

MORPHOLOGICAL MISMATCH IN ELLIPSIS

NOT ALL FORMS OF MORPHOLOGICAL MISMATCH ARE
ACCEPTABLE IN VERB-PHRASE ELLIPSIS

By

TIFFANY DESCHAMPS, M.Sc., B.A.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Master of Science

McMaster University

© Copyright by Tiffany Deschamps, August 2012

MASTER OF SCIENCE (2012)
(Linguistics and Languages)

McMaster University
Hamilton, Ontario

TITLE: Not All Forms of Morphological Mismatch Are Acceptable in Verb-Phrase Ellipsis
AUTHOR: Tiffany Deschamps, M.Sc., B.A. (McMaster University)
SUPERVISOR: Dr. Ivona Kučerová and Dr. Victor Kuperman
NUMBER OF PAGES: xiii, 91

Abstract

Not All Forms of Morphological Mismatch Are Acceptable in Verb-Phrase Ellipsis

Tiffany Deschamps
Master of Science
Department of Linguistics and Languages
McMaster University
2012

The Recycling Hypothesis of verb-phrase ellipsis states that elided verb phrases with non-parallel antecedents are interpreted by reconstructing the appropriate verb phrase structure using the information available in the antecedent (Arregui, Frazier, Clifton, & Moulton, 2006). The hypothesis predicts that structurally more complex antecedents will involve more complicated reconstruction operations, which will lower the acceptability of the sentences. The experiments reported in this thesis tested two underlying assumptions of the Recycling Hypothesis as well as one prediction that follows from the proposal. First, the hypothesis assumes that elided verb phrases with parallel antecedents are interpreted by copying the structure of the antecedent into the ellipsis site (Frazier & Clifton, 2001). Second, Arregui et al. (2006) argued that changes in verbal morphology were “really easy (p. 242)” to recover from, suggesting that verbal morphology is not a factor in determining parallelism between the antecedent and elided verb phrases. Results from three written survey experiments in which participants were asked to judge the acceptability of verb-phrase ellipsis with matching or non-matching verbal morphology contradicted these assumptions. Morphologically more complex antecedents were rated less acceptable than simpler antecedents, regardless of whether the antecedent morphology matched the morphology on the elided verb phrase. The fact that verbal morphology affected acceptability ratings suggests that this factor plays a critical role in determining parallelism in ellipsis. Furthermore, the fact that parallel antecedents patterned with non-parallel antecedents suggests that the two must be processed in a similar fashion. Finally, if more complex antecedents require more complicated reconstruction operations, it might be predicted that word-by-word reading times at the ellipsis site should be correlated with the level of difficulty (Gibson, 1998). One self-paced reading experiment using the same materials showed no such correlation. These results are discussed with reference to two other psycholinguistic theories of verb-phrase ellipsis comprehension.

Acknowledgements

Given that this is my second master's thesis, I never expected it to be more difficult than my first. If it were not for the support of my supervisors, Dr. Ivona Kučerová and Dr. Victor Kuperman, completing this project would have been impossible. Their time, patience, guidance, and encouragement proved to be invaluable contributions to this thesis. For this, I sincerely thank them both.

Similarly, I could not have completed this project without the unconditional and unwavering support of my family and friends. My parents – Angus Deschamps, Lorrie Boundy, Gary Boundy, and Jennifer Milton – have seen me through three post-secondary degrees, and would proudly see me through three more. I thank them for their limitless support and encouragement. I would also like to thank my brother and sisters – Angus Deschamps Jr., Melissa Boundy, and Michelle Boundy – for their continued support. I am blessed with support from so many wonderful people, and it constantly amazes me. A mass, impersonal thanks to all of the friends who have carried me through this project feels inadequate; however, to list them all would be an enormous task, and I would inevitably forget some very important people. That said, I would like to specifically thank Anna Kata, Andrea Unrau, and Angela Harrison for their contributions. Anna spent a great deal of time and effort editing the many pages that follow, which was certainly an arduous task for someone with no linguistics background. I am very grateful for her assistance. Andrea and Angela were primarily responsible for holding me up when I was about to fall down (sometimes literally, in Andrea's case). Morning coffees with Angela would help me start the day with a good attitude, while evening beers with Andrea renewed my spirits at the end of a hard day. For this, I am truly thankful.

Contents

1	What Is Ellipsis? Evidence from Linguistics	1
1.1	General Overview	1
1.2	Evidence for Unpronounced Syntactic Structure	3
1.3	Defining “Parallelism”	6
1.3.1	Syntactic Parallelism	7
1.3.1.1	Deletion Theories	7
1.3.1.2	Copy Theories	9
1.3.2	Non-Syntactic Parallelism	12
1.3.3	Merchant’s e-GIVENness	14
1.4	Non-Parallel Ellipsis	17
1.4.1	Voice Mismatch	17
1.4.2	Morphological Mismatch	18
1.5	Accommodating Non-Parallelism	23
1.6	What is the problem? (Summary)	25
2	How do we Process Ellipsis? Evidence from Psycholinguistics	27
2.1	General Overview	27
2.2	Testing ”Parallelism”	28
2.2.1	Parallelism and Deep vs. Surface Anaphors	28
2.2.2	Parallelism and Strict vs. Sloppy Identity Readings	32

2.3	Dealing with Non-Parallelism	34
2.3.1	Cue-Based Direct-Access Memory Retrieval Hypothesis	34
2.3.2	Structure Search Hypothesis	38
2.3.3	Recycling Hypothesis	40
2.4	What is the problem? (Summary)	44
3	Does the Recycling Hypothesis Work?	47
3.1	Review of Theory and Problem	47
3.2	Experiment 1A: Ellipsis in Isolation	49
3.2.1	Motivation	49
3.2.2	Methodology	50
3.2.2.1	Participants	50
3.2.2.2	Stimuli	50
3.2.2.3	Procedure	50
3.2.3	Results	51
3.2.4	Discussion	52
3.3	Experiment 1B: Ellipsis with Context	54
3.3.1	Motivation	54
3.3.2	Methodology	54
3.3.2.1	Participants	54
3.3.2.2	Stimuli	54
3.3.2.3	Procedure	54
3.3.3	Results	55
3.3.4	Discussion	56
3.4	Experiment 1C: Non-Elided Controls	57
3.4.1	Motivation	57
3.4.2	Methodology	57

3.4.2.1	Participants	57
3.4.2.2	Stimuli	58
3.4.2.3	Procedure	58
3.4.3	Results	58
3.4.4	Discussion	59
3.5	Experiment 2: Self-Paced Reading	60
3.5.1	Motivation	60
3.5.2	Methodology	61
3.5.2.1	Participants	61
3.5.2.2	Stimuli	62
3.5.2.3	Procedure	62
3.5.3	Results	62
3.5.3.1	Acceptability Judgments Reaction Times	63
3.5.3.2	Reading Times	64
3.5.4	Discussion	64
3.6	What do we know now? (Summary)	65
4	What Have We Learned About Ellipsis?	67
4.1	General Overview	67
4.2	Is morphological assignment cost-free?	67
4.3	Is Copy α (Frazier & Clifton, 2001) Cost-Free?	73
4.4	Do Real-Time Comprehension Measures Show Effects of Recycling?	75
4.5	Which theory is best?	76
A	Appendix A	79

List of Tables

3.1	Experiment 1A: Percent ‘Acceptable’ responses	52
3.2	Experiment 1A: Logistic regression summary	52
3.3	Experiment 1B: Percent ‘Acceptable’ responses	55
3.4	Experiment 1B: Logistic regression summary	55
3.5	Experiment 1C: Percent ‘Acceptable’ responses	59
3.6	Experiment 1C: Logistic regression summary	59
3.7	Experiment 2: Percent ‘Acceptable’ responses	62
3.8	Experiment 2: Logistic regression summary	62
3.9	Experiment 2: Mean judgment reaction times	63
3.10	Experiment 2: Regression summary (reaction times)	63
3.11	Experiment 2: Mean reading times	64
3.12	Experiment 2: Regression summary (reading times)	64

List of Figures

2.1 Hypothetical SAT functions (Martin & McElree, 2008)	35
---	----

Chapter 1

What Is Ellipsis? Evidence from Linguistics

1.1 General Overview

A verb phrase is said to be elided when it is missing from the linguistic signal but its meaning can be recovered from the context. The sentence in (1) represents a standard case of verb-phrase ellipsis. (Note: Here, and throughout, I mark elided material with ~~crossed out font~~.)

- (1) a. John went to the store, and Mary did too.
b. John went to the store, and Mary did ~~went to the store~~ too.

This sentence is interpreted such that both John and Mary went to the store, despite the fact that the string *Mary went to the store* does not overtly exist in the signal. To be clear, it is traditionally thought that the clause *Mary did too* represents an unpronounced verb phrase that inherits its meaning from the overt verb phrase in the preceding clause.

What is interesting about verb-phrase ellipsis is that it represents a linguistic phenomenon where there is meaning but no sound. Sentences with elided verb phrases are grammatical despite the fact that the sound-meaning correspondence rules inherent to language break down. What is even more interesting is the fact that the meanings assigned to these types of sentences are very consistent. In fact, all native speakers of English will assign the same interpretation to sentences such as that in (1). This is a relatively unique phenomenon. Imagine a scenario in which you and your friends have gone hiking and a moose crosses your path. One of your friends points and says “Moose!” It is nearly impossible to be sure exactly what this exclamation means because there are too many possible interpretations. Your friend may be expressing his excitement at seeing a moose for the first time, or expressing fear that the moose might decide to charge toward you. He may be implying that you should stop walking in its direction or that you should

stop talking so that the moose does not get scared and run away. The exact interpretation that the speaker intended is not entirely clear. Verb-phrase ellipsis constructions are not so flexible, which offers an interesting puzzle for linguists and psycholinguists alike.

Fundamentally, the puzzle of ellipsis is twofold. The first question is one of licensing: Under what circumstances is verb-phrase ellipsis allowed to occur? The second question is one of interpretation: How does the elided verb phrase receive its interpretation, and where does this meaning come from? To a certain extent, the answers to these questions go hand-in-hand. It is generally accepted that verb-phrase ellipsis is licensed whenever the context allows recovery of the elided verb phrase. This recovery process is thought to rely on the existence of another verb phrase in the discourse (the antecedent verb phrase) that is in some way “parallel” with the elided verb phrase. Thus, the presence of one verb phrase is thought to be enough to license the ellipsis of a repeated verb phrase. Furthermore, the antecedent verb phrase is thought to be responsible for providing the elided verb phrase with its interpretation in some way. In other words, the elided verb phrase typically shares its meaning with the verb phrase that licensed the ellipsis.

The trick, then, is determining how to draft the definition of parallelism. Early theories (e.g., Hankamer & Sag, 1976; Ross, 1969; Sag, 1976) proposed a strictly syntactic definition of parallelism, such that the (unpronounced) structure of the elided verb phrase must match the structure of the antecedent verb phrase. Meanwhile, other theories (e.g., Culicover & Jackendoff, 2005; Dalrymple, Shieber, & Pereira, 1991) rejected the syntactic definition of parallelism in favour of a semantic definition. Semantic theories of verb-phrase ellipsis posit that there is no syntactic structure underlying the ellipsis, and that ellipsis is interpreted by applying the semantics of the antecedent clause to the elided clause. Still other theories (e.g., Merchant, 2001) suggested that both syntactic and semantic properties are relevant for explaining verb-phrase ellipsis phenomena.

However, the idea of parallelism can lead to some difficulties. For example, parallelism does not account for the acceptability of sentences like (2) and (3), in which the structure of the elided verb phrase does not match the structure of the antecedent, ultimately violating any definition of parallel surface structure. Nor does it account for the fact that the elided verb phrase in (4) can receive multiple interpretations. If the two verb phrases are identical, why doesn't the elided verb phrase automatically receive the same meaning as the antecedent?

- (2) a. The garbage had to be taken out, but Bill refused to.
 b. The garbage had to be taken out, but Bill refused to ~~take the garbage out~~.
- (3) a. John understands the situation and surely Peter should.
 b. John understands the situation and surely Peter should ~~understand the situation~~.
- (4) a. John likes his job, and Bill does too.
 b. John likes his job, and Bill does ~~likes his job~~ too.
 i. = ... Bill likes John's job.
 ii. = ... Bill likes Bill's job.

Instances of verb-phrase ellipsis without parallel antecedents, like those in (2) and (3), are the focus of this thesis. Generally, the research question can be stated as: How do we comprehend elided verb phrases when they are not parallel to their antecedents? Specifically, the experiments presented in this thesis focus on the types of sentences in (3), where the only difference between the two verb phrases is the morphology on the verb. Are sentences like these always easy to comprehend? If not, how can we account for the different levels of complexity?

This first chapter reviews the linguistic evidence used to describe verb-phrase ellipsis. The layout of this chapter is as follows: Section 1.2 provides evidence that the ellipsis site hosts unpronounced syntactic structure. Section 1.3 describes several theories of parallelism in detail. Section 1.4 discusses cases of non-parallel ellipsis while Section 1.5 offers some explanation of how the system might overcome issues of non-parallelism. Section 1.6 provides a summary of the evidence and builds the foundation for the research question pursued here.

Chapter 2 fully develops the current research question after reviewing the psycholinguistic research surrounding verb-phrase ellipsis. Chapter 3 reports a series of experiments designed to investigate how individuals comprehend ellipsis when the verbal morphology differs between the antecedent and elided verb phrases. Chapter 4 discusses the implications of the results reported in Chapter 3.

1.2 Evidence for Unpronounced Syntactic Structure

A great deal of evidence suggests that there is unpronounced syntactic structure in the ellipsis site. More accurately, the evidence shows that the elided material is sensitive to the same kinds of syntactic constraints as overt material – a fact that is unexpected if there is no syntactic structure in the ellipsis site. This section will review three pieces of evidence that support the existence of unpronounced syntactic structure, though there are several others (see Merchant, 2009, for a review). The first piece of evidence shows that verb-phrase ellipsis constructions are sensitive to syntactic locality domains, or syntactic island constraints. The final two pieces of evidence come from instances of sluicing, which is the label assigned to constructions in which an entire clause has been elided, leaving only a stranded *wh*-phrase. These examples show that sluicing constructions share syntactic properties with standard *wh*-questions. Specifically, sluiced clauses are (i) subject to case-matching effects, where the stranded *wh*-phrase must carry the same case as its correlate in the antecedent clause (Ross, 1969), and (ii) sensitive to the preposition-stranding generalization; that is, only languages that allow preposition stranding in standard *wh*-questions will allow it under sluicing (Merchant, 2001). All of these facts rely on the presence of syntactic structure, and thus support the hypothesis that the ellipsis site hosts unpronounced syntactic material.

Syntactic islands are constituents that do not permit internal elements to be extracted to positions outside of the constituent. Ross (1967) introduced several types of island

constraints on movement, two of which will be applied here. First, verb-phrase ellipsis is sensitive to Ross' Complex Noun Phrase Constraint, which blocks movement of elements within a relative clause to positions outside of the noun phrase that it modifies. The examples in (5) show how this applies to verb-phrase ellipsis.

- (5) a. John wants to take a job that offers competitive benefits.
 b. *I don't know [which benefits]_i John wants to take [_{NP} a job [_{CP} that offers [*t_i*]]].
 c. *John wants to take a job that offers competitive benefits, but I don't know [which benefits]_i John does ~~wants to take~~ [_{NP} a job [_{CP} that offers [*t_i*]]].

(5a) represents a grammatical sentence with an embedded subject-relative clause. (5b) shows that extracting the embedded object of *offers* to a matrix clause position results in ungrammaticality. (5c) further shows that this ungrammaticality holds even when the verb phrase containing the complex noun phrase has been elided.

Verb-phrase ellipsis is also sensitive to the Coordinate Structure Constraint (Ross, 1967), which forbids the movement of an element when it has been joined with another by coordination. This is shown in (6).

- (6) a. John likes his job and his house.
 b. * [What]_i does John like [*t_i*] and his house?
 c. * John likes his job and his house, but I don't know [what]_i Bill does ~~like~~ [*t_i*] ~~and his house~~

Again, (6a) shows a grammatical sentence with two coordinated noun phrases. (6b) shows that the sentence becomes ungrammatical when *wh*-movement targets an element inside only one of the conjuncts, and (6c) shows that it remains ungrammatical even when both conjuncts have been elided. The fact that verb-phrase ellipsis constructions show the same pattern of behaviour as fully-articulated sentences with respect to syntactic islands suggests that the verb-phrase ellipsis constructions indeed host unpronounced syntactic structure.

As stated previously, the next two pieces of evidence supporting the existence of syntax in the ellipsis site come from sluicing. The sentence in (7) is a representative example of sluicing: a *wh*-phrase stands in for an entire interrogative clause. Ross (1969) observed that, like cases of verb-phrase ellipsis, the missing clause appears to have been elided, and typically shares (approximately) the same meaning as a nearby clause. In this case, the interpretation given to (7a) implies that the speaker knows that John went to one of the stores in town, but does not know which store John went to.

- (7) a. John went to the store, but I don't know which (store).
 b. John went to the store, but I don't know [_{CP} [which (store)]_i ~~John went to~~ [*t_i*]].

Ross observed that sluiced clauses are generally understood as elided interrogative clauses, and proposed that sluiced clauses possess unpronounced syntactic structure that resembles the structure of standard *wh*-questions. For this to be true, he expected to see that sluiced clauses share syntactic properties with standard *wh*-questions. Indeed, he showed that there is evidence to support this hypothesis. First, he observed that the *wh*-words that head the sluiced clause must carry the same case as its correlate in the antecedent clause. This fact is difficult to see in English, due to its relatively impoverished inflectional system. However, it is quite easy to see in languages such as German, which utilize a rich inflectional system. Ross offered the German examples in (8) to show this case-matching effect. The verb *schmeicheln* assigns dative case to its object *jemandem* in the first clause of (8a), and dative case is the only case permitted on the stranded *wh*-phrase; all other case markings generate an ungrammatical sentence. A similar phenomenon can be seen in (8b) in which the verb *loben* assigns accusative case to its object; in this case, the stranded element must carry accusative case.

- (8) a. Er will jemandem schmeicheln, aber sie wissen nicht, { *wer
 he wants someone.DAT flatter but they know not who.NOM
 /*wen /wem }
 who.ACC who.DAT
 ‘He wants to flatter someone, but they don’t know who.’
- b. Er will jemanden loben, aber sie wissen nicht, { *wer /wen
 he wants someone.ACC praise but they know not who.NOM who.ACC
 /*wem }
 who.DAT
 ‘He wants to flatter someone, but they don’t know who.’
- (9) a. { *wer /*wen /wem } will er schmeicheln?
 who.NOM who.ACC who.DAT want he flatter
 ‘Who does he want to flatter?’
- b. { *wer /wen /*wem } will er loben?
 who.NOM who.ACC who.DAT want he praise
 ‘Who does he want to flatter?’

When you compare the contrasts in (8) to the standard *wh*-questions in (9), you can see that the *wh*-words in the standard questions must also carry the same object case that would be assigned by the verb in a declarative sentence. A typical analysis of the case effects in (9) proposes that the *wh*-phrase is originally generated in the object position, where it will receive the appropriate case assignment, then moves to the front of the sentence to form a question. That the stranded *wh*-phrases in the sluicing examples in (8) show a similar pattern of behaviour suggests that they may also be originally generated in the object position of the unpronounced syntactic structure in the ellipsis site, as Ross proposed.

The final piece of evidence explored in this section is the distribution of preposition stranding in sluicing. The availability of preposition stranding is language-dependent –

that is, it is observed in some languages but not others. Merchant (2001) formulated the preposition-stranding generalization which states that only languages that permit stranded prepositions in standard *wh*-questions will permit stranded prepositions in sluicing. This generalization can be observed in examples (10) and (11), but see Merchant (2001) for a more complete set of examples. Languages like English, which permit preposition stranding in standard questions (10a) also allow it in sluicing (10b). In contrast, languages like German do not allow preposition stranding in standard *wh*-questions (11a) or in sluicing (11b).

- (10) a. Who was he talking with?
 b. Peter was talking with someone, but I don't know (with) who.
- (11) a. *Wem hat sie mit gesprochen?
 who has she with spoken
 b. Anna hat mit jemandem gesprochen, aber ich weiß nicht, *(mit) wem.
 Anna has with someone spoken, but I know not with who

The fact that sluicing patterns with *wh*-questions with regards to case-matching and the language-specificity of preposition stranding suggests that these two linguistic phenomena share a common syntactic structure, and that this structure is unpronounced in sluiced clauses.

In this section, it was observed that verb-phrase ellipsis constructions are sensitive to locality conditions on *wh*-movement, including the Complex Noun Phrase Constraint and the Coordinate Structure Constraint (Ross, 1967). It was also observed that sluicing constructions (in which an entire interrogative clause has been elided) behave exactly like standard *wh*-questions with respect to case assignment (Ross, 1969; Merchant, 2001) and preposition-stranding phenomena (Merchant, 2001). All of these facts are consistent with the hypothesis that there is unpronounced syntactic structure underlying the ellipsis site, and are unexpected under any hypothesis that assumes otherwise.

1.3 Defining “Parallelism”

As outlined in Section 1.1, verb-phrase ellipsis is thought to be licensed by the presence of a parallel verb phrase. Defining “parallelism” in this context has long been a challenge for linguists, and a variety of theories now exist in the literature. Many theories suggest that the relevant parallelism domain operates at the level of syntax. These syntactic parallelism theories are described in Section 1.3.1. Other theories suggest that parallelism should be defined at some post-syntactic level, such as semantics or discourse levels. These non-syntactic theories are discussed in Section 1.3.2. The currently accepted definition of parallelism is addressed in Section 1.3.3. This theory claims that neither syntax nor semantics alone is sufficient for defining the relevant parallelism domain. Instead, both levels are relevant for determining whether verb-phrase ellipsis can

occur. Before discussing a phenomenon that poses difficulties for many of these theories, this section concludes with a brief review.

1.3.1 Syntactic Parallelism

Section 1.2 provided evidence to support the hypothesis that the ellipsis site hosts unpronounced syntactic structure. These facts led to theories of verb-phrase ellipsis which propose a syntactic parallelism constraint: verb phrases may be elided only if there exists another (usually overt) verb phrase with an identical syntactic structure. While the theories discussed in this section agree that elided verb phrases require a syntactically parallel antecedent, they can be divided into two camps based on how they believe the elided verb phrase acquires its structural content. One camp believes that the unpronounced syntactic structure exists in the ellipsis site throughout the entire derivation, but is deleted at a later stage just before pronunciation. The other camp believes that the ellipsis site is inherently empty until the structure of the antecedent verb phrase has been copied into the empty position at some stage preceding interpretation. In this section, deletion theories are described first, followed by a discussion of the copy theories.

1.3.1.1 Deletion Theories

Ross (1969) proposed that ellipsis could be explained by a deletion operation that proceeds under the condition that the elided surface structure is parallel to the surface structure of another constituent in the sentence. He applied this theory to sluicing, showing that the sluiced clause must contain syntactic structure, as evidenced by the case-matching effects shown in (8) in Section 1.2 and the selectional restrictions that determine the distribution of sluicing. The examples in (13) show that sluicing is restricted to only those clauses which would otherwise host a full interrogative clause (12).

- (12) a. I wonder [_{CP[+Q]} how many men she's inviting.]
 b. *I wonder [_{NP} those old men].
- (13) a. She says she's inviting some men – I wonder how many men.
 b. She says she's inviting some men – I wonder [_{CP} [how many men]_i ~~she's inviting~~ [_{t_i}]].
 c. *She says she's inviting some men – I wonder those old men.

Ross went on to show that the relevant parallelism condition operates at surface structure using evidence from pied piping. Pied piping results from *wh*-extraction targeting a full prepositional phrase rather than only the embedded *wh*-phrase. Pied piping is realized as a preposition that has been fronted along side the targeted *wh*-phrase, and can be seen in both sluicing (15) and standard questions (14).

- (14) a. John went to the store with someone.
 b. With whom did John go to the store?
 c. Who did John go to the store with?
- (15) a. John went to the store with someone, but I don't know with whom.
 b. John went to the store with someone, but I don't know who.

Observing that sluicing and standard questions show the same patterns of pied piping, Ross concluded that sluiced interrogatives must undergo the same question formation operations as standard questions. Assuming that these operations result in surface structure, he suggested that deletion must operate on surface structure, and is only licensed when the elided surface structure is parallel to some antecedent surface structure.

Sag (1976), however, showed that surface parallelism accounts make certain incorrect predictions, particularly with regards to semantic factors. First, he argued that these accounts predict that the scope distribution possibilities in sentences like those in (16) will not be affected by ellipsis. In this example, Sag observed that the first clause may receive both a surface scope reading (such that the same individual hit every other person) and an inverse scope reading (such that every person was hit, but not necessarily by the same individual), despite the fact that the second clause must receive a surface scope reading (because *Bill* is not a quantificational object).

- (16) Someone hit everyone, and then Bill hit everyone.
 a. = $[_{CP} [\exists x] [\forall y] [_{VP} x \text{ hit } y]]$ (*surface scope*)
 b. = $[_{CP} [\forall y] [\exists x] [_{VP} x \text{ hit } y]]$ (*inverse scope*)

In contrast, when the verb phrase in the second clause is elided, Sag observed that the interpretation of the first clause is restricted to only the surface scope reading, as seen in (17). According to Sag, the surface structures for (16) and (17) are identical, so surface parallelism theories would predict that (17) would show the same scope distribution patterns, contrary to fact.

- (17) Someone hit everyone, and then Bill did.
 a. = $[_{CP} [\exists x] [\forall y] [_{VP} x \text{ hit } y]]$ (*surface scope*)
 b. \neq $[_{CP} [\forall y] [\exists x] [_{VP} x \text{ hit } y]]$ (*inverse scope*)

Sag further argued that surface parallelism incorrectly predicts that (18a) is grammatical in the context in (18b) because the elided verb phrase in *What did Bill?* is parallel to the verb phrase in *What did Harry take a picture of?* on the surface.

- (18) a. *What did Bill?
 b. A: What did Harry take a picture of?
 B: An elephant.

- A: What did Bill?
 B: A tiger.

