., 1994). My results are wholly consistent with these studies, because they show that the linguistic positions of subject and focus are associated with salient referents. However, I have gone further: The corpus analysis shows that subjects and foci are also associated with a tendency to refer more often to their referents, which I have argued is an index of how important they are to the discourse. This suggests that salience has its roots in the way people use language: the things people refer to more often are the things represented with greater activation, and are thus viewed as more salient.

By anchoring salience in patterns of languages use, my approach has a second advantage: it does not rely on a rule positing that subjects and foci are salient positions. Previous accounts have treated the salience of subjects and other positions as a characteristic of the positions themselves (e.g., Gordon et al., 1993; McDonald and MacWhinney, 1995; Stevenson et al., 1994). This has been used to explain why pronouns and other less-specified forms of reference can be used, through the stipulation that less-specified forms are used when the referent is more salient. This rule-based approach is also seen in Centering Theory (e.g., Gordon et al., 1993; Grosz et al., 1995), which uses subject position as a categorical marker that one entity is the "backward-looking center", or the most accessible referent. The same claim holds for the salience of sentence and discourse topics (e.g., Ariel, 1990; Gernsbacher, 1990). In contrast, I have suggested that the referents of subjects and foci are more salient than other referents because speakers favor these positions for referents that they are likely to refer to again later.

A third advantage of my account is that it looks to general properties of language use to suggest how subjects and foci mediate the mental representations of the speaker and listener. My approach builds on two traditions in psycholinguistics. One is H. Clark's (1996) theory of language use as a joint activity, which focuses on the goal-driven, communicative nature of language. The other is the constraint-based models developed by MacDonald, Tanenhaus, and colleagues (e.g., MacDonald et al., 1994; Trueswell and

Tanenhaus, 1994; Tanenhaus and Trueswell, 1995), which suggest that people keep track of how language is used, and use their knowledge of usage patterns during language processing. Both of these traditions consider linguistic competence to be the direct result of experience with language in the social and pragmatically situated uses of day-to-day life.

Finally, my account has the advantage of explaining the results of the experiment and the corpus analysis just presented. The rating experiment revealed that subjects and foci are viewed similarly in terms of reference form: readers preferred passages that used pronouns to refer to the referents of both subjects and foci, but preferred names to refer to other referents. The corpus analysis showed that subjects and foci share a second characteristic, a relatively high probability that their referents will be mentioned in the following sentence. The activation account has suggested how these patterns relate to one another. The association of Subjecthood and Focus with Reference Continuation is stored as a part of the categories subject and focus, and this allows discourse participants to use these positions as tools to coordinate the cognitive status of the discourse referents.