UNIVERSITY OF CALGARY

The Role of Syntactic Form in Incremental Referential Interpretation

by

Jodi Edwards

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF ARTS

DEPARTMENT OF LINGUISTICS

CALGARY, ALBERTA

JULY, 2003

UNIVERSITY OF CALGARY

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled, "The Role of Syntactic Form in Incremental Referential Interpretation", submitted by Jodi Edwards in partial fulfillment of requirements for the degree of Master of Arts.

____

Supervisor, Dr. Craig G. Chambers, Departments of Linguistics and Psychology

____

Dr. Hotze Rullmann, Department of Linguistics

____

Dr. Chris Sears, Department of Psychology

July 29/03

Date

Abstract

There is increasing evidence that semantic interpretation proceeds incrementally. As a noun phrase (e.g., *the big red triangle*) unfolds in time, each successive word restricts attention to candidates whose properties match the available information. However, not all properties can be expressed by prenominal adjectives like *big / red* - some require postnominal constructions (e.g., the square *with the diamonds*). Experiments 1 and 2 investigated whether these 'codability' differences are considered when evaluating candidates for an unfolding noun phrase. Eye movements were monitored as participants heard instructions to manipulate objects. Results demonstrated that the form typically used for a property plays a predictive role in referential processing. Experiment 3 investigated if codability influences the resolution of syntactic ambiguities. Results did not confirm this prediction. Overall, the experiments suggest that listeners assign interpretations by reasoning about the speaker's communicative intentions, including those underlying the use of particular syntactic forms for referring expressions.

Acknowledgements

I greatly appreciate the guidance provided by my supervisor, Dr. Craig Chambers, over the two years it took to complete this thesis. His knowledge, patience, and outstanding research abilities were essential to this project and my success as a graduate student. I would also like to thank the members of my committee, Dr. Hotze Rullmann and Dr. Chris Sears. Their helpful comments greatly improved the quality of my work. I am especially grateful to my husband, Chris Baskin, for all of his patience and support throughout this process. Finally, I would like to thank Kristinn Meagher, Cory Gerritsen and Valerie San Juan for all of their help in running these experiments.

TABLE OF CONTENTS

Approval Page....ii
Abstract....iii
Acknowledgements....iv
Table of Contents....v
List of Figures....vii

CHAPTER ONE: INTRODUCTION....1
Structural Information and Real-Time Language Interpretation....1
Referential Interpretation....5
Current Research....10

CHAPTER TWO: EFFECTS OF FORM IN ON-LINE REFERENTIAL INTERPRETATION....12
Codability....13
Production Study....14
Method....16
Participants....16
Materials and Design....16
Procedure....16
Results....17
Discussion....17
Experiment 1....19
Predictions....20
Method....22
Participants....22
Materials and Design....22
Visual Counterbalancing in Filler Trials....23
Linguistic Counterbalancing in Filler Trials....24
Procedure....24
Results....25
Discussion....27

CHAPTER THREE: THE TIME-COURSE OF FORM EFFECTS IN ON-LINE REFERENTIAL INTERPRETATION....31
Verb Effects in On-Line Referential Interpretation....32
Experiment 2....34
Predictions....37
Method....37
Participants....37
Materials and Design....38
Visual Counterbalancing in Filler Trials....39
Linguistic Counterbalancing in Filler Trials....39
Procedure....39
Results....40

Discussion........42

CHAPTER FOUR: THE EFFECT OF FORM IN SYNTACTIC AMBIGUITY RESOLUTION........45
Experiment 3........47
Predictions........50
Method........51
Participants........51
Materials and Design........51
Visual Counterbalancing in Filler Trials........52
Linguistic Counterbalancing in Filler Trials........53
Procedure........54
Results........54
Discussion........56

CHAPTER FIVE: GENERAL DISCUSSION
Referential Processing........60
Syntactic Processing........61
Underlying Mechanism........62
Activation in an Associative Network........63
Pragmatic Inference Mechanism........65
Future Directions: Further Evaluation of a Pragmatics-Based Explanation........66

REFERENCES........70

APPENDIX A: Materials for Experiment 1........75

APPENDIX B: Materials for Experiment 2........78

APPENDIX C: Materials for Experiment 3........80

LIST OF FIGURES

FIGURE 1.0.: Syntactic Representations for the Sentences in Example (2)……………

FIGURE 1.1.: Sample Critical Display from Eberhard et al. (1995)……………………

FIGURE 1.2.: Sample Critical Display from Sedivy et al., (1999)……………………

FIGURE 1.3.: Sample Critical Display from Altmann & Kamide, (1999)………………

FIGURE 2.0.: Sample Critical Display from Pilot Study………………………………

FIGURE 2.1.: Sample Critical Display from Experiment 1……………………………

FIGURE 2.2.: Mean Fixation Time to Competitor in 2000 ms Window…………………

FIGURE 2.3.: Mean Fixation Time to Competitor in 2000 ms Window…………………

FIGURE 3.0.: Sample Referential Scenario………………………………………………

FIGURE 3.1.: Sample Critical Display from Experiment 2……………………………

FIGURE 3.2.: Mean Fixation Time to Competitor in 2000 ms Window…………………

FIGURE 3.3.: Mean Fixation Time to Competitor in 2000 ms Window…………………

FIGURE 4.0.: Sample Displays from Spivey, et al., (2002)……………………………

FIGURE 4.1.: Sample Critical Display from Experiment 3……………………………

FIGURE 4.2.: Mean Fixation Time to False Goal in 1500 ms Window…………………..

FIGURE 5.0.: Schematic of Associative Network Model……………..…………………..

Introduction

In normal conversation between native English speakers, listeners perceive an average of 180 to 200 words per minute. In order to interpret spoken language in such a limited amount of time, listeners must perform multiple tasks simultaneously. Some of the processes involved in sentence comprehension include: distinguishing speech from ambient noise; identifying individual speech sounds and lexical items; computing syntactic and semantic representations for the linguistic input; and relating this information to the real world or discourse context. The central question in research on sentence processing has been how these various tasks are achieved as language unfolds in real time.

1.0. Structural Information and Real-Time Language Interpretation

Over the years, a considerable amount of attention has been devoted to investigating how the *form* or *structure* of a sentence, as opposed to its meaning, shapes the course of interpretation. Some of the earliest investigations on this topic involved experiments where listeners reported the location of 'clicks' in a recording of continuous speech. Listeners reported the clicks as occurring at syntactic clause boundaries, even when the clicks occurred clause medially (e.g., Bever, 1970; Bever & Hurtig, 1975). These findings provided support for the notion that syntactic constituents function as 'processing units' in on-line interpretation, i.e., linguistic information must be assembled into grammatical constituents during the interpretation process. This idea was extended into a theory of comprehension where the first step in processing is to construct a complete syntactic analysis of a sentence. However, recent findings indicate that many aspects of sentence interpretation are in fact more incremental and continuous than this theory would suggest (e.g., Frazier, 1987; Marslen-Wilson, 1975; Marslen-Wilson, Brown & Tyler, 1988; Steedman, 1989; Tanenhaus, et al., 1995).

Although incremental processing of this sort might appear to be advantageous for the efficiency of the comprehension system, one disadvantage is that the processing of a word or phrase may begin before sufficient linguistic information has become available to allow the intended interpretation to be identified. As a result, much of the linguistic input that a listener receives during on-line processing is highly indeterminate. Consider the sentence in (1) (from Frazier, 1989).

(1) The girl knew the answer was correct.

Although this sentence, as a whole, is unambiguous, certain phrases within the sentence may be temporarily consistent with multiple interpretations as the sentence unfolds in time. For example, at the point in the unfolding utterance where the listener perceives the noun phrase *the answer*, there are at least two possible interpretations for the available sentence information. This noun phrase may be interpreted as the direct object of the preceding main clause verb *knew*, or it may be interpreted as the subject noun phrase in an upcoming sentence complement. It is only once the verb phrase *was correct* is heard that the listener will have enough linguistic information to resolve the ambiguity and select the latter interpretation.

Throughout the 1980's, much of the research in language processing focused on the on-line interpretation of these structurally indeterminate sentences. Several reading time studies demonstrated that people display consistent "preferences" for resolving the indeterminacy in a particular way (e.g., Ferreira & Clifton, 1986; Frazier, 1987; Rayner et al., 1983). For example, people are slower to read the (underlined) disambiguating region in (2b) compared to (2a), revealing an expectation for the phrase *the answer* to function as a direct object, rather than as the subject of a sentence complement.

(2) a. The girl knew the answer by heart.

b. The girl knew the answer was correct.

These findings suggested that listeners make on-line commitments to the structure of the sentence they are hearing, even though these commitments may prove to be incorrect. This, in turn, raises the question of why listeners choose particular structural analyses.

According to a number of researchers, decisions about ambiguous grammatical relationships are based on considerations of syntactic complexity at the point where the ambiguity is encountered (e.g., Frazier & Fodor, 1978; Kimball, 1973). For example, the syntactic representation for the sentence complement interpretation of the unfolding sentence *The girl knew the answer...,* shown in (2b) above, is structurally more complex than the direct object interpretation, (2a), (see Figure 1.0. below).

Figure 1.0. Syntactic representations for the sentences in example (2).

(a) Direct Object Structure

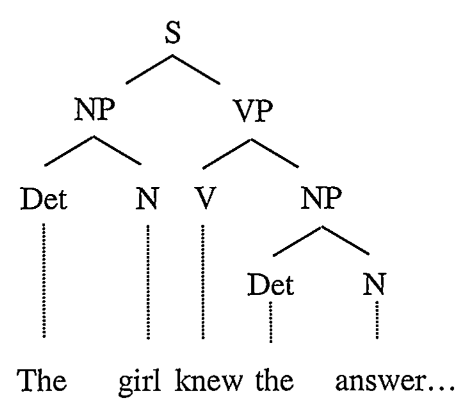

(b) Sentence Complement Structure

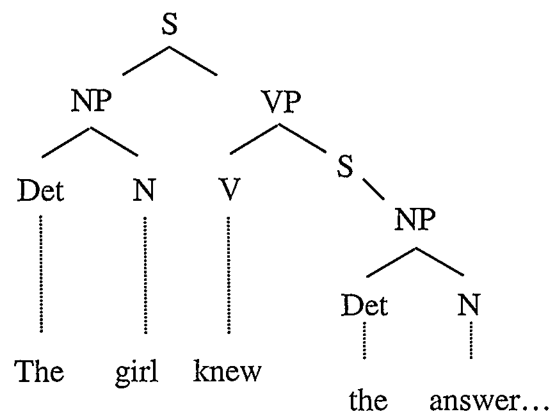

As depicted by Figure 1.0., the sentence complement interpretation of the ambiguous phrase *the answer*, in (b), requires a more complex structural representation than the direct object interpretation, (a). In (a), the noun phrase *the answer* is directly associated with an existing verb phrase. However, in (b), it occupies the subject position in a sentential complement phrase, which requires the construction of an additional sentence level 'node'.

If processing mechanisms build only the simplest grammatical structure necessary at the point where the ambiguity is encountered, then it should follow that the structure in (a) will be selected because it involves less structural complexity (i.e., fewer nodes to create) in its representation. So, when the phrase *by heart* is eventually encountered, it is easily accommodated into the current syntactic representation. However, when the phrase *was correct* is encountered, parts of the syntactic structure must be eliminated and reconstructed (i.e., creating a structure like (b)), thereby increasing reading times. As a result, the sentence in (b) is considered to be 'less preferred' for interpretation, as it involves the computation of a more complex structural representation.

The consistent finding that ambiguity in these 'less preferred' structures took readers longer to resolve provided support for a model of comprehension where the source of the information used in real time interpretation was a consideration of linguistic form. In this model, decisions that enable syntactic processing to be continuous and incremental are based on considerations about the evolving constituent structure of the unfolding utterance. An important assumption of this approach was that these form-based factors are given priority in the course of on-line processing. Thus, by this account, form plays a predictive role in syntactic decisions. However, recent studies of syntactic ambiguity resolution have challenged some aspects of this proposal. Details of this debate will be discussed in more depth in Chapter 4.

1.1. Referential Interpretation

Referential interpretation, namely the process of mapping a linguistic expression to a specific entity, or entities, in a given referential domain, can also be understood as a problem of dealing with temporary indeterminacies as language unfolds. For example, as a listener hears *The collector found the old rare manuscript in an antique shop*, the semantic information in the unfolding referential expression (e.g., 'the old...') may apply to the properties of multiple entities in the referential domain (e.g., other old things in the context). At first glance, it seems that only when the noun phrase is complete does the listener have enough information to identify the intended referent. As a result, most traditional theories of referential interpretation assume that the processing system requires listeners to wait until sufficient information becomes available before attempting to map the linguistic information to the world. However, recent studies of reference resolution show increasing evidence for the *incremental* processing of referential information in real time interpretation (e.g., Altmann & Kamide, 1999; Chambers, et al., 2002; Eberhard, et al., 1995; Tanenhaus, et al., 1995; Sedivy et al., 1999). However, in contrast to investigations of incremental processing at other levels of representation, these studies have shown little evidence that *form* plays a predictive role in this process.

Many of these studies have adopted an experimental paradigm that allows for a very fine-grained analysis of interpretive processes. Tanenhaus et al., (1995) developed a visual world eye-tracking technique that uses real world visual domains to examine on-line spoken language comprehension. This methodology enables researchers to map the time course of the unfolding linguistic information onto participants' fixations on objects in the referential domain. They found that eye movements to a target object were closely time-locked to the point in the sentence when a referring expression becomes unambiguous. This provided an experimental technique for the investigation of the information sources involved in referential processing.

Using this visual-world eyetracking paradigm, Eberhard et al. (1995) investigated the on-line interpretation of complex noun phrases, using constructions where the head noun was always modified by two prenominal adjectives (e.g, the *starred yellow* square). They predicted that, if referential interpretation proceeds incrementally, then each successive word in the noun phrase should be evaluated against referential candidates in the relevant contextual domain. If so, the time course of processing should reflect the number of objects in the display whose properties are consistent with the linguistic information as the utterance unfolds (see Figure 1.1.).

Figure 1.1. Sample critical display from Eberhard et al., (1995).