Sag (1976) proposed that these examples could be explained by a deletion account if the parallelism condition is moved to the level of logical form (LF). To be clear, he suggested that deletion operates on surface structures, but is licensed only when the LF structure of the elided verb phrase is parallel to the LF structure of the antecedent. Here, Sag defines parallelism as ‘alphabetic variance’. For one structure to be an alphabetic variant of another, the two structures must be identical up to the identity of any shared variables. Sag applied this proposal to the scope restrictions seen in (17). His analysis is shown in (19). The elided verb phrase can take only the surface scope reading (19a), which leaves only the surface scope reading of the first clause (19b-i) available as an antecedent (because the inverse scope reading (19b-ii) is not an alphabetic variant).

- (19) Someone hit everyone, and then Bill did.
- a. ...and then Bill did ~~hit everyone~~.
 - i. = $[_{CP} [\text{Bill}]_x [\forall y] [_{VP} x \text{ hit } y]]$
 - ii. $\neq [_{CP} [\forall y] [\text{Bill}]_x [_{VP} x \text{ hit } y]]$
 - b. Someone hit everyone...
 - i. = $[_{CP} [\exists x] [\forall y] [_{VP} x \text{ hit } y]]$
 - ii. $\neq [_{CP} [\forall y] [\exists x] [_{VP} x \text{ hit } y]]$

Sag also showed how this explains (18b). His analysis is presented in (20).

- (20) a. What did Harry take a picture of?
 $[_{CP} [\text{what}]_w [\text{Harry}]_x [_{VP} x \text{ take a picture of } w]]$
- b. What did Bill ~~take a picture of~~?
 $[_{CP} [\text{what}]_y [\text{Bill}]_z [_{VP} z \text{ take a picture of } y]]$

Sag suggested that each occurrence of a *wh*-phrase introduces a unique variable into the LF structure. Therefore, the reason why *What did Harry take a picture of?* is not a good antecedent for *What did Bill?* is because the *wh*-phrase in the elided sentence is associated with a variable that is distinct from the one bound by the *wh*-phrase in the LF for the antecedent sentence. From this, Sag concluded that deletion is not licensed in this context because the parallelism condition of alphabetic variance is not met.

1.3.1.2 Copy Theories

As just described, Sag (1976) observed several semantic problems for deletion accounts that rely on a surface parallelism condition. However, he argued that the deletion operation inherent to these accounts could be maintained if the parallelism condition was shifted to the level of logical form. Several authors approached these problems from a

different direction, claiming that the structure of the elided phrase was in fact empty until it reached LF, where it would acquire content via a mechanism that copies the LF structure of the antecedent into the empty ellipsis site (Chung, Ladusaw, & McCloskey, 1995; Fiengo & May, 1994; Williams, 1977).

Williams (1977) argued that this LF-copy account was superior to a deletion account. Fundamental to Williams' argument was a distinction between Sentence Grammar and Discourse Grammar. He suggested that Sentence Grammar consisted of transformation rules that derive surface structures from deep structures, and a set of semantic interpretation rules that operate on surface structures to derive LF structures. On the other hand, Discourse Grammar is comprised of a set of semantic interpretation rules that are different from those in Sentence Grammar in that they only operate on LF structures. Under this distinction, he proposed that the verb-phrase ellipsis rule was a part of the Discourse Grammar, and that this rule stipulated that the LF structure of the antecedent must be copied into the empty category representing the elided information. The reason this copy approach was superior to a deletion approach, he claimed, was because deletion violated the principle of 'strict utterance'. This principle states that all Discourse Grammar rules must apply after all Sentence Grammar rules have been applied. Williams proposed that verb-phrase deletion rules violate this principle because the DO-support operation (a Sentence Grammar rule) is expected to follow deletion under most standard deletion accounts. This is shown in (21), where DO-support applies only after the verb phrase has been deleted.

- (21) a. A: Who left?
 B: John left.
 b. A: Who left?
 B: John ~~left~~. (*deletion*)
 c. A: Who left?
 B: John did ~~left~~. (*DO-support*)

In contrast, Williams suggested that the rule that derives ellipsis is a Discourse Grammar rule, and will necessarily apply after DO-support, as shown in (22).

- (22) a. A: Who left?
 B: John did [_{VP} e].
 b. A: Who left?
 B: John did [_{VP} ~~leave~~]. (*LF copy*)

Here, the structure of the ellipsis site is empty, and only receives content after the verb-phrase ellipsis rule has been applied.

Chung et al. (1995) extended this theory to sluicing, which offers the intuitive derivation shown in (23).

- (23) Joan ate dinner with someone, but I don't know (with) who(m).

- a. $\exists x$. Joan ate dinner with x , but I don't know (with) who $[_{IP} e]$.
- b. $\exists x$. Joan ate dinner with x , but I don't know (with) who $[_{IP} \exists x . \text{Joan ate dinner with } x]$.

However, they acknowledged that there are cases of sluicing in which the *wh*-remnant does not have an overt correlate in the antecedent. In this type of situation, simply copying the LF structure of the antecedent into the empty ellipsis site will not generate the correct interpretation. This is because there will be no variable present in the structure for the *wh*-remnant to bind, as shown in (24).

- (24) Joan ate dinner, but I don't know with whom.
- a. $[_{IP} \text{Joan ate dinner}]$, but I don't know with whom $[_{IP} e]$.
- b. $[_{IP} \text{Joan ate dinner}]$, but I don't know with whom $[_{IP} \text{Joan ate dinner}]$.
- (25) Joan ate dinner, but I don't know with whom.
- a. $[_{IP} \text{Joan ate dinner}]$, but I don't know with whom $[_{IP} e]$.
- b. $[_{IP} \text{Joan ate dinner}]$, but I don't know with whom $[_{IP} \exists x . \text{Joan ate dinner with } x]$.

To remedy this problem, Chung et al. claim that an appropriate position will be 'sprouted' at LF, resulting in the derivation in (25). Here, when the LF structure is copied into the empty ellipsis site and there is no free variable to bind, the system will construct a variable of the appropriate type, thus rescuing the derivation.

The copy mechanism has also been used to explain why sentences like those in (4), (26), can receive more than one interpretation. The ambiguity in these sentences is tied to the pronoun in the elided clause - does it refer to the subject of the antecedent clause (the strict reading in (26a)) or to the subject of the elided clause (the sloppy reading in (26b))? Given that the elided material receives its interpretation from the antecedent, it is typically assumed that this ambiguity originates in the antecedent clause.

- (26) John likes his job, and Bill does ~~like his job~~ too.
- a. = ... Bill likes John's job.
- b. = ... Bill likes Bill's job.

Fiengo and May (1994) suggested that the ambiguity follows from the type of dependency relation the pronoun holds with its referent. According to their theory, if the pronoun represents an independent occurrence (or α -occurrence), then the pronoun behaves as a typical referential noun phrase. If the pronoun represents a dependent occurrence (or β -occurrence), it only receives its referent via coindexation with another pronoun. The theorists claimed that these indexical dependencies are aspects of the LF structure, and thus assume an LF identity condition. They proposed that the system necessarily copies the dependency relation associated with the pronoun when it copies the LF structure from the antecedent to the ellipsis site.

- (27) Max saw his mother and Oscar did too.
- a. Max₁ [_{VP} saw his₁^α mother] and Oscar₂ [_{VP} ~~saw his₁^α mother~~].
 - b. Max₁ [_{VP} saw his₁^β mother] and Oscar₂ [_{VP} ~~saw his₂^β mother~~].

In the case of α -occurrences, they claimed that the system is instructed to copy the entire occurrence, including the index. In the case of β -occurrences, the system is instructed to copy only the dependency relation. In these cases, the pronoun will receive its index from a referential expression that sits in the same configuration with the pronoun as that in the antecedent. This can be observed in (27). The strict reading in (27a) arises from copying an antecedent with an α -occurrence pronoun. In contrast, the sloppy reading in (27b) arises from copying an antecedent with a β -occurrence pronoun.

1.3.2 Non-Syntactic Parallelism

The last two sections presented two types of syntactic theories of verb-phrase ellipsis. The first type proposed that the syntactic structure in the ellipsis site exists from the beginning of the derivation, but is deleted prior to pronunciation. The second type claimed that the ellipsis site is inherently empty until the antecedent structure is copied into the site at some point before an interpretation is assigned. These syntactic theories typically run into problems when addressed from a semantic perspective. For example, verb-phrase ellipsis shares many properties with ‘do-X’ anaphora (28), which cannot be explained exclusively by a deletion approach (Culicover & Jackendoff, 2005; Hankamer & Sag, 1976).

- (28) a. Robin read the newspaper today, but Leslie didn’t do so.
 b. Robin smokes a pipe after breakfast, and Leslie does so during breakfast.

For this reason, a number of theorists have opted to abandon the idea that ellipsis hosts syntactic structure in favour of theories that are strictly semantic (Culicover & Jackendoff, 2005; Dalrymple, Shieber, & Pereira, 1991) or that appeal to the greater discourse structure (Hardt, 1999; Kehler, 2000).

Hankamer and Sag (1976) argued for two different types of verb-phrase anaphora: verb-phrase ellipsis which they classified as ‘surface anaphora’, and ‘do it’ anaphors which they termed ‘deep anaphora’. Despite a number of superficial similarities between the two forms of anaphora, the authors observed that surface anaphors (29) allow an overt pronoun to refer to an unpronounced referent in the ellipsis site (29b), but deep anaphors do not (30).

- (29) a. I’ve never ridden a camel, but Ivan has ridden a camel_i, and he says it_i stank horribly.
 b. I’ve never ridden a camel, but Ivan has, and he says it_i stank horribly.

- (30) a. * Jack didn't cut Betty with a knife – Bill did it, and *it* was rusty.
 [*it* = the knife]
 b. * Jack didn't cut Betty with a knife – Bill did so, and *it* was rusty.

Based on the contrast observed with respect to this 'missing antecedent' phenomenon, Hankamer and Sag argued that the verb-phrase ellipsis examples in (29) must possess (deleted) syntactic structure which contains a noun phrase in the appropriate position to bind the pronoun. As such, these surface anaphors must contain syntactic structure, and their interpretation is derived from the syntax. In contrast, the overt verbal pro-form (*it* or *so*) in the 'do X' examples in (30) offer no such syntactic structure, leaving the pronoun without a referent. Because these deep anaphors contain no syntactic structure, the authors argue that they must be interpreted via pragmatic control – that is, their interpretation is derived from the context. Thus, Hankamer and Sag proposed a combined syntactic-semantic account of verb-phrase anaphora.

Culicover and Jackendoff (2005) advocated a purely semantic account of verb-phrase ellipsis, which places the mechanism of interpreting ellipsis outside of language proper and relies on more general cognitive principles. They suggested that meaning is represented in the Conceptual Structure, which they claim is a "central system of the mind (p 20)". Thus, the Conceptual Structure is a cognitive system that is not language-specific, and is the system that is fundamentally responsible for cognitive reasoning. In terms of deriving an interpretation for verb-phrase ellipsis, Culicover and Jackendoff proposed a mechanism called 'indirect licensing'. They argued that the antecedent clause (notated with superscript ANT) activates a proposition in the Conceptual Structure. Under this proposal, verb-phrase ellipsis is indirectly licensed by this activated proposition in the following way: The elided clause (notated with superscript IL) contains a constituent that has been orphaned (ORPH), which represents information in the elided proposition that differs from some target in the antecedent. Upon encountering the orphan, the system is instructed to search for an active proposition in the Conceptual Structure and match the orphan with a target within the proposition. Refer to (31) for an illustration.

- (31) [John^{TARGET} went to the store]^{ANT}, and [Mary^{ORPH} did]^{IL} too.

Upon encountering *Mary*, the system searches the Conceptual Structure for a proposition containing an appropriate target element. There, it will find *John went to the store*, match the orphan *Mary* with *John*, and derive an interpretation for the ellipsis site.

Hardt (1999) proposed a similar process that operates at the discourse level. He argued that all propositions can interact with the discourse in two ways. One way is to update the discourse structure with the new information provided by the proposition. The other way is to use the current discourse structure as a context for interpreting the incoming information. He showed how this can explain the difference between strict and sloppy identity readings, like those in (26), repeated here in (32).

- (32) John likes his job, and Bill does ~~like his job~~ too.

- a. = ... Bill likes John's job.
- b. = ... Bill likes Bill's job.

He argued that the discourse structure centres around a single element, and the rest of the discourse provides information about this central element. In terms of sentences like that in (32), Hardt proposed that the antecedent *John likes his job* creates a discourse structure that is centred on *John*. Upon encountering the elided clause, the system can maintain the current discourse centre (i.e., keep *John* as the central element) or it can change the discourse such that the central element becomes *Bill*. Hardt proposed that this is exactly how the strict and sloppy identity readings come about. When the system maintains *John* as the centre of the discourse, the elided pronoun will refer back to John – deriving the strict reading in (32a). However, if the system changes the discourse to centre around *Bill*, then the pronoun will refer to Bill and produce the sloppy reading in (32).

1.3.3 Merchant's e-GIVENNESS

Merchant (2001) offered the most comprehensive account of ellipsis to date. His theory builds upon the arguments of several authors who note that ellipsis must satisfy a condition on the distribution of focus (Rooth, 1992; Schwarzschild, 1999; Tancredi, 1992). Rooth (1992) and Tancredi (1992) observed that verb-phrase ellipsis shares characteristics with phonologically reduced verb-phrases, in which redundant information is pronounced with a lower pitch accent. In particular, the researchers acknowledged that the two phenomena are in free variation (33). (Note: I mark phonologically reduced material with *italic font*.) For this reason, Rooth (1992) suggested that the verb phrases in both cases must satisfy the focus condition in (34).

- (33)
- a. Abby was reading the book while [Ben]_F was ~~reading the book~~.
 - b. Abby was reading the book while [Ben]_F *was reading the book*.
 - c. Abby called Chuck an idiot after [Ben]_F did ~~call Chuck an idiot~~.
 - d. Abby called Chuck an idiot after [Ben]_F *called Chuck an idiot*.

- (34) **Rooth's Focus Condition:** An XP α in XP_E can be deleted only if there is an XP_A, where XP_A either is or implies an element of XP_E.

In particular, Rooth argued that the antecedent verb phrase and the elided/reduced verb phrase must belong to the same set of focus alternatives, which is a set of propositions of the same form. Under this proposal, the set of propositions that make up the focus value of the elided verb phrase in (33) takes the form [$\lambda x.x$ was reading the book], and it can be seen that both the antecedent (*Abby was reading the book*) and the elided (*Ben was reading the book*) verb phrases belong to this set.

However, Rooth (1992) also observed that there are cases in which phonological reduction is permitted but verb-phrase ellipsis is not (35).

- (35) a. *Abby was reading the book while [Ben]_F was ~~reading~~.
 b. Abby was reading the book while [Ben]_F *was reading*.
 c. *Abby called Chuck an idiot after [Ben]_F did ~~insult him~~.
 d. Abby called Chuck an idiot after [Ben]_F *insulted him*.

Thus, Rooth argued that ellipsis is subject to an additional syntactic parallelism constraint, such that the syntax of the elided constituent must be identical to that of the antecedent. This constraint accounts for the acceptability of ellipsis in (33) because the elided constituent is parallel to the antecedent. It also accounts for the unacceptability of ellipsis in (35) because the elided constituent is in fact non-parallel to the antecedent.

Merchant (2001) rejected Rooth's (1992) syntactic parallelism condition, however, as it does not account for the acceptability of sluicing with implicit correlates (36) or sluicing of infinitival clauses with finite antecedents (37). Despite the fact that the elided constituents are not parallel to their antecedents, these sentences remain grammatical.

- (36) a. Abby was reading, but I don't know what ~~Abby was reading~~ *t*.
 b. Ben called - guess when ~~Ben called~~ *t* !

- (37) Decorating for the holidays is easy if you know how!
 a. ≠ * ...how [decorating for the holidays]
 b. = ...how [to decorate for the holidays]

Although Merchant (2001) believed that elided constituents contain fully derived syntactic structures (this accounts for the evidence presented in Section 1.2), and that this syntactic structure is deleted at the level of phonological form (PF), he argued in favour of a strictly semantic licensing condition on ellipsis. In particular, he offered a revised focus condition (38) which elided verb phrases must satisfy before deletion. He adopted and revised Schwarzschild's (1999) GIVENNESS condition (which offers an analysis of (33) that is roughly analogous to Rooth's (1992) analysis, described above), and developed the notion of e-GIVENNESS (39). Essentially, Merchant argued that a constituent can only be elided if it has an overt antecedent, and both constituents (antecedent and elided) mutually entail each other.

- (38) **Merchant's Focus Condition:** An XP α can be deleted only if α is e-GIVEN.

- (39) **e-GIVENNESS:** An expression E counts as e-GIVEN iff E has a salient antecedent A and, modulo \exists -type shifting,
 a. A entails F-clo(E), and

b. E entails F-clo(A).

- (40) **F-closure:** The F-closure of α , written $F\text{-clo}(\alpha)$, is the result of replacing F-marked parts of α with \exists -bound variables of the appropriate type (modulo \exists -type shifting).

In addition, Merchant proposed that deletion is licensed by a feature located on the head preceding the ellipsis site. He called this the E-feature, and suggested that the E-feature is responsible for informing PF to ignore the syntactic structure in the complement of its host. He further suggested that the E-feature was encoded with its own semantic information (41), which is designed to enforce the focus condition in (38). (41) is a partial identity function that assigns the ellipsis site the interpretation of some proposition p , provided that p is e-GIVEN.

- (41) $\llbracket E \rrbracket = \lambda p : p \text{ is e-GIVEN} . p$

One of the advantages of Merchant's proposal is the idea that the E-feature is encoded with both the licensing and interpretation requirements for ellipsis. That is, his theory assumes that the licensing and interpretation requirements are intrinsically linked – an assumption that previous theories lacked.

The paradigm in (42) shows how Merchant's (2001) e-GIVENNESS proposal accounts for the verb-phrase ellipsis in (33c) while (43) excludes the ungrammatical (35c). In (42), the two verb-phrases mutually entail each other, satisfying the e-GIVEN requirement. This is not true in (43), however: while VP_A entails VP_E (because calling someone an idiot entails insulting him/her), the reverse is not true (insulting someone does not necessarily entail calling him/her an idiot).

- (42) Abby called Chuck an idiot after Ben did ~~call Chuck an idiot~~.

- a. $VP_A' = \exists x.x \text{ called Chuck an idiot}$
- b. $F\text{-clo}(VP_E) = \exists x.x \text{ called Chuck an idiot}$
 $\therefore VP_A \text{ entails } F\text{-clo}(VP_E)$
- c. $VP_E' = \exists x.x \text{ called Chuck an idiot}$
- d. $F\text{-clo}(VP_A) = \exists x.x \text{ called Chuck an idiot}$
 $\therefore VP_E \text{ entails } F\text{-clo}(VP_A)$

- (43) *Abbey called Chuck an idiot after Ben did ~~insult him~~.

- a. $VP_A = \exists x . x \text{ called Chuck an idiot}$
- b. $F\text{-clo}(VP_E) = \exists x . x \text{ insulted Chuck}$
 $\therefore VP_A \text{ entails } F\text{-clo}(VP_E)$
- c. $VP_E = \exists x . x \text{ insulted Chuck}$
- d. $F\text{-clo}(VP_A) = \exists x . x \text{ called Chuck an idiot}$
 $\therefore VP_E \text{ does not entail } F\text{-clo}(VP_A)$

The paradigms in (42) and (43) show how Merchant's e-GIVENNESS condition can account for the same examples Rooth (1992) explained ((33c) and (35c)) without appealing to a syntactic parallelism condition. Thus, Merchant's theory avoids the various problems associated with syntactic parallelism constraints (see, for example, Sag's semantic problems in Section 1.3.1.1), while still accounting for the evidence suggesting that the elided constituent contains syntactic structure (see Section 1.2).

1.4 Non-Parallel Ellipsis

As discussed in the previous section, verb-phrase ellipsis is typically licensed only in the presence of a parallel antecedent verb phrase. However, ellipsis is sometimes acceptable even in cases where the antecedent verb phrase is not parallel to the elided verb phrase. As mentioned in Section 1.1, the experiments reported in this thesis were designed to test comprehension of verb-phrase ellipsis with non-parallel antecedents. This section describes two forms of acceptable non-parallelism in ellipsis: (i) voice mismatch (Section 1.4.1) and (ii) morphological mismatch (Section 1.4.2). Instances of verb-phrase ellipsis with non-parallel morphology are the focus of the experiments described below.

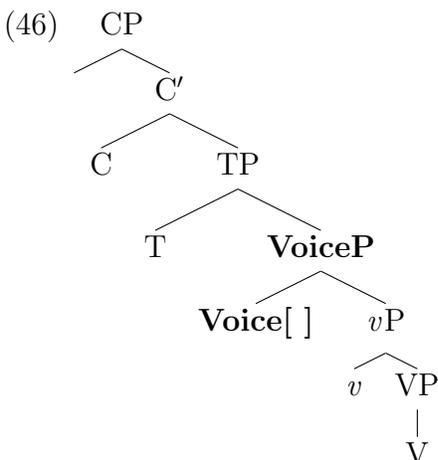
1.4.1 Voice Mismatch

It is widely accepted that verb-phrase ellipsis permits voice mismatch between the antecedent and elided verb phrases (44) but sluicing (45) does not (e.g., Merchant, 2008a; 2012). In (44a,b), the antecedent verb phrase is in active form while the elided verb phrase is passive (the reverse is true for (44c,d)). (45a) exemplifies an ungrammatical instance of an active antecedent for a passive sluiced clause (again, the reverse is true for (45b)).

- (44) a. The janitor must remove the trash whenever it is apparent that it should be ~~removed~~.
- b. Actually, I have implemented it [\bar{a} computer system] with a manager, but it doesn't have to be ~~implemented with a manager~~.
- c. This problem was to have been looked into, but obviously nobody did ~~look into this problem~~.
- d. This information could have been released by Gorbachev, but he chose not to ~~release the information~~.
- (45) a. * Joe was murdered, but we don't know who ~~murdered Joe~~.
- b. * Someone murdered Joe, but we don't know who by ~~Joe was murdered~~.

Merchant (2008a; 2012) has provided the predominant explanation for this distinction, though it should be noted that a number of researchers have adopted this account with

some modifications (e.g., Baltin, 2012; Tanaka, 2011). Following Collins (2005), Merchant assumed that there is a Voice head external to the verbal projection, as in (46). He also suggested that the antecedent and elided constituents are required to satisfy a syntactic featural identity condition – that is, the features within the structure must be identical across the two constituents.



Merchant argued that the distinction between verb-phrase ellipsis and sluicing could be explained by the level of the structure targeted for deletion by the E-feature. For verb-phrase ellipsis, he suggested that the E-feature is located directly on the Voice head and licenses the deletion of the *vP* projection. Deletion at the *vP* level does not delete the Voice information, so it need not satisfy the feature identity condition on ellipsis. In contrast, the E-feature is located on the C head in sluicing, licensing deletion at the TP level. In this case, the Voice information is deleted, so it must satisfy the feature identity condition by maintaining consistent active or passive voice across constituents. This explains why the sentences in (45) are ungrammatical: the value on the Voice head in the two clauses differ, violating the feature identity condition.

1.4.2 Morphological Mismatch

Verb-phrase ellipsis of main verbs is generally shown to be resilient to mismatches between the antecedent and elided verb phrases in terms of morphological markings (e.g., Potsdam, 1997), as shown in (47).

- (47)
- a. Jack fell down and Jill did ~~fall down~~ too.
 - b. I didn't touch the TV set, but Percy might have ~~touched the TV set~~.
 - c. If you haven't told them yet, you really should ~~tell them~~.
 - d. A: Are they arguing?
B: Yes, they always do ~~argue~~.
 - e. John slept, and Mary will ~~sleep~~ too.

However, this is not true of ellipsis of auxiliary verbs, as shown by the contrast between cases in which the elided auxiliary matches the antecedent (48) and cases in which the elided auxiliary differs from the antecedent (49).

- (48) a. John will be happy about the result and Mary will ~~be happy about the result~~ too.
 b. He is being audited by the IRS because his company is ~~being audited by the IRS~~.
- (49) a. *I am confused about ellipsis and you will ~~be confused about ellipsis~~ too.
 b. *Mary is an auto mechanic and her daughter wants to ~~be an auto mechanic~~ also.

Several theories of morphology have been developed to account for this pattern of data, four of which are discussed below. The first three to be described all suffer from (near) fatal limitations; the fourth – the Distributed Morphology framework – offers the most successful explanation, and will be adopted in this thesis.

One early approach to verbal morphology – Affix Hopping (Chomsky, 1957) – proposed that verbs and auxiliaries are uninflected when they are inserted into the structure. Tense morphemes are generated under the T head and adjoin to the verb by moving downward, as seen in (50b) for the simple sentence in (50a).

- (50) a. The boy liked the girl.
 b.
-

This type of theory could presumably account for the examples of morphological mismatch on the main verbs in (47). Prior to the Affix Hopping procedure, the two verb phrases are identical; thus, one would simply posit that the parallelism requirement must be satisfied prior to Affix Hopping. However, this theory incorrectly predicts that the examples in (49) will be grammatical.

The lexicalist approach to verbal morphology (Chomsky, 1995; Potsdam, 1997) proposed that verbs and auxiliaries are inserted into the structure in their fully inflected forms. The lexicon stores instances of each morphological form of a verb, so the trick is simply to select the correct form from the lexicon. This account would be promising

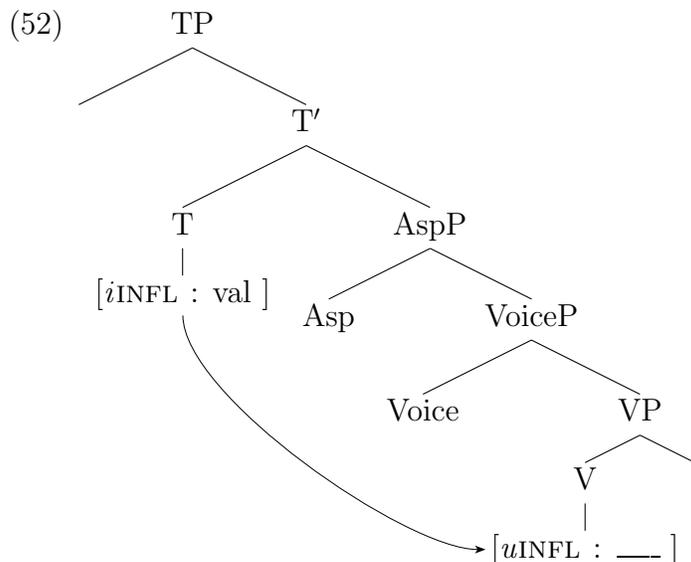
as an explanation for the examples in (47) and (49). However, the analysis described violates the parallelism requirement on ellipsis – if the verbal items are inserted in their inflected form, then there is no stage in the derivation at which the elided verb phrase is parallel to the antecedent.

Lasnik (1995) proposed a hybrid approach to morphology, such that main verbs like the ones in (47) are inserted into the structure uninflected and undergo Affix Hopping procedures while auxiliary verbs (49) are inserted in their fully inflected forms (see also Omaki, 2007; 2009). This account could provide an explanation for the full paradigm described: morphological mismatch is permitted under ellipsis of main verbs because the elided and antecedent verb phrases are parallel prior to the application of Affix Hopping; morphological mismatch is not permitted under ellipsis of auxiliary verbs because the parallelism requirement is not satisfied.

- (51) a. Mary will [_{VP} leave] and John already has [-en [_{VP} ~~leave~~]].
- b. * John will not [_{VP} enter the competition], but Peter is [-ing [_{VP} ~~enter the competition~~]].

However, this theory also suffers from some difficulties. One such difficulty is that this analysis leads to structures that violate the Stray Affix Filter (which states that morphological affixes must be overtly realized at PF). As shown in (51), grammaticality is maintained when certain affixes are stranded (51a), but this is not true for all affixes (51b). Crucially, Lasnik’s hybrid approach provides no method of explaining this inconsistency.

Finally, the Distributed Morphology framework (DM) claims that morphological structure is fundamentally syntactic structure (Embick & Noyer, 2006; Halle & Marantz, 1993; Marantz, 1997). That is, the syntactic structure encodes a system of features which, when fed into the morphological component of the grammar, produces what eventually becomes realized as morphology. In other words, morphology is encoded in the syntax, but is interpreted post-syntactically. DM distinguishes between two types of morphemes: Roots are morphemes that represent lexical items, whose full (bare) forms are present in the structure throughout the entire derivation of the sentence. Abstract morphemes represent functional categories, and are encoded as features in the syntax that eventually receive phonological properties at PF. The mechanism that supplies abstract morphemes with their phonological properties is called Vocabulary Insertion, which relies on a ‘vocabulary’ of (what roughly equates to) phonological rules which determine the phonological realization of a given abstract morpheme based on its syntactic environment. Crucially, morphology as it applies to verbal inflection is not overtly realized until PF, despite the fact that the inflectional features are present throughout the derivation.

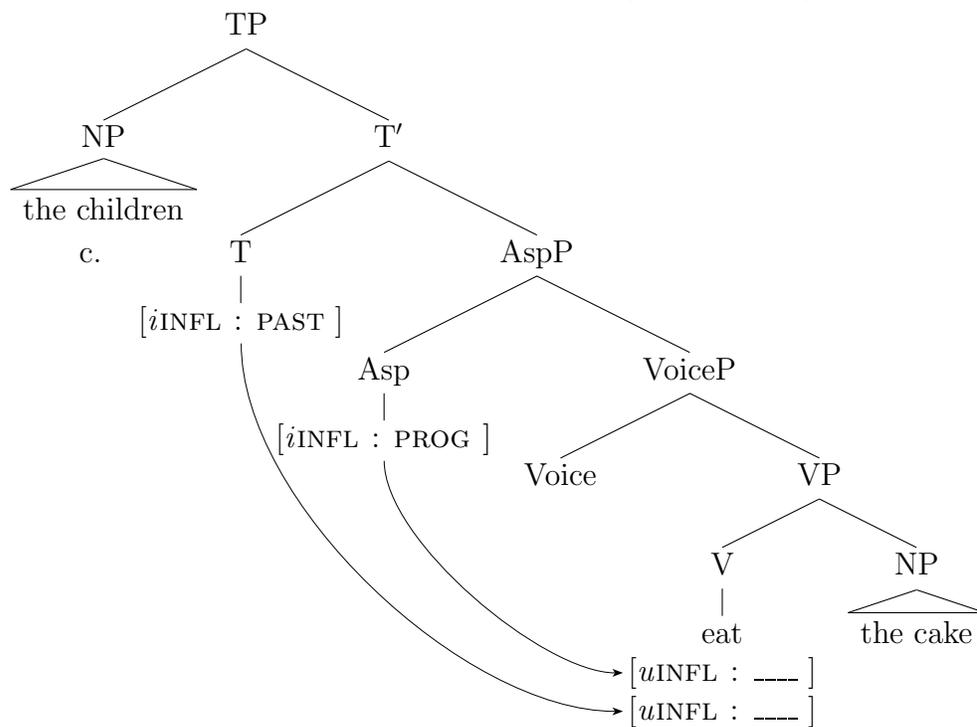
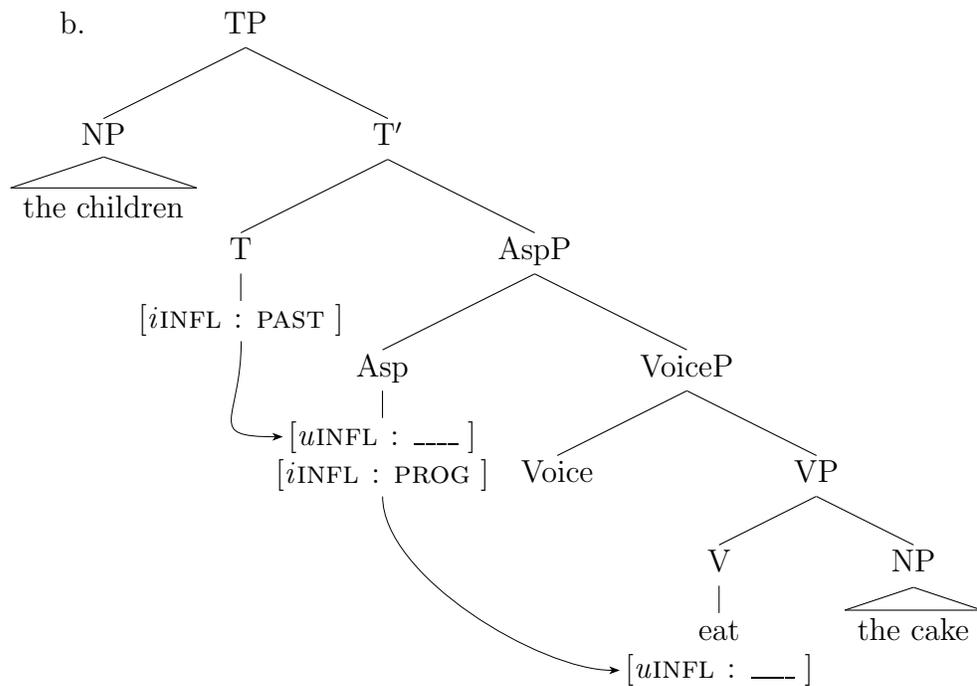


Working within a DM framework, Bjorkman (2011) proposed that verbal inflection is governed by an Agree relation that allows downward feature valuation. For the simplest form of tense assignment, in a structure such as the one in (52), she argued that the interpretable inflection feature (*i*INFL) on the T head provides a value for the uninterpretable feature (*u*INFL) on V. She further argued that the T head would assign its value to any intervening head that hosts an additional *u*INFL feature, and the intervening head would then pass the value along to V.

Crucially, Bjorkman proposed that auxiliaries *be* and *have* are the overt realization of INFL features that are blocked from reaching the verb. In other words, the presence of *be* or *have* represents a last-resort repair mechanism that functions to realize stranded morphological features (much like the traditional formulation of DO-support). She demonstrated how this works for the past progressive sentence in (53b), ruling out the possibility that both PAST and PROG features will be realized on the main verb (53c).

Bjorkman argued that an *i*INFL feature must provide a value to the nearest *u*INFL feature. She further argued that any functional head that hosts a *u*INFL feature simultaneously hosts an *i*INFL feature, which either carries its own value or adopts the value assigned from above. When the intervening *i*INFL feature carries a distinct value, the corresponding *u*INFL feature is left stranded on the intervening functional head and is blocked from being realized on the main verb. Bjorkman proposed that these are just the situations in which *be* and *have* occur.

(53) a. The children were eating the cake.



The contrasts in (47) and (49) can be accounted for in the Distributed Morphology framework quite easily. Morphological mismatch is permitted in ellipsis of main verbs because the $i\text{INFL}$ feature on T and $u\text{INFL}$ feature on V maintain a local Agree relation even under ellipsis. This morphological mismatch effect can be accounted for using a similar analysis to the one Merchant provides for the voice mismatch effects in the previous section. The $i\text{INFL}$ feature on T in the elided clause need not match the feature

in the antecedent clause because ellipsis targets a node lower in the structure, and so the parallelism requirement does not include the T head. In contrast, morphological mismatch is illicit in ellipsis of auxiliary verbs because these auxiliaries represent a last-resort repair operation that permits the realization of stranded INFL features. In (49), the inflectional features in the ellipsis site are stranded on a functional head, but deletion blocks the repair operation. Thus, the features cannot be realized at all, and the derivation fails. Note, however, that targeting a lower projection with ellipsis will allow the repair operation to proceed, and saves the derivation (54).

- (54) a. I am confused about ellipsis and (after reading this paper) you will be ~~confused~~
~~about ellipsis~~ too.
- b. Mary is an auto mechanic and her daughter wants to be ~~an auto mechanic~~
also.

In summary, this section described how four theories of verbal morphology attempt to explain the morphological mismatch paradigm seen in (47) and (49). Three of the four theories suffered from very serious problems: the Affix Hopping approach incorrectly predicted the sentences in (49) to be as acceptable as those in (47), lexicalism abandoned the widely accepted parallelism requirements on ellipsis, and Lasnik's (1995) hybrid theory had undesirable consequences outside of the paradigm considered here. In the end, the best account of the data came from the Distributed Morphology framework, particularly under the model put forward by Bjorkman (2011).

1.5 Accommodating Non-Parallelism

Section 1.3 showed that verb-phrase ellipsis is constrained by some form of parallelism, as a general rule. Section 1.4 contradicted this, showing that there are grammatical cases of verb-phrase ellipsis with non-parallel antecedents. The question that falls out of this paradox is: How do comprehenders deal with cases of non-parallelism? More specifically, what do comprehenders do to accommodate speakers who produce non-parallel ellipses?

When two people participate in a conversation, they develop a collection of information (more accurately: propositions) that they both assume or know to be true. This collection of propositions is called the 'common ground' (Stalnaker, 1970; 2002). Each sentence in the conversation updates the common ground with new information. All propositions within the common ground are said to be GIVEN; that is, the truth of each proposition is taken for granted because it is already known to be true by both participants. When a speaker produces a sentence for which the interpretation depends on the truth of some additional unspoken proposition, the speaker is said to be 'presupposing' the truth of the unspoken proposition. If this proposition already belongs to the common ground, this presupposition is relatively harmless; the listener must simply integrate the new material relative to the GIVEN proposition. However, if the unspoken proposition does not already belong to the common ground, the listener has two options: accommo-

date the speaker or let the conversation fail (von Stechow, 2000; 2006). Consider (55) as an illustration.

(55) I am sorry that I am late. I had to take my daughter to the doctor.

(55) presupposes the fact that the speaker has a daughter. According to von Stechow (2000; 2006), accommodation in these types of cases can be defined as a process of tacitly updating the common ground to include the presupposed proposition before integrating the asserted proposition (here, *I am late because I took my daughter to the doctor*).

In terms of ellipsis, only GIVEN information can be elided (Merchant, 2001; Rooth, 1992; Schwarzschild, 1999; Tancredi, 1992). When the elided verb phrase has a parallel antecedent, the ellipsis is relatively harmless; the listener must simply interpret the verb phrase relative to the existing verb phrase. However, in the case of non-parallel ellipsis, the listener is forced to accommodate the speaker. According to Fox (1999), a speaker may only produce a non-parallel ellipsis sentence if the antecedent verb phrase entails an ‘accommodation sentence’ that functions as an appropriate antecedent for the elided verb phrase. From the perspective of the listener, Fox argued that the accommodation process is governed by an economy condition. That is, he claimed that accommodation sentences are not always available to the listener, because the listener does not engage in accommodation processes without motivation. The only time accommodation is possible is when the elided clause contains ‘accommodation-seeking material’ – pronounced (deaccented) material that is absent in the overt antecedent. This accounts for the paradigm in (56)

- (56) a. * John proved that I’m innocent. [Fingerprints]_F did, too.
 b. * Fingerprints proved that I’m innocent. [John]_F did, too.
 c. John proved that I’m innocent. *Fingerprints that* [Bill]_F *presented* did, too.

All three of the examples in (56) contain verb-phrase ellipsis with non-parallel antecedents (because *John* and *fingerprints* have different thematic roles). However, only (56c) is grammatical because it is the only example that provides a trigger for accommodation - i.e., it is the only example that includes accommodation-seeking material. Fox concluded that, without this trigger, the listener has no reason to perform the accommodation procedure, and the sentence is deemed ungrammatical (56a,b).

Interestingly, Fox’s (1999) economy condition on accommodation competes with Merchant’s (2008b) economy condition on ellipsis, MaxElide. MaxElide is a speaker-oriented economy-driven principle which prefers to elide the largest possible constituent (see also Parker & Seely, 2010; Takahashi & Fox, 2005). Merchant introduced MaxElide to account for the data in (57) which show that verb-phrase ellipsis is not permitted in environments where sluicing may occur – namely those environments that involve *wh*-extraction out of syntactic islands.

- (57) a. Sue criticized someone, but I don’t know who ~~Sue criticized~~ *t*.

- b. *Sue criticized someone, but I don't know who she did ~~criticize~~ *t*.

Merchant's (2008b) MaxElide prefers that the speaker elide the largest possible constituent, saving him/her the effort of pronouncing redundant material. Meanwhile, Fox's (1999) economy condition requires that some material be present in the signal in order to trigger the necessary accommodation processes.

1.6 What is the problem? (Summary)

This chapter reviewed several theories of how to account for the preference for parallelism in verb-phrase ellipsis (Section 1.3), provided evidence that verb-phrase ellipsis can remain grammatical even with non-parallel antecedents (Section 1.4), and briefly discussed some methods of accommodating non-parallelism in verb-phrase ellipsis (Section 1.5). Having covered this material, this chapter sets the stage for the coming chapters. Chapter 2 reviews the psycholinguistic research conducted to explore the comprehension of verb-phrase ellipsis. Chapter 3 presents a series of experiments designed to investigate how verb-phrase ellipsis comprehension proceeds when the morphology on the elided verb phrase is not parallel to the morphology on the antecedent verb phrase. Chapter 4 discusses the consequences of the results for the psycholinguistics and linguistics domains.

Chapter 2

How do we Process Ellipsis? Evidence from Psycholinguistics

2.1 General Overview

Psycholinguistic approaches to ellipsis comprehension can be divided into two categories: (i) those that attempt to test linguistic theories and predictions and (ii) those that attempt to explain the processes that underlie comprehension performance. Section 2.2 reviews literature from the first category. Specifically, Section 2.2.1 describes a series of experiments designed to test the theory that VP anaphors come in two varieties: surface anaphors that require a syntactically parallel antecedent and deep anaphors that can be interpreted based on information in the discourse (Hankamer & Sag, 1976). Hankamer and Sag (1976) originally proposed the distinction between the two types of anaphora; their proposal is described in Section 1.3.2 of the previous chapter. Sag and Hankamer (1984) later predicted that surface anaphors would be sensitive to syntactic manipulations made in the antecedent, but deep anaphors would not. The psycholinguistic experiments reviewed in this section test this prediction. Section 2.2.2 reviews a series of experiments designed to explore the availability of strict and sloppy identity readings during the on-line comprehension of elided pronouns. As discussed in the previous chapter, many linguistic theories predict that both interpretations can often be assigned to these types of sentences. The psycholinguistic evidence reviewed below attempts to address whether both readings are automatically available at the ellipsis site, or if one is preferred over the other.

Section 2.3 reviews three psycholinguistic models of ellipsis comprehension. These models attempt to explain the comprehension processes that underlie both parallel and non-parallel ellipsis by integrating linguistic theory with psycholinguistic evidence, to varying degrees. Section 2.3.1 describes a model that rejects the memory-search approach in favour of cue-based direct-access retrieval from memory (McElree, 2003; Martin & McElree, 2008; 2009; 2011). This model assumes that retrieval cues are provided by

linguistic variables (namely, syntax and semantics), leaving room for a full integration of linguistics and psycholinguistics. In Section 2.3.2, I describe a model that attempts to fully integrate linguistic theory and psycholinguistic theory. Kim et al. (2011) first develop the grammar that they believe feeds into the parser, then go on to explain how the parser uses heuristics inherent in the grammar to constrain its search through the possible structures held in memory. Finally, Section 2.3.3 outlines a model called the Recycling Hypothesis (Arregui et al., 2006) that predicts that ellipsis comprehension follows one of two procedures. Ellipsis with syntactically parallel antecedents is interpreted by copying the antecedent structure into the ellipsis site. On the other hand, non-parallel antecedents must undergo a reconstruction process during which the parser recycles the information available in the antecedent in order to build the correct ellipsis structure.

Section 2.4 summarizes the literature reviewed in this chapter and develops the research questions explored in this thesis. To anticipate: The experiments reported in Chapter 3 were designed to test the underlying assumptions of the Recycling Hypothesis, which are outlined in detail in Section 2.3.3. In Section 2.4, I review these assumptions, explain their limitations, and propose a method for testing them.

2.2 Testing "Parallelism"

2.2.1 Parallelism and Deep vs. Surface Anaphors

Early psycholinguistic investigations into VP ellipsis tested predictions made by Sag and Hankamer (1984). A detailed description of this proposal is presented in Section 1.3.2 of the previous chapter. To recap, Hankamer and Sag (1976) distinguished between two types of VP anaphors: surface anaphors (1a) that required a syntactically parallel antecedent, and deep anaphors (1b) whose antecedent could be pragmatically controlled.

- (1) a. Someone needed to create this example, so Tiffany did.
- b. Someone needed to create this example, so Tiffany did it.

Sag and Hankamer (1984) extended this proposal and suggested that surface anaphors are interpreted by copying the syntactic structure of their antecedents from short-term memory, and further noted that this syntactic representation has a short lifespan in the memory system. Given their dependence on syntactic structure, the interpretation of surface anaphors should be affected by syntactic factors. For example, they predict that the size of the antecedent should affect comprehension because copying larger antecedent structures should take more time and/or effort. They also predicted that the distance between the antecedent and the ellipsis site would affect the interpretation of surface anaphors. Increasing the distance between antecedent and anaphor by introducing intervening material should lower comprehension, assuming that new information interferes with/displaces old information in short-term memory. In contrast, deep anaphors should

not be affected by antecedent size or distance because their interpretation does not depend on syntactic structure. Instead, deep anaphors can receive their interpretation from discourse models, which are assumed to have a longer lifespan in short-term memory.

Murphy (1985) conducted three self-paced reading experiments whose results contradicted the predictions outlined above. First, he showed that the length of the antecedent affected the interpretation of both types of anaphors. He embedded sentence pairs like those in (2) into short paragraphs, and found that the sentences containing the anaphors took longer to read when they referred to a larger antecedent (2b), regardless of the type of anaphor.

- (2) a. Jimmy swept the floor. Later, his uncle did/did it too.
b. Jimmy swept the tile floor behind the chairs free of hair and cigarettes. Later, his uncle did/did it too.

He also showed that both types of anaphors were read more slowly when the anaphor and its antecedent were separated by an intervening sentence, and when the syntactic form of the antecedent did not match the intended form of the ellipsis site. Essentially, he showed that not only were surface anaphors affected by syntactic variables, but deep anaphors were as well, in contrast to Sag and Hankamer's (1984) predictions.

In a later study, Murphy (1990) examined the effect of distance between antecedent and anaphor in ellipsis comprehension. In a self-paced reading experiment, he found that sentences with long-distance antecedents took longer to read than those with near antecedents, regardless of the type of anaphor in the ellipsis site. In a separate experiment using a simple acceptability judgment task, he further found that sentences with near antecedents were rated more acceptable than far antecedents, independent of anaphor type, confirming the previous result. However, results showed that surface anaphors were judged more slowly with far antecedents, while deep anaphors were actually judged somewhat more quickly with far antecedents. Murphy suggested that the different patterns of results found in these two experiments were a function of the tasks involved. He argued that judgment processes are sensitive to grammatical rule violations and whether these violations are permitted in the context – thus, distant antecedents are equally acceptable for both deep and surface anaphors, but the parser takes longer to determine this for surface anaphors. On the other hand, reading time measures were not sensitive to the ability to violate grammatical rules, and thus showed the same pattern of results for both types of anaphors. Taken together, the results from Murphy (1985) and Murphy (1990) provide evidence against the predictions advanced by Sag and Hankamer (1984).

In contrast, Tanenhaus and Carlson (1990) provided strong evidence in support of Sag and Hankamer's (1984) hypotheses. Participants judged the acceptability of ellipsis sentences with and without parallel antecedents. In one experiment, they compared active and passive structures (3) as antecedents for active ellipsis sites. Importantly, active and passive sentences describe essentially the same event, but use different syntactic structure (read: word order). They predicted that surface anaphors would be sensitive to the

active-passive distinction because their interpretation depends on finding an antecedent with an identical syntactic structure. On the other hand, deep anaphors should not be affected by this distinction, because both of the events described will take the same form in the discourse model regardless of the syntactic structure used.

- (3) a. Someone had to take out the garbage, but Bill refused to / to do it.
 b. The garbage had to be taken out, but Bill refused to / to do it.

Indeed, the results showed that surface (but not deep) anaphors were more likely to be judged acceptable when the form of the antecedent matched the intended form of the ellipsis site. However, the speed with which participants judged the sentences showed parallelism effects for both surface and deep anaphors – participants were slower to judge sentences with non-parallel antecedents regardless of the type of anaphor in the ellipsis site. Similar results were found in an experiment comparing verb phrase antecedents to nominalised antecedents. Taken together, these results show that although parallelism does not affect the acceptability ratings of deep anaphors, it does affect the speed with which these ratings are made. The authors took this as evidence to support a distinction between deep and surface anaphors, but suggest that non-parallel antecedents slow the interpretation of deep anaphors because they encode discourse structures that require more effort to access.

Mauner, Tanenhaus, and Carlson (1995) reanalyzed the data from Tanenhaus and Carlson (1990) and found that the results described above were modulated by the form of the passive sentence structure. When the passive antecedent contained an agentive *by*-phrase (4a), the speed of judging the acceptability of both anaphors was slowed. When the passive antecedent did not have its *by*-phrase (4b), only the speed of judging surface anaphors was slowed.

- (4) a. The kitten needs to be fed by someone. Joey forgot to again.
 b. The kitten needs to be fed. Joey forgot to again.
 c. Someone needs to feed the kitten. Joey forgot to again.

Mauner et al. conducted an acceptability judgment experiment comparing the short passive antecedents (4b) to standard active antecedents (4c), and showed that non-parallel antecedents resulted in lower acceptability ratings and slower judgment speeds for surface (but not deep) anaphors. In a separate experiment comparing the long passive antecedents (4a) to standard active antecedents (4c), results showed that non-parallel antecedents resulted in lower acceptability ratings and lower judgment speeds for both types of anaphors. These results corroborate the results found in the reanalysis of the original data. The authors maintain that these results support the hypotheses put forth by Sag and Hankamer (1984), but suggest that these results should be explained by some (undefined) extra-syntactic factor.

Garnham and Oakhill (1987) provided one possible explanation for what extra-syntactic factor may be behind the variation in acceptability ratings observed in Mauner et al.

(1995). Garnham and Oakhill argued that memory for verbatim surface structure information is relatively low (despite evidence suggesting otherwise, Bock, 1986; and see Branigan, 2007, for a review). As such, they predicted that other factors might interact with syntactic parallelism to derive the correct interpretation of the ellipsis site. Specifically, they suggested that one of these factors might be pragmatic plausibility. They conducted a self-paced reading time experiment in which participants read short paragraphs like those in (5). One version of the paragraph described a highly plausible scenario (5a), while the other described a less plausible scenario (5b). Each paragraph ended with a yes/no question that the participant was expected to answer. The question was designed to measure which interpretation the reader assigned to the preceding elliptical sentence.

- (5) a. It had been a busy morning in the hospital. The elderly patient had been examined by the doctor (during the ward round). The child had too.
- i. Did the doctor examine the child?
 - ii. Did the child examine the elderly patient?
- b. It had been a busy morning in the hospital. The elderly patient had been examined by the doctor (during the ward round). The nurse had too.
- i. Did the doctor examine the nurse?
 - ii. Did the nurse examine the elderly patient?

In both plausible and implausible versions of the paragraphs, the correct answer required the participants to assign the syntactically parallel antecedent *had been examined by the doctor* to the ellipsis site – thus, the correct answer was that the doctor examined the child/nurse. However, it is more plausible for a doctor to examine a child than it is for a doctor to examine a nurse. The researchers predicted that participants would be more likely to choose the correct answer when it described a plausible scenario, as in (5a), compared to when it was implausible (5b). The results showed that participants answered the questions faster and more accurately when the correct answer was plausible, and they made greater use of this plausibility strategy when the antecedent was further from the ellipsis site. A similar plausibility effect was found in the reading time data. The authors took these results to suggest that readers prefer to find a syntactically parallel antecedent, and will spend extra time looking for one when it is implausible. However, readers also made use of plausibility cues in assigning interpretations to ellipsis sites.

To summarize, early experimental work provided mixed support for the distinction between deep and surface anaphors, as proposed by Hankamer and Sag (1976). Murphy (1985; 1990) showed that deep anaphors were affected by syntactic factors, contrary to the proposal that only surface anaphors would show this effect. Meanwhile, Tanenhaus and colleagues (1990; 1995) showed that only surface anaphors were sensitive to syntactically non-parallel antecedents, receiving lower acceptability ratings in these conditions. However, evidence showing that mismatching antecedents could also affect the speed with which deep anaphors were interpreted suggested that syntactic factors might interact with extra-linguistic information. Garnham and Oakhill (1987) provided evidence to

support this claim, showing that the likelihood of choosing an interpretation consistent with a parallel antecedent depended on the plausibility of the event being described. Thus, there is no unequivocal psycholinguistic evidence to support or refute Hankamer and Sag's (1976) distinction between deep and surface anaphors.