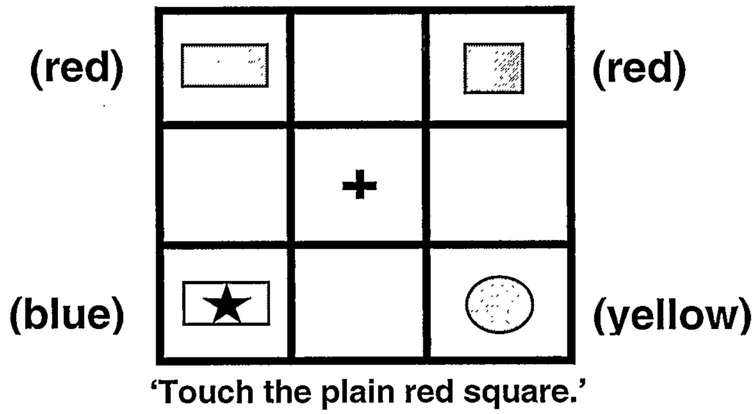

The results indicated that eye movements to the target referent were in fact closely time-locked to the point in the instruction where the disambiguating information was given. For example, when listeners heard the adjective 'plain', they reduced their consideration of the object in the display that was not plain (e.g., the starred rectangle). Then, when they heard the adjective 'red', they narrowed the domain even further and fixated only either of the two objects in the display that displayed this property (e.g., the red square and the red rectangle). Finally, when they heard the disambiguating noun 'square', they eliminated the final referential candidate (e.g., the red rectangle) and individuated the target referent (e.g., the red square).

The data from this study showed that referential interpretation is continuous and incremental. The semantic concept associated with each unfolding word is used to successively refine the set of potential referential candidates to the intended target. This provided evidence to indicate that processing is not delayed until the entire noun phrase is heard.

Using an experimental design similar to the Eberhard et al. study, Sedivy et al. (1999) investigated whether immediate interpretation would still be observed for *scalar* adjectives (e.g., the *tall* glass). As argued by Frazier (1999), Sedivy et al. raised the possibility that scalar adjectives could not be interpreted immediately because their meaning is partially dependent on the following head noun, shown in (3).

(3) a. the tall….glass

b. the tall…building

The dimension conveyed by the adjective *tall* is much different in relation to the noun *glass*, as in (3a), than *building* in (3b). Despite this apparent dependency, Sedivy et al. reported that participants did in fact use the adjective information immediately in the process of referential interpretation, by assuming that the adjective provided contrastive information. For instance, if two objects in the referential domain displayed the property of being tall (e.g., a tall glass and a pitcher), yet one of the objects was accompanied by a smaller contrasting object (e.g., a small glass), then participants began to individuate the target referent (*the tall glass*) on the basis of the adjective alone, prior to the presentation of the head noun (see Figure 1.2.).

Figure 1.2. Sample critical display from Sedivy et al., (1999).

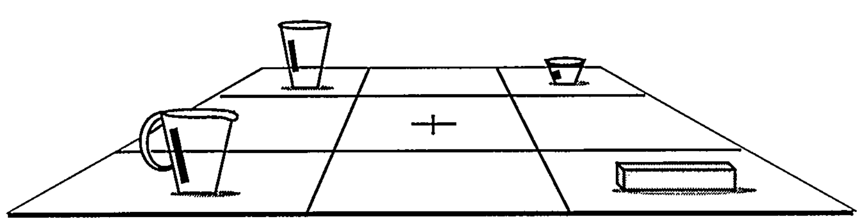

'Pick up the tall glass.'

These findings provided further evidence for the relative unimportance of form in incremental referential interpretation, demonstrating that, even when the interpretation of a unit is arguably formally dependent on a following unit, processing occurs continuously.

In addition, evidence for continuous referential processing is not even limited to particular *kinds* of syntactic forms. Altmann & Kamide (1999) found that the semantic properties of verbs can also be used to rapidly eliminate referential competitors during on-line comprehension. Given the sentence "the boy ate the cake", participants began to individuate the display object compatible with the action of eating (e.g., *the cake*) at the point of the verb, prior to when the noun phrase was heard (see Figure 1.3.).

Figure 1.3. Sample critical display from Altmann & Kamide, (1999).

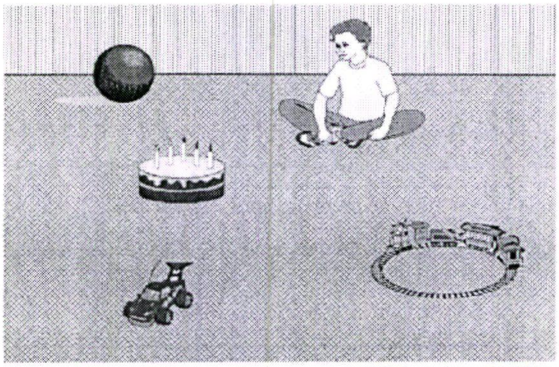

'The boy will eat the cake.'

Chambers et al. (2002) found a similar result for the interpretation of spatial prepositions (e.g., Put the cube *inside* the can). The results from their study indicated that the semantic properties of spatial prepositions immediately constrained the consideration of compatible referential candidates, again before the noun phrase was encountered. The fact that processing of this type is observed outside the noun phrase is inconsistent with many theoretical and experimental models of reference that assume referential interpretation does not begin until the end of a noun phrase. For example, in Discourse Representation Theory (DRT) (e.g., Kamp & Reyle, 1993), referential mapping is initiated only when a particular triggering condition, namely a noun phrase, is encountered in the input.

These studies provide increasing evidence to suggest that aspects of linguistic form, such as apparent dependencies within the constituent, or even the identity of constituents themselves, have little importance in real time referential interpretation. Rather, it appears that the mechanisms underlying referential processing make use of *semantic*, or conceptual, information to incrementally map the unfolding utterance to the referential world, on a word-by-word basis. Thus, in contrast to early models of sentence comprehension, these studies show very little support for the predictive value of linguistic form.

1.2. Current Research

As described above, recent research demonstrates that interpretive processes for referential interpretation are incremental and continuous. Evidence from these studies suggests that the 'referential processor' makes use of the available conceptual information to successively constrain the referential domain for the incoming linguistic input. Importantly, it also appears that this process is relatively unaffected by the structural realization of this information.

These findings suggest that considerations of linguistic form play a fairly minimal role in real time interpretation. However, in this study, I examine the possibility that issues of form *do* play a role in comprehension, just not in the way that these issues have traditionally been conceptualized (i.e., in terms of delaying processing until an entire constituent is encountered.) Instead, I propose that the 'referential processor' utilizes knowledge of the syntactic forms that are typically used to express the properties of referential entities during on-line interpretation. For example, it is possible that listeners make use of the fact that some properties (e.g., colour) usually occur in the form of an adjective before the noun, in order to identify objects in the referential domain.

The purpose of the current study was to examine whether listeners consider linguistic form during on-line referential interpretation. To investigate this issue, I conducted a series of eye-tracking experiments, manipulating the properties of potential referential candidates in a visual display. Chapter 2 provides evidence that listeners do consider the linguistic form of object properties, indicating that syntactic structure plays a role in on-line referential processing. Chapter 3 describes an experiment designed to more carefully index the time course of this process and Chapter 4 describes an experiment designed to investigate whether this process can ultimately constrain other interpretive tasks, such as the identification of grammatical relationships. Chapter 5 provides a summary of the results and discusses the implications of these results for the architecture of the sentence processing system.

Effects of Form in On-line Referential Interpretation

As described in Chapter 1, recent studies of reference resolution have provided considerable evidence for rapid and continuous semantic interpretation. Evidence of this sort challenges the traditional view that semantic interpretation occurs only at the boundaries of certain syntactic constituents, for example after the entire noun phrase has been heard. In contrast, processing appears to be relatively insensitive to characteristics of the linguistic form used to express semantic information.

In the following, I explore the idea that real time referential interpretation is in fact sensitive to aspects of linguistic form. However, this sensitivity is not reflected in interpretation being delayed until constituent boundaries. Instead, I propose that incremental referential interpretation is influenced by the relationship between the semantic properties of the denoted object and the syntactic constructions typically used to express these properties.

As a starting point, recall the study by Eberhard et al. (1995), which examined the referential interpretation of complex noun phrases (e.g., Touch *the plain red square*). In this experiment, the number of "competitor" objects, i.e. objects sharing similar semantic properties with the intended referent, was varied. However, it is important to note that this study, and most studies of referential interpretation conducted to date, used target and competitor objects whose semantic properties are expressed with prenominal adjectives (i.e., adjectives occurring before the main noun). As a result, these findings did not take into account any variation in the different kinds of structures used to denote different classes of semantic properties.

2.0. Codability

For any given entity, there are typically a number of different forms that a speaker can use to refer to the entity. The range of available forms is often referred to as the 'codability' of a given object. However, in general, there are certain forms that are more conventional than others and, as a result, used more frequently in natural speech. For example, given an entity such as a dog, there are a variety of linguistic expressions that a speaker may use to refer to this entity, such as *four-legged mammal*, or *canine*, but the intuition is that the term used most often would be *dog* or *puppy*. The same holds true for the *syntactic structure* used to realize semantic properties of entities. For example, a property, such as *furry*, may be realized in a 'prenominal' construction, i.e. as an adjective before the noun (e.g, the *furry* dog), or it may be realized in a 'postnominal' syntactic construction, i.e. in a modifier phrase after the noun (e.g., the dog *that is furry*). Again, there are differences in the frequency with which some structures are used for a given property. For instance, some properties, such as colour and dimension, are typically expressed prenominally (e.g., the *green* square/the *tall* glass), whereas other properties seem to be most naturally expressed with a postnominal preposition phrase (e.g., the square *with the arrows*). Moreover, there are many semantic properties that are not restricted to one specific syntactic construction and may be realized either prenominally or postnominally (e.g., the *striped* square/the square *with the stripes*). Thus, for a given object property, certain forms are associated with particular denotative functions. Prenominal forms are most naturally used to denote colour and, to denote stripedness, either a prenominal or postnominal form is used. Throughout this study, I will refer to these relationships as 'form-function mappings' for object properties. In addition, I will refer to the listener's knowledge of the form-function mappings for a particular property as knowledge of the 'codability' for that property.

The purpose of the experiments described in this chapter was to investigate if knowledge of the codability of object properties is considered during the course of referential interpretation. To investigate this issue, Experiment 1 tests whether differences in the codability of properties for objects in a given domain is reflected in listeners' ability to eliminate referential candidates as a noun phrase unfolds. However, prior to Experiment 1, it was necessary to first identify a set of appropriate object properties that could be easily manipulated in an experimental setting.

2.1. Production Study

Although most native English speakers have fairly strong intuitions about the linguistic forms typically used to encode semantic properties, it is important to establish that these can be revealed in the type of setting used in the experiments that follow. To do this, I conducted a pilot production study, in which participants viewed a display with two objects and were asked to name one of the objects in a way that clearly distinguished it from the other object in the display. This design ensured that participants would produce names that included descriptions of object properties. The goal of the experiment was to confirm that these descriptions reflect the use of particular syntactic forms to describe particular semantic properties. It was predicted that speakers would refer to objects with a colour property using a prenominal syntactic structure and objects with pattern properties using postnominal preposition phrases. For those objects displaying patterns that may be easily denoted by either form, it was predicted that that there would be a weaker bias to consistently use a particular syntactic alternative.

Each display contained two objects: a target object and a competitor object. On critical trials, the target object and the competitor object were the always the same shape, but differed in terms of some visual property (see Figure 2.0.).

Figure 2.0. Sample critical display from pilot study.

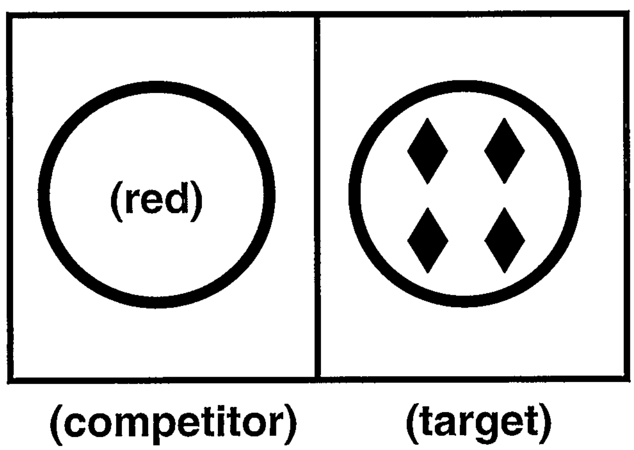

Target objects in this study had distinguishing properties that varied across three conditions, according to the syntactic form that would intuitively be used to express these properties. The *prenominal* condition contained objects that differed in colour, and which are likely distinguished using a noun phrase construction that contains a prenominal adjective (e.g., the *green* square). The *postnominal* condition contained objects with visual patterns that would typically be denoted using a noun phrase construction containing a postnominal prepositional phrase (e.g., the square *with the arrows*). Finally, the *either* condition contained objects with visual patterns that can be expressed either as prenominal adjectives, or using postnominal preposition phrases (e.g., the *striped* square/the square *with the stripes*).

2.1.1. Method

Participants

The participants in the production study were undergraduate students at the University of Calgary who participated in exchange for bonus credit in a psychology course. There were 14 participants in this experiment. All participants had normal or corrected-to-normal vision and reported that English was their native language.

Materials and Design

The visual materials for this experiment consisted of displays presented on a computer screen. There were 6 critical trials in each of the three semantic property conditions, yielding 18 critical trials. The experiment also included 18 filler trials, in which the target object and the competitor object were different shapes and displayed different semantic properties. The semantic property that distinguished the target object was also varied across filler trials.

Procedure

Participants were seated in front of a computer screen and told that they would view a visual display containing two objects. They were instructed that they would be told to name one of the objects out loud, in a way that would clearly distinguish it from the other object in the display, using a description that was as natural as possible. On each trial, participants received a spatial instruction (e.g., left or right) to indicate which object in the display to name. Participants were given eight practice trials to ensure they understood the procedure prior to beginning the experiment. All responses were manually recorded by the experimenter. The experimental session lasted approximately 30 minutes, after which participants were debriefed and given their participation credit.

2.1.2. Results

The percentage of responses containing the predicted form for the semantic property of the target was calculated and averaged over participants. Results showed that, for the prenominal condition, speakers named the target object using a prenominal syntactic structure on 100% of the trials. For the postnominal condition, 93% of the target objects were named using a postnominal prepositional phrase. Finally, in the either condition, speakers named 25% of target objects using a prenominal modifier and 74% of targets named using a postnominal prepositional phrase.

2.1.3. Discussion

As predicted, in the prenominal condition, the property of colour was consistently expressed in a prenominal syntactic construction. The strength of the form-function mapping for the postnominal condition was also fairly robust, with the majority of patterns, such as arrows, denoted by postnominal prepositional phrases. Finally, as predicted, there was less consistency in the mappings for patterns, such as stripes, from the either condition. Some of the trials were named using a prenominal construction, but the majority of trials of this kind were named using a postnominal prepositional phrase.

Most generally, the results indicate that the type of distinguishing property for a referent (e.g., colour vs. another visual pattern) is reflected in the syntactic form that the speaker selects in creating a description of the referent, confirming general intuitions. This outcome is consistent with most models of language production, which assume that speakers make choices among alternative syntactic representations, in the course of producing a grammatical utterance (e.g., Ferreira, 1996; Ferreira & Dell, 2000; Pickering et al., 2000; Pickering, Branigan & McLean, 2002). However, what is most important for current purposes is that robust effects of codability in language production have been observed using a referential environment of the type that will be used in all experiments to follow.

With this in mind, I now return to the question of whether the codability of semantic properties plays a role in the on-line *interpretation* of referential expressions. To date, the studies that are most relevant to this question are those that investigate how a particular *class* of noun phrases influences processing difficulty. For example, Gordon and Scearce (1995), among others, found that readers were slower to interpret a coreferential noun phrase when it is a repeated name (4b), rather than a pronoun (4a).

(4) a. Fred went to the store. He bought some milk.

b. Fred went to the store. Fred bought some milk.

While this result shows that the choice of a particular form influences referential processing, this choice involves deciding among alternative lexical items that differ in the semantic information that they convey, and not a decision about the syntactic frame used for a given lexical item. As a result, these studies provide little insight into whether codability would constrain reference resolution as listeners hear an expression made up of a *sequence* of words unfolding in real time, and consequently have little relevance for questions about syntactic form.

2.2. *Experiment 1*

Experiment 1 was designed to investigate if the codability of object properties constrains incremental referential interpretation. Following Eberhard et al. (1995), I examined the interpretation of modified noun phrases in visual contexts containing 'competitor' objects that could function as referential competitors for the intended referent. However, unlike previous studies, this experiment involved a codability manipulation, where the distinguishing properties of competitors were varied according to the syntactic form naturally used to express these properties.

Participants were presented with visual displays that each contained a target object, a competitor object, and several unrelated objects (see Figure 2.1.).

Figure 2.1. Sample critical display from Experiment 1.

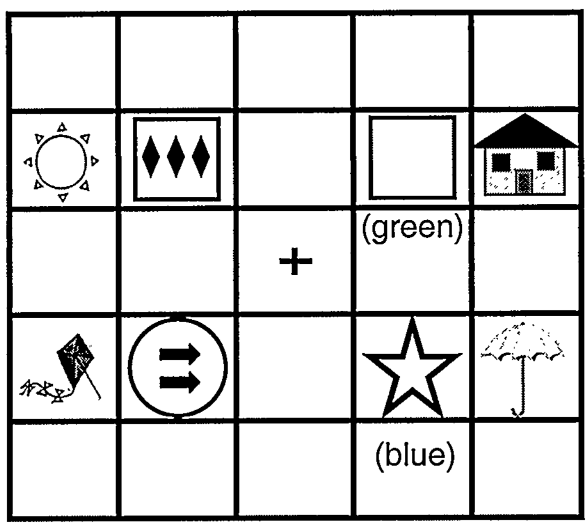

'Click on the square with the diamonds.'

The distinguishing property of the target object was always a visual pattern that would be most naturally described in a postnominal phrase (e.g., the square *with the diamonds*). In this experiment, the distinguishing properties of the competitor objects were varied across conditions that included: (i) a *prenominal* condition, where the distinguishing property of the competitor would typically be expressed prenominally (e.g., the *green* square); (ii) a *postnominal* condition, where the competitor property would be most naturally expressed with a postnominal prepositional phrase (e.g., the square *with the arrows*); and, (iii) an *either* condition, where the competitor property could be naturally described by either type of construction (e.g., the *striped* square / the square *with the stripes*). To provide a baseline, a control condition was also included, where the competitor object was a different shape than the target, but had a distinguishing property of the same type as the target (e.g., the *circle* with the arrows).

Using a head-mounted eye-tracking paradigm (e.g., Tanenhaus, et al., 1995), participants' eye movements were monitored as they heard instructions to manipulate objects on a computer screen. The critical measure in this experiment was the degree to which the competitor object attracts fixations as the instruction unfolds in time. Fixations to the competitor indicate the extent to which the listener considers this object as a potential referent.

2.2.1. Predictions

If differences in property codability are considered during real time referential interpretation, the competitor manipulation should vary the point in time where the competitor object can be excluded from consideration. The earliest "point of disambiguation" in the input occurs for the control condition, because here the competitor is a different shape than the target (e.g,. a *circle* with diamonds). Thus, given the instruction 'Click on the square with the diamonds', listeners have enough linguistic information to disambiguate between the target and the competitor object simply upon hearing the noun 'square'. As a result, the prediction is that the proportion of fixations to the competitor object will be lowest for this condition.

By contrast, the point of disambiguation for the postnominal condition occurs very late in the input string. In this condition, the competitor object is the same shape as the target, and the information that distinguishes the two is expressed at the end of a modifying postnominal phrase. As a result, the listener will not be able to determine which of these two objects is the intended target until the end of the instruction (i.e., at the noun 'diamonds'). Thus, this condition should show a significantly greater proportion of fixations to the competitor object, because it 'competes' for referential consideration for a longer period of time as the instruction unfolds.

The most critical prediction from the 'codability' hypothesis is the result for the prenominal condition. In this condition, the point of disambiguation could potentially be as early as in the control condition. This is because the distinguishing property of the competitor object (i.e., colour) would be most naturally expressed using a prenominal adjective (e.g., the *green* square). Thus, once listeners hear the noun 'square' (with no prenominal modification), they could potentially eliminate the competitor as a candidate referent for the unfolding noun phrase. If this object was the intended referent, the adjective 'green' should have been encountered before this point. Thus, if listeners can completely eliminate the competitor object on this basis, then the competitor object should be fixated to the same extent in this condition as in the control condition.

Finally, in the *either* condition, it is predicted that the proportion of fixations to the competitor should fall in between the results for the prenominal and postnominal condition. This is because the distinguishing property of the competitor object could be expressed either prenominally or postnominally (e.g., the *striped* square/the square *with the stripes*). Consequently, the strength of the relationship between the property and a prenominal construction is weaker, and should provide only a moderate or weak basis for eliminating the competitor when the noun is heard.

2.2.2. Method

Participants

The participants in Experiment 1 were undergraduate students at the University of Calgary who participated in exchange for bonus credit in a psychology course. There were 24 participants in this experiment. All participants had normal or corrected-to-normal vision and reported that English was their native language. None of the participants in Experiment 1 had participated in the production study.

Materials and Design

The visual materials for this experiment consisted of a 5 x 5 grid display, presented on a computer screen. Each display contained from three to five objects that were moved by the participant over the course of the trial. Each display also contained four fixed pictures, one in each corner of the grid, that were used as reference points for the goal location of the target object (see Figure 2.1. above). For each display in the experiment, participants heard two pairs of auditory instructions to move objects around in the grid (e.g., 'Click on the X with the Y. Now place it above the Z') (see Appendix A). Each pair was preceded by an initial instruction to 'Look at the cross.'

There were a total of 60 displays in the experiment, of which 24 were used in critical trials. Each critical trial display contained a target object, a competitor, and two unrelated objects. The competitor was manipulated across four conditions, as described above, such that all trials occurred in a given competitor condition. To ensure that any effects obtained in the critical trials were not due the particular arrays of objects in a given condition, four versions of the experiment were created by cycling each array of objects through the four experimental conditions. Equal numbers of participants were assigned to each of the four versions of the experiment.

Visual Counterbalancing in Filler Trials

In addition to the 24 critical trials, 36 fillers were included in the experiment. The fillers were used to prevent participants from employing strategies based on characteristics of the critical trials.

In the visual displays used for three of the critical trial conditions, the target object and competitor object were always the same shape. Thus, the listener might develop a bias to expect instructions to refer to one of the two similar items in the display. To counteract this strategy, 12 of the filler displays contained a pair of objects of the same shape that were never referred to, while another 12 contained a set of three objects of the same shape. Further, on critical trials, the target and competitor object differed by only one visual property (e.g., target = plain square with diamonds, competitor = green square with no other pattern). Thus, some of the filler trials contained objects with more than one distinguishing property (e.g., a blue square with arrows). Also, to provide some variation about the number of objects in the display, twelve filler trials contained three objects while another twelve contained five objects. A final consideration was the configuration of the objects within the displays. Across trials, the position of the objects was counterbalanced such that the intended target and competitor objects occurred in a different starting positions on the grid.

Linguistic Counterbalancing in Filler Trials

Of the 120 total instructions (2 instructions per display), 72 referred to a target referent using modified nouns and 36 did not contain modified nouns. Of the 72 trials with modified nouns, 36 trials had targets modified by a prenominal adjective and 36 had targets modified postnominally. This counteracts any bias toward expecting the referent to be an object described by a noun with postnominal modification. The fillers with 'prenominal' instructions included 12 in which the prenominal modifier could have occurred in postnominal position (e.g., the *striped* square).

Procedure

Participants were seated in front of a computer screen and told that they would receive auditory instructions, presented through the computer speakers, to click on an object in a grid display on the computer screen and place it in a different area of the grid. They were told that each display would contain two sets of instructions to complete before forwarding to the next display. Participants were then fitted with an ASL Model 501 head-mounted eye tracking device that consists of a monocle, an eye camera and a scene camera mounted to a headband. Participants then completed a brief calibration procedure to determine eye position relative to the scene environment. Participants were given four practice trials to ensure they understood the experimental procedure prior to beginning the experiment. Each trial was recorded on a video cassette recorder. The experimental session lasted approximately 30 minutes, after which participants were debriefed and given their participation credit.

2.2.3. Results

Data were analyzed using a frame-by-frame playback of the video recording. A time code for the onset and offset of certain word boundaries in the critical instructions was recorded and mapped onto the visual recording of participants' eye fixations during this time period. Eye movements to the competitor object were coded beginning with the first fixation following the onset of the initial word in the instruction and ending with the initiation of movement of the computer mouse toward the target object. Trials in this experiment that were unable to be properly coded due to an uninterpretable track were excluded from the analysis (1.4%).

First, I will consider evidence for a 'codability' effect. This evidence would constitute a linearly increasing time spent fixating the competitor object across the prenominal, either, and postnominal conditions. The mean fixation time data indicates that listeners showed a consistently increasing tendency to fixate the competitor object across the competitor conditions (see Figure 2.2.).

Figure 2.2. Mean Fixation Time to Competitor in 2000 ms Window

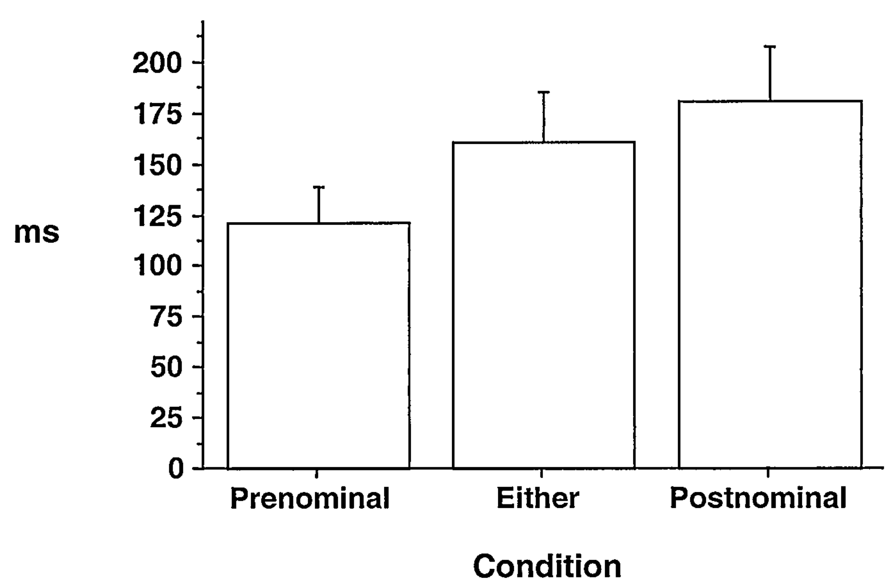

As predicted, the shortest mean time spent fixating the competitor object occurred in the prenominal condition. The postnominal condition showed the longest time spent fixating the competitor object and the mean for the either condition fell in between the prenominal and postnominal conditions.

In order to determine if this increase in the mean fixation time to the competitor object across conditions was significant, the results were also evaluated by using a repeated-measures analysis of variance (ANOVA) to test for a linear trend. The version of the experiment was included as a between subjects factor in this analysis and in all subsequent analyses. Unless otherwise stated, the effect of version was not significant. The effect of linear trend was significant: ($F(1, 20) = 5.44$, $p < .05$) and the effect of quadratic trend did not reach significance ($F(1, 20) = 0.21$, $p = .65$). This outcome provides statistical evidence for a codability effect in this study and confirms the prediction that listeners were using knowledge of the syntactic form used to denote object properties during the course of referential interpretation.

Next, I considered the issue of whether the competitor was *completely* excluded from consideration as the head noun (e.g., *square*) was heard because of the failure to hear a prenominal adjective. If so, the competitor should be considered to the same extent in the prenominal condition as in the control condition (where the competitor is an object of a different shape and therefore should receive little consideration.) The relevant data are shown in Figure 2.3., showing that listeners had fixated the competitor object longer in the prenominal condition than in the control condition. A planned pairwise contrast indicated that the 54.2 ms difference between the prenominal condition and the control condition was marginally significant ($F(1,60) = 3.81$, $p =$ .06).

Figure 2.3. Mean Fixation Time to Competitor in 2000 ms Window

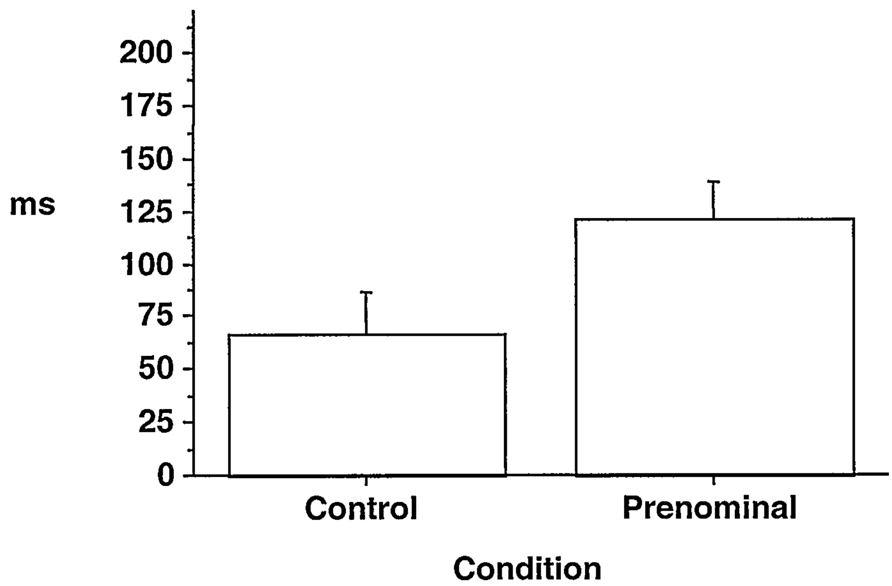

Given this marginally significant result, it would be premature to conclude that listeners can categorically exclude the competitor from consideration in the prenominal condition. Rather, there is some evidence (although relatively weak) that the competitor in the prenominal condition is considered to some degree.

2.2.4. Discussion

The results of Experiment 1 demonstrated that fixation times for the competitor object varied according to the distinguishing properties of the competitor in the display. The fact that the prenominal condition showed significantly shorter fixation times for the competitor object is consistent with the proposal that listeners reduced their consideration of this object when they heard the main noun because they failed to hear a prenominal adjective. Interestingly, this outcome suggests that the absence of linguistic information (e.g., failing to hear an adjective) can be used as a cue to the identity of the intended referent as the expression unfolds. The competitor was considered longest in the postnominal condition, which is expected because the information necessary to individuate the target referent occurs very late in the input. Interestingly, consideration of the competitor in the 'either' condition fell in between the prenominal and postnominal conditions. Overall, the gradually increasing fixation times for the competitor across these conditions appears to reflect the probabilistic nature of form-function mappings for object properties. The more likely it is that the distinguishing property of the competitor would be expressed as a prenominal adjective, the less often the competitor is considered a less likely referential candidate for the unfolding referential expression (which does not contain a prenominal adjective).

An additional question was whether fixations to the competitor object would be the same for both the prenominal and control condition, suggesting that the failure to hear a prenominal adjective categorically ruled out the competitor referent. The data show that this prediction was not fully confirmed. Fixations to the competitor were marginally longer in the prenominal condition than in the control condition. Thus, it seems that the competitor object was still being considered to some degree even though the listener had failed to hear a prenominal adjective, which would presumably be necessary for the description of this object.

Two interpretations of this result are available. One possible interpretation is that the consideration of form-function mappings does not play a role in the earliest moments of referential processing. Perhaps codability information is relatively difficult to compute and exerts an effect somewhat later in the course of referential interpretation. Thus, upon hearing 'square' in the prenominal condition, listeners momentarily consider the green competitor square before using codability considerations to eliminate it.

An alternative explanation of the results is that there is a confound in Experiment 1 that delays the earliest point at which codability information could be effectively used. A closer inspection of the displays used in this experiment reveals the possibility that listeners could have expected *spatial* information to distinguish the intended object (e.g., click on the square *on the left/on the right*), and not simply property descriptions. If this were the case, considerations of codability could not be effective in eliminating candidate referents until the preposition was heard. For example, given the instruction 'click on the square with the diamonds', it was initially predicted that the listener would be able to eliminate a competitor in the prenominal condition (e.g., the green square) once they had heard the word *square*. However, the prenominal competitor could in fact still be referred to using a postnominal spatial modifier. It is only upon hearing the preposition *with* that the listener would know that a spatial description will not be produced (because the preposition in a spatial description would most likely need to be *beside, on,*etc.). Thus, the competitor in the prenominal condition could 'compete' with the intended target referent for several words in the referring expression before it was eliminated by codability considerations. Importantly, on this account, codability information can constrain processing immediately, but its effect isn't observed until the point where it could logically come into play.

It is impossible to determine which of these two interpretations is correct from the data in Experiment 1. Moreover, this debate is difficult to resolve using experimental paradigms in which participants view objects in a visual display. This is because there is always a strong possibility for spatial descriptors to be used to distinguish objects from one another. It is possible, however, to attempt to resolve this issue using a more indirect strategy. This can be achieved by comparing the time course of the "codability effect" with the time course of effects that are known to constrain the earliest moments of referential processing. One effect of this type is the use of *verb* information to limit the set of referential candidates in an unfolding utterance (e.g., Altmann & Kamide, 1999). Chapter 3 describes an experiment that makes this comparison in order to evaluate whether codability information constrains referential processing from the earliest possible moments or whether its use is slightly delayed.

The Time-Course of Form Effects in On-line Referential Interpretation

Experiment 1 provided evidence to indicate that form-function mappings for object properties are considered in the course of on-line referential interpretation. However, the fact that listeners in the prenominal condition were not able to fully eliminate the competitor at the head noun makes the time course of this effect somewhat unclear. The precise time course of the codability effect is important for understanding the underlying architecture of the language processing system. In particular, it raises the possibility that considerations of syntactic form are computed somewhat later in the course of processing and do not constrain the earliest moments of referential interpretation. It is therefore necessary to further examine this effect and establish the time course for the use of codability information with greater precision.

As described in the previous chapter, the fact that listeners might expect spatial modifiers to distinguish the intended referent means that elimination of the competitor in the prenominal condition would occur at the preposition following the noun, rather than at the noun itself. Because spatial descriptions are almost always a possibility in the visual world eyetracking paradigm, it is difficult to create an experimental design that can determine with precision whether codability considerations can eliminate a competitor as soon as the head noun is heard (i.e. as soon as the listener can determine that a prenominal modifier was not uttered). As a result, it necessary to find another way of indexing the time course of the codability effect. In Experiment 2, I attempt to do this by directly comparing the effect of syntactic form to verb effects that have been established to occur early in comprehension.

3.0. Verb effects in On-line Referential Interpretation

As mentioned in Chapter 1, Altmann & Kamide (1999) found that verb information immediately constrained the domain of interpretation for linguistic reference. In their experiment, participants were presented with sentences from two different verb conditions (e.g., the boy will *move* the cake or the boy will *eat* the cake), while viewing a visual display containing a number of objects (see Figure 1.3.). Importantly, in the *eat* condition, the only object in the display that afforded the action denoted by the verb (e.g., *eating)* was the referent of the direct object noun phrase (e.g., *the cake*). The authors found that, in this condition, listeners began identifying the intended referent upon hearing the verb. In the *move* condition, they did not identify the target until the noun was heard. These results provide evidence that verb information is rapidly used in the course of referential interpretation. Importantly, the immediate use of verb information has also been observed in other domains, such as syntactic ambiguity resolution (e.g., Britt, 1994), establishing that this is a highly constraining information source.

How can this phenomenon be used to investigate the time course of the use of codability information? First, it is necessary to modify the experimental design such that the verb information can be used to differentiate the intended referent from a competitor of the same type when the noun is heard. For example, imagine a referential scenario with three things: an opened box, a closed box, and an opened jar (see Figure 3.0.).

Figure 3.0. Sample referential scenario

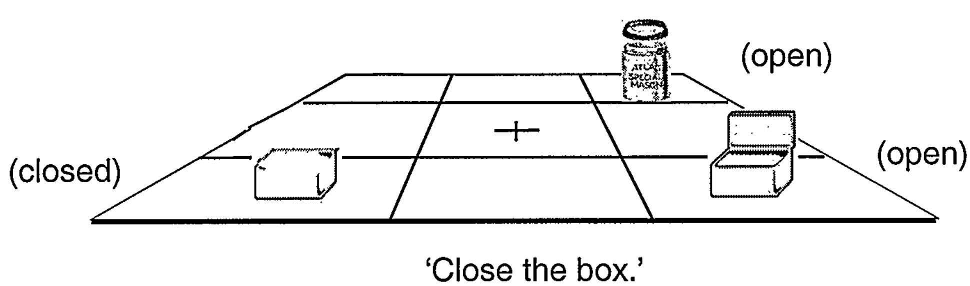

When the listener hears a verb in an unfolding sentence (e.g., *close* the box), this verb information allows the listener to narrow the referential domain to two objects, i.e., the opened box and the opened jar. When the listener hears the word 'box' later on, he/she has good evidence that it is the opened box that is being referred to. Importantly, even though the other box in the domain, i.e., the closed box, is also compatible with the description 'box', information from the verb indicates that this object is not the intended referent.

Experiment 2 uses this type of scenario to investigate the time course of the use of codability information. Like the above example, this experiment includes a condition where intended referent is identified using semantic information associated with the verb. For example, a potential scenario for this condition would include: a closed jar; an opened jar; a change purse; a rubber ducky. Once the listener hears the verb in an unfolding instruction (e.g., *open* the jar…), the referential domain can be narrowed to two objects, i.e., the closed jar and the closed change purse. When the listener hears the word *jar*, he/she has enough evidence to determine that the target referent is the closed jar.

In addition, Experiment 2 also includes a 'prenominal' condition similar to Experiment 1, where the failure to hear a prenominal adjective can potentially eliminate the competitor object from consideration. The referential scenario for this condition would be similar to before, except that the two jars are: a closed jar and a closed small jar. In this case, once the listener hears a verb in an unfolding sentence (e.g., *open* the jar…), this verb information allows the listener to narrow the referential domain to 3 objects, i.e., the closed jar, the closed small jar and the closed change purse. Then, when the listener hears the word 'jar', he/she has evidence to indicate that the target referent is not the small jar, based on the fact that a prenominal modifier was not used.

The difference between these two conditions is the type of information that the listener can use to eliminate the competitor jar and individuate the target referent. In the first condition, the listener may use information from the verb (e.g., *close*) to eliminate the competitor (e.g., the jar that is already opened) from consideration. However, in the second condition, the listener may use information about the syntactic form typically used to express the property of the competitor (i.e., in this case, a small jar, which would be distinguished using a prenominal adjective) to eliminate the competitor from consideration. The critical test of the time course for the use of codability is to compare the pattern of responses that occur when the noun is heard in the first condition to the pattern of responses that occur when the noun is heard in the second condition. Since it has been established that verb effects are early and highly constraining (e.g., Altmann & Kamide, 1999), this comparison will determine if codability information is also computed in the earliest moments of processing.

3.1. Experiment 2

Participants' eye movements were monitored as they heard instructions to manipulate objects in a visual display using real objects (e.g., Open the jar with the strawberries) (see Figure 3.1. below). Unlike Experiment 1, the verb in the instruction varied from trial to trial. On critical trials, the distinguishing property of the target object was expressed using a postnominal preposition phrase (e.g., Open the jar *with the strawberries*). As before, displays contained "competitor" objects whose properties

Figure 3.1. Sample critical display from Experiment 2

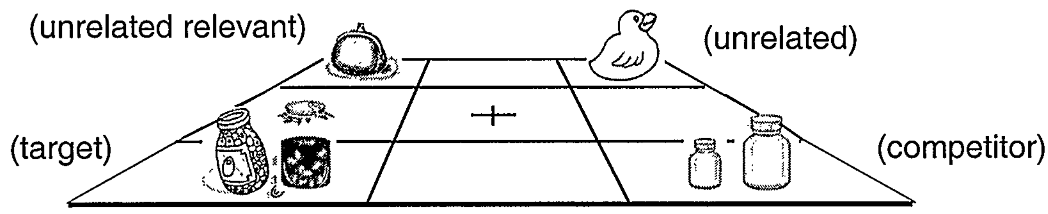

varied across four experimental conditions. In the *prenominal* condition, the distinguishing property of competitor objects was relative size, which is typically denoted using a prenominal scalar adjective (e.g., the *tall/short* jar). The use of scalar adjectives instead of colour adjectives in this experiment was designed to test the generality of the codability effect. Because dimensional descriptions (e.g., big/small) are always relative to another object, the target and competitor objects in this experiment were always paired with another object of the same class in the same display region. For example, the target jar was paired with another jar that would be most naturally described using postnominal modification, and the competitor jars might consist of two jars varying in height, enabling the use of terms such as 'big' or 'small' to individuate them.

In the *postnominal afforded* condition, the distinguishing property of the competitors would be most naturally expressed using a postnominal preposition phrase (e.g., the jar *with the pickles*). In this condition, the competitors 'afforded' the action denoted by the verb, e.g., they were able to undergo the described action (e.g,. the competitor jar was not opened). In the *postnominal unafforded* condition, the distinguishing property of the competitor objects would again be most naturally expressed using a postnominal preposition phrase, but these objects could not afford the action denoted by the verb, (e.g., these jars already had their lids removed, and thus could not be 'opened'). Thus, given an instruction such as *Open the jar with the strawberries*, the listener can distinguish the intended referent when the noun *jar* is encountered. It is important to note that, in this condition, although the competitor jars are not afforded for the action denoted by the verb 'open', an unafforded competitor is not necessarily *categorically* eliminated upon hearing the verb. This is because listeners can in fact manipulate the competitor to make it compatible with the denoted action. Consider, for example, the instruction 'Pop the balloon with the happy faces' (and see other materials listed in Appendix B). In the postnominal unafforded condition, the unafforded competitor object is a deflated balloon, however, it is possible that the listener could manipulate the object (i.e., blow up the balloon) and then complete the required action. Thus, it is likely that there will still be some consideration of the competitor until additional individuating evidence is encountered. Finally, there was an *early* condition, to provide an independent test of the rapid use of verb information. In this condition, the competitor object was completely unrelated to the target and also could not afford the action denoted by the verb (e.g., a set of building blocks). Thus, in this condition, the listener could eliminate this object from consideration at the verb *open.*

As before, the dependent variable for this experiment is the time spent fixating the competitor object, indicating the extent to which listeners considered the competitor object to be a potential referent for the unfolding description. By manipulating the semantic properties of the competitor objects, it was possible to compare the linguistic codability effect from Experiment 1 to a verb effect similar to those found in previous studies. In this experiment, the prenominal condition and the postnominal afforded condition represent the conditions designed to replicate the codability effect from

Experiment 1. These conditions provide a test of whether the listener can use information about the relationship between linguistic form and semantic properties in order to disambiguate between the target object and the competitor objects. As before, the difference between these conditions is that, for the prenominal condition, the listener receives information that could be used for disambiguation early in the input string, whereas, for the postnominal afforded condition the information occurs later in the instruction. These two conditions will be evaluated alongside the *early* condition, where consideration of the competitor should be minimal. Finally, by directly comparing postnominal unafforded condition to the prenominal condition, it is possible to carefully evaluate how 'early' codability information is used.

3.1.1. Predictions

If this experiment successfully replicates the codability effect from Experiment 1, it is predicted that the time spent fixating the competitor should increase in a generally linear fashion in the early, prenominal, and postnominal conditions. In addition, if information about form-function relationships for object properties is applied in the earliest moments of referential interpretation, consideration of the competitor in the prenominal condition should be comparable to the postnominal unafforded condition.

3.1.2. Method

Participants

The participants in Experiment 2 were 24 undergraduate students at the University of Calgary who participated in exchange for bonus credit in a psychology course. All participants had normal or corrected-to-normal vision and reported that English was their native language.

Materials and Design

The visual materials for this experiment consisted of a variety of real objects placed in four quadrants on a tabletop. The tabletop display was marked with a central fixation cross. On critical trials the quadrants contained: a pair of target objects; a pair of competitor objects; an unrelated relevant object; and an unrelated irrelevant object (see Figure 3.0. above). As explained above, this experiment required the use of *pairs* of objects in order to make it more plausible for the listener to expect that a scalar adjective (e.g., small vs. tall) would be used to describe the competitor in the prenominal scalar condition.

There were a total of 16 critical trials, 4 in each of the conditions described above. Half of the critical displays used a verb that required an instrument to perform the action (e.g., the verb *unlock* required a key). On these trials, instrument objects were placed with every item in the display that could afford the action denoted by the verb.

For each display, the participants heard two pairs of instructions, preceded by an instruction to 'Look at the cross'. The form of the first instruction in each critical pair was "Verb the X with the Y." (see Appendix B). The second instruction in the critical pair was a filler instruction. In all of the critical instructions, the target object was referred to using a postnominal prepositional phrase. Every critical instruction used a different verb. As with Experiment 1, to ensure that results could not be attributed to the particular arrays of objects occurring in a given experimental condition, four versions if experiment were created by cycling the experimental trials through the four experimental conditions. An equal number of participants was assigned to each version of the experiment.

Visual Counterbalancing in Filler trials

In addition to the critical trials, 20 filler trials were created. As with the experimental trials, filler displays also contained two pairs of a given object type. However, to counteract any expectation that one of these objects would always be referred to, the fillers included trials where the similar items were never referred to in the instructions. Half of the critical trials used a verb that required an instrument to complete the action. To eliminate any bias that the instruction would require an action involving instruments, half of the filler displays also contained objects with instruments, but the filler instructions did not require an instrument to complete the described action.

Linguistic Counterbalancing in Filler trials

To address other linguistic contingencies that might lead to response biases, the number of modified vs. unmodified instructions was roughly equated. 41 of the 72 instructions were designed to refer to a target noun with some form of modification and 31 instructions referred to targets that were unmodified. Of the 41 instructions with modified targets, 21 instructions had targets modified by a prenominal adjective and 20 had targets modified postnominally. As in Experiment 1, the quadrant position of object types was counterbalanced across trials.

Procedure

The procedure for Experiment 2 was very similar to Experiment 1, with the following exceptions. Participants were seated in front of a tabletop and told that they would receive auditory instructions, presented through computer speakers, to perform an action on an object in the grid display. Participants were given two practice trials to ensure they understood the experimental procedure prior to beginning the experimental trials. The experimental session lasted approximately 1 hour, after which participants were debriefed and given their participation credit.

3.1.3. Results

Data were analyzed using the same procedure as in Experiment 1. Eye movements to the competitor object were coded beginning with the first fixation following the onset of the initial word in the instruction and ending with the initiation of the required action. Trials in this experiment that were unable to be properly coded due to an uninterpretable track were excluded from the analysis (0.5%).

The mean time spent fixating the competitor object across the early, prenominal, and postnominal afforded conditions is shown in Figure 3.2. As expected, there was an increasing trend to fixate the competitor object across these conditions.

Figure 3.2. Mean Fixation time to Competitor in 2000 ms Window

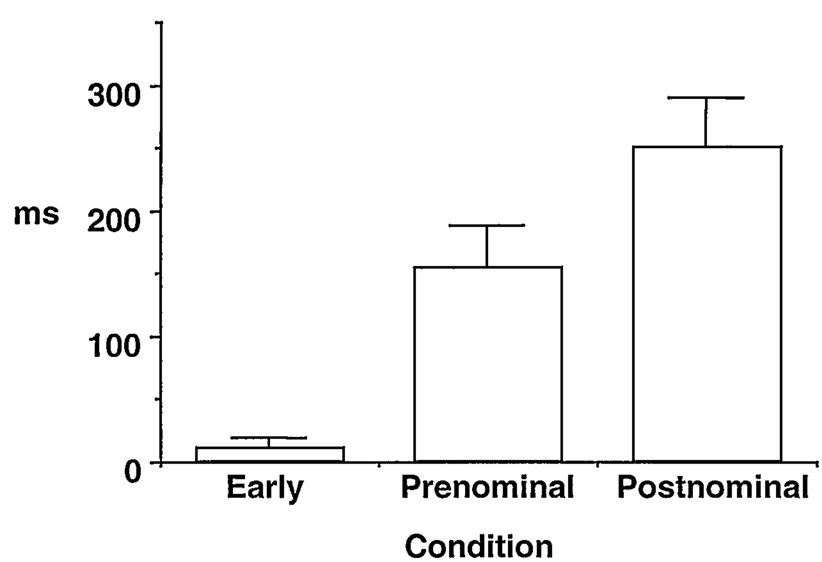

The shortest time spent fixating the competitor object occurred in the early condition, indicating that listeners used verb information at the beginning of the instruction to immediately eliminate the control competitor (e.g., set of blocks), since this object was unrelated to the verb action, i.e., a set of blocks cannot be opened. In the prenominal condition, listeners spent a longer time fixating the competitor object compared to the early condition. Importantly, the results showed that the fixation times for this condition were shorter than in the postnominal afforded condition. This suggests that listeners were able to use codability information to reduce their consideration of the competitor object when the noun was heard. In the postnominal afforded condition, listeners did not receive the information necessary to eliminate the competitor (e.g., the jar with the blueberries) until the very end of the sentence. As a result, listeners continued to fixate on this object until late in the input. The pattern of data in Experiment 2 appears to replicate the codability effect from Experiment 1.

As in Experiment 1, the pattern of fixation times across conditions was evaluated using an ANOVA to test for the effect of linear trend. The effect of linear trend was significant ($F(1, 20) = 52.98, p < .001$), and the effect of quadratic trend did not reach significance ($F(1, 20) = .70, p = .41$).

Next, I considered the issue of whether the competitor was eliminated from consideration to the same extent in the prenominal condition as in the postnominal unafforded condition. If information about form-function relationships for object properties is applied in the earliest moments of referential interpretation, than the time spent fixating the competitor object in the prenominal condition should be comparable to the postnominal unafforded condition (see Figure 3.3.).

Figure 3.3. Mean Fixation Time to Competitor in 2000 ms Window

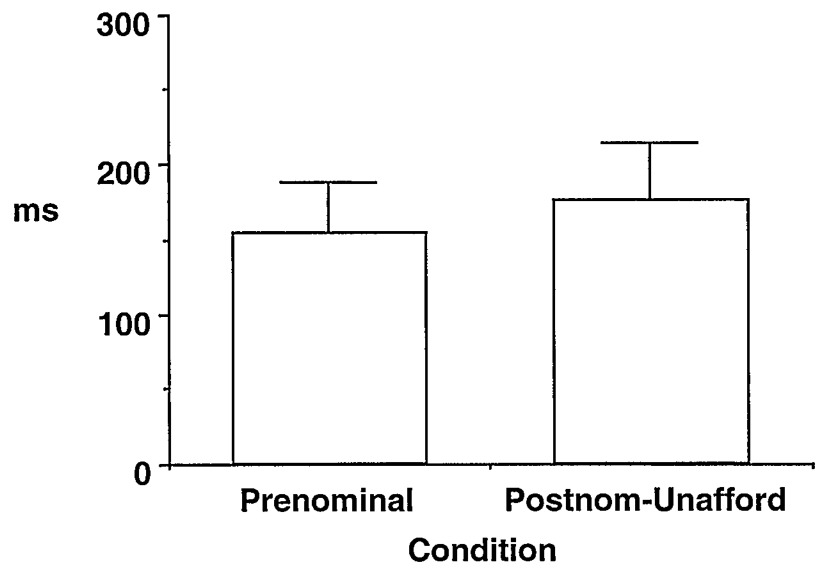

The data show very similar fixation times for the two conditions, with a difference of less than 12 ms. A planned pairwise comparison was then conducted. The results confirmed that there were no significant differences in the time spent fixating the competitor object in the prenominal condition compared to the postnominal unafforded condition ($F(1,60) = .33, p = .57$). This supports the prediction that the time course of reference resolution is the same for these conditions and suggests that codability information constrains even the earliest moments of processing.

3.1.4. Discussion

The results of Experiment 2 replicated the codability effect from Experiment 1. The time spent fixating the competitor object increased across the early, prenominal, and postnominal afforded conditions, indicating that listeners used knowledge of the syntactic structure used to refer to the competitor properties influenced the consideration of this object.

There were no significant differences in the time spent fixating the competitor in the prenominal condition and the postnominal unafforded condition. This result suggests that the use of codability information in referential interpretation occurs as quickly as the use of verb information, which has independently been shown to have immediate effects in semantic processing. The fact that there was no detectable advantage for the postnominal unafforded condition is interesting considering that listeners could have potentially used the verb information in this condition to eliminate the competitor object. Overall, the results from this comparison contribute to the growing evidence that many of the processes that were traditionally thought to be computationally demanding can occur quickly enough to constrain the earliest moments of processing.

Taken together, Experiments 1 and 2 provide evidence that knowledge about form-function mappings for object properties constrains referential processing and that the this information can be used very quickly. A further test of the 'robustness' of the codability effect can be conducted by examining whether the effect has consequences for other aspects of on-line processing. One possible candidate is *syntactic processing*. Recent studies have shown that reference plays an important role in syntactic ambiguity resolution (e.g., Altmann & Steedman, 1988; Spivey-Knowlton & Tanenhaus, 1994; Spivey et al., 2000). In particular, these studies show that the number of the candidates available for reference can influence expectations about the constituent structure of an unfolding utterance i.e., whether an unfolding phrase is understood as a modifier or not. Since the results of Experiments 1 and 2 have shown that codability information constrains the consideration of possible referents, it is also possible that it may have consequences for syntactic processing. This would provide an even more stringent test of whether codability is computed quickly in the course of language processing. This is because syntactic ambiguity resolution involves the coordination of multiple complex processes, of which reference resolution is only one component. Chapter 4 describes an experiment designed to determine if codability influences the determination of grammatical relationships in syntactically ambiguous sentences.

The Effect of Form in Syntactic Ambiguity Resolution

Throughout this study, I have been looking at the role of linguistic form in reference resolution. Experiments 1 and 2 have shown that form-function relationships for object properties constrain the course of on-line referential interpretation. Experiment 3 extends this analysis of codability effects in reference processing to syntactic ambiguity resolution.

As discussed in Chapter 1, much of the evidence from early studies of syntactic ambiguity resolution provided support for a model of comprehension where initial syntactic decisions were based on structural criteria (e.g., Frazier & Fodor, 1978; Frazier & Rayner, 1982). However, more recent studies of syntactic processing have proposed that these decisions reflect referential factors (e.g., Altmann & Steedman, 1988; Crain & Steedman, 1985; Spivey-Knowlton & Tanenhaus, 1994). Consider one well-studied syntactic ambiguity involving a temporarily ambiguous prepositional phrase, as shown in (6).

(6) Put the apple on the towel in the box.

As this sentence unfolds in time, the prepositional phrase *on the towel* is temporarily ambiguous between two interpretations: (i) a 'goal' interpretation, where the prepositional phrase specifies the intended location for the apple, or (ii) a 'modification' interpretation, where the prepositional phrase specifies properties of the apple to be moved. Although a modifier interpretation for (6) is syntactically more complex, and, as a result, 'less-preferred' according to traditional structure-based theories of ambiguity resolution, it would be required in a referential context where the intended referent is not uniquely identifiable, i.e., when there is more than one apple. As a result, researchers began to consider the possibility that there was a relationship between the referential context and the resolution of syntactically ambiguous sentences.

In one study, Spivey et al. (2002) examined the on-line interpretation of instructions such as that in (6). Two different referential contexts were used (see Figure 4.0.).

Figure 4.0. Sample displays from Spivey et al., (2002).

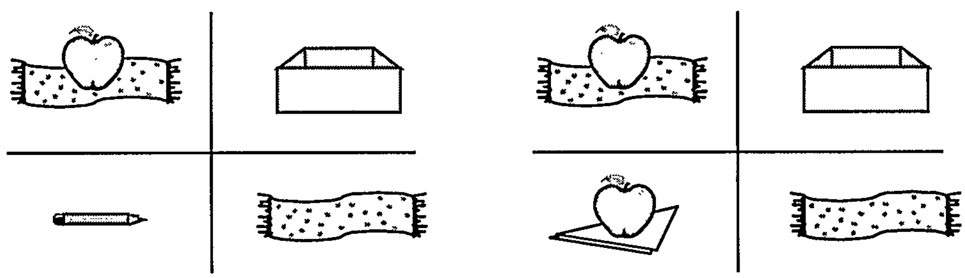

In the *one-referent* context, the corresponding visual display contained an apple on a towel, an empty towel, a pencil, and a box. In the *two-referent context*, the display contained an apple on a towel, another apple on a napkin, an empty towel, and a box. Spivey et al. found that, in the one referent condition, when listeners heard the ambiguous prepositional phrase, they had an overwhelming bias to fixate on the empty towel. This suggests that, in this condition, listeners were interpreting the prepositional phrase as a goal argument. However, in the two-referent context, listeners were no more likely to fixate the empty towel than when unambiguous instructions were used (e.g., put the apple *that's* on the towel in the box). This suggests that the difficulty associated with the 'less-preferred' interpretation, i.e., the modifier interpretation, disappeared in the two-referent context, where listeners parsed the ambiguous prepositional phrase as a modifier used to distinguish between the two potential referents, i.e., the two apples.

Evidence of this sort indicates that reference has implications for aspects of language processing other than the interpretation of noun phrases. In particular, it appears that reference plays a role in the on-line identification of grammatical relationships. Experiments 1 and 2 provided evidence to suggest that knowledge of the codability of object properties constrains reference resolution. A further test of whether this information constrains real time interpretation is to determine if it affects referential effects on syntactic ambiguity resolution.

4.0. Experiment 3

Following Spivey et al. (2002), the experimental materials for Experiment 3 consisted of 'put' instructions containing a temporarily ambiguous prepositional phrase, as shown in (7) below.

(7) Put the apple on the plate in the sandbucket.

A control version of the instructions was also included (e.g., put the apple *that's* on the plate in the sandbucket) to provide a baseline measure for the interpretation of unambiguous constructions. Spivey et al. observed that the interpretation of the first prepositional phrase (e.g., on the towel) varied according to whether the visual context supported a goal interpretation (one-referent) or a modification interpretation (two-referents). However, in Experiment 3, the visual contexts were always one-referent contexts and thus always supported a goal interpretation of the ambiguous preposition phrase. For example, the display corresponding to (7) contained a target object, a set of 'false goal' objects, a 'goal' object and an unrelated object (see Figure 4.1.).

Figure 4.1. Sample Critical Display from Experiment 3.

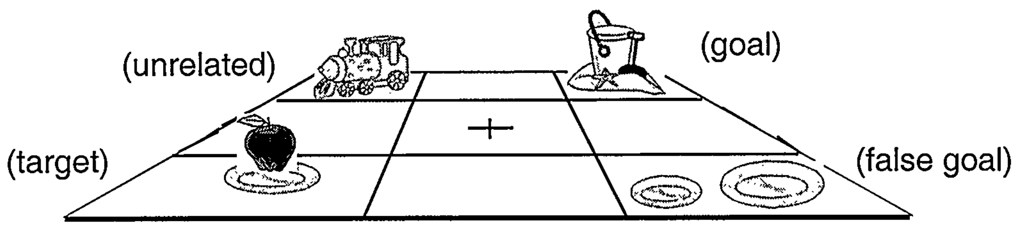

'Put the apple on the plate in the sandbucket.'

Thus, the listener should be likely to fixate 'false goal' objects (e.g, the plates), in the course of interpreting the ambiguous version of the sentence. It is not until the sentence final information is encountered (e.g., the preposition phrase *in the sandbucket*) that the listener can revise his/her interpretation of the first prepositional phrase. Unlike the Spivey et al. study, Experiment 3 did not include two-referent contexts. Two-referent contexts allow the listener to resolve the prepositional phrase ambiguity based on information from the visual domain, but the goal of Experiment 3 was to determine if listeners could identify the grammatical relationships based strictly on an additional type of linguistic information. Thus, by creating a visual domain that only supports a goal interpretation (one-referent context), listeners are required to rely on properties of the available linguistic information for disambiguation.

Thus in addition to the ambiguity manipulation, Experiment 3 also involved a codability manipulation, where the distinguishing properties of the 'false goal' were varied according to the syntactic form used to express these properties. In the *prenominal* condition, the false goal objects differed in size, and would typically be distinguished with prenominal adjectives (e.g., the *big* plate / the *small* plate). As in Experiment 2, this experiment required the use of *pairs* of objects in order to make it more plausible for the listener to expect that a scalar adjective may be used in the prenominal condition. In the *postnominal* condition, the distinguishing properties of the false goal objects would be most naturally expressed with postnominal prepositional phrases (e.g., the plate *with the flowers* / the plate *with the hearts*). Thus, there were four experimental conditions in this experiment: (i) ambiguous instruction-prenominal, (ii) ambiguous instruction- postnominal, (iii) unambiguous instruction- prenominal, and (iv) unambiguous instruction- postnominal.

By manipulating the distinguishing properties of the false goal objects, it is possible to determine if listeners used knowledge of the codability of these properties in their interpretation of the ambiguous prepositional phrase. For example, when the listener hears 'put the apple…', he/she will be able to identify the intended apple, i.e., the only apple in the display. Then, when the listener hears 'on the…', he/she will presumably interpret this unfolding phrase as specifying a goal argument, indicating the intended destination of the apple. Next, the listener will hear the noun 'plate', and will identify the 'false goal' objects, i.e., the plates, as candidate locations. Importantly, if listeners are using knowledge of property codability, it is at this point in the instruction where this information could begin to exert an influence on interpretation. In the prenominal condition, the listener may realize that, if one of the plates was the intended goal, then he/she should have already heard a prenominal modifier at this point, identifying which plate is the destination. If so, the listener may begin to reconsider the analysis of this phrase as a goal. In contrast, in the postnominal condition, there is nothing until the final prepositional phrase is encountered to indicate that the first prepositional phrase does not specify the intended location.

4.0.1. Predictions

The critical measure in this experiment was the time spent fixating the 'false goal' objects after the noun in the first prepositional phrase was heard (e.g., *plate* in (7)). This measure reflects the extent to which listeners considered the grammatical role of the first prepositional phrase to that of a goal, rather than a modifier.

As described above, if information about the linguistic form used to denote certain properties can be computed fast enough to constrain syntactic ambiguity resolution, the time spent fixating the 'false goal' in the conditions with ambiguous instructions should be shorter in the prenominal condition than in the postnominal condition. This is because the distinguishing property of the 'false goal' objects (i.e., size) would be most naturally expressed using a prenominal adjective (e.g., the *big/small* plate). Thus, once listeners hear the noun 'plate' (with no prenominal modification), they should be able to reduce their consideration of these objects and may begin to revise their assumption that *on the plate* denotes where the apple is to be put. By contrast, in the postnominal condition, the properties of the 'false goal' objects would be most naturally expressed postnominally (e.g., the plate *with the flowers/with the hearts*). As a result, the listener will not be as likely to eliminate these objects from consideration until at least the preposition following the noun 'plate' (e.g., *in,* in the sentence 'Put the apple on the plate *in* the sandbucket'). Only at this point does the listener begin to obtain evidence that the preposition phrase *on the towel* denotes some property of the apple, and not the intended location of the apple.

The unambiguous conditions were included to provide a baseline to measure interpretation when there is early linguistic evidence that the first prepositional phrase is a modifier. Because the listener heard *that's* prior to the preposition phrase, he/she should not initially misinterpret this phrase and, as a result, a relatively short fixation time to the false goal objects should result.

4.0.2. Method

Participants

The participants in Experiment 3 were undergraduate students at the University of Calgary who were paid in exchange for participation. There were 28 participants in this experiment. All participants had normal or corrected-to-normal vision and reported that English was their native language.

Materials and Design

The visual materials for this experiment consisted of a variety of real objects placed on a tabletop display. Each display contained several objects that were manipulated over the course of the trial. For each display in the experiment, participants heard two pairs of auditory instructions to move objects around in the grid (e.g., 'Put the X on the Y in the Z. Now put the X in the Y.') (see Appendix C). Each pair was preceded by an initial instruction to 'Look at the cross.'

There were a total of 32 displays in the experiment, of which 16 were critical displays. Each critical trial contained a target object, a set of 'false goal' objects, a 'goal' object, and several unrelated objects (see Figure 4.0. above). As described above, two versions of the 16 critical instructions were created (e.g., syntactically ambiguous vs. syntactically unambiguous). In addition, the properties of the 'false goal' objects were manipulated across two conditions, as described above (e.g., prenominal vs. postnominal properties). To ensure that any effects obtained in the experimental trials were not due to the particular arrays of objects occurring in a given condition, four versions of the experiment were created by cycling each array of objects through the experimental conditions. Equal numbers of participants were assigned to each of the four versions of the experiment.

Visual Counterbalancing in Filler Trials

In addition to the 16 critical trials, 16 filler trials were included in the experiment. The fillers were used to prevent participants from employing strategies based on characteristics of the critical trials.

Critical displays were always one-referent displays, i.e., there was always only one target item (e.g., one apple). Thus the listener might develop a bias to consider the ambiguous prepositional phrase to be a goal argument, based on the fact that there was only one potential target referent and, thus, there was no need for modification to distinguish between two potential referents. To ensure that the prepositional phrase modifying the target referent wasn't always infelicitous, eight of the filler trials contained two potential target referents. In critical displays, the 'false goal' was always a set of objects and the goal object was always a single item. Thus, four of the filler trials contained instructions modelled after the critical instructions except that the display contained a 'false goal' with only one object and another 16 of the filler displays contained a pair of 'goal' objects. A final consideration was the configuration of the objects within the displays. Across trials, the position of the object types was counterbalanced such that the intended target, 'false goal' and 'goal' objects occurred in a different starting positions on the grid.

Linguistic Counterbalancing in Filler Trials

There were 64 total instructions (two instructions per display), 16 of which were critical instructions. On critical trials, the critical instruction always came first and the second instruction was always a filler. In each list, eight of the critical instructions were syntactically ambiguous (e.g., *put the apple on the plate in the sandbucket*) and eight were syntactically unambiguous (e.g., *put the watch that's on the book in the teacup*). Recall that the syntactically ambiguous critical instructions always referred to a target referent using a postnominal preposition phrase modifier, i.e., the ambiguous prepositional phrase, which was then followed by the destination (goal) phrase. In order to prevent participants from developing an expectation for a specific syntactic structure, 16 filler instructions also contained a sequence of two prepositional phrases. However, in these instructions, the first phrase denoted the intended destination and the second phrase modified the first. (e.g., *Put the candle in the bowl on the frisbee*, where the display contained two bowls, one of which was on a frisbee). Half of these were syntactically ambiguous and half were unambiguous (e.g., *Put the candle in the bowl that's on the frisbee*). A variety of syntactic structures for the instructions was also achieved by including syntactically unambiguous instructions containing only a single prepositional phrase. Four of these contained an unmodified target referent (e.g., *Put the bubbles on* *the chocolate bar*) and four contained a prenominally modified target referent (e.g., *Put the small teddy bear on the soapdish*).

Procedure

The procedure for this experiment was similar to the procedure in Experiments 1 and 2. Here, participants were given two practice trials to ensure they understood the experimental procedure prior to beginning the experiment. The experimental session lasted approximately 45 minutes, after which participants were debriefed and given their payment.

4.0.3. Results

Data were analyzed using the same procedure as previous experiments. Eye movements to display objects were coded beginning with the first fixation following the onset of the initial word in the instruction and ending with the initiation of required action. The mean time spent fixating the 'false goal' region within a 1500 msec window from the onset of the noun in the first preposition phrase (e.g., on the *plate*) is reported in Figure 4.2. The end point of this window corresponds to the point where most new fixations were finished and the participant had begun reaching for the target object. Trials in this experiment that were unable to be properly coded due to an uninterpretable track were excluded from the analysis (1.1%).

Figure 4.2. Mean Fixation Time to False Goal in 1500 ms Window

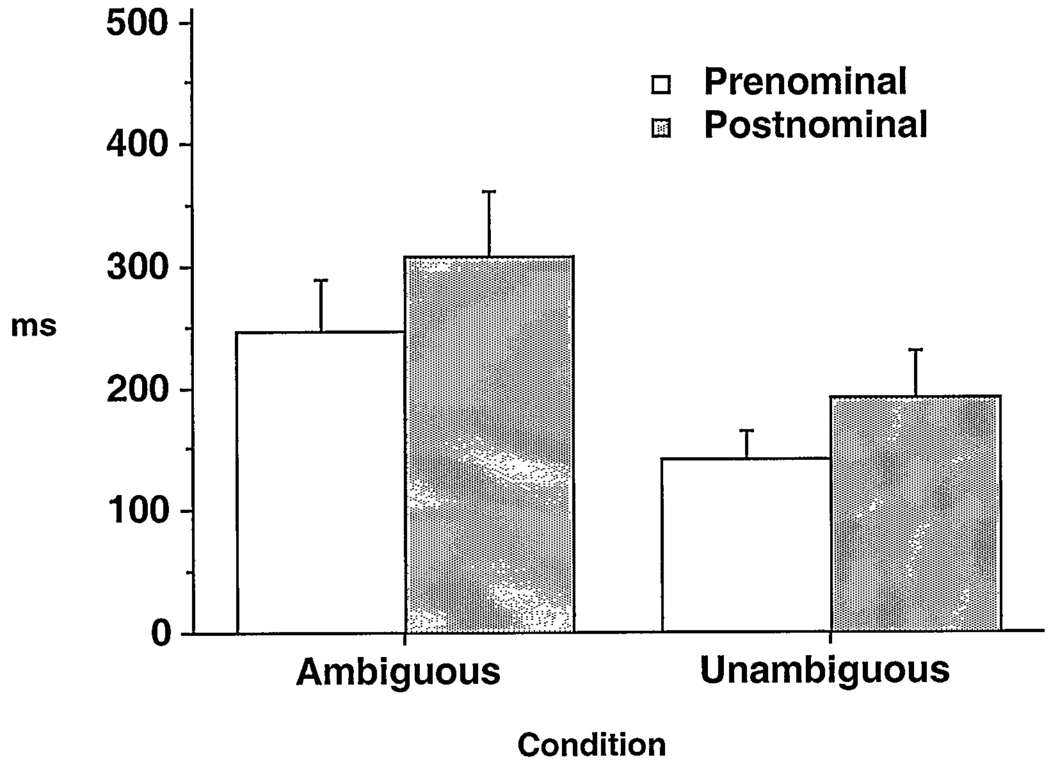

Most generally, the data clearly show that the false goal objects were fixated least in the unambiguous conditions. On average, the false goal region was fixated 111 ms longer in the ambiguous conditions than in the unambiguous conditions. There is some hint that the time spent fixating the false goal region also varied by property condition. The data showed this region was fixated 50 ms longer in the postnominal condition with unambiguous instructions, and 62 ms longer in the ambiguous instructions.

The data were submitted to a 2 (instructions: ambiguous/unambiguous) x 2 (false goal property: prenominal/postnominal) repeated-measures ANOVA. There was a main effect of ambiguity ($F(1, 24) = 6.39, p < .05$), confirming that the 'false goal' regions was fixated significantly longer in the ambiguous conditions than in the unambiguous conditions. Neither the main effect of false goal property nor the interaction was significant: $F(1, 24) = 2.84, p = .10$; $F(1, 24) = 2.04, p = .14$. Thus, in contrast to the predictions, listeners did not reduce their consideration of the ‘false goal’ objects faster in the prenominal condition than the postnominal condition when ambiguous instructions were used.

4.0.4. Discussion

The results of Experiment 3 showed that fixations to the false goal objects varied according to whether the prepositional phrase in the instruction was linguistically ambiguous. The data indicated that listeners were more likely fixate on the false goal objects if the instruction was ambiguous than if it was unambiguous. This suggests that, in the ambiguous instruction condition, listeners were misinterpreting the initial prepositional phrase as a goal argument. However, unlike Experiments 1 and 2, there was no evidence that underlying knowledge of the syntactic form that is typically used to express a property influenced the course of language processing. The data showed that there were no significant differences in listeners’ ability to eliminate the ‘false goal’ objects from consideration in the prenominal property condition compared to the postnominal property condition. This suggests that knowledge of the codability of object properties was not used quickly enough to reduce the listener’s interpretation of the ambiguous prepositional phrase as a goal argument.

There are a number of potential reasons for why a codability effect was not observed in this experiment. Although Experiments 1 and 2 demonstrated that form-function mappings are computed rapidly enough to constrain reference resolution, it is possible that the computation of this information is not quite fast enough to influence syntactic ambiguity resolution. Syntactic processing is generally assumed to be a computationally complex component of sentence comprehension. In particular, there are a range of information sources, such as intonation, lexical biases, plausibility, and discourse information, which must be coordinated in parallel to identify grammatical relationships. One possibility is that the language processing system simply cannot compute information about form-function relationships quickly enough to produce a codability effect in this context. However, the lack of an effect in this experiment does not preclude the possibility that codability information does play a role in syntactic processing, it merely demonstrates that this effect is not robust enough to constrain the earliest moments of syntactic processing.

In addition to computational load, another explanation for why the codability effect was not detected in this experiment is based on the strength of the information that causes the listener to initially adopt a goal analysis of the ambiguous prepositional phrase. To begin with, recall that the instructions all contained the verb 'put'. The argument structure of a verb contains information about the kinds of syntactic phrases that occur, either obligatorily, or optionally, with the verb in a sentence (e.g., Chomsky, 1965). The argument structure for the verb 'put' is shown in (8) below.

(8) put: <agent, theme, location>

It is important to note that the verb 'put' has an obligatory 'location' argument. This indicates that the verb 'put' has different requirements than other verbs (e.g., send), such that it must occur with a phrase specifying 'location' in order to be grammatical. MacDonald et al., (1994) argued that both lexical and syntactic ambiguities are governed by the same kinds of knowledge representations. Thus, the argument structure for the verb 'put' may create such a strong bias to interpret the initial prepositional phrase as a goal argument that information about form-function mappings was not able to override this bias. Moreover, the referential context in this experiment also supported a goal interpretation for the initial prepositional phrase. In particular, all critical displays contained only one target referent. As a result, the various informational cues that bias the interpretation of the ambiguous phrase towards a goal (rather than modifier) interpretation could be almost overwhelming. Thus, it is possible that the codability information, while available in principle, is simply not strong enough to compete strongly with the informational constraints that lead the listener to analyse the ambiguous phrase as specifying a goal.

There are also some methodological considerations that may have made the codability effect difficult to detect in this experiment. The displays in this experiment were much more complex than the referential contexts used in previous studies (e.g., Spivey et al., 2000). This increase in visual complexity may have attenuated the discernable differences in fixations patterns to the various regions in the display and, in turn, reduced the ability to detect an effect. A final issue is that only a short stretch of speech separated the points where codability information could reduce consideration of the false goals and where the rest of the unfolding instruction could reduce consideration. For example, if codability information is used, reanalysis of the ambiguous phrase in the prenominal condition could be begin when the listener hears the noun in the first prepositional phrase (e.g., put the apple on the *plate* in the sandbucket). However, a listener in the postnominal condition could possibly begin to reanalyze the syntactic role of the ambiguous phrase upon hearing the preposition in the second prepositional phrase (e.g., put the apple on the plate *in* the sandbucket), since this word indicates that additional information is coming that could also specify the intended location for the apple. The closeness of these two points in the speech stream may have made it difficult to detect a difference between the two conditions. In contrast, in Experiments 1 and 2, the relevant points in the speech stream were more separated (e.g., with the instruction *click on the square with the diamonds*, the relevant points are *square* and *diamonds).*

Although the data from this experiment did not show a statistically significant main effect of 'false goal' property, there was a slight numerical difference (i.e., a 62.2 ms effect) in the time spent fixating the false goal objects between the ambiguous prenominal condition and the ambiguous postnominal condition. This can be interpreted as limited evidence that listeners can begin to revise syntactic hypotheses on the basis of codability information. However, it is unclear why this difference should be observed for both the ambiguous and unambiguous conditions. It is clear that additional experimentation is required to establish whether the codability of object properties can affect listeners' on-line decisions about the grammatical role of ambiguous constituents.

General Discussion

The series of experiments conducted in this study examined the role of syntactic form in real time interpretation. In general, this study provides evidence that knowledge of the syntactic form used to denote semantic properties of objects is considered during the course of incremental referential interpretation, but does not constrain the interpretation of syntactic relationships.

Referential Processing

The results from Experiment 1 show a that the form of the modifier used to express object properties is taken into account during the referential interpretation of complex noun phrases. For example, in a display with two squares, listeners rapidly reduced their consideration of a square whose distinguishing property was colour (e.g,. a green square), when the noun 'square' was heard. This is because a prenominal adjective should have been encountered at that point. Listeners were slower to eliminate competitors with properties that can be denoted either prenominally or postnominally (e.g., an object that could be referred to as *the striped square/the square with the stripes*), and were even slower when competitor properties required postnominal modification (e.g., a square with an arrow pattern). This outcome suggests that the relative strength of the mapping between the form of the modifier and the property type is reflected in listener's use of this information during processing.

Experiment 2 provided further support for the syntactic form effect using dimension, rather than colour, as a property for which listeners expect prenominal modification. This experiment also explicitly examined the time course of the effect, showing that 'prenominal' referential competitors (e.g., a tall jar) were eliminated as quickly as when these competitors could be eliminated on the basis of the preceding verb. This suggests that information about form-function relationships for object properties is computed very early in processing and immediately constrains reference resolution. This set of experiments provides new and compelling evidence that the syntactic structure used to realize semantic property information plays a predictive role in interpretation of referring expressions.

Syntactic Processing

The goal of Experiment 3 was to evaluate whether information about the syntactic form used to describe object properties also constrains syntactic ambiguity resolution. In particular, this study investigated whether listeners could begin to revise the syntactic role assigned to an unfolding phrase when objects denoted by that phrase had properties that would normally be described with a different construction. In contrast to Experiments 1 and 2, a codability effect was not detected. Despite clear evidence that form-function mappings for object properties play an important role in reference resolution, the results of this experiment suggest that this information is not strong enough to influence the course of syntactic ambiguity resolution.

The main objective of this series of experiments was to assess the role of form in referential processing. As discussed in Chapter 1, early models of sentence comprehension appealed to elements of linguistic form to explain how sentence information is organized as it is encountered. One example was the claim that information must be first assembled into syntactic constituents. Additional evidence came from early studies of syntactic ambiguity, proposing that listeners predicted the syntactic role of an ambiguous constituent based on structural criteria (e.g., Frazier, 1987;

Frazier & Rayner, 1989). However, research in referential processing has shown little evidence that the form in which linguistic information is encountered has any real significance for the way in which referential processing proceeds (e.g., Altmann & Kamide, 1999; Chambers, et al., 2002; Eberhard, et al., 1995; Sedivy, et al, 1999; Spivey, et al., 2000). Despite this fact, the current results do show how considerations of form can affect on-line referential interpretation. In particular, the findings illustrate that listeners consider the type construction typically used to express object properties, as they map modifiers encountered in speech, to candidates in the referential world. Form can therefore be used alongside semantic information to make sophisticated predictions about the objects the speaker intends to refer to. In the following, I will consider some of the implications that these results have for the development of an explicit account of the mechanisms underlying form-function mappings in comprehension.

5.0. Underlying Mechanism

The syntactic form effect found in this study is not straightforwardly captured in any existing models of referential interpretation. Most generally, evidence for the use of form-function mappings in referential interpretation seems most consistent with an interactive approach, where multiple information sources are continuously coordinated during the course of on-line comprehension. However, the relationship between form and function has been relatively unexplored in area of research on comprehension. In contrast, there has been much more consideration of this issue in the language production literature. For example, Bierwisch and Schreuder (1992) proposed an interactive model of lexical selection for noun phrases involving, among other things, mappings from conceptual information to linguistic form. This kind of framework can be used as a starting point for explaining the basis for the codability effect observed in Experiments 1 and 2.

5.1. Activation in an Associative Network

Associative network models, such as those used in the production literature mentioned above, provide one type of interactive processing architecture that may be useful in characterizing the codability effect. In this framework, bi-directional connections among units at different levels of representation capture correspondences among these units. The strength of the correspondences (or *associations)* is reflected in the relative weighting of the association. One factor thought to influence the weighting of an association is the frequency with which the two units are activated together. One possibility is that mappings between object properties and the linguistic forms used to express them can be captured as connections between representational nodes at the conceptual level of a network and nodes at the level of syntactic form (see Figure 5.0.). In this type of system, the weighting of a particular connection determines the threshold for activating associated nodes.

Figure 5.0. Schematic of Associative Network Model.

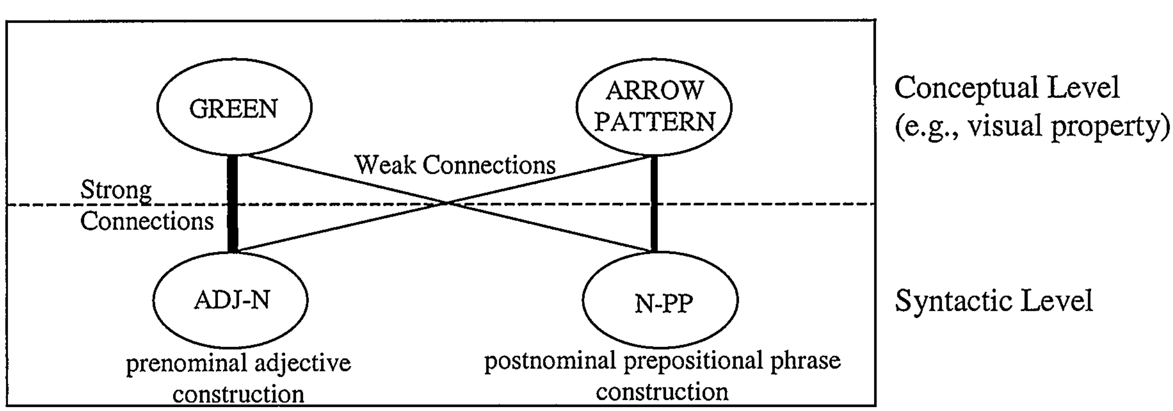

Figure 5.0. depicts a representation of a small portion of a network, showing connections between visual properties such as colour (e.g., green) and patterns (e.g., arrows), and the syntactic structures that denote these properties (e.g,. prenominal and postnominal constructions). It is important to note that the connections occurring directly between a particular semantic property and the syntactic structure typically used to express this property are very strong, whereas the connections between this property and other syntactic structures are much weaker. For those semantic properties that may be expressed either prenominally or postnominally (e.g, stripedness), the strength connections between the property node and the two nodes for prenominal and postnominal syntactic structures would be more equal.

In this model, the idea is that when a prenominal adjective is heard, the property concepts associated with that node are strongly activated. In turn, attention could be directed to referential candidates who possess those properties. However, when an adjective is not heard, the property concepts associated with that node are not activated. Thus, the failure to hear a prenominal adjective would result in a failure to direct attention to candidates whose properties would be described prenominally.

It would be necessary to address a number of additional issues to establish whether this kind of model could adequately simulate the 'codability effect' (e.g., whether inhibitory links would be necessary to 'suppress' activation of alternative candidates), but it is clear this approach provides a general way to capture how knowledge of property codability could influence on-line interpretation using existing frameworks of mental representation and process. For current purposes, an important feature of this type of model is that the connections between nodes are static, reflecting learned associations based on the most commonly occurring patterns. As a result, the associations will not vary with moment-to-moment changes in the context.

5.2. Pragmatic Inference Mechanism

An alternative to the associative network approach is the possibility that the use of codability information during comprehension is based on a mechanism of pragmatic inference. This type of approach proposes that listeners continuously assign interpretations to utterances by reasoning about the speaker's communicative intentions, including the intention underlying the use of a particular linguistic form. For example, a different communicative effect can be produced in the assertion of the sentence in (8a) compared to (8b), even though they have similar semantic content.

(8) a. Fred is tall.

b. Fred is not short.

Consider also the well-known example of a professor's response when asked about a job candidate, shown in (9).

(9) She is punctual and makes a great apple pie.

The fact that the professor did not comment on more relevant talents of the job candidate conveys a certain meaning above and beyond that which the professor actually conveyed. These examples illustrate a particular type of *implicature*, namely how a listener's knowledge of the *alternative* forms and information that were *not* linguistically expressed influence interpretation (e.g., see Grice (1961) for a detailed account of conversational implicature). Thus, it is likely that listeners must continuously evaluate the alternative forms or information that could have been expressed in an utterance in order to understand the intended meaning.

The proposal here is that the results from the prenominal conditions in Experiments 1 and 2 can be explained in terms of this type of inferencing mechanism. In particular, the listener may be able to eliminate a referential competitor whose distinguishing property would typically be expressed prenominally (e.g., a green square), based on the listener's reasoning that, if the speaker had intended to refer to this object in this specific context, then he/she would have produced a prenominal adjective. If this explanation is correct, then the results from Experiments 1 and 2 illustrate that this kind of pragmatic reasoning can occur fast enough to constrain the process of relating linguistic information to real-world referents. In addition, a pragmatically-based process may also offer a potential explanation for why a codability effect was not observed in Experiment 3. Although this type of pragmatic reasoning may be a relatively fast process, it may be not quite fast enough to influence interpretation when the task of identifying referents is occurring in parallel with other types of complex processes, as is the case during syntactic ambiguity resolution. (Note that the associative network proposal described in 5.1 does not provide any straightforward reason for why an effect was not found in Experiment 3, although, as stated in 4.0.4., some methodological factors may have limited the ability to detect an effect).

5.2.1. Future directions: Further evaluation of a pragmatics-based explanation

It is relevant to note that some independent arguments for a pragmatic inferencing account of the codability effect come from observations in linguistic theory. One point is that intuitions about whether a given semantic property may be naturally expressed as a prenominal adjective often depend on how 'intrinsic' the property is to the object. Thus, the reason for why properties such as colour and dimension are typically expressed prenominally is that these tend to be stable, characterizing attributes of an object. Properties expressed postnominally, such as fine-grained visual patterns, tend to be less integral to an object. For example, they are perhaps more likely to change over time or contexts, or may be perceived as an accidental or relatively unimportant characteristic of the object.

This generalization receives support from an examination of the expression of properties that can occur in both postnominal or prenominal position, e.g., *the corner table/ the table in the corner* (Quirk & Greenbaum, 1973). A table located in the corner of a restaurant is more likely to be referred to as a 'corner table' than a table located in the corner of a furniture store. This is presumably because the location of a table in a restaurant is more permanent and, thus, is considered to be an inherent property of the table. A related type of intuition holds when considering the types of stimulus objects used in the above experiments. For example, although *striped* can be denoted either prenominally or postnominally, the exhaustivity of the object property, i.e., whether the object is completely striped or just has a single stripe on it, seems to affect intuitions about whether a prenominal adjective sounds natural[1]. When the object has only minimal striping, the prenominal construction sounds somewhat odd. These observations illustrate that the relationship between a property type and the linguistic form used to describe it are based on *situation-specific* factors.

---

[1] This observation is credited to Dr. Elizabeth Ritter

Importantly, this outcome poses a problem for an "associative network model" explanation of the codability effect. In the associative network, the connections between nodes are static, such that there is no way to switch the weightings of the connections in response to a particular situation. As a result, this type of model cannot account for any variability in form-function mappings. It is likely then, that form-function mappings for object properties cannot be adequately modelled as weighted connections among nodes at different levels of representation. A more plausible theory is that the syntactic expression of object properties may be determined based on a mechanism of pragmatic reasoning about how integral certain properties are to the entities they describe.

One way to test this latter proposal is to conduct an experiment, similar to Experiment 1, using objects with distinguishing properties denoted by a variety of syntactic structures. However, in this experiment, filler trials would involve a speaker referring to the objects in the display with infelicitous descriptions. For example, for a particular target referent in a filler trial (e.g., a square with arrows), the speaker could refer to the object using an atypical syntactic structure (e.g., the arrowed square), or with an atypical description of the property (e.g., the square with horizontal lines with triangles on the ends). However, the experimental trials would have the same structure as in the original experiment, e.g., a 'typical' description of the target object would be used. If pragmatic inference underlies the 'codability' effect observed in Experiments 1 and 2, the infelicitous filler items should eliminate the effect. The reason is that the listener must assume that the speaker is producing descriptions that convey meaning in appropriate ways in order to gauge the speaker's reasons for selecting one expression over another. If the speaker's utterances are perceived as erratic and uncooperative from the perspective of the listener, then the reasoning process that underlies the codability effect would be disturbed. If, on the other hand, the mechanism underlying the codability effect involves a static network with weighted connections, than the codability effect should obtain on experimental trials, regardless of whether the speaker is infelicitous on filler trials.

In sum, while more research is necessary to determine the precise nature of the mechanisms that underlie on-line referential comprehension, the results of the above experiments highlight the fact that considerations of syntactic form can play a predictive role in referential processing. These outcomes create new opportunities for researchers in sentence processing to explore the particular relationships between meaning and structure in natural language and examine the role of pragmatic inferences in on-line processing.

References

Altmann, G., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition, 73,* 247-264.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30,* 191-238.

Bierwisch, M., & Schreuder, R. (1992). From concepts to lexical items. *Cognition, 42,* 23-60.

Bever, T.G. (1970). The cognitive basis for linguistic structure. In J.R. Hayes (Ed.), *Cognition and the Development of Language.* New York: Wiley.

Bever, T.G., & Hurtig, R.R. (1975). Detection of a non-linguistic stimulus is poorest at the end of a clause. *Journal of Psycholinguistic Research, 4,* 1-7.

Britt, M.A (1994). The interaction of referential ambiguity and argument structure in the parsing of prepositional phrases. *Journal of Memory and Language, 33*, 251-283.

Chambers, C.G., Tanenhaus, M.K., Eberhard, K.M., Filip, H., & Carlson, G.N. (2002). Circumscribing referential domains during real-time language comprehension. *Journal of Memory and Language, 47,* 30-49.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

Clark, H., Schreuder, R., & Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning & Verbal Behavior, 22,* 245-258.

Clifton, C., Ferreira, F. 1989). Ambiguity in context. *Language and Cognitive Processes, 4,* 77-104.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In D.R. Dowty, L. Kartunnen, & A. Zwicky (Eds.), *Natural language parsing: Psychological, computational and theoretical perspectives.* Cambridge, UK: Cambridge University Press.

Eberhard, K.M., Spivey-Knowlton, M.J., Sedivy, J.C., & Tanenhaus, M.K. (1995). Eye movements as a window into real-time spoken language processing in natural contexts. *Journal of Psycholinguistic Research, 24,* 409-436.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25*, 348-368.

Ferreira, V. (1996). Is it better to give than to donate? Syntactic flexibility in language production. *Journal of Memory and Language, 35,* 724-755.

Ferreira, V. & Dell, G. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology, 40,* 296-340.

Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and Performance XII: The Psychology of Reading.* London: Erlbaum.

Frazier, L. (1989). Against lexical generation of syntax. In W. Marslen-Wilson (Ed.), *Lexical Representation and Process.* Cambridge, MA: MIT Press.

Frazier, L. (1999). *On Sentence Interpretation.* Dordrecht: Kluwer Academic Publishers.

Frazier, L & Fodor, J.D. (1978). The sausage machine: A new two-stage parsing model. *Cognition, 6,* 291-325.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology, 14,*178-210.

Grice, H.P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society, Sup. Vol. 35,* 121-152.

Kamp, H., & Reyle, U. (1993). *From discourse to logic. Introduction to modeltheoretic semantics of natural language, formal logic, and discourse representation theory.* Dordrecht: Kluwer Academic Publishers.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition, 2,* 15-47.

MacDonald, M.C., Pearlmutter, N.J., & Seidenberg, M.S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*, 676-703.

Marslen-Wilson, W. (1975). Sentence perception as an interactive parallel process. *Science, 189,* 226-228.

Marslen-Wilson, W., Brown, C., & Tyler, L.K. (1988). Lexical representations in language comprehension. *Language and Cognitive Processes, 3,* 1-16.

Newman, J., & Dell, G. (1978). The phonological nature of phoneme monitoring: A critique of some ambiguity studies. *Journal of Verbal Learning & Verbal Behavior, 17(3),* 359-374.

Pickering, M., Branigan, H.P., & McLean, J. (2002). Constituent structure is formulated in one stage. *Journal of Memory & Language*, 46, 586-605.

Pickering, M., Branigan, H.P., Cleland, A., & Stewart, A.J. (2000). Activation of syntactic information during language production. *Journal of Psycholinguistic Research, 29,* 205-216.

Quirk, R. & Greenbaum, S. (1973). *A university grammar of English.* Harlow: Longman.

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

Rayner, K., Pollatsek, A., & Binder, K.S. (1998). Phonological codes and eye movements in reading. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24(2),* 476-497.

Sedivy, J.C., Tanenhaus, M.K., Chambers, C.G., & Carlson, G.N. (1999). Achieving incremental interpretation through contextual representation. *Cognition, 71,* 109-147.

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

Spivey-Knowlton, M.J., & Tanenhaus, M.K. (1994). Referential context and syntactic ambiguity resolution. In C. Clifton, L. Frazier, & K. Rayner (Eds.), *Perspectives on sentence processing.* Hillsdale, NJ: Erlbaum.

Spivey, M.J., Tanenhaus, M.K., Eberhard, K.M., & Sedivy, J.C. (2002). Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution. *Cognitive Psychology, 45,* 447-481.

Steedman, M. (1989). Grammar, interpretation, and processing from the lexicon. In: W. Marslen-Wilson, (Ed.), *Lexical Representation and Process.* Cambridge, MA: MIT Press.

Stemberger, J.P., Elman, J., & Haden, P. (1985). Interference between phonemes during monitoring: Evidence for an interactive activation model of speech perception. *Journal of Experimental Psychology: Human Perception & Performance, 11(4),* 475-489.

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

Tanenhaus, M.K., Spivey-Knowlton, M.J., Eberhard, K.M., & Sedivy, J.C. (1995). Integration of visual and linguistic information during spoken language comprehension, *Science, 268,* 1632-1634.

Trueswell, J.C., Tanenhaus, M.K., & Garney, S.M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language, 33,* 285-318.

Appendix A: Materials for Experiment 1

Experimental Trials

Click on the star with the arrows.
Now place it below the cross.
Click on the circle.
Now place it above the house.

Click on the circle with the trees.
Now place it below the house.
Click on the star.
Now place it in the bottom left corner.

Click on the square with the lightbulbs.
Now place it above the cross.
Click on the star.
Now place it below the house.

Click on the triangle with the diamonds.
Now place it above the cross.
Click on the square.
Now place it below the sun.

Click on the star with the lightbulbs.
Now place it above the kite.
Click on the square.
Now place it below the house.

Click on the triangle with the arrows.
Now place it below the sun.
Click on the rectangle.
Now place it in the bottom right corner.

Click on the circle with the keys.
Now place it above the house.
Click on the triangle.
Now place it in the bottom left corner.

Click on the triangle with the moons.
Now place it above the umbrella.
Click on the circle
Now place it below the cross.

Click on the square with the keys.
Now place it in the top left corner.
Click on the rectangle.
Now place it above the umbrella.

Click on the triangle with the lightbulbs.
Now place it below the cross.
Click on the square.
Now place it above the sun.

Click on the rectangle with the keys.
Now place it above the sun.
Click on the square.
Now place it in the bottom left corner.

Click on the square with the moons.
Now place it in the top right corner.
Click on the rectangle.
Now place it beside the cross.

Click on the circle with the diamonds.
Now place it above the umbrella.
Click on the square.
Now place it beside the cross.

Click on the circle with the arrows.
Now place it above the kite.
Click on the star.
Now place it in the bottom left corner.

Click on the circle with the lightbulbs.
Now place it below the house.
Click on the rectangle.
Now place it in the top right corner.

Click on the triangle with the keys.
Now place it in the top right corner.
Click on the star.
Now place it below the sun.

Click on the rectangle with the moons.
Now place it above the house.
Click on the triangle.
Now place it below the cross.

Click on the square with the diamonds.
Now place it above the kite.
Click on the triangle.
Now place it beside the sun.

Click on the square with the arrows.
Now place it below the cross.
Click on the rectangle.
Now place it in the bottom left corner.

Click on the rectangle with the arrows.
Now place it above the kite.
Click on the triangle.
Now place it below the umbrella.

Click on the star with the trees.
Now place it in the top left corner.
Click on the circle.
Now place it above the kite.

Click on the square with the trees.
Now place it in the top right corner.
Click on the triangle.
Now place it below the umbrella.

Click on the star with the lightbulbs.
Now place it above the house.
Click on the square.
Now place it below the cross.

Click on the rectangle with the lightbulbs.
Now place it below the sun.
Click on the circle.
Now place it above the house.

Appendix B: Materials for Experiment 2

Experimental Trials

Pop the balloon with the happy faces.
Now touch the green book.

Close the bag with the smarties.
Now pick up the red scissors.

Tie the ribbon with the diamonds.
Now pick up the ice cube tray.

Stamp the card with the bells.
Now place the watch on top of the paper.

Knock over the glass with the straw.
Now touch the rolling pin.

Fold the towel with the fish.
Now touch the bubbles.

Tear the paper with the numbers.
Now place the stapler on the plastic bag.

Clean the plate with the flowers.
Now switch the glass and the kitten.

Cover the doll with the pigtails.
Now touch the spatula.

Cut the cupcake with the racecar.
Now pick up the rubber ducky.

Hang the postcard with the mountain.
Now turn over the newspaper clipping.

Open the jar with the strawberries.
Now switch the water bottle and the flashlight.

Paint the house with the window.
Now pick up the green drum.

Unzip the case with the horses.
Now place the yellow mug below the cross.

Unlock the box with the sun.
Now place the apple beside the small box.

Unwrap the present with the balloons.
Now switch the candle and the chocolate bar.

Appendix C: Materials for Experiment 3

Experimental Trials

Put the doll (that's) on the towel in the basket.
Now put the small crayon with the doll.

Put the mug (that's) on the postcard in the pot.
Now put the pig in the mug.

Put the apple (that's) on the plate in the sandbucket.
Now put the sandbucket with the train.

Put the watch (that's) on the book in the teacup.
Now put the spatula on the small book.

Put the scissors (that are) on the card in the vase.
Now put the dishsoap on the card with the hearts.

Put the knife (that's) on the jar in the newspaper.
Now put the large envelope with the knife.

Put the spoon (that's) on the pencilcase in the shoe.
Now put the stamp pad in the {pencilcase with the dinosaurs / large pencilcase}.

Put the hammer (that's) on the kite in the glass.
Now put the glass with the drum.

Put the blocks (that are) on the napkin in the tin.
Now put the trumpet on the {large napkin / napkin with the unicorn}.

Put the lemon (that's) on the ribbon in the measuring cup.
Now put the car with the lemon.

Put the yo-yo (that's) on the can in the bowl.
Now put the large doll in the bowl.

Put the stamp (that's) on the present in the bag.
Now put the bag with the toothbrush.

Put the flowers (that are) on the paper in the ice cream pail.
Now put the ketchup with the flowers.

Put the glue (that's) on the coaster in the sock.
Now put the sock on the paints.

Put the straw (that's) on the box in the wine bottle.
Now put the bottle with the rubber ducky.

Put the pencil (that's) on the placemat in the ketchup bottle.
Now put the small duck with the ketchup.