2.2.2 Parallelism and Strict vs. Sloppy Identity Readings

Several studies have provided evidence that the conjunction used to conjoin the antecedent clause and the ellipsis clause may play a critical role in the on-line comprehension of elided verbs (Callahan, Shapiro, & Love, 2010; Shapiro & Hestvik, 1995; Poirier, Wolfinger, Spellman, & Shapiro, 2010). Shapiro and Hestvik (1995) attempted to chart the time-course of comprehending ellipsis constructions containing reflexive pronouns, which can be interpreted in two ways. Using (6) as an example, the strict identity reading results from assigning the interpretation where the fireman defends the policeman. In contrast, the sloppy identity reading leads to the interpretation where the fireman defends himself. (A detailed discussion of the linguistic properties associated with strict and sloppy identity readings is presented in Sections 1.3.1.2 and 1.3.2, above.)

- (6) a. The policeman defended himself, and the fire [1] man did [2] too, according to someone [3] who was there.
 b. The policeman defended himself because the fire [1] man did, [2] according to someone [3] who was there.

To assess the time-course of ellipsis comprehension, the researchers employed the cross-modal lexical decision task. This task requires participants to listen to sentences while simultaneously making decisions about visually presented words. These words were presented at one of three points in the sentence (as indicated in (6)). The researchers predicted that if the strict identity reading was assigned automatically, then words related to the subject of the first clause (in this case, the policeman) should be recognized more quickly at the ellipsis site than unrelated words. If this effect – called ‘priming’ – occurred at the ellipsis site (but not before), then it would reflect the fact that information from the first clause was reactivated at the ellipsis site, suggesting that the participant had activated the strict identity reading. The results did show this effect, but only when the ellipsis and antecedent clauses were conjoined using the coordinate conjunction, *and*. When the clauses were conjoined using the subordinate conjunction, *because*, reactivation of the first clause information was not observed until after the ellipsis site. The authors suggested that subordinate conjunctions introduce additional semantic processing that delays reactivation of first clause information. Specifically, they noted that the parser must compute the cause-effect relation between the antecedent and ellipsis clauses when they are conjoined by *because*. Thus, they concluded that the strict identity reading is retrieved automatically at the ellipsis site when only structural information is involved and delayed when additional semantic factors emerge. Callahan et al. (2010) provided further support for this claim, showing that first-clause subject activation rapidly decayed

within the first clause, but was reactivated at the coordinate conjunction and remained active throughout the clause containing the ellipsis. They concluded that the activation at the conjunction site suggests that *and* creates an expectation of syntactic parallelism, and thus elicits the reactivation of information from the first clause.

Shapiro, Hestvik, Lesan, and Garcia (2003) tested the strength of this proposed parallelism expectation using verbs that were inherently reflexive. These verbs do not allow for a strict identity reading of the ellipsis clause because it is impossible to assign the reading where the pilot asserts the optometrist in (7).

- (7) a. The optometrist who had signed the release form asserted herself, and the pilot who needed to pass the training exam did too, according to the receptionist.

Therefore, if the lexical properties of the verb are consulted on-line during ellipsis comprehension, then there should be no prediction for parallelism, and thus no priming effect at the ellipsis site for words related to the first-clause subject. Surprisingly, they found activation of first clause material in the ellipsis site, suggesting that the strict identity reading was available. They took the results as evidence that syntactic information is consulted before lexical or probabilistic information.

Poirier, Walenski, and Shapiro (2012) found evidence to support this conclusion. They conducted a similar experiment using unaccusative (8a) and unergative (8b) verbs, which have contrasting argument structures, despite showing similar surface structures. Subjects of unaccusative verbs are thought to be generated in the object position before moving to their surface subject position, leaving a post-verbal trace. In contrast, subjects of unergative verbs are thought to be generated in their surface position and leave no such trace.

- (8) a. The dog disappeared in the crowded street fair, and the child with the blue jumpsuit on did too, much to the family's dismay.
 b. The musician winked at the cute bartender, and the tourist with a carefree smile did too, while people were entering the club.

The researchers predicted that first-clause information would be activated for both verb types, due to the parallelism expectation created by the coordinate conjunction. However, only unaccusative verbs leave a trace that must be bound by the ellipsis clause subject, therefore only unaccusative verbs should show activation of the ellipsis clause subject at the ellipsis site. These are exactly the results that the researchers found, providing further support for the hypothesis that the lexical properties of conjunctions can bias the parser toward a preference for parallelism, but the lexical properties of verbs cannot reverse this preference.

In summary, it appears that some types of lexical information can influence the real-time comprehension of ellipsis. For example, Shapiro and Hestvik (1995) showed that the

conjunction used to conjoin the ellipsis clause with the antecedent clause played a significant role, leading them to suggest that the strict identity readings of reflexive pronouns are automatically activated at the ellipsis site, but can be delayed by additional semantic processing. Callahan et al. (2010) confirmed this result, and suggested that the coordinate conjunction, *and*, biases the parser to an expectation of parallelism. Shapiro et al (2005) and Poirier et al. (2012) constrained this bias to strictly syntactic parallelism, showing that strict identity readings are available even for verbs that forbid such interpretations. Additionally, only verbs that selected for an (underlying) object produced first clause activation at the ellipsis site. Thus, it seems that only lexical information that provides an indication of what the upcoming syntactic structure might look like will influence ellipsis comprehension.

2.3 Dealing with Non-Parallelism

2.3.1 Cue-Based Direct-Access Memory Retrieval Hypothesis

One psycholinguistic theory of sentence processing predicts that comprehension is constrained by the same set of general cognitive processes that govern memory performance (Lewis & Vasishth, 2005; Lewis, Vasishth, & Van Dyke, 2006; McElree, 2000; McElree, Foraker, & Dyer, 2003). Early evidence that memory retrieval is critical to sentence comprehension came from studies using cross-modal lexical priming tasks (Nicol & Swinney, 1989).

- (9) a. The policeman saw the boy₁ who₁ the witness had accused *t*₁ at the bus stop.

While listening to sentences that contained embedded clauses (9), participants were faster to recognize a visually presented word associated with *boy* at the trace site immediately following the embedded verb *accused* than at the position immediately preceding the verb. These results indicate that memory representations are active at trace sites (but not before), suggesting that the trace site cues the retrieval of the item from memory.

Based partly on this result, McElree (2000) and McElree et al. (2003) suggested that syntactic and semantic information are responsible for providing retrieval cues to access the relevant memory representations. Furthermore, they argued that these cues provide direct access to the items in memory – contrary to Kim, Kibele, Runner, and Halle’s (2011) proposal that item retrieval involves a labourious search process, described below. This direct-access proposal is crucial, because these theorists believe that the capacity of the short-term memory store is very small, holding only one item at a time. This limited storage space makes it impossible to hold several items or even multiple instances of the same item in an active state, because incoming information will displace the old material. At the same time, McElree and colleagues predicted that search processes are effortful and time-consuming, and increasing the amount of information in the search domain should

increase the length of time it takes to retrieve the relevant memory representation. The absence of such an effect supports the hypothesis that items are directly accessed in memory. They used a speed-accuracy tradeoff (SAT) procedure to test this prediction. Unlike standard measures of processing speed (e.g., reading time), the SAT procedure allows researchers to measure the speed and accuracy of performance on a given task simultaneously, but examine the effects of each factor independently. By examining the effects of speed and accuracy independently, researchers can draw clear and precise conclusions about the factors involved in performing the task at hand. Specifically, the SAT technique derives a function that represents how accuracy changes over time. Variables that affect performance accuracy will be represented by changes in a horizontal asymptote, while variables that affect processing speed will be represented by changes in the function's intercept and/or slope (see Figure 2.1)

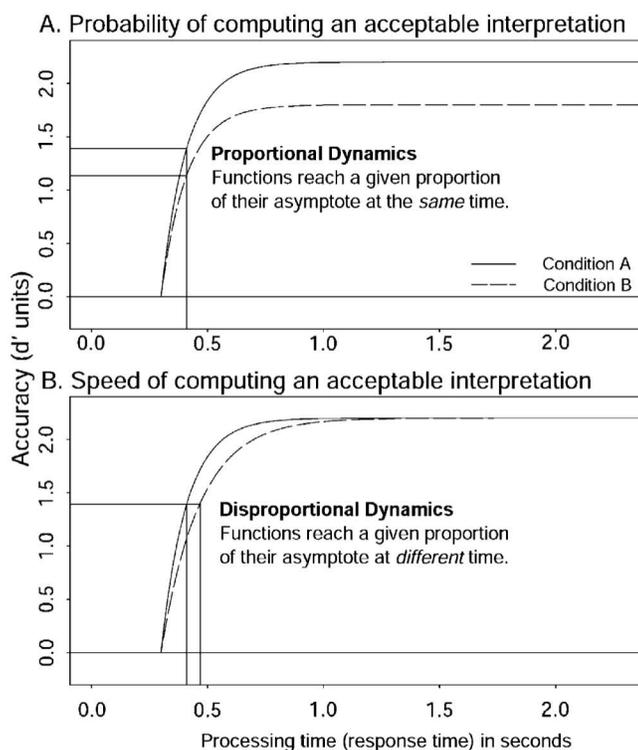


Fig. 1. Hypothetical SAT functions illustrating two conditions that differ by SAT asymptote only (A) or SAT rate (B). The intersection of the horizontal and vertical lines shows the point in time (abscissa) when the functions reach two-thirds of their respective asymptote (ordinate). When dynamics are proportional (A), the functions reach the two-thirds point at the same time.

Figure 2.1: Martin and McElree's (2008) hypothetical speed-accuracy tradeoff functions, which illustrate the independent effects of accuracy and speed of retrieval.

For the current example, if item retrieval in sentence comprehension involves an effortful search process, and introducing additional information into the search domain affects the speed of this search process, then the intercept and/or the slope of the function should change as a function of the amount of interfering information. Conversely, if sentence comprehension involves cue-dependent direct-access retrieval from memory, additional information will degrade the quality of the retrieval cues, rendering them less effective. Degraded retrieval cues should only interfere with the probability (read: accuracy) of retrieving an item from memory; they should not affect the speed of the retrieval process. Thus, this interference effect will manifest itself as differences in the functions asymptote. Low probability of retrieval can reflect either of two possible effects of the degraded retrieval cues. First, it may reflect the fact that the intervening material caused so much interference that the memory representation is not accessed at all. Alternatively, it may reflect the fact that the degraded retrieval cue initially accessed the wrong memory representation, forcing reanalysis processes to take place.

McElree (2000) initially tested the predictions outlined above by introducing relative clauses of various lengths between a noun phrase and the trace site that it must bind. In particular, he presented cleft sentences (10) at a controlled rate of 250 ms per word, and asked participants to make an acceptability judgment in response to a tone that sounded at one of six time points after the onset of the final word (50, 300, 500, 800, 1200, or 3000 ms).

- (10) a. This was the book that the editor admired/*amused.
 b. This was the book that the editor who the receptionist married admired/*amused.
 c. This was the book that the editor who the receptionist who quit married admired /*amused.

He then plotted the rate of acceptable responses as a function of response latency, which produced an SAT function. If the memory representation of *the book* is retrieved at the trace site using a search mechanism, the speed of this search process should be slower when there is more information between the noun phrase and final verb, as represented by changes in the SAT intercept/slope. However, the results showed no difference at the intercept or slope. Instead, the results showed a difference at the asymptote, indicating that the intervening information lowered the probability of retrieving the memory representation, but did not affect the amount of time it took to make the retrieval. This suggests that sentence comprehension of filler-gap dependencies such as those in cleft sentences is guided by a cue-based direct-access mechanism, not a time-consuming search process.

McElree et al. (2003) replicated these results showing that the complexity of the intervening embedded clause consistently affected the accuracy of memory retrieval, but not the speed of retrieval, regardless of the nature of the complexity. In addition, they extended the results to show that it takes longer to fill two gaps in double object constructions like (11a,b) compared to single gaps in intransitive constructions (11c,d) overall, but that the amount of intervening material still only affected the probability of retrieving

the items from memory.

- (11) a. This is the album₁ that the stamps₂ were difficult to mount t_2 in t_1 .
 b. This is the album₁ that the stamps₂ which obviously angered the fussy collector were difficult to mount t_2 in t_1 .
 c. This is the album₁ that the collector found difficult to spread open t_1 .
 d. This is the album₁ that the customer who obviously angered the fussy collector found difficult to spread open t_1 .

The authors argued that resolving double dependencies takes longer than single dependencies because order information becomes important in these cases. It had been shown previously that serial order information is retrieved from memory using a search mechanism (McElree & Doshier, 1993), so the authors took these results as evidence to support their hypothesis. When required to fill two gaps as in (11a,b), the parser must perform a search operation to ensure the arguments are retrieved in the correct positions, but the intervening material does not affect the search process, only the success of retrieving the correct items.

Most importantly for the topic discussed in this thesis, Martin and McElree (2008; 2009; 2011) provided evidence that ellipsis comprehension can be explained under the same model. Extrapolating from the predictions described above, if accessing antecedents for ellipsis sites involves a search process, then the amount of information between the antecedent and ellipsis site should slow the retrieval speed. On the other hand, if only accuracy of retrieval is affected, this suggests that the antecedent is directly accessed using the cues provided by the ellipsis site. Indeed, Martin and McElree (2008) showed that the distance between the antecedent and the ellipsis site affects only the asymptote of the SAT function. Furthermore, they found no evidence that more complex antecedents affect the speed or accuracy of antecedent access, suggesting that ellipsis comprehension does not involve a copy mechanism (contra Frazier & Clifton, 2001).

Martin and McElree (2009) reasoned that search processes can, in principle, be performed by moving backward through the most recent material to the least recent material, or by moving forward starting with the least recent material. They recognized that the stimuli used in their previous experiment only targeted backward search processes, so their previous results may have reflected the fact that their stimuli were not sensitive to the type of search process that is necessary for comprehension. Thus, they replicated their previous ellipsis experiment using materials with additional information before the antecedent (12a) compared to the standard materials with additional information after the antecedent (12b).

- (12) a. Sometime in the early morning yesterday, Claudia filed a complaint. Ron did too.
 b. Claudia filed a complaint sometime in the early morning yesterday. Ron did too.

This manipulation essentially creates a difference in the locality of the antecedent relative to the ellipsis site. The results showed that only the probability of retrieving the antecedent was affected by this locality manipulation. Since there was no effect on the speed of retrieval, the authors could safely reject the hypothesis that forward search processes are used in sentence comprehension. Furthermore, results showed that the probability of retrieving the antecedent was lower when the additional information was presented between the antecedent and ellipsis site. This provides further evidence to support the hypothesis that antecedent access is guided by cue-based direct-access retrieval mechanisms: When the additional material intervenes between the antecedent and ellipsis site, this material interferes with the quality of the memory representation for the antecedent, which lowers the effectiveness of the retrieval cue provided at the ellipsis site.

To recap, McElree (2000) showed that trace sites in cleft sentences were less likely to be resolved when there was intervening material between the cleft and the trace, but the time it took to resolve the dependency stayed constant. This suggested that long-distance dependencies were recovered by retrieving the relevant material from memory using direct-access retrieval cues, not a labourious search process. Later, Martin and McElree (2008; 2009; 2011) showed similar results using elided sentences when they introduced intervening material between the antecedent and ellipsis site. This suggests that ellipsis comprehension may be governed by direct-access of an antecedent in memory, based on retrieval cues provided in the syntax and/or semantics of the sentence.

2.3.2 Structure Search Hypothesis

In contrast to the previous proposal, Kim et al. (2011) advanced a model of ellipsis comprehension that argues that antecedents are accessed from memory through a search procedure that is guided by linguistically motivated heuristics. Assuming a grammar that contains the operations of Merge, Move, and Delete combined with the linguistic principles MaxElide and Canonical Representation, Kim et al. claimed that they can account for the variability in acceptability of ellipsis constructions with and without parallel antecedents. In particular, they claimed that Delete is licensed only when a syntactically identical antecedent is available in the finished derivation. That is, they favour a strict syntactic identity condition on ellipsis licensing. On this account, the parser interprets ellipsis constructions by searching through all possible structures to find the one that contains an antecedent that matches the syntactic structure intended in the ellipsis site. To account for cases of acceptable ellipsis constructions with non-parallel antecedents, they suggested that this search process is constrained by the principles of MaxElide and Canonical Representation. MaxElide is a principle that suggests that the largest constituent must be deleted, whenever possible. The parser uses this principle as a heuristic, which constrains the search space to only those syntactic structures that delete the largest possible constituent. Canonical Representation is a principle that suggests there is an overall preference for sentences that follow the canonical word order. In English, this word order is subject-verb-object (SVO). Sentences that violate this canonical word order (e.g. passive constructions) are still grammatical, but dispreferred.

Thus, Canonical Representation is a heuristic that will guide the parser toward the possible structures that obey the canonical word order before addressing those that violate it. Together, these heuristics constrain the search space to only the relevant preferred structures. From here, the parser searches through the structures to choose the one with the best possible antecedent. It is through this process that Kim et al. can explain the graded acceptability cline found in mismatch ellipsis cases.

Kim et al. (2011) presented a series of experiments designed to test the various components of their model. Their first experiment used cases of voice mismatch ellipsis (13) as an initial test of the parsing heuristics: MaxElide and Canonical Representation.

- (13)
- a. Jill betrayed Abby, and Matt did ~~betray Abby~~ too.
 - b. Jill betrayed Abby, and Matt was ~~betrayed by Jill~~ too.
 - c. Abby was betrayed by Jill, and Matt was ~~betrayed by Jill~~ too.
 - d. Abby was betrayed by Jill and Matt did ~~betray Abby~~ too.

They claimed that MaxElide prefers parallel ellipsis because parallel antecedents permit the deletion of the largest possible constituent. Thus, they predicted that ellipsis with parallel antecedents will show higher ratings of acceptability than non-parallel antecedents. They further predicted that this match effect will occur only in ellipsis constructions because MaxElide only applies to cases of ellipsis. To be clear, they predicted that non-elided sentences should not show the same preference for parallelism as elided sentences. On the other hand, Canonical Representation prefers active sentence constructions compared to passive because active sentences conform to the canonical SVO word order. Furthermore, Canonical Representation is not an ellipsis-specific principle, so the preference for active sentences should be observed in both elided and non-elided sentences. Results from this experiment showed that parallel clauses produced higher acceptability ratings, but only when the second clause contained an ellipsis, confirming the predictions based on MaxElide. Results further showed that sentences with active antecedents were rated more acceptable than sentences with passive antecedents, regardless of whether the second clause contained an ellipsis, thus confirming the predictions based on Canonical Representation. Two other experiments using adjectivizations and nominalizations as antecedents showed similar results, lending further support to the hypothesis that ellipsis comprehension involves a search through memory for the correct syntactically parallel antecedent.

To summarize, Kim et al. (2011) suggested that ellipsis comprehension is governed by a search process that is guided by the grammar. Specifically, the parser uses built-in grammatical principles – MaxElide and Canonical Representation – to constrain the set of possible structures to those that permit deletion of the largest possible constituent and (preferentially) conform to canonical word order. Non-parallel ellipsis constructions violate these preferences, which forces the search space to become larger. As the search space increases, retrieval of the correct structure becomes more difficult, which results in lower acceptability due to more difficult comprehension processes.

2.3.3 Recycling Hypothesis

In this section, I outline the Recycling Hypothesis of verb-phrase ellipsis (Arregui et al., 2006) in detail. First, I give a general description of the hypothesis. Second, I provide a summary of the underlying assumptions, along with the data that motivate those assumptions. Third, I describe the Recycling Hypothesis as well as present the experimental evidence supporting this theory.

The Recycling Hypothesis of verb-phrase ellipsis is a model that aims to explain how humans process ellipsis in the context of a non-parallel antecedent. Specifically, the model suggests that the parser will use the information available in the non-parallel antecedent to build the correct syntactic structure for the ellipsis site. However, this model makes critical assumptions about how we process ellipsis with parallel antecedents. The most fundamental assumption that this model makes is that the ellipsis site hosts unpronounced syntactic structure. Frazier and Clifton (2005) provided evidence that this is the case (see also Section 1.2). They embedded antecedents within a conjoined verb phrase, and showed that when the antecedent was in the nearby conjunct, ellipsis sentences were rated more acceptable and read more quickly than if the antecedent was in the far conjunct. They also reported that sluiced sentences were processed differently depending on if the *wh*-remnants targeted arguments or adjuncts. Both of these results are unexpected under the hypothesis that there is no syntactic structure in the ellipsis site.

The second most important assumption underlying the Recycling Hypothesis regards how this unpronounced syntactic structure is created. Does the parser engage in costly syntactic construction processes, or is there a potentially cost-effective copy mechanism that can be used? Frazier and Clifton (2001) suggested the latter, introducing the mechanism that they call Copy α . They presented experimental evidence that structural complexity did not affect the length of time it took readers to process ellipsis sentences. More convincingly, they referred to results from Frazier and Clifton (2000), where it was reported that there was no difference in reading times for the second sentence (*Tina did too*) in the mini discourses in (14), despite the fact that the antecedent in (14b) is much larger than the antecedent in (14a).

- (14) a. Sarah left her boyfriend last May. Tina did ~~leave her boyfriend last May~~ too.
 b. Sarah got up the courage to leave her boyfriend last May. Tina did ~~get up the courage to leave her boyfriend last May~~ too.

The authors reasoned that if the parser engaged in reconstruction processes at the ellipsis site, it should take more effort to reconstruct a larger, more complex antecedent, which should result in longer reading times. As evidenced, this was not the case, and Frazier and Clifton took these results to favour the existence of an essentially cost-free copy mechanism. This copy mechanism is assumed by Arregui et al. (2006), where the same group of researchers seek to explain how the parser comprehends ellipsis constructions when faced with non-parallel antecedents.

A third assumption underlying the Recycling Hypothesis is one regarding the role of extra-syntactic information in facilitating (or harming) processing performance. Specifically, the Recycling Hypothesis (implicitly) assumes that discourse processing plays a crucial role in ellipsis comprehension. Frazier and Clifton (2005; 2006; and others) directly investigated the role of discourse structure in ellipsis comprehension. Frazier and Clifton (2005) proposed the Main Assertion Hypothesis, which states that information in a new sentence is preferentially related to information from the main clause of the preceding sentence. In terms of ellipsis comprehension, this predicts that ellipsis that spans across sentences (15a) will preferentially choose the main verb phrase of the preceding clause as its antecedent, even when the embedded verb phrase offers a plausible antecedent. In contrast, ellipsis that has an antecedent within the same sentence (15b) will preferentially choose the embedded verb phrase as its antecedent.

- (15) a. John said that Fred went to Europe. Mary did too.
b. John said that Fred went to Europe and Mary did too.

The researchers showed that ellipsis constructions like those in (15a) received significantly more main verb interpretations than embedded verb phrase interpretations, while sentences like those in (15b) showed the opposite pattern. The authors took these results as evidence to support the Main Assertion Hypothesis, as well as the role of discourse information in ellipsis comprehension. In particular, they argued that sentence boundaries such as those in (15a) close the syntactic structure of the first sentence, which makes it unavailable for building syntactic relations with the new sentence. This is where the discourse processor comes in – it will look back to the information in the previous sentence to find the appropriate antecedent for the ellipsis site. There is no longer any way to read the detailed syntactic structure of the sentence, so the discourse processor instead chooses the most salient piece of information in the discourse: the main clause of the sentence. Notice that this explanation indirectly introduces the idea of constructing the correct syntactic structure at the ellipsis site. Indeed, this may be viewed as a precursor to the Recycling Hypothesis.

The Recycling Hypothesis (Arregui et al., 2006) proposed that the preferred method of comprehending elided sentences is to copy the structure of a parallel antecedent from a previous clause into the ellipsis site. In the absence of a parallel antecedent, the parser is thought to use the available information to build the correct structure for the ellipsis site. To do this, they suggested that the parser will recycle the material provided in the non-parallel antecedent, and perform the syntactic operations required to build the ellipsis structure that yields the correct interpretation. Thus, the Recycling Hypothesis predicts that the acceptability of elided sentences with non-parallel antecedents will depend on three factors: (i) what material is available in the antecedent clause, (ii) the operations that must be performed to build the correct structure, and (iii) the availability of additional information to guide the parser to the correct interpretation.

To test the predictions made by the Recycling Hypothesis, Arregui et al. (2006) asked participants to rate the acceptability of sentences like those in (16), predicting that more

complicated antecedents will result in a higher number of operations required to build the ellipsis structure. Compared to (16a), the sentences in (16b-d) represent elided sentences with increasingly complex antecedents. According to Arregui et al., the embedded verb phrase in (16b), *seeing the comet*, requires the parser to search for the antecedent in an unexpected position. In (16c), the antecedent verb is in the expected position, but it also holds the trace of the object that has been topicalized – thus, the parser must seek out the verb’s object and then combine it with the verb to build the ellipsis structure. The sentence in (16d) becomes even more complex because the relevant verb is embedded within an adjective, *unseeable*. Here, the parser must first break apart the adjective to retrieve the verb, then seek out its object, and then combine the verb with the object to build the ellipsis structure.

- (16) a. None of the astronomers saw the comet, but John did.
 b. Seeing the comet was nearly impossible, but John did.
 c. The comet was nearly impossible to see, but John did.
 d. The comet was nearly unseeable, but John did.

The Recycling Hypothesis predicts that these sentences should decrease in acceptability in a step-wise fashion from (16a) to (16d). This is in fact what the researchers found. Sentences like those in (16a) were most likely to be rated acceptable, while those in (16d) were least likely to be judged acceptable, with (16b) and (16c) taking distinct positions in between. The authors took this result as evidence that the acceptability of verb phrase ellipsis depends on rebuilding the structure underlying the ellipsis site.

In a similar experiment, Arregui et al. (2006) asked participants to rate the acceptability of sentences containing a verbal gerund antecedent (17a,b) and those containing a nominal gerund antecedent (17c,d). Verbal gerunds are distinct from nominal gerunds in the following way: Verbal gerunds take the structure of standard verb phrases, but take on the distributional properties of noun phrases. Nominal gerunds take the structure of standard noun phrases, but are derived from verb phrases and so possess the properties of verbal argument structure.

- (17) a. Singing the arias tomorrow night will be difficult, but Mary will.
 b. Singing the arias slowly tomorrow night will be difficult, but Mary will.
 c. Tomorrow night’s singing of the arias will be difficult, but Mary will.
 d. Tomorrow night’s slow singing of the arias will be difficult, but Mary will.

The Recycling Hypothesis predicts that (17a,b) will be judged as more acceptable than (17c,d) because there are additional syntactic operations required for building the ellipsis site in the presence of nominal gerund antecedents because they do not already have the structure of standard verb phrases, the way verbal gerunds do. Furthermore, the Recycling Hypothesis also predicts that (17b) will be more acceptable than (17a) because the adverb, *slowly*, will guide the parser toward the correct ellipsis structure (i.e., a verb phrase); in contrast, (17d) should be worse than (17c) because the adjective, *slow*, should

instead guide the parser away from the correct ellipsis structure (i.e., toward a noun phrase). The results showed that verbal gerunds made better antecedents than nominal gerunds, but the modifiers in (17b,d) did not modulate the acceptability. Like the results from the previous experiment, these results are in line with the prediction that the complexity of the syntactic operations that must be performed will affect acceptability of ellipsis sentences. However, these results failed to support the prediction that additional information will help guide the parser toward the correct structure. Nevertheless, the authors took these results to support the Recycling Hypothesis.

In a final experiment, Arregui et al. (2006) explored the role of the speaker in creating ellipsis constructions that are acceptable to the listener. They suggested that mismatch ellipsis cases occur when there is more than one way of expressing a proposition. Specifically, mismatch ellipsis results from the speaker misremembering the original paraphrase. They put forward the Systematic Paraphrase Hypothesis, which claims that when there is more than one possible paraphrase, the speaker is more likely to (mis)remember the simpler version. Thus, mismatch ellipsis constructions should be more acceptable when the ellipsis structure is simpler than the preceding antecedent. Active sentence structure is simpler than passive structure, so this hypothesis predicts that sentences with a passive antecedent followed by an active ellipsis site (18a,b) should be more acceptable than sentences with an active antecedent followed by the more complex passive ellipsis (18c,d). The researchers also predicted that ellipsis constructions will be more acceptable when they contain a presupposition trigger (e.g., *already*, *too*) as in (18a,c). They suggested that the presence of this presupposition trigger indicates to the listener that the speaker intended for there to be an appropriate antecedent for the ellipsis site. To test these hypotheses, the researchers asked their participants to rate the acceptability of the types of sentences presented in (18), as well as versions that used *too* as the sentence-final presupposition trigger.

- (18)
- a. The dessert was praised by the customer after the critic did already.
 - b. The dessert was praised by the customer and the critic did.
 - c. The customer praised the dessert after the appetizer was already.
 - d. The customer praised the dessert and the appetizer was.

The relevant results showed that sentences with the simpler ellipsis structure (18a,b) were more acceptable than those with the more complex ellipsis structure (18c,d), which the authors took as evidence to support the Systematic Paraphrase Hypothesis. However, it could be the case that active structures are simply easier for the listeners to rebuild (assuming the Recycling Hypothesis is correct). The results also showed that sentences that contained a presupposition trigger (18a,c) were more acceptable than those without (18b,d), which the authors claimed support their hypothesis about the speaker's intention in using these types of lexical items. However, there may be other ways to explain these results. Indeed, what is easy for the speaker to produce is not always easy for the listener to comprehend (see Keysar, 2007, for a discussion). Furthermore, the presence of a presupposition trigger can sometimes make mismatch ellipsis constructions worse.

Taking Arregui et al.'s own stimuli from their first experiment as an example, we can see that the presence of the presupposition trigger in (19b,c) actually causes the generally acceptable (19a) to become (at least marginally) ungrammatical.

- (19)
- a. None of the astronomers saw the comet, but John did.
 - b. *None of the astronomers saw the comet, but John did too.
 - c. *None of the astronomers saw the comet, and John did too.
 - d. None of the astronomers saw the comet, but John did already.
 - e. None of the astronomers saw the comet, and John did already.

To summarize, Arregui et al. (2006) argued that parallel ellipsis comprehension is relatively easy because it simply involves copying the antecedent structure into the ellipsis site. However, comprehending ellipsis in the context of non-parallel antecedents is not so simple. In this case, the parser must refer to the information in the antecedent, and manipulate it in such a way that it fits into the ellipsis site. This suggests that non-parallel ellipsis constructions will be more or less acceptable depending on (i) what information the parser can access in the antecedent, (ii) what operations must be performed on that information to build the correct ellipsis structure, and (iii) the usefulness of any additional information in guiding the parser through these steps. To support their hypothesis, they offered evidence from several experimental questionnaires. First, they found that antecedents that required more complex operations to fit into the ellipsis site were less acceptable than antecedents that required less complex operations. Second, they replicated their first experiment showing that verbal gerunds made better antecedents for VP ellipsis than nominal gerunds, but found no evidence that additional information affected acceptability judgments (contrary to their predictions). Finally, they reported that ellipsis constructions that differed from their antecedents only by voice (e.g., active vs. passive) were more acceptable when the ellipsis site contained the easier active structure relative to the antecedent's more difficult passive structure. There are some key issues with this hypothesis and the data that have been used to support it. The experiments reported in this thesis were designed to address these concerns. The concerns will be discussed in detail in the next section, and the experiments will be described in Chapter 3.

2.4 What is the problem? (Summary)

The series of experiments reported in Chapter 3 were designed to test the predictions made by the Recycling Hypothesis (Arregui et al., 2006), described in the previous section. To recap, the Recycling Hypothesis assumes that elided constituents are interpreted by copying the syntactic structure of an identical antecedent. In the absence of an identical antecedent, the parser uses the information available in the provided antecedent to rebuild the correct structure for the ellipsis site. As such, the Recycling Hypothesis predicts that the acceptability of ellipsis constructions with non-parallel antecedents will be

degraded to varying degrees depending on three factors: (i) what information is available in the antecedent, (ii) the type of syntactic operations that must be performed on that material in order to derive the correct underlying structure for the ellipsis site, and (iii) what additional information is available in the context to guide the parser toward the correct interpretation. Arregui et al. attempted to provide support for these predictions, but – in my view – fell short, for a variety of reasons. An initial concern is the appropriateness of the stimuli they used in their experiments. The stimuli they used in their experiments (20) were built on the assumption that every syntactic operation carries the same cognitive cost and that these costs are additive.

- (20) a. None of the astronomers saw the comet, but John did.
 b. Seeing the comet was nearly impossible, but John did.
 c. The comet was nearly impossible to see, but John did.
 d. The comet was nearly unseeable, but John did.

Furthermore, their stimuli are confounded with the assumption that differences in verbal morphology carry no cognitive cost at all. The researchers provide no evidence to support this assumption, but simply state that “the class of really easy recycling operations seems to be those that do not involve structural alterations at all, for example, changing the features on a verb phrase from +tense to –tense. (p. 242)”

Secondly, their results are not consistent with all of their predictions. For example, when Arregui et al. (2006) introduced additional information into the context of their stimuli with the intent of guiding the parser toward a specific interpretation, they found no effect on acceptability ratings – a result that completely contradicts their third prediction. Finally, as discussed earlier, the results from their final experiment do not so clearly support their conclusions about the role of the speaker’s intention with ellipsis constructions. Presupposition triggers do not always make sentences easier to comprehend. Furthermore, what may be easier for the speaker to produce is not necessarily easy for the listener to comprehend.

The experiments described in this thesis were crafted primarily to address the first concern: the stimuli that Arregui et al. (2006) used in their experiments were inadequate. As an initial investigation into this claim, I used stimuli like those in (21).

- (21) a. None of the astronomers would see the comet, but John did / would.
 b. None of the astronomers would get to see the comet, but John did / would.
 c. None of the astronomers would be seeing the comet, but John did / would.

In the antecedent clause, only the morphology on the verb is manipulated. If changes in verbal morphology carry no cognitive costs, as Arregui et al. assume, then there should be no difference in acceptability across the three antecedent types. Intuitively, the sentences in (21) vary in their acceptability, with (21a) being fairly acceptable, and (21c) being wildly unacceptable. Thus, I predict that acceptability for these sentences

will be modulated by the morphology on the antecedent verb phrase. If this prediction is borne out, it will indicate that Arregui et al.'s assumption was incorrect, and that the changes in verbal morphology seen across their stimuli may have played a confounding role in their results.

Another shortcoming in the stimuli used by Arregui et al. (2006) is that they did not provide appropriate baseline measures – that is, they did not compare the non-parallel antecedents in (20) to ellipsis constructions in which the same antecedents were in fact parallel antecedents. Thus, I manipulated the modal verb located in the ellipsis site to allow for such comparisons. When the modal verb *would* was present in the ellipsis site, it is possible for the parser to simply copy and paste the entire antecedent – verbal morphology and all – into the ellipsis site. However, when the modal verb *did* was located in the ellipsis site, it was necessary for the parser to reconstruct the verbal morphology to fit the simple past tense form on the modal. As Arregui et al. note, “the processor attempts to create a matching antecedent only when it cannot find an already matching one (p. 242).” If this is the case, then I predict that ellipsis with *would* will be uniformly acceptable because a matching antecedent is available, and the content of the antecedent should not affect the copy mechanism used for matching antecedents (Frazier & Clifton, 2001). If, instead, these matching antecedent clauses show the same pattern of results as the mismatching antecedent cases, this could mean one of two things: (i) Arregui et al.'s assumption that the processor only rebuilds the ellipsis site when there is no matching antecedent present is incorrect, or (ii) Frazier and Clifton's (2001) Copy α mechanism does not exist.

Chapter 3

Does the Recycling Hypothesis Work?

3.1 Review of Theory and Problem

This chapter describes four experiments that were designed to test the predictions made by the Recycling Hypothesis (Arregui et al., 2006). The Recycling Hypothesis is described in detail in Section 2.3.3 of the previous chapter. To review, the Recycling Hypothesis aims to explain how ellipsis constructions with non-parallel antecedents are comprehended. The hypothesis assumes that ellipses with parallel antecedents are comprehended by copying the syntactic structure of the antecedent into the ellipsis site. It further assumes that this procedure is used whenever it is available; that is, it is the parser's preferred treatment of ellipsis constructions. When a parallel antecedent is not available, the Recycling Hypothesis claims that the parser uses the information available in the antecedent to rebuild the correct ellipsis structure. Specifically, the parser recycles the material in the antecedent and performs the necessary syntactic operations on the antecedent structure to construct the correct interpretation of the ellipsis site. Under this hypothesis, the model makes the following empirical predictions:

- (1) Ellipsis with non-parallel antecedents will vary in acceptability depending on:
 - a. what information is available in the antecedent
 - b. the number and type of syntactic operations that must be performed on the antecedent material to build the correct ellipsis site
 - c. what additional information is available in the context to guide the parser toward the correct interpretation

Arregui et al. (2006) used stimuli like those in (2) to test these predictions. Compared to (2a), the sentences in (2b-d) contain antecedents of increasing complexity as determined by the complexity of the operations that are assumed to be necessary to rebuild the

ellipsis site.

- (2) a. None of the astronomers saw the comet, but John did.
- b. Seeing the comet was nearly impossible, but John did.
- c. The comet was nearly impossible to see, but John did.
- d. The comet was nearly unseeable, but John did.

They reported a grammaticality judgment experiment that showed that the sentences in (2) decreased in acceptability as the complexity of the antecedent increased. From this, they concluded that the predictions in (1a,b) were borne out.

At first blush, this result may seem fairly clear and straightforward. However, there are a number of methodological issues that must be addressed, and these are the focus of the experiments presented below. First, the sentences in (2) do not form a minimal set – each sentence is substantially different from any other sentence in the set. Thus, it is not possible to make any clear conclusions about the stimuli. Secondly, there are a number of confounding factors that the researchers fail to take into account. Specifically, they assumed that each syntactic operation comes with its own cognitive cost, and that these cognitive costs are additive in nature. This is not necessarily the case, as models of parallel sentence processing are quick to point out (Marslen-Wilson, 1975; Rumelhart & McClelland, 1987). Furthermore, they assumed that certain syntactic operations carry no cognitive costs at all; namely, they did not consider the effect that differences in verbal morphology may have on ellipsis comprehension. Third, the authors did not provide appropriate baseline comparisons; they only ever measured the acceptability of the antecedents in (2) when they are in non-parallel contexts, never in parallel contexts. Therefore, they did not explicitly provide evidence that these antecedents can actually be simply copied into the ellipsis site, as they assumed. The three survey experiments reported in Sections 3.2-3.4 attempt to address these concerns.

Finally, Arregui et al. (2006) provide no traditional on-line measure of processing difficulty, such as word-by-word reading times. Word-by-word reading time measures have been said to be correlated with the amount of effort involved in integrating information into a sentence (e.g., Gibson, 1998). Frazier and Clifton (2001) argued that parallel antecedents are copied into the ellipsis site via a cost-free copying mechanism. In contrast, Arregui et al. proposed that nonparallel antecedents must undergo a costly reconstruction operation. In the case of parallel antecedents, the complexity of the antecedent structure should not affect the copy procedure; indeed, Frazier and Clifton showed that complexity did not affect reading times at the ellipsis site. On the other hand, antecedent complexity should play a critical role in the reconstruction of nonparallel antecedents. If the Recycling Hypothesis is to be believed, one should find that reading times at the ellipsis site should depend on the complexity of the nonparallel antecedents. Section 3.5 reports a self-paced reading experiment that sought to address these predictions.

3.2 Experiment 1A: Ellipsis in Isolation

3.2.1 Motivation

The three survey experiments reported here were designed to test the assumption that the syntactic operations responsible for altering verbal morphology between the antecedent and ellipsis site carry no cognitive costs. Arregui et al. (2006) assumed that this is the case, but provided no control studies to show that this assumption is valid. Furthermore, their stimuli (2) were created in such a way that verbal morphology was a confounding factor in interpreting their results. Each antecedent in their comparison set contained a different form of verbal morphology, in addition to the manipulations that the researchers were specifically investigating. To address the issue of morphological complexity, the experiments reported here held the structural properties of the antecedent constant, but manipulated the verbal morphology across sentences. Examples of the stimuli used in these experiments can be found in (3), below. If differences in verbal morphology do not carry different cognitive costs, then there should be no difference in acceptability ratings due to antecedent type. If results indicate that acceptability ratings depend on antecedent type, then we must reject the assumption that the syntactic operations responsible for altering verbal morphology carry no cognitive costs.

Another aim of the experiments reported below was to compare the same antecedents under parallel and non-parallel conditions. Frazier and Clifton (2001) suggested that parallel antecedents are interpreted via a copy mechanism, such that the structure of the antecedent is copied into the ellipsis site. They showed that the size of the antecedent (i.e., the material within the antecedent) did not affect the length of time it takes to perform this copy operation. Arregui et al. (2006) took these results for granted, and assumed that this operation is performed whenever a parallel antecedent is available. However, they did not show evidence that their highly complex antecedents can be targeted by the copy mechanism. That is, they did not compare their non-parallel antecedents to conditions under which the same antecedents served as parallel antecedents. Therefore, it is unclear whether their results reflect an ellipsis-specific non-parallelism effect or simply a generalized morphosyntactic complexity effect. The stimuli used in the experiments reported in this thesis attempted to correct for this by manipulating the modal verb that represents the ellipsis site. In all antecedents, the modal verb would precede the verb phrase. Thus, when the same modal verb was used in the ellipsis site, the parallel antecedent reading was available. When the modal verb did not appear in the ellipsis site, this forced participants to rebuild the correct simple past morphology for the elided structure. If a copy mechanism is used whenever a parallel antecedent is available (Arregui et al., 2006, p. 242), and the material to be copied does not affect the copy procedure (Frazier and Clifton, 2001), then the parallel antecedent conditions should be rated as more acceptable and should show no effect of morphology on the acceptability ratings. If the parallel antecedent conditions are not rated more highly than the non-parallel conditions, then we must conclude that the parallelism manipulation was ineffective as the preference for parallel antecedents is widely documented in the

literature (e.g., Arregui et al., 2006; Mauner et al., 1995; Tanenhaus & Carlson, 1990; Dickey & Bungler, 2011). However, if our parallel antecedent conditions are susceptible to the morphology manipulation, then we must reconsider the hypothesis that parallel antecedents are copied into the ellipsis site. At the very least, we must conclude that the copy mechanism is sensitive to complexity effects. More strongly, we might conclude that the copy mechanism is not used at all.

3.2.2 Methodology

3.2.2.1 Participants

Participants were recruited from the undergraduate participant pool in the Department of Linguistics at McMaster University. In total, 115 students completed the survey. All 115 respondents were compensated with partial course credit. Forty-three students were non-native English speakers and were excluded from the analysis. Data from the remaining 72 native English speakers are presented here. There were 66 females and eight males. Participants ranged from 17 to 31 years in age, with an average of 19.75 years.

3.2.2.2 Stimuli

Test sentences were adapted from those used by Arregui et al. (2006). A sample set can be found in (3). Items from all 16 sets appear in Appendix A. Each set contained six sentences, created by crossing three antecedent types with two ellipsis types. Antecedents could be active (*would see the comet*), infinitive (*would get to see the comet*), or progressive (*would be seeing the comet*). Whether the antecedent was parallel or non-parallel was determined by the modal preceding the ellipsis site: *would* indicated parallel antecedents, while *did* indicated non-parallel antecedents. Six presentation lists were created with 16 test sentences (one from each set) and 32 filler sentences. All filler sentences were well-formed and contained no instances of ellipsis.

- (3)
- a. None of the astronomers would see the comet, but John did/would.
 - b. None of the astronomers would get to see the comet, but John did/would.
 - c. None of the astronomers would be seeing the comet, but John did/would.

3.2.2.3 Procedure

Participants completed the experiment via a survey hosted on the SurveyMonkey website (<http://www.surveymonkey.com/>). Each of the six presentation lists was presented in a separate version of the survey. Participants were instructed to read each sentence and judge whether it was acceptable or unacceptable. Following Arregui et al. (2006), I

defined unacceptable sentences as those sentences that violated the conventional rules of English – all other sentences were to be judged as acceptable whether or not they were insightful, interesting, true, or elegant.

3.2.3 Results

An ordinary logistic regression model was fitted to the data. Logistic regression models are designed specifically for analyzing categorical data, and have been reported to be a more accurate method of analysis for this type of data than standard analysis of variance measures (Jaeger, 2008). Logistic regression models the likelihood that an event will occur (typically defined as success or failure) given a set of predictor variables. In the current experiment, the model predicts the occurrence of an Acceptable response (i.e., success) or that of an Unacceptable response (i.e., failure) in the sentences containing verb-phrase ellipsis, given the type of antecedent and the type of ellipsis clause. Logistic regression analysis returns a log-odds coefficient value for each predictor in the model. The model's regression coefficient estimates the direction and magnitude of the relationship between the value of the predictor and the likelihood of an Acceptable rating. A significant coefficient suggests that the predictor is reliably associated with the likelihood of the event. The regression coefficients are expressed in units of log-odd ratio. The odds ratio is a function of probability: $\text{odds} = p/(1 - p)$, where p is the events probability. Thus, a probability of 0.5 translates into the odds ratio of 1, such that both outcomes are equally likely to occur. The odds ratio is further (natural) log-transformed to achieve the distributional properties desirable for regression analyses.

The model contrasts one predictor variable against another predictor variable that has been assigned as a baseline. All predictors in the model were categorical variables, so the coefficients reflected the effects of contrasts on the likelihood of an Acceptable rating. A positive coefficient suggests that the probability of a successful event is higher under the specified condition compared to a baseline, while a negative coefficient suggests that the probability of a successful event occurring is lower under the same conditions. For each contrast, I report the log-odds coefficient (β) and its level of significance (SE = standard error) as well as the difference in odds between conditions ($e\beta$). This $e\beta$ indicates the likelihood that one condition will lead to an Acceptable response compared to another. That is, if $e\beta = 1.5$ when comparing Condition A to a baseline Condition B, then a successful event is 1.5 times more likely to occur in Condition A than Condition B.

A total of 1152 acceptability judgment responses were collected: 386 active antecedent type judgment responses (197 with parallel ellipsis types, 189 with nonparallel ellipsis types), 377 infinitive antecedent type judgment responses (190 with parallel ellipsis types, 187 with nonparallel ellipsis types), and 389 progressive antecedent type judgment responses (192 with parallel ellipsis types, 197 with nonparallel ellipsis types). Table 3.1 presents the percentage of Acceptable responses for each condition. Active antecedent types were judged acceptable more often than infinitive antecedent types, which were

Table 3.1: Experiment 1A: Percent ‘Acceptable’ responses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	61.42	51.32
Infinitive (<i>get to see</i>)	47.89	41.18
Progressive (<i>be seeing</i>)	42.19	26.90

Table 3.2: Experiment 1A: Summary of logistic regression analysis.

Predictors	Coefficient	Std Error	Wald’s Z	Odds	
Intercept	0.03383	0.11953	0.283	1.03	$p = 0.7$
Antecedent:Infinitive	-0.48207	0.14680	-3.284	0.62	$p < 0.01$
Antecedent:Progressive	-0.90755	0.14900	-6.091	0.40	$p < 0.001$
Ellipsis:Parallel	0.45058	0.12124	3.717	1.57	$p < 0.001$

judged acceptable more often than progressive antecedent types (active, 56.4%; infinitive, 44.6%; progressive, 34.6%). Parallel antecedent types were judged acceptable more often than nonparallel antecedent types (parallel, 50.5%; nonparallel, 39.8%)

A logistic regression model predicted the probability of an Acceptable response given the factors of antecedent type (active vs. infinitive vs. progressive) and ellipsis type (parallel vs. nonparallel). Table 3.2 presents a summary of the analysis. This analysis revealed a main effect of antecedent type and a main effect of ellipsis type, but no interaction ($\chi^2(2) = 1.94$, $p = 0.38$). The main effect of antecedent type ($\chi^2(2) = 38.16$, $p < 0.001$) confirmed the presence of a stepwise pattern of acceptability: infinitive antecedent types were less likely to be judged acceptable than active antecedent types ($\beta = -0.48$, $SE = 0.15$, $p < 0.01$, $e\beta = 0.62$), and progressive antecedent types were less likely be judged acceptable than both infinitive antecedent types ($\beta = -0.42$, $SE = 0.15$, $p < 0.01$, $e\beta = 0.65$) and active antecedent types ($\beta = -0.91$, $SE = 0.15$, $p < 0.001$, $e\beta = 0.40$). This suggests that the acceptability of ellipsis sentences is partially determined by the complexity of the antecedent verb phrase. The main effect of ellipsis type ($\chi^2(1) = 13.91$, $p < 0.001$) showed that parallel ellipsis types were more likely to be judged acceptable than nonparallel antecedent types ($\beta = 0.45$, $SE = 0.12$, $p < 0.001$, $e\beta = 1.57$). This suggests that ellipsis sentences are more acceptable when the antecedent verb phrase and the ellipsis verb phrase are identical.

3.2.4 Discussion

This experiment sought to examine the assumptions underlying the Recycling Hypothesis (Arregui et al., 2006). Primarily, it tested the assumption that the syntactic operations responsible for assigning verbal morphology carry no cognitive costs. In addition, it tested two assumptions inherent in the Copy α principle (Frazier & Clifton, 2001), which claims that parallel antecedents are copied into the ellipsis site. The first assumption predicts that this operation is used any time a parallel antecedent is avail-

able. The second assumption predicts that the effectiveness of the copy operation is not affected by the complexity of the antecedent to be copied.

The results showed that parallel antecedents were somewhat more acceptable than non-parallel antecedents, which coincides with previous evidence suggesting a preference for parallel antecedents (e.g., Arregui et al., 2006; Mauner et al., 1995; Tanenhaus & Carlson, 1990; Dickey & Bungler, 2011). More importantly, the results indicated that the morphological form of the antecedent verb phrase affected the acceptability of non-parallel antecedents. Specifically, active antecedents were rated the most acceptable, while progressive antecedents were rated the least acceptable. This result is unpredicted under the hypothesis that there is no cognitive cost associated with switching the morphology on the antecedent verb phrase to fit the ellipsis site. On the contrary, this result suggests that there are different cognitive costs associated with stripping the antecedent of its verbal morphology and applying the simple past tense required by the ellipsis site. Furthermore, this result suggests that verbal morphology may have been a confounding factor in the results that Arregui et al. (2006) found in support of the Recycling Hypothesis.

The results also showed that the morphological form of the antecedent affected judgments of parallel antecedents. These results are unpredicted under the hypothesis that parallel antecedents are automatically copied into the ellipsis site, and the material inside the antecedent does not affect the application of this copy operation (Frazier & Clifton, 2001). At the very least, these results suggest that the material inside the antecedent can, indeed, affect the application of the copy operation. Specifically, the copy mechanism appears to be sensitive to the morphology on the antecedent verb phrase. Interestingly, morphological form affected parallel and non-parallel antecedents equally (as indicated by the lack of a significant interaction between antecedent type and ellipsis type). In fact, parallel antecedents showed the same pattern of results as non-parallel antecedents: active antecedents were rated the best, while progressive antecedents were rated the worst. This similarity suggests that parallel and non-parallel antecedents were processed in the same way. However, it is unclear what this similarity means for the status of the copy mechanism: Does the copy operation apply in both parallel and non-parallel ellipsis environments? Or, does this result suggest that there is in fact no copy mechanism, and the Recycling Hypothesis applies to both parallel and non-parallel ellipsis cases? These questions are addressed in more detail in Section 3.6, and at length in Chapter 4.

One caveat to the results reported in this experiment is that the acceptability ratings for all of the sentences were very low. In fact, the most acceptable sentence type was judged to be acceptable only 61% of the time. These low acceptability ratings may reflect participants difficulty with judging ellipsis sentences out of context. Thus, the experiment reported in Section 3.3 presented ellipsis sentences with a preceding context sentence with the aim of increasing the acceptability ratings.

3.3 Experiment 1B: Ellipsis with Context

3.3.1 Motivation

The primary aim of the experiment reported in this section was to replicate the results from the previous experiment, while simultaneously increasing the acceptability ratings of the sentences. To increase the acceptability of the sentences, I paired the elided sentences from Experiment 1A with a preceding context sentence. If the results from the previous experiment prove to be replicable, we can conclude that these results are robust. If the preceding context sentences improve the acceptability of the elided sentences, then we can conclude that the low ratings found in the previous experiment were due to the difficulty with interpreting ellipsis sentences out of context.

3.3.2 Methodology

3.3.2.1 Participants

Participants were recruited from the McMaster University community. Forty-one participants completed the survey as volunteers or for course credit. One participant was a non-native English speaker and was excluded from the analysis. Data from the remaining 40 native English speakers are presented here. There were 33 females and eight males. Participants ranged from 18 to 56 years in age, with an average of 23.68 years.

3.3.2.2 Stimuli

The stimuli used in this experiment were the same as for Experiment 1A. However, instead of presenting each sentence in isolation as in Experiment 1A, each sentence in Experiment 1B was presented with a preceding context sentence. For example, the context sentence for the sample stimuli in (3) was: *The local newspaper reported that a comet might pass through tonight's meteor shower.* All context sentences are presented with their corresponding stimuli sets in Appendix A.

3.3.2.3 Procedure

The procedure for this experiment was the same as for Experiment 1A, with one exception. In this experiment, stimuli were presented in two-sentence texts, and participants were instructed to judge the acceptability of the second sentence. All other elements of the experiment stayed the same.

Table 3.3: Experiment 1B: Percent ‘Acceptable’ responses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	61.90	53.21
Infinitive (<i>get to see</i>)	57.01	42.45
Progressive (<i>be seeing</i>)	35.58	19.27

Table 3.4: Experiment 1B: Summary of logistic regression analysis.

Predictors	Coefficient	Std Error	Wald’s Z	Odds
Intercept	0.02537	0.15998	0.159	1.03 $p = 0.08$
Antecedent:Infinitive	-0.32395	0.19675	-1.646	0.72 $p = 0.09$
Antecedent:Progressive	-1.30806	0.20926	-6.251	0.27 $p < 0.001$
Ellipsis:Parallel	0.57525	0.16651	3.455	1.78 $p < 0.001$

3.3.3 Results

As in Experiment 1A, an ordinary logistic regression model was fitted to the data. For each contrast, I report the log-odds coefficient (β) and its level of significance (SE = standard error) as well as the difference in odds between conditions ($e\beta$).

A total of 640 acceptability judgment responses were collected: 214 active antecedent type judgment responses (105 with parallel ellipsis types, 109 with nonparallel ellipsis types), 213 infinitive antecedent type judgment responses (107 with parallel ellipsis types, 106 with nonparallel ellipsis types), and 213 progressive antecedent type judgment responses (104 with parallel ellipsis types, 109 with nonparallel ellipsis types). Table 3.3 presents the percentage of Acceptable responses for each condition. Active antecedent types were judged acceptable more often than infinitive antecedent types, which were judged acceptable more often than progressive antecedent types (active, 57.7%; infinitive, 49.8%; progressive, 27.4%). Parallel antecedent types were judged acceptable more often than nonparallel antecedent types (parallel, 51.5%; nonparallel, 38.3%)

A logistic regression model predicted the probability of an Acceptable response given the factors of antecedent type (active vs. infinitive vs. progressive) and ellipsis type (parallel vs. nonparallel). Table 3.4 presents a summary of the analysis. This analysis revealed a main effect of antecedent type and a main effect of ellipsis type, but no interaction ($\chi^2(2) = 1.32, p = 0.52$). The main effect of antecedent type ($\chi^2(2) = 44.42, p < 0.001$) reflects the fact that progressive antecedent types were less likely to be judged acceptable than both infinitive antecedent types ($\beta = -0.98, SE = 0.21, p < 0.001, e\beta = 0.37$) and active antecedent types ($\beta = -1.31, SE = 0.21, p < 0.001, e\beta = 0.27$). There was a non-significant trend indicating that infinitive antecedents were less likely to be judged acceptable than active antecedents ($\beta = -0.32, SE = 0.20, p = 0.09, e\beta = 0.72$). This suggests that the acceptability of ellipsis sentences is partially determined by the complexity of the antecedent verb phrase. The main effect of ellipsis type ($\chi^2(1) = 12.08, p < 0.001$) reflects the fact that parallel ellipsis types were more likely to be

judged acceptable than nonparallel ellipsis types ($\beta = 0.58$, $SE = 0.17$, $p < 0.001$, $e\beta = 1.78$). This suggests that ellipsis sentences are more acceptable when the antecedent verb phrase and the elided verb phrase are identical.

3.3.4 Discussion

The results from this experiment showed the same pattern of results as Experiment 1A. Parallel antecedents were rated more highly than non-parallel antecedents. Furthermore, both antecedent types showed varying degrees of acceptability depending on the morphological form of the antecedent. For both parallel and non-parallel antecedents, active antecedents were most acceptable and progressive antecedents were least acceptable. This result suggests that the findings from the previous experiment are reliable, and provide further evidence for the conclusion that the syntactic operations involved in changing morphological properties between the antecedent and ellipsis site do in fact carry cognitive costs.

Interestingly, the preceding context sentence did not improve acceptability ratings. In fact, a comparison between the two experiments showed no significant differences (all $ps < 0.1$). This result suggests that the low acceptability ratings found in the previous experiment were not due to the difficulty associated with interpreting elided sentences out of context. One might question whether the low acceptability ratings might be improved by presenting the same stimuli auditorily. Several studies have shown that prosody can affect the interpretation given to an elided sentence (Carlson, 2001; Carlson, Dickey, Frazier, & Clifton, 2009; Frazier, Clifton, & Carlson, 2007). Carlson et al. (2009) investigating the effect of prosody on the interpretation of sluiced sentences in which the *wh*-phrase could be correlated with either the subject or object of the antecedent clause (4). In a written questionnaire, through which participants received no indication of prosodic structure, the researchers found that participants were more likely to choose (4b) as the more appropriate interpretation – that is, participants were more likely to correlate the remnant *wh*-phrase with the object in the antecedent clause.

- (4) The captain talked with the co-pilot, but we couldn't find out who else.
- a. We couldn't find out who else talked with the co-pilot. (subject antecedent)
 - b. We couldn't find out who else the captain talked with. (object antecedent)

In a second experiment, the researchers presented the stimuli auditorily, and manipulated which antecedent noun phrase received focal stress: the subject, the object, both, or neither (stress was placed on the verb, here). They found that sentences with focal stress on the subject were more likely to receive subject-antecedent interpretations (4a) than the three other conditions, suggesting that focal stress can bias individuals to particular interpretations.

The current experiments did not investigate elided sentences with potentially ambiguous antecedents, so Carlson and colleague's (2001; 2009; 2007) results do not serve

to provide explicit predictions for the current set of stimuli. However, the results suggest that ellipsis comprehension can be manipulated using prosody. Other research suggests that prosody may also affect sentence acceptability: In an auditory sentence-shadowing task, researchers have found that grammatical sentences are easier to shadow than ungrammatical sentences, but only if the grammatical sentences are presented with appropriate prosodic structure (Miller & Isard, 1963; Martin, 1968; as cited in Cutler, Oahan, & van Donselaar, 1997). Together, these results suggest that the acceptability and interpretation of elided sentences may be improved by auditory presentation with appropriate prosodic cues.

3.4 Experiment 1C: Non-Elided Controls

3.4.1 Motivation

One question that remains is how the antecedents used in the previous experiments behave in non-elided contexts. Non-elided constructions are a standard baseline measure for determining whether the effects observed in elided sentences are due to general sentence processing constraints, or whether they are due to ellipsis-specific properties. If elided and non-elided constructions show the same pattern of results under the same conditions, this suggests that the effect can be explained in terms of some general processing constraint. However, if results show that only elided constructions are affected by some manipulation (here, the morphological form of the antecedent), then the effect can be said to be ellipsis-specific and results from the interaction of the manipulation with the ellipsis context.

The experiment described in this section examined whether the effects observed in the previous experiments were caused by general or ellipsis-specific processing constraints by measuring the acceptability of the same stimuli in non-elided constructions. Examples of these sentences can be found in (5), below. If the effect of antecedent morphology observed in the previous experiments is caused by a general processing constraint, then these non-elided sentences should show a similar pattern of results to the elided constructions in the first two experiments. However, if the effect of antecedent morphology observed above is caused by an ellipsis-specific processing constraint, then the results should show no difference in acceptability among the sentences.

3.4.2 Methodology

3.4.2.1 Participants

Participants were recruited from Amazon Mechanical Turk community. In total, 53 individuals completed the survey. All 53 respondents were compensated with \$0.005 per

sentence judged. 2 individuals were non-native English speakers and were excluded from the analysis. Data from the remaining 51 native English speakers are presented here. There were 28 females and 23 males. Participants ranged from 18 to 60 years in age, with an average of 33.53 years.

3.4.2.2 Stimuli

The stimuli used in this experiment were non-elided versions of the stimuli used for Experiment 1B.

- (5)
- a. None of the astronomers would see the comet, but John did/would see the comet.
 - b. None of the astronomers would get to see the comet, but John did/would get to see the comet.
 - c. None of the astronomers would be seeing the comet, but John did/would be seeing the comet.

3.4.2.3 Procedure

The procedure for this experiment was the same as for Experiment 1B, except for the website that hosted the survey. Participants completed the experiment via a survey hosted on the Amazon Mechanical Turk website (<https://www.mturk.com/mturk/>). The Amazon Mechanical Turk provides a reliable method of conducting survey-based experiments. Several experiments have compared using Amazon Mechanical Turk with lab-based procedures, and have found excellent compatibility of results (Gibson, Piantadosi, & Fedorenko, 2011; Mason & Suri, 2012; Munro, Bethard, Kuperman, Lai, Melnick, Potts, Schnoebelen, & Tily, 2010; Schnoebelen & Kuperman, 2010; Snow, OConnor, Jurafsky, & Ng, 2008; Sprouse, 2011). All other elements of the experimental procedure stayed the same.

3.4.3 Results

As in the previous two experiments, an ordinary logistic regression model was fitted to the data. For each contrast, I report the log-odds coefficient (β) and its level of significance (SE = standard error) as well as the difference in odds between conditions ($e\beta$).

A total of 816 acceptability judgment responses were collected: 271 active antecedent type judgment responses (134 with parallel ellipsis types, 137 with nonparallel ellipsis types), 274 infinitive antecedent type judgment responses (138 with parallel ellipsis types,

Table 3.5: Experiment 1C: Percent ‘Acceptable’ responses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	70.90	75.91
Infinitive (<i>get to see</i>)	85.51	72.06
Progressive (<i>be seeing</i>)	69.53	75.52

Table 3.6: Experiment 1C: Summary of logistic regression analysis.

Predictors	Coefficient	Std Error	Wald’s Z	Odds
Intercept	1.14788	0.19980	5.745	3.15 $p < 0.001$
Antecedent:Infinitive	-0.20050	0.27647	-0.725	0.82 $p = 0.4$
Antecedent:Progressive	-0.02110	0.27884	-0.076	0.98 $p = 0.9$
Ellipsis:Parallel	-0.25757	-0.27584	-0.934	0.77 $p = 0.3$
Infinitive:Parallel	1.08514	0.41362	2.624	2.96 $p < 0.01$
Progressive:Parallel	-0.04414	0.38832	-0.114	0.96 $p = 0.9$

136 with nonparallel ellipsis types), and 271 progressive antecedent type judgment responses (128 with parallel ellipsis types, 143 with nonparallel ellipsis types). Table 3.5 presents the percentage of Acceptable responses for each condition. Infinitive antecedent types were judged acceptable more often than active and progressive antecedent types (active, 73.4%; infinitive, 78.8%; progressive, 72.5%). Nonparallel antecedent types were judged acceptable more often than parallel antecedent types (parallel, 75.3%; nonparallel, 74.5%)

A logistic regression model predicted the probability of an Acceptable response given the factors of antecedent type (active vs. infinitive vs. progressive) and ellipsis type (parallel vs. nonparallel). Table 3.6 presents a summary of the analysis. This analysis revealed a significant interaction, but no main effects. The significant antecedent type x ellipsis type interaction ($\chi^2(2) = 9.53$, $p < 0.01$) reflects the fact that the infinitive-parallel condition was more likely to be judged acceptable than all other conditions ($\beta = 2.09$, $SE = 0.41$, $p < 0.01$, $e\beta = 2.96$). The main effect of antecedent type was not significant ($\chi^2(2) = 3.31$, $p = 0.19$), nor was the main effect of ellipsis type ($\chi^2(1) = 0.08$, $p = 0.77$). This suggests that although the acceptability of nonelided sentences is not determined by antecedent type and ellipsis type alone, these factors may interact to influence the acceptability as seen in the infinitive-parallel condition.

3.4.4 Discussion

The experiment reported in this section aimed to show that the effects observed in the previous experiments were caused by an ellipsis-specific processing constraint, by ruling out the possibility of a general processing constraint. Results showed that the morphological form of the antecedent affected the judgments of non-elided sentences only in the context of infinitive-parallel antecedents. Infinitive-parallel antecedents were judged more

acceptable than all of the other sentence types, for which there were no differences in acceptability judgments. It is unclear why the infinitive-parallel antecedents were judged to be most acceptable. In the previous experiments, active-parallel antecedents were most acceptable in ellipsis contexts. Therefore, one would predict that active-parallel antecedents would have an advantage in non-elided contexts as well. This does not appear to be the case. However, the other five conditions showed approximately equal ratings of acceptability, suggesting that they were all processed in generally the same way. I take these results to suggest that infinitive-parallel antecedents were also processed in a similar fashion. Thus, the results indicate that the effects of morphological form observed in the two previous experiments were caused by an ellipsis-specific processing constraint, not a general one.

These results contrast with Dickey and Bungler's (2011) finding that sluiced clauses and non-elided controls are equally sensitive to non-parallelism, suggesting that parallelism is in fact a general linguistic preference, not an ellipsis-specific one. The different findings between their experiment and the current experiments might be explained by three factors. First, Dickey and Bungler used sluicing instead of verb-phrase ellipsis – as discussed in Section 1.4.2, sluicing and verb-phrase ellipsis are remarkably similar, but do deviate from each other in some cases. Whether or not the elided sentences pattern the same way as non-elided sentences may be one of those cases. Secondly, they used a different type of parallelism manipulation. They manipulated whether or not the sluiced *wh*-remnant had an overt correlate in the antecedent, which represents a strictly syntactic manipulation that creates fundamentally different syntactic structure. It is possible that changing the syntactic structure creates a stronger manipulation than simply modifying morphosyntactic features (as was done in the current experiments), and this allowed the researchers to observe a preference for parallelism in non-elided sentences. Finally, the difference between the two experiments may be a result of different measurement tools: Dickey and Bungler measured on-line reading times, while the current experiment used off-line acceptability judgments. It might be the case that their on-line measures offer a more sensitive tool for evaluating whether parallelism is an ellipsis-specific preference.

3.5 Experiment 2: Self-Paced Reading

3.5.1 Motivation

The primary goal of the self-paced reading experiment reported here was to assess the effect of antecedent complexity on the on-line comprehension of non-parallel antecedents. Arregui et al. (2006) reported a series of grammaticality judgment tasks in support of their Recycling Hypothesis, which predicts that the number and difficulty of the syntactic operations required to build the ellipsis site will affect comprehension. However, they did not report any experiments using traditional on-line measures of processing difficulty (e.g. reading time). Gibson (1998) suggests that reading time correlates with processing

difficulty, such that increases in processing difficulty result in longer reading times. Furthermore, he suggests that these longer reading times should be observed at the point in the sentence where the difficulty occurs. Therefore, if processing non-parallel antecedents is thought to be more difficult than parallel antecedents because it requires rebuilding the structure underlying the ellipsis site, then longer reading times should be observed at the ellipsis site for non-parallel antecedents when compared to parallel antecedents. Furthermore, if antecedent complexity affects non-parallel ellipsis comprehension, then longer reading times should be observed at the ellipsis site for more complex antecedents compared to simpler antecedents.

If the parser rebuilds the ellipsis structure in non-parallel antecedents, as suggested by the Recycling Hypothesis, then the time it takes to read the ellipsis site should vary depending on the complexity of the antecedent. The survey experiments reported above have shown that parallel antecedents were more acceptable than non-parallel antecedents, suggesting that non-parallel ellipsis is the more complex construction. Thus, one might predict that reading times at the ellipsis site will be longer in the context of non-parallel antecedents. Furthermore, the previous experiments also showed that antecedent morphology affects acceptability of elided constructions such that active antecedents are most acceptable and progressive antecedents are least acceptable, with infinitive antecedents ranking somewhere in between. Therefore, reading times should be longest at the ellipsis site under progressive antecedent conditions, and shortest under active antecedent conditions.

If results show that reading times at the ellipsis site are indeed modulated by the effects of parallelism and/or antecedent morphology, this will support the hypothesis that the parser rebuilds the structure of the ellipsis site in the case of non-parallel antecedents. However, if results show no difference in reading times at the ellipsis site, this will suggest that processing elided constructions with non-parallel antecedents occurs in much the same time frame as with parallel antecedents. If this is the case, it may be possible to conclude that non-parallel ellipsis does not involve rebuilding the elided structure. Alternatively, it may indicate that even parallel ellipsis comprehension involves the rebuilding operations suggested by the Recycling Hypothesis.

3.5.2 Methodology

3.5.2.1 Participants

Participants were recruited from the undergraduate participant pool in the Department of Linguistics at McMaster University. In total, 33 students participated in the experiment. All 33 respondents were compensated with course credit. Four students were non-native English speakers and were excluded from the analysis. Data from the remaining 29 native English speakers are presented here. There were 24 females and five males. Participants ranged from 18 to 60 years in age, with an average of 22.03 years.

3.5.2.2 Stimuli

The stimuli used in this experiment were the same as for Experiment 1B.

3.5.2.3 Procedure

Participants completed this experiment in individual 45-minute sessions. As in the previous experiments, each of the six presentation lists was presented in a separate version of the experiment. Sentences were presented one word at a time in the centre of the screen, with each subsequent word replacing the word before it. Participants were instructed to read each sentence at a comfortable pace by pressing the space bar to view each new word. After reading each sentence, participants judged whether it was acceptable or unacceptable. As in the previous experiments, I defined unacceptable sentences as those that violated the conventional rules of English – all other sentences were to be judged as acceptable whether or not they were insightful, interesting, true, or elegant.

3.5.3 Results

Again, an ordinary logistic regression model was fitted to the data. For each analysis, I report the log-odds coefficient (β) and its level of significance (SE = standard error) as well as the difference in odds between conditions ($e\beta$).

In total, 464 acceptability judgment responses were collected: 155 active antecedent type judgment responses (76 with parallel ellipsis types, 79 with nonparallel ellipsis types), 156 infinitive antecedent type judgment responses (78 with parallel ellipsis types, 78 with nonparallel ellipsis types), and 153 progressive antecedent type judgment responses (78 with parallel ellipsis types, 75 with nonparallel ellipsis types). Table 3.7 presents the percentage of Acceptable responses for each condition. Active antecedent

Table 3.7: Experiment 2: Percent ‘Acceptable’ responses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	52.63	60.76
Infinitive (<i>get to see</i>)	55.13	52.56
Progressive (<i>be seeing</i>)	32.05	33.33

Table 3.8: Experiment 2: Summary of the logistic regression analysis.

Predictors	Coefficient	Std Error	Wald’s Z	p	Odds Ratio
Intercept	0.31966	0.18743	1.706	1.38	$p > 0.08$
Antecedent:Infinitive	-0.11764	0.22829	-0.515	0.89	$p > 0.61$
Antecedent:Progressive	-0.99404	0.23671	-4.199	0.37	$p < 0.001$
Ellipsis:Parallel	-0.09557	0.19035	-0.502	0.91	$p > 0.62$

types were judged acceptable slightly more often than infinitive antecedents, while progressive antecedents were least likely to be judged acceptable (active, 56.7%; infinitive, 53.9%; progressive, 32.7%). Nonparallel ellipsis types were slightly more likely to be judged acceptable than parallel ellipsis types (parallel, 46.6%; nonparallel, 48.9%).

A logistic regression model predicted the probability of an Acceptable response given the factors of antecedent type (active vs. infinitive vs. progressive) and ellipsis type (parallel vs. nonparallel). Table 8 presents a summary of the analysis. This analysis revealed a significant main effect of antecedent type, but no effect of ellipsis type ($\chi^2(1) = 0.25$, $p = 0.62$) and no interaction ($\chi^2(2) = 0.92$, $p = 0.63$). The main effect of antecedent type ($\chi^2(2) = 21.60$, $p < 0.001$) reflects the fact that progressive antecedent types were less likely to be judged acceptable than both active antecedent types ($\beta = -0.99$, $SE = 0.24$, $p < 0.001$, $e\beta = 0.37$) and infinitive antecedent types ($\beta = -0.88$, $SE = 0.24$, $p < 0.001$, $e\beta = 0.42$). Active and infinitive antecedents were equally likely to be judged acceptable ($\beta = -0.12$, $SE = 0.23$, $p = 0.61$, $e\beta = 0.89$). This suggests that the acceptability of ellipsis sentences is partially determined by the complexity of the antecedent verb phrase. These results differ from the previous experiments using elided sentences, where there was a significant difference between all three antecedent conditions. In the previous experiments, active antecedents were most likely to elicit an acceptable judgment, while progressive antecedents were least likely. This difference is likely due to the different methodology used in this experiment, compared to the previous experiments. This point is further discussed in Section 3.5.4.

Table 3.9: Experiment 2: Mean reaction time for the acceptability judgment task, as measured in milliseconds. Standard deviation in parentheses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	1281 (847)	1457 (1158)
Infinitive (<i>get to see</i>)	1200 (833)	1322 (930)
Progressive (<i>be seeing</i>)	1351 (883)	1300 (838)

Table 3.10: Experiment 2: Summary of regression analysis for judgment reaction times.

Predictors	Coefficient	Std Error	t	p
Intercept	6.99910	0.07599	92.110	$p < 0.001$
Antecedent:Infinitive	-0.02545	0.10862	-0.234	$p = 0.815$
Antecedent:Progressive	-0.02391	0.11032	-0.217	$p = 0.829$
Ellipsis:Parallel	-0.02456	0.10903	-0.225	$p = 0.822$
Infinitive:Parallel	-0.04472	0.15365	-0.291	$p = 0.771$
Progressive:Parallel	0.06807	0.15679	0.434	$p = 0.664$

3.5.3.1 Acceptability Judgments Reaction Times

Table 3.9 reports the mean reaction time for the acceptability judgment task. The data were fitted to an ordinary regression model. Table 3.10 presents a summary of

Table 3.11: Experiment 2: Mean final-word reading times for each sentence, as measured in milliseconds. Standard deviation in parentheses.

Ellipsis Type	Antecedent Type	
	Parallel (<i>would</i>)	Non-Parallel (<i>did</i>)
Active (<i>see</i>)	1337 (1013)	1605 (1134)
Infinitive (<i>get to see</i>)	1594 (1236)	1446 (1021)
Progressive (<i>be seeing</i>)	1362 (989)	1529 (1193)

Table 3.12: Experiment 2: Summary of regression analysis for final-word reading times.

Predictors	Coefficient	Std Error	t	
Intercept	1482.4	112.5	13.182	$p < 0.001$
Antecedent:Infinitive	-125.6	158.4	-0.793	$p = 0.428$
Antecedent:Progressive	-149.1	161.6	-0.923	$p = 0.357$
Ellipsis:Parallel	-232.2	159.6	-1.455	$p = 0.296$
Infinitive:Parallel	236.7	226.3	1.046	$p = 0.296$
Progressive:Parallel	161.1	227.6	0.708	$p = 0.480$

this analysis. The analysis revealed no difference in reaction times between any of the conditions. This is possibly due to the high variability (SD > 800 ms) in reaction times among the participants.

3.5.3.2 Reading Times

Table 3.11 reports the mean reading times for the final word of each sentence, would in parallel ellipsis type sentences, did in nonparallel ellipsis type sentences. Table 3.12 presents a summary of this analysis. The analysis revealed no difference in reading times between any of the conditions. Again, this is possibly due to the high variability in participants' reading times (SD > 900 ms.)

3.5.4 Discussion

The self-paced reading experiment presented in this section was designed to determine whether on-line measures of sentence processing were sensitive to the processes involved in rebuilding the elided structure when there is no parallel antecedent. Acceptability judgment responses showed that active and infinitive antecedents were more acceptable than progressive antecedents, independent of parallelism effects. There was also no difference in acceptability between parallel and non-parallel antecedent conditions. These results differ from the results reported in the survey experiments described earlier in this chapter. Specifically, the previous experiments showed that all three morphological forms differed from each other with respect to their effects on acceptability ratings active antecedents were most acceptable, progressive antecedents were least acceptable, and infinitive antecedents ranked somewhere in the middle. Earlier survey experiments also

showed that parallel antecedents were more acceptable than non-parallel antecedents. One explanation for the difference between this self-paced reading experiment and the survey experiments reported previously may be the on-line nature of the self-paced reading task. While completing the survey experiments, participants could take as long as they desired to judge the acceptability of the sentences. In contrast, participants were forced to make acceptability judgments within five seconds after they finished reading the sentence. This time constraint may have influenced the acceptability ratings in the self-paced reading experiment.

Results also showed no difference in either of the on-line processing measures. There was no difference in judgment speed among the sentences, and there was no difference in the reading times at the ellipsis site. These results indicate that the relative complexity of the antecedent types did not affect the length of time it took participants to process these sentences, suggesting that the parser may not actively rebuild the elided structure at the ellipsis site. If this is true, then these results serve as evidence against the Recycling Hypothesis (Arregui et al., 2006). However, this explanation is speculative as it relies on a null result. One alternative explanation of these findings is that the processing measures employed in this experiment were simply not sensitive to the complexity manipulations. Another explanation may be that the variation in complexity created by different verbal morphology may not be strong enough to reliably produce an effect on reading times. These possibilities are briefly discussed in the next section, and outlined more explicitly in Chapter 4.

3.6 What do we know now? (Summary)

The experiments reported in this thesis were designed to test two of the assumptions underlying the Recycling Hypothesis (Arregui et al., 2006). This hypothesis attempts to explain the comprehension of non-parallel cases of ellipsis, suggesting that the parser reuses the material from the antecedent in order to build the syntactic structure underlying the ellipsis site. The experiments primarily address the assumption that the syntactic operations required for assigning morphological form carry no cognitive costs. This assumption predicts that variations in antecedent morphology will have no effect on the acceptability of elided sentences. These experiments also address the assumption that the parser automatically copies parallel antecedents into the ellipsis site when they are available, and thus only rebuilds the elided structure when no parallel antecedent is available. This assumption predicts that any effect that antecedent morphology may have on the acceptability of elided sentences should only be seen under non-parallel antecedent conditions. Under parallel antecedent conditions, the entire antecedent verb phrase (morphology and all) should be copied into the ellipsis site. However, the antecedent verb phrase must be stripped of its morphological properties in non-parallel antecedent conditions, as this is the only way to prepare the verb phrase to be assigned the morphology required by the ellipsis site. Therefore, only judgments of non-parallel ellipsis should be sensitive to morphology manipulation.

Results from three survey experiments and one self-paced reading experiment show that the syntactic operations associated with assigning morphological properties carry cognitive costs, and the costs vary depending on the morphological form being analyzed. Antecedents with active morphology were generally rated more acceptable than antecedents with infinitive morphology, and antecedents with progressive morphology were rated the worst overall. This effect was observed for both parallel and non-parallel antecedents, suggesting that whatever operations are responsible for interpreting ellipsis, those operations are the same for both parallel and non-parallel antecedents. The broader implications of these results are discussed in Chapter 4.

Chapter 4

What Have We Learned About Ellipsis?

4.1 General Overview

This thesis examined the Recycling Hypothesis (Arregui, 2006) by testing two key underlying assumptions as well as a prediction that follows from the hypothesis. The Recycling Hypothesis is described in detail in Chapter 2. The experiments designed to test the hypothesis are reported in Chapter 3. Firstly, these experiments tested the assumption that the syntactic operations that are responsible for assigning verbal morphology are relatively easy (Arregui et al., 2006, p. 242) or cost-free. Secondly, they tested the assumption that ellipsis constructions with parallel antecedents are comprehended via a cost-free copy mechanism (Frazier & Clifton, 2001), by which the syntactic structure of the antecedent is copied into the ellipsis site. Finally, this thesis tested the prediction that on-line reading time measures should be sensitive to complexity effects associated with reconstructing the structure of the ellipsis site via recycling. Sections 4.2 through 4.4 address these questions in turn. Section 4.5 concludes the chapter, and situates the findings from this thesis within the models of ellipsis comprehension described in Section 2.3. The current findings seem to run counter to the predictions made by the Recycling Hypothesis, but can they be explained by Kim et al.'s (2011) memory search process theory or by Martin and McElree's (2008; 2009; 2011) cue-dependent direct-access memory retrieval account?

4.2 Is morphological assignment cost-free?

Arregui et al. (2006) claimed that “the class of really easy recycling operations seems to be those that do not involve structural alterations at all, for example, changing the features on a verb phrase from +tense to –tense (p. 242)”. Unfortunately, they provided

no evidence to support this claim and created stimuli (1) in which they did not control for the morphological form of the verb phrase.

- (1) Stimuli from Arregui et al. (2006)
 - a. None of the astronomers saw the comet, but John did.
 - b. Seeing the comet was nearly impossible, but John did.
 - c. The comet was nearly impossible to see, but John did.
 - d. The comet was nearly unseeable, but John did.

- (2) Stimuli from the Current Experiments
 - a. None of the astronomers would see the comet, but John did / would.
 - b. None of the astronomers would get to see the comet, but John did / would.
 - c. None of the astronomers would be seeing the comet, but John did / would.

Specifically, they assumed that stripping the verb phrase of its simple past tense morphology in (1a) is equivalent to stripping the verb phrase of its gerundive morphology in (1b), but never showed that this is true. In fact, they do not even mention this operation as one that must be performed in order to comprehend these sentences. One of the aims for the experiments reported in Chapter 3 was to test the assumption that morphological assignment is easy. If Arregui et al.'s assumption is correct, the three forms of verbal morphology in (2) – active (2a), infinitive (2b), and progressive (2c) – should have shown similar acceptability ratings under the same ellipsis contexts. In fact, two online survey experiments and one self-paced reading experiment showed that active morphology is more acceptable than infinitive morphology under ellipsis, and both are better than progressive morphology. These results run counter to the assumption that Arregui et al. make about morphological assignment. To be clear, (re-)assigning verbal morphology is not necessarily easy, and different forms of morphology seem to cause more problems than others.

Hyönä and Vainio (2002) provide further evidence that there may be cognitive costs associated with inflectional morphology. These authors conducted research on Finnish, which is a so-called free word order language with an intricate morphological system. In fact, Finnish morphology allows for highly inflected verbs – called *converbs* – to take the place of entire clauses. Converbs represent individual lexical items that encode all the required information for a full clause. In an experiment measuring eye movements, the researchers found that *converb* constructions elicited longer gaze durations than full clauses, suggesting that the *converbs* were more difficult to process. Additionally, Hyönä, Vainio, and Laine (2002) investigated the effect of morphological complexity on Finnish noun phrases. Unlike English, in which the role of each participant in an event is determined by a fixed word order (subject-verb-object), Finnish uses affixes to identify the role of each participant in an event. This allows for variation in word order because the role assigned to a participant is marked directly on the lexical item and is, thus, independent of word order. The researchers showed that the morphological complexity effect

for inflected noun phrases typically found in standard lexical decision tasks disappeared when the items were embedded into sentence contexts. When the noun phrases were presented in isolation, the inflected items took longer to recognize than the uninflected items. When the same noun phrases were presented as part of a sentence, the two sets of items elicited comparable gaze durations and lexical decision times. One might argue that the removal of the morphological complexity effect implies that nominal morphology does not influence participants' reading abilities. However, it is important to note that there may be additional contextual cues that facilitate processing of nominal morphology in full sentences, which may alleviate the difficulties found in isolated presentation. Taken together, these results suggest that assigning morphological structure to a sentence may not be as simple as Arregui et al. (2006) assumed.

Interestingly, Hyönä et al. (2002) note that the inflected noun phrases used in their experiments would typically only be seen in isolation when used as part of an elided construction (e.g., a response to a question). As such, these items – and perhaps the Finnish language in general – may prove to be a useful tool for investigating the effects of morphological complexity on ellipsis comprehension. This may be particularly relevant for cases of mismatch ellipsis, such as those explored in this thesis and in Arregui et al. (2006). To explore the issue, one might make the antecedent and ellipsis verb phrases more or less similar by manipulating the number of morphemes that differ between the two. Presumably, increasing the number of morphemes that differ between the antecedent and ellipsis verb phrases should increase the complexity of the ellipsis construction. This might be reflected in decreased acceptability judgments, as observed in the experiments reported here, or as increases in reading times, as observed in the experiments conducted by Hyönä and colleagues (2001; 2002). Clearly, this would be an interesting and potentially fruitful line of research to pursue.

One last question that these results raise is: How is the information encoded in verbal morphology used during parsing? That is, what role does it play in the process of deriving a syntactic structure for a sentence? How can it be applied to the results found in Chapter 3? There are several possible answers to these questions, three of which will be discussed in this section. It is important to note that all of these hypotheses are consistent with the results described above. Further research is required to determine if any are viable explanations.

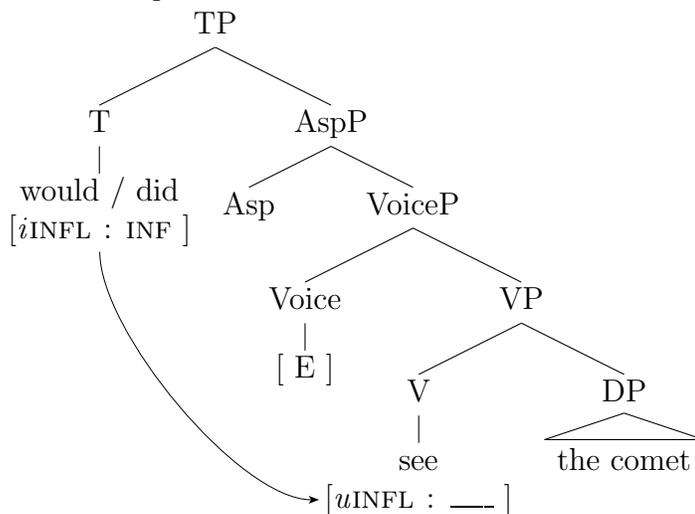
One hypothesis, derived from the model of morphology advanced by Bjorkman (2011) and described in Section 1.4.2, explains morphological complexity in terms of the number of INFL features that must be valued. In particular, this hypothesis predicts that sentences with a higher number of uninterpretable (*u*INFL) features will be more complex, morphologically speaking. Recall that Bjorkman's model operates within the Distributed Morphology framework, which argues that morphological features are encoded in the syntactic structure and only receive phonological realization at the level of Phonological Form (PF). Although these features do not receive phonological properties until PF, the Agree relation required for feature valuation must be satisfied in the syntax – in fact, the acceptability of a sentence depends on all uninterpretable features being valued in the

syntax. Thus, when the comprehender reaches an ellipsis site, he/she must determine whether or not the elided *uINFL* features have been properly checked. This would involve a rather complicated guessing task, through which the comprehender must first determine what *INFL* features exist in the structure and then establish whether the *uINFL* features have been valued. This can be rather difficult because the comprehender is typically not given very many hints as to what the underlying structure of the ellipsis site is.

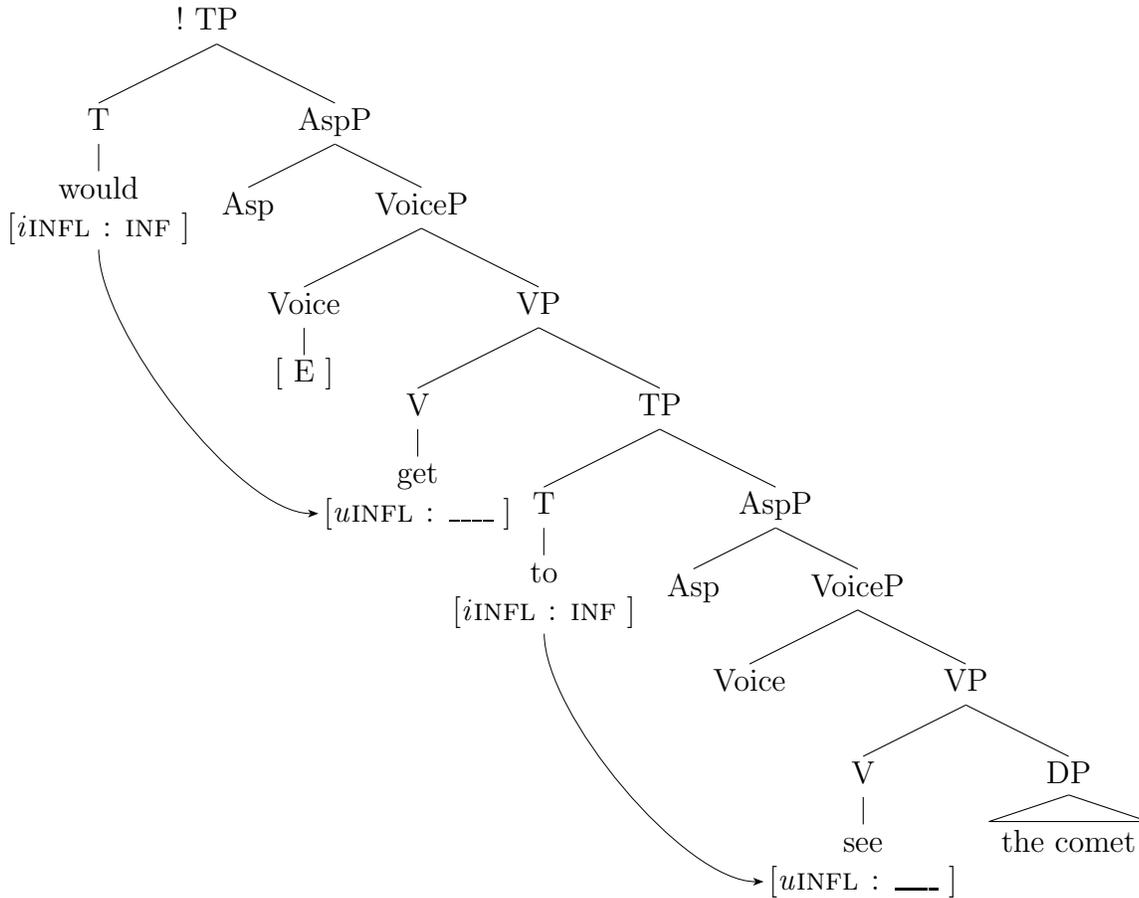
This account can readily explain the morphological complexity patterns observed in Chapter 3. The phrase structure trees below represent the elided structure for the parallel versions of the active (3), infinitive (4), and progressive (5) verb phrases used in the current experiments. In the active ellipsis site (3), the comprehender must first determine that a *iINFL* feature exists on *T*, and ensure that it is in the appropriate Agree relation to *V* in order to value the *uINFL* feature. This process involves assessing the value of only one *uINFL* feature, and should be relatively easy. In principle, the same process should apply for the non-parallel ellipsis sites.

Compare this with the infinitive ellipsis site (4). Here, the comprehender must determine the value of two *uINFL* features. First, it must ensure that the *uINFL* feature on *get* can be valued by the first *iINFL* feature. Then, he/she must discover the second *iINFL* feature, and ensure that it can provide a value for the *uINFL* feature on *see*. This process requires the comprehender to assess an additional *uINFL* feature, relative to the active sentence, and should therefore show more difficulty. Note that the two *iINFL* features in (4) share the same value, *INF*. Now, observe the progressive ellipsis site in (5). This tree also requires that two *uINFL* features be valued, but not by the same *iINFL* feature. Thus, the progressive sentences should be more difficult to process than the infinitival sentences because the comprehender must discover that the second *iINFL* feature is not the same as the first, and proceed accordingly.

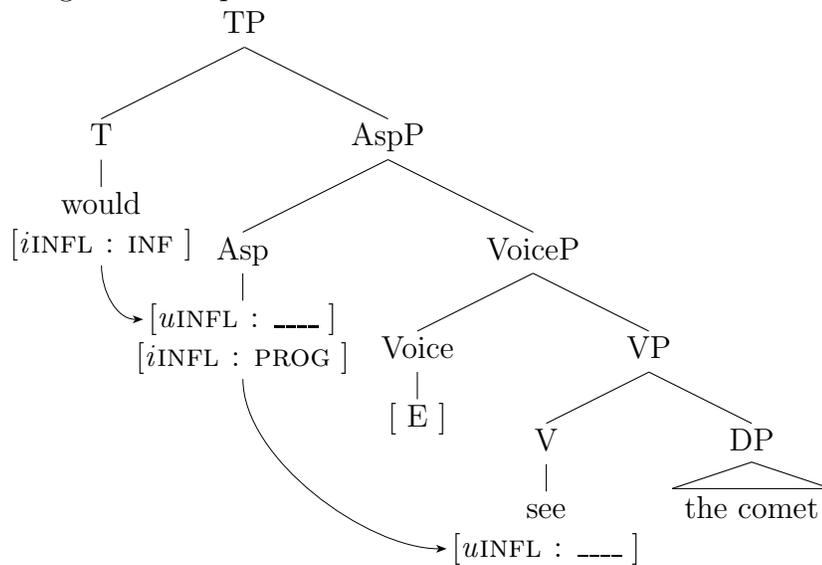
(3) Active Ellipsis Site:



(4) Infinitive Ellipsis Site:



(5) Progressive Ellipsis Site:



First, he/she must determine that the $u\text{INFL}$ feature in Asp will be assigned the INF value from above. Then, he/she will discover the second $i\text{INFL}$ feature with the PROG

value, and must determine whether it can provide a value for the *uINFL* feature on *see*. It could be the case that simply switching from one feature value to another creates the difficulty, or it could be the case that PROG features involve more ‘effort’, in some sense. In either case, Bjorkman’s model of morphology predicts that progressive morphology is more difficult than infinitive morphology, which in turn is more difficult than active morphology. The results from the current experiments show that acceptability judgments decrease as the level of morphology increases. Thus, the current findings are compatible with Bjorkman’s model.

A second hypothesis implicates the relative frequency with which these verbal forms occur in ellipsis. In fact, a substantial body of evidence shows that, in general, high frequency items show a range of superiority effects relative to low frequency items (see Ellis, 2002, for a comprehensive review). It may be the case that progressive verbs rarely occur in ellipsis, thus explaining their very low acceptability ratings. Similarly, it could be the case that infinitive verbs occur less frequently than finite active verbs, which would explain the result showing that infinitive ellipsis is less acceptable than active ellipsis. However, this hypothesis is purely speculative. These predictions would need to be confirmed with corpus frequency data for this hypothesis to have any weight.

Finally, verbal morphology may be used as a pruning mechanism that facilitates the parsing process by eliminating parses that differ from the antecedent. Standard chart parsers, such as the Earley model (Earley, 1970; Hale, 2001; Levy, 2008), generate a set of all possible parses for a given input using a top-down, left-to-right prediction algorithm. Upon reaching the ellipsis site, the parser must select a parse that matches the one generated for the antecedent. To do this, it will retrieve the parse for the antecedent and compare it to each individual prediction in the set generated for the ellipsis clause. If a direct match is not made immediately, the parser is forced to search through all of the possible alternative parses. This comparison process has the potential to be incredibly costly. To circumvent this cost, the parser may use morphological information as a pruning tool to narrow the space within which the parser must search for the correct representation. That is, the parser may use the information encoded in verbal morphology to discard the parses that are morphologically too distinct from the antecedent. When the correct parse for the elided verb phrase is in fact not very distinct from the antecedent (e.g., present vs. past tense verbs), this mechanism may in fact lower the cost associated with searching for a mismatching parse. As the parser begins the comparison process by looking at the least different parses first, it will quickly reach the correct parse. However, if the correct parse is highly distinct from the antecedent (e.g., active vs. progressive verbs), this pruning mechanism may in fact discard the correct parse, resulting in an incomplete parse, or ungrammatical sentence.

To summarize, the evidence suggests that morphological processing may not be as easy as Arregui et al. (2006) claimed. Results from the experiments reported in Chapter 3 showed that different forms of verbal morphology behave differently under conditions of ellipsis. Elided verbs with active morphology were more acceptable than elided infinitival verbs, and elided progressive verbs were worst overall. Importantly, these differences

in verbal morphology were consistent across parallel and non-parallel ellipsis contexts. Furthermore, Hyönä and colleagues (2001; 2002) have shown that morphology can be costly even in agglutinative languages such as Finnish. Highly inflected verbs (i.e., *converbs*) were shown to be more difficult to process than the full subclauses they represent, suggesting that inflectional morphology is more complex than a string of unbound morphemes. This increase in complexity may create complications with the parsing process, resulting in degraded or ungrammatical parses that are too costly to repair.

4.3 Is Copy α (Frazier & Clifton, 2001) Cost-Free?

Frazier and Clifton (2001) suggested that parallel ellipsis constructions were comprehended by copying the underlying structure of the antecedent into the ellipsis site. They proposed that this procedure involved a cost-free copy mechanism, which they dubbed Copy α , which shares features of the copy mechanisms described in Section 1.3.1.2. To support their hypothesis, they showed that the size and/or complexity of the antecedent did not increase the length of time it took to read the ellipsis region of a sentence. Arregui et al. (2006) assumed that this hypothesis is true, and did not show how their antecedents behave under parallel ellipsis conditions (because their manipulations do not lend themselves to such an investigation). However, they state that the processor attempts to create an antecedent only when it cannot find an already matching one (p. 242). This means that when there is a matching antecedent available, the processor will always choose to copy it into the ellipsis site over attempting to build a different one. The results from two online survey experiments reported in Chapter 3 contradict this hypothesis. By using the same modal verb (i.e., *would*) in both the antecedent and ellipsis clauses on half of the trials, the same antecedents could be judged under both parallel and non-parallel ellipsis conditions. Under the hypothesis that the content of the antecedent does not affect comprehension of parallel ellipsis, one would predict that there would be no effect the verbal morphology manipulation under conditions of parallel ellipsis. However, Experiment 1A and Experiment 1B showed that – even under conditions of parallel ellipsis – active morphology offered a better antecedent than infinitive morphology, and both were better than progressive morphology. The self-paced reading experiment took a slight departure from this pattern, showing that elided verbs active and infinitive morphology were equally acceptable, but remained consistent in showing that elided verbs with progressive morphology were the worst overall.

These results speak directly to the validity of the Copy α mechanism (Frazier & Clifton, 2001) and call into question Arregui et al.'s (2006) claims that recycling is a last resort. With regards to Copy α , these results suggest that it is at least sometimes sensitive to the material inside the antecedent. Assuming that the mechanism exists, this is the only way to explain the current findings. Frazier and Clifton (2001) manipulated the size of the antecedent under the assumption that larger antecedents would be more difficult to copy into the ellipsis site, and found no difference between smaller and larger antecedents. If we wish to hold on to the Copy α notion, we must conclude that it

may not be sensitive to the size of the antecedent to be copied, but is quite sensitive to the morphological content of the antecedent. However, it is unclear whether we should maintain the Copy α mechanism as an explanatory tool. Under the Recycling Hypothesis, parallel and non-parallel antecedents are processed in inherently different ways. Parallel antecedents are copied into the ellipsis site, while non-parallel antecedents are rebuilt from the material that is present in the antecedent. Arregui et al. (2006) further claim that the recycling processes used for non-parallel antecedents are only performed when no parallel antecedent can be located. The results reported above suggest something quite different: the acceptability of both parallel and non-parallel ellipsis constructions was affected by variations in verbal morphology. Both forms of ellipsis showed the exact same pattern of results, which suggests that parallel and non-parallel antecedents are in fact processed in much the same way. So, do parallel antecedents undergo the same recycling processes as non-parallel antecedents? Or are non-parallel antecedents somehow copied into the ellipsis site the way parallel antecedents are? Neither option seems particularly appealing. Although adopting a modified version of the Copy α mechanism might explain the behaviour of the minimally different antecedents used in the current experiments, it is less clear how it would account for the behaviour of the maximally different antecedents used by Arregui and her colleagues. Similarly, modifying the Recycling Hypothesis to explain parallel ellipsis comprehension might account for the current results, but it would not be able to account for Frazier and Clifton's (2001) results as larger syntactic structures should indeed take longer to rebuild than smaller structures. How do we comprehend ellipsis, then? Section 4.5 considers two possible alternatives.

It is also important to note that the findings reported here are consistent with the idea that there may be an overall preference for parallel ellipsis (e.g., Arregui et al., 2006; Mauner et al., 1995; Tanenhaus & Carlson, 1990; Dickey & Bunker, 2011). One possible explanation for this parallelism preference is that the comprehension of parallel ellipsis sites is facilitated through syntactic priming. Syntactic priming is a phenomenon showing that recent experience with a particular syntactic structure facilitates current processing of an identical structure. Bock (1986) showed that participants were more likely to produce a passive sentence structure to describe a picture after having recently repeated a passive sentence than if they had recently repeated an active sentence. Similarly, participants were more likely to use double object constructions after repeating a double object construction compared to when they repeated alternatives with prepositional phrases (Bock, 1986; Bresnan, Cueni, Nikitina, & Baayen, 2007). Branigan, Pickering, and McLean (2005) further showed that syntactic priming holds for comprehension as well. There is potential for extending this hypothesis to the ellipsis domain – repetition of parallel verb structure should facilitate processing of the elided construct. Furthermore, the expectation of parallelism built by conjunctions (Shapiro & Hestvik, 1995) might strengthen this priming effect in ellipsis. Thus, it may prove beneficial to investigate the link between syntactic priming and ellipsis comprehension.

To summarize, the findings reported in this thesis showed that parallel and non-parallel forms of ellipsis are similarly sensitive to variations in verbal morphology. This suggests that both types are processed in similar ways, leading us to reject Copy α (as

well as the other copy theories in Section 1.3.1.2) as an explanation for parallel ellipsis and the Recycling Hypothesis as an explanation for non-parallel ellipsis. Nevertheless, there appears to be a preference for parallel over non-parallel ellipsis, which might be explained by syntactic priming.

4.4 Do Real-Time Comprehension Measures Show Effects of Recycling?

Arregui et al. (2006) provided no real-time measures of reading comprehension to support their Recycling Hypothesis. Word-by-word reading time measures reflect the amount of effort involved in integrating a word into a sentence (Gibson, 1998), such that more difficult integration processes are associated with longer reading times. If the Recycling Hypothesis is correct, reading times should increase for ellipsis constructions as the complexity of the operations required to build the ellipsis site increase. Experiment 2 tested this prediction, and found no evidence to support it. Reading times at the ellipsis site did not differ as a function of antecedent complexity. The fact that there were no differences between conditions suggests that the processor may not rebuild the ellipsis site. If this is the case, we must reject the Recycling Hypothesis. However, this conclusion is speculative as it relies on a null result that may be explained by other factors.

First, it is possible that the effect of morphology found here is too small to be observed in reading times. The sentences used in these experiments were minimally different from each other; these small differences may produce an effect that is too small to observe. However, this seems somewhat unlikely given the robustness of the acceptability judgment responses. Alternatively, it is possible that the measures used in this experiment are not sensitive to the verbal morphology manipulation. While reading time measures generated from button presses is a standard measure of reading difficulty, it is possible that other measures may be more sensitive to the manipulations reported above. As such, it may prove beneficial to perform similar experiments using other techniques, such as eye-tracking methodology which records regressive fixations (a likely index of the difficulty of reconstructing elided structures). Finally, the null result may reflect methodological complications inherent in the experiment. The primary region of interest was the word that represented the ellipsis site, which was always the final word of the sentence. Evidence suggests that words are fixated longer when they end a sentence compared to when they are in the middle of a sentence, suggesting that there may be something special about sentence-final semantic integration (Kuperman, Dambacher, Nuthmann, & Kliegl, 2010; Rayner, Sereno, Morris, Shmauder, & Clifton, 1989). The results from this experiment showed particularly high variability in reading times at this location, so it is possible that the expected verbal morphology effect was masked by any sentence wrap-up processes that the participants might have been engaged in.

4.5 Which theory is best?

The two previous sections offered evidence against the Recycling Hypothesis. Section 4.3 argued that the results from the experiments described in Chapter 3 suggest that parallel and non-parallel ellipsis constructions are comprehended in a similar manner. It proved impossible to adopt a modified version of Frazier and Clifton's (2001) Copy α mechanism to explain non-parallel ellipsis comprehension, because it would not account for the results found by Arregui et al. (2006). Similarly, it was impossible to adopt a modified version of the Recycling Hypothesis that would apply to parallel ellipsis comprehension, as this hypothesis would predict that larger syntactic structures would take longer to rebuild, contrary to evidence from Frazier and Clifton (2001). Section 4.4 further argued that the experiments reported in Chapter 3 found no on-line evidence to support the hypothesis that the processor rebuilds the ellipsis structure at the ellipsis site. Therefore, it seems unlikely that the Recycling Hypothesis is the correct theory for describing ellipsis comprehension. The final section in this chapter considers whether the other theories described in Section 2.3 might account for the current findings.

Kim et al. (2011) proposed that verb-phrase ellipsis comprehension is governed by a memory search process that is guided by grammatically motivated heuristics, MaxElide and Canonical Representation. MaxElide favours parallel ellipsis constructions because these constructions permit the deletion of the largest possible constituent. Thus, this theory predicts that the parallel ellipsis sentences in this experiment will be more acceptable than the non-parallel ellipsis sentences. On the other hand, Canonical Representation favours sentences that conform to standard SVO word order (in English). This heuristic makes no predictions for the stimuli used in the current experiment, as all sentences conformed to standard English word order. The researchers argued that the parser uses these heuristics to constrain the search space to only those structures that conform to the canonical word order and permit the deletion of the largest possible constituent. For the stimuli used in these experiments, the parser would constrain the search space to only those sentences that permit deletion of the largest constituent. It is not entirely clear whether this theory can account for the current results. The predictions derived from the MaxElide principle appear to be borne out: the results showed that parallel ellipsis was more acceptable than non-parallel ellipsis. As Canonical Representation makes no predictions for these experiments, the results neither support nor refute the existence of this principle. Most importantly, Kim et al.'s proposal makes absolutely no predictions about the morphology manipulation used in these experiments. Therefore, it is unclear whether this theory can account for the current findings. For this model to work, we would have to adopt a grammatically motivated heuristic that identifies the preferred verbal morphology. Furthermore, this heuristic would necessarily apply to elided and non-elided sentences alike (because verbal morphology is present in both elided and non-elided sentences). Such a heuristic seems unlikely, given the results reported here, which showed that elided sentences were sensitive to changes in verbal morphology (Experiments 1A and 1B) but non-elided sentences were not (Experiment 1C). Therefore, it seems unlikely that Kim et al.'s memory search model of ellipsis comprehension could

adequately explain the current results.

Martin and McElree (2008; 2009; 2011) provided a more promising proposal in which they claimed that ellipsis comprehension involves directly accessing antecedents from memory based on the cues provided at the ellipsis site. These researchers directly reject the possibility of a search process by showing that various complexity manipulations consistently affect the accuracy of retrieval, but never the speed of retrieval. The current experiments employed a morphological complexity manipulation; thus, Martin and McElree would predict that this manipulation should affect the accuracy of retrieval, but not the speed of retrieval. The results from the current set of experiments demonstrate exactly this pattern of results. The verbal morphology manipulation affected acceptability ratings (read: accuracy of retrieval) of elided constructions in Experiments 1A, 1B, and 2. However, there was no evidence that the speed of retrieval was affected: Experiment 2 showed that there was no difference in judgment speed or reading times at the ellipsis site. Thus, the results from the current experiments are compatible with Martin and McElree's model of ellipsis comprehension. Given that the current findings serve to reject the Recycling Hypothesis, and do not seem to be compatible with Kim et al.'s (2011) memory search model, it seems most reasonable to adopt Martin and McElree's direct-access memory retrieval model.

Appendix A

Materials used in the reported experiments. Note: Experiment 1A did not make use of the context sentences, and Experiment 1C used the non-elided versions. The non-elided structure is included in parentheses at the end of each sentence.

(1) Set 1

Context: The local newspaper reported that a comet might pass through tonight's meteor shower.

- a. None of the astronomers would see the comet, but John did (see the comet).
- b. None of the astronomers would get to see the comet, but John did (see the comet).
- c. None of the astronomers would be seeing the comet, but John did (see the comet).
- d. None of the astronomers would see the comet, but John would (see the comet).
- e. None of the astronomers would get to see the comet, but John would (get to see the comet).
- f. None of the astronomers would be seeing the comet, but John would (be seeing the comet).

(2) Set 2

Context: Across town, the police rushed to the scene of the crime.

- a. Few people would hear the distant siren, but Jane did (hear the distant siren).
- b. Few people would get to hear the distant siren, but Jane did (hear the distant siren).
- c. Few people would be hearing the distant siren, but Jane did (hear the distant siren).
- d. Few people would hear the distant siren, but Jane would (hear the distant siren).
- e. Few people would get to hear the distant siren, but Jane would (get to hear the distant siren).
- f. Few people would be hearing the distant siren, but Jane would (be hearing the distant siren).

(3) Set 3

Context: The trapped man covered his face from the flames and waved his flashlight in order to be found.

- a. No firement would detect the faint signal, but the dogs did (detect the faint signal).
- b. No firement would get to detect the faint signal, but the dogs did (detect the faint signal).
- c. No firement would be detecting the faint signal, but the dogs did (detect the faint signal).
- d. No firement would detect the faint signal, but the dogs would (detect the faint signal).
- e. No firement would get to detect the faint signal, but the dogs would (get to detect the faint signal).
- f. No firement would be detecting the faint signal, but the dogs would (be detecting the faint signal).

(4) Set 4

Context: In between performances, the lion sat on the platform outside of his cage.

- a. Almost nobody would approach the lion, but the trainer did (approach the lion).
- b. Almost nobody would get to approach the lion, but the trainer did (approach the lion).
- c. Almost nobody would be approaching the lion, but the trainer did (approach the lion).
- d. Almost nobody would approach the lion, but the trainer would (approach the lion).
- e. Almost nobody would get to approach the lion, but the trainer would (get to approach the lion).
- f. Almost nobody would be approaching the lion, but the trainer would (be approaching the lion).

(5) Set 5

Context: The teacher walked between the desks during the math test.

- a. Few students would solve the calculus problems, but Erica did (solve the calculus problems).
- b. Few students would get to solve the calculus problems, but Erica did (solve the calculus problems).
- c. Few students would be solving the calculus problems, but Erica did (solve the calculus problems).
- d. Few students would solve the calculus problems, but Erica would (solve the calculus problems).
- e. Few students would get to solve the calculus problems, but Erica would (get to solve the calculus problems).

- f. Few students would be solving the calculus problems, but Erica would (be solving the calculus problems).

(6) Set 6

Context: Today, the restaurant staff met with the new hotel management team.

- a. Almost nobody would like the owner, but the waitress did (like the owner).
- b. Almost nobody would get to like the owner, but the waitress did (like the owner).
- c. Almost nobody would be liking the owner, but the waitress did (like the owner).
- d. Almost nobody would like the owner, but the waitress would (like the owner).
- e. Almost nobody would get to like the owner, but the waitress would (get to like the owner).
- f. Almost nobody would be liking the owner, but the waitress would (be liking the owner).

(7) Set 7

Context: Many tourists were seen exploring the waterfront walking trail.

- a. Very few visitors would find Princess Point, but Pierre did (find Princess Point).
- b. Very few visitors would get to find Princess Point, but Pierre did (find Princess Point).
- c. Very few visitors would be finding Princess Point, but Pierre did (find Princess Point).
- d. Very few visitors would find Princess Point, but Pierre would (find Princess Point).
- e. Very few visitors would get to find Princess Point, but Pierre would (get to find Princess Point).
- f. Very few visitors would be finding Princess Point, but Pierre would (be finding Princess Point).

(8) Set 8

Context: Jim was talking loudly and making crude jokes.

- a. Few people would tolerate Jim's antics, but Sally did (tolerate Jim's antics).
- b. Few people would get to tolerate Jim's antics, but Sally did (tolerate Jim's antics).
- c. Few people would be tolerating Jim's antics, but Sally did (tolerate Jim's antics).
- d. Few people would tolerate Jim's antics, but Sally would (tolerate Jim's antics).
- e. Few people would get to tolerate Jim's antics, but Sally would (get to tolerate Jim's antics).
- f. Few people would be tolerating Jim's antics, but Sally would (be tolerating Jim's antics).

(9) Set 9

Context: The scraggly dog with the missing eye was named Fiona.

- a. Few people would love Fiona, but Max did (love Fiona).
- b. Few people would get to love Fiona, but Max did (love Fiona).
- c. Few people would be loving Fiona, but Max did (love Fiona).
- d. Few people would love Fiona, but Max would (love Fiona).
- e. Few people would get to love Fiona, but Max would (get to love Fiona).
- f. Few people would be loving Fiona, but Max would (be loving Fiona).

(10) Set 10

Context: This month, Joe defended his title of heavyweight champion.

- a. Few wrestlers would beat Joe, but Sam did (beat Joe).
- b. Few wrestlers would get to beat Joe, but Sam did (beat Joe).
- c. Few wrestlers would be beating Joe, but Sam did (beat Joe).
- d. Few wrestlers would beat Joe, but Sam would (beat Joe).
- e. Few wrestlers would get to beat Joe, but Sam would (get to beat Joe).
- f. Few wrestlers would be beating Joe, but Sam would (be beating Joe).

(11) Set 11

Context: A tornado touched down near the large forested area this morning.

- a. Almost nobody would predict the direction of the tornado, but Channel 7 did (predict the direction of the tornado).
- b. Almost nobody would get to predict the direction of the tornado, but Channel 7 did (predict the direction of the tornado).
- c. Almost nobody would be predicting the direction of the tornado, but Channel 7 did (predict the direction of the tornado).
- d. Almost nobody would predict the direction of the tornado, but Channel 7 would (predict the direction of the tornado).
- e. Almost nobody would get to predict the direction of the tornado, but Channel 7 would (get to predict the direction of the tornado).
- f. Almost nobody would be predicting the direction of the tornado, but Channel 7 would (be predicting the direction of the tornado).

(12) Set 12

Context: Timothy dropped his dirty gym socks behind the couch.

- a. Almost nobody would identify the source of the odor, but Lisa did (identify the source of the odor).
- b. Almost nobody would get to identify the source of the odor, but Lisa did (identify the source of the odor).
- c. Almost nobody would be identifying the source of the odor, but Lisa did (identify the source of the odor).
- d. Almost nobody would identify the source of the odor, but Lisa would (identify the source of the odor).

- e. Almost nobody would get to identify the source of the odor, but Lisa would (get to identify the source of the odor).
- f. Almost nobody would be identifying the source of the odor, but Lisa would (be identifying the source of the odor).

(13) Set 13

Context: The boss was caught having an affair with his assistant.

- a. Few people would mention the scandal, but Tom did (mention the scandal).
- b. Few people would get to mention the scandal, but Tom did (mention the scandal).
- c. Few people would be mentioning the scandal, but Tom did (mention the scandal).
- d. Few people would mention the scandal, but Tom would (mention the scandal).
- e. Few people would get to mention the scandal, but Tom would (get to mention the scandal).
- f. Few people would be mentioning the scandal, but Tom would (be mentioning the scandal).

(14) Set 14

Context: A journalist attempted to uncover illegal practices in the corporate office.

- a. Almost no one would notice the hidden camera, but the secretary did (notice the hidden camera).
- b. Almost no one would get to notice the hidden camera, but the secretary did (notice the hidden camera).
- c. Almost no one would be noticing the hidden camera, but the secretary did (notice the hidden camera).
- d. Almost no one would notice the hidden camera, but the secretary would (notice the hidden camera).
- e. Almost no one would get to notice the hidden camera, but the secretary would (get to notice the hidden camera).
- f. Almost no one would be noticing the hidden camera, but the secretary would (be noticing the hidden camera).

(15) Set 15

Context: Fifty people applied for the open salesperson position.

- a. Almost nobody would verify the applicants' credentials, but Higgins did (verify the applicants' credentials).
- b. Almost nobody would get to verify the applicants' credentials, but Higgins did (verify the applicants' credentials).
- c. Almost nobody would be verifying the applicants' credentials, but Higgins did (verify the applicants' credentials).
- d. Almost nobody would verify the applicants' credentials, but Higgins would (verify the applicants' credentials).

- e. Almost nobody would get to verify the applicants' credentials, but Higgins would (get to verify the applicants' credentials).
- f. Almost nobody would be verifying the applicants' credentials, but Higgins would (be verifying the applicants' credentials).

(16) Set 16

Context: City council ignored the residents who complained about the increasing number of traffic accidents.

- a. Few people would think about the scale of the problem, but Ben did (think about the scale of the problem).
- b. Few people would get to think about the scale of the problem, but Ben did (think about the scale of the problem).
- c. Few people would be thinking about the scale of the problem, but Ben did (think about the scale of the problem).
- d. Few people would think about the scale of the problem, but Ben would (think about the scale of the problem).
- e. Few people would get to think about the scale of the problem, but Ben would (get to think about the scale of the problem).
- f. Few people would be thinking about the scale of the problem, but Ben would (be thinking about the scale of the problem).

Bibliography

- Arregui, A., Clifton Jr, C., Frazier, L., and Moulton, K. (2006). Processing elided verb phrases with flawed antecedents: The recycling hypothesis. *Journal of Memory and Language*, 55:232–246.
- Baltin, M. (2012). Deletion versus pro-forms: An overly simple dichotomy? *Natural Language and Linguistic Theory*, 30:381–423.
- Bjorkman, B. (2011). *BE-ing Default: The Morphosyntax of Auxiliaries*. PhD thesis, Massachusetts Institute of Technology.
- Bock, K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, 18:355–387.
- Branigan, H. (2007). Syntactic priming. *Language and Linguistics Compass*, 1:1–16.
- Branigan, H., Pickering, M. J., and McLean, J. F. (2005). Priming prepositional-phrase attachment during comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31:468–481.
- Bresnan, J., Cueni, A., Nikitina, T., and Baayen, H. (2007). Predicting the dative alternation. In Boume, G., Kraemer, I., and Zwarts, J., editors, *Cognitive Foundations of Interpretation*, pages 69–94. Royal Netherlands Academy of Science, Amsterdam.
- Callahan, S., Shapiro, L., and Love, T. (2010). Parallelism effects and verb activation: The sustained reactivation hypothesis. *Journal of Psycholinguistic Research*, 39:101–118.
- Carlson, K. (2001). The effects of parallelism and prosody in the processing of gapping structures. *Language and Speech*, 44:1–26.
- Carlson, K., Dickey, M. W., Frazier, L., and Clifton Jr, C. (2009). Information structure expectations in sentence comprehension. *The Quarterly Journal of Experimental Psychology*, 62:114–139.
- Chomsky, N. (1957). *Syntactic Structures*. Mouton de Gruyter, Berlin, Germany.
- Chomsky, N. (1995). *The Minimalist Program*. MIT Press, Cambridge, MA.
- Chung, S., Ladusaw, W., and McCloskey, J. (1995). Sluicing and logical form. *Natural Language Semantics*, 3:239–282.

- Collins, C. (2005). A smuggling approach to the passive in English. *Syntax*, 8:81–120.
- Culicover, P. and Jackendoff, R. (2005). *Simpler Syntax*. Oxford University Press, Oxford.
- Cutler, A., Oahan, D., and van Donselaar, W. (1997). Prosody in the comprehension of spoken language: A literature review. *Language and Speech*, 40:141–201.
- Dalrymple, M., Shieber, S., and Pereira, F. (1991). Ellipsis and higher-order unification. *Linguistics and Philosophy*, 14:399–452.
- Dickey, M. W. and Bungler, A. C. (2011). Comprehension of elided structure: Evidence from sluicing. *Language and Cognitive Processes*, 26:63–78.
- Earley, J. (1970). An efficient context-free parsing algorithm. *Communications of the Association for Computing Machinery*, 13:94–102.
- Ellis, N. (2002). Frequency effects in language processing. *Studies in Second Language Acquisition*, 24:143–188.
- Embick, D. and Noyer, R. (2006). Distributed morphology and the syntax/morphology interface. In Ramchand, G. and Reiss, C., editors, *The Oxford Handbook of Linguistic Interfaces*, pages 289–324. Oxford University Press, Oxford.
- Fiengo, R. and May, R. (1994). *Indices and Identity*. MIT Press, Cambridge, MA.
- Fox, D. (1999). Focus, parallelism, and accommodation. In Matthews, T. and Strolovitch, D., editors, *Proceedings of the 9th Semantics and Linguistics Theory Conference*, pages 70–90. University of California, Santa Cruz.
- Frazier, L. and Clifton Jr, C. (2000). On bound variable interpretations: The lf-only hypothesis. *Journal of Psycholinguistic Research*, 29:125–140.
- Frazier, L. and Clifton Jr, C. (2001). Parsing coordinates and ellipsis: Copy α . *Syntax*, 4:1–22.
- Frazier, L. and Clifton Jr, C. (2005). The syntax-discourse divide: Processing ellipsis. *Syntax*, 8:121–174.
- Frazier, L. and Clifton Jr, C. (2006). Ellipsis and discourse coherence. *Linguistics and Philosophy*, 29:315–346.
- Frazier, L., Clifton Jr, C., and Carlson, K. (2007). Focus and VP ellipsis. *Language and Speech*, 50:1–21.
- Garnham, A. and Oakhill, J. (1987). Interpreting elliptical verb phrases. *Quarterly Journal of Experimental Psychology*, 39A:611–627.
- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68:1–76.

- Gibson, E., Piantadosi, S., and Fedorenko, K. (2011). Using Mechanical Turk to obtain and analyze English acceptability judgments. *Language and Linguistics Compass*, 5:509–524.
- Hale, J. (2001). A probabilistic early parser as a psycholinguistic model. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Halle, M. and Marantz, A. (1993). Distributed Morphology and the pieces of inflection. In Hale, K. and Keyser, S. J., editors, *The View From Building 20: Essays in Linguistics in Honor of Sylvain Bromberger*. MIT Press, Cambridge, MA.
- Hankamer, J. and Sag, I. (1976). Deep and surface anaphora. *Linguistic Inquiry*, 7:391–428.
- Hardt, D. (1999). Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy*, 22:185–219.
- Hyönä, J. and Vainio, S. (2001). Reading morphologically complex clause structures in Finnish. *European Journal of Cognitive Psychology*, 13:451–474.
- Hyönä, J., Vainio, S., and Laine, M. (2002). A morphological effect obtains for isolated words but not for words in a sentence context. *European Journal of Cognitive Psychology*, 14:417–433.
- Jaeger, F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59:434–446.
- Kehler, A. (2000). Coherence and the resolution of ellipsis. *Linguistics and Philosophy*, 23:533–575.
- Keysar, B. (2007). Communication and miscommunication: The role of egocentric processes. *Intercultural Pragmatics*, 4:71–84.
- Kim, C., Kobele, G., Runner, J., and Hale, J. (2011). The acceptability cline in VP ellipsis. *Syntax*, 14:318–354.
- Kuperman, V., Dambacher, M., Nuthmann, A., and Kliegl, R. (2010). The effect of word position on eye-movements in sentence and paragraph reading. *The Quarterly Journal of Experimental Psychology*, 63:1838–1857.
- Lasnik, H. (1995). Verbal morphology: Syntactic structures meets the Minimalist Program. In Campos, H. and Kempchinsky, P., editors, *Evolution and Revolution in Linguistic Theory: Essays in Honor of Carlos Otero*, pages 251–275. Georgetown University Press, Washington, DC.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106:1126–1177.
- Lewis, R. and Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29:375–419.

- Lewis, R., Vasishth, S., and Dyke, J. V. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Sciences*, 10:447–454.
- Marantz, A. (1997). No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. In Dimitriadis, A., Siegel, L., Surek-Clark, C., and Williams, A., editors, *University of Pennsylvania Working Papers in Linguistics*, volume 4.2, pages 201–225. University of Pennsylvania.
- Marslen-Wilson, W. (1975). Sentence perception as an interactive parallel process. *Science*, 18:226–228.
- Martin, A. and McElree, A. (2008). A content-addressable pointer mechanism underlies comprehension of verb-phrase ellipsis. *Journal of Memory and Language*, 58:879–906.
- Martin, A. and McElree, B. (2009). Memory operations that support language comprehension: Evidence from verb-phrase ellipsis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35:1231–1239.
- Martin, A. and McElree, B. (2011). Direct-access retrieval during sentence comprehension: Evidence from sluicing. *Journal of Memory and Language*, 64:327–343.
- Martin, J. G. (1968). Temporal word spacing and the perception of ordinary, anomalous, and scrambled strings. *Journal of Verbal Learning and Verbal Behaviour*, 7:154–157.
- Mason, W. and Suri, S. (2012). Conducting behavioural research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44:1–23.
- Maurer, G., Tanenhaus, M. K., and Carlson, G. N. (1995). A note on parallelism effects in processing deep and surface verb-phrase anaphora. *Language and Cognitive Processes*, 10:1–12.
- McClelland, J. and Rumelhart, D. (1986). *Parallel Distributed Processing*. MIT Press, Cambridge, MA.
- McElree, B. (2000). Sentence comprehension is mediated by content-addressable memory structures. *Journal of Psycholinguistic Research*, 29:536–571.
- McElree, B. and Doshier, B. A. (1993). Serial retrieval processes in the recovery of order information. *Journal of Experimental Psychology: General*, 122:291–315.
- McElree, B., Foraker, S., and Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, 48:67–91.
- Merchant, J. (2001). *The Syntax of Silence: Sluicing, Islands, and the Theory of Ellipsis*. Oxford University Press, Oxford.
- Merchant, J. (2008a). An asymmetry in voice mismatches in VP-ellipsis and pseudogapping. *Linguistic Inquiry*, 39:169–179.

- Merchant, J. (2008b). Variable island repair under ellipsis. In Johnson, K., editor, *Topics in Ellipsis*, pages 132–153. Cambridge University Press, Cambridge.
- Merchant, J. (2012). Voice and ellipsis. Unpublished manuscript, University of Chicago, Chicago, IL.
- Merchant, J. (to appear). Ellipsis. In Alexiadou, A., Kiss, T., and Butt, M., editors, *Handbook of Contemporary Syntax*. Mouton de Gruyter, Berlin, Germany, 2 edition.
- Miller, G. A. and Isard, S. D. (1963). Some perceptual consequences of linguistic rules. *Journal of Verbal Learning and Verbal Behaviour*, 2:217–228.
- Munro, R., Bethard, S., Kuperman, V., Lai, V. T., Melnick, R., Potts, C., Schnoebelen, T., and Tily, H. (2010). Crowdsourcing and language studies: The new generation of linguistic data. In *Proceedings of the NAACL Workshop on Creating Speech and Language Data With Amazon's Mechanical Turk*, pages 122–130.
- Murphy, G. (1985). Processes of understanding anaphora. *Journal of Memory and Language*, 24:290–303.
- Murphy, G. (1990). Interpretation of verb phrase anaphora: Influences of task and syntactic context. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology*, 42:675–692.
- Nicol, J. and Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research*, 18:5–19.
- Omaki, A. (2007). Revisiting revived syntactic structures: The extended hybrid approach to verbal morphology. In Omaki, A., Ortega-Santos, I., Sprouse, J., and Wagers, M., editors, *University of Maryland Working Papers in Linguistics*, volume 16, pages 89–110, College Park, MD.
- Omaki, A. (2009). Verbal morphology: Return of the affix hopping approach. In *Proceedings of NELS 38*, volume 2, pages 193–204, Amherst, MA.
- Parker, D. and Seely, D. MaxElide and its domain of application. Talk given at the 85th Annual Meeting of the Linguistic Society of America, Baltimore, Maryland.
- Poirier, J., Walenski, M., and Shapiro, L. (2012). The role of parallelism in the real-time processing of anaphora. *Language and Cognitive Processes*, 27:868–886.
- Poirier, J., Wolfinger, K., Spellman, L., and Shapiro, L. (2010). The real-time processing of sluiced sentences. *Journal of Psycholinguistic Research*, 39:411–427.
- Potsdam, E. (1997). English verbal morphology and VP ellipsis. In *Meeting of the North Eastern Linguistics Society*. GLSA.
- Rayner, K., Sereno, S. C., Morris, R. K., Schmauder, A. R., and Clifton Jr, C. (1989). Eye movements and on-line language comprehension processes. *Language and Cognitive Processes*, 4:S121–S149.

- Rooth, M. (1992). Ellipsis redundancy and reduction redundancy. In an A Hestvik, S. B., editor, *Proceedings of the Stuttgarter Ellipsis Workshop*.
- Ross, J. (1967). *Constraints on Variables in Syntax*. PhD thesis, Massachusetts Institute of Technology.
- Ross, J. (1969). Guess who? In Binnick, R., Davison, A., Green, G., and Morgan, J., editors, *Papers from the 5th Regional Meeting of the Chicago Linguistic Society*, pages 252–286, Chicago. Chicago Linguistic Society.
- Sag, I. (1976). *Deletion and Logical Form*. PhD thesis, Massachusetts Institute of Technology.
- Sag, I. and Hankamer, J. (1984). Toward a theory of anaphoric processing. *Linguistics and Philosophy*, 7:325–345.
- Schnoebelen, T. and Kuperman, V. (2010). Using Amazon Mechanical Turk for linguistic research. *Psihologija*, 43:441–464.
- Schwarzschild, R. (1999). Givenness, AvoidF, and other constraints on the placement of accent. *Natural Language Semantics*, 7:141–177.
- Shapiro, L. and Hestvik, A. (1995). On-line comprehension of VP-ellipsis: Syntactic reconstruction and semantic influence. *Journal of Psycholinguistic Research*, 24:517–532.
- Shapiro, L. P., Hestvik, A., Lesan, L., and Garcia, A. R. (2003). Charting the time-course of VP-ellipsis sentence comprehension: Evidence for an initial and independent structural analysis. *Journal of Memory and Language*, 49:1–19.
- Snow, R., O’Connor, B., Jurafsky, D., and Ng, A. Y. (2008). Cheap and fast – but is it good? Evaluating non-expert annotations for natural language tasks. In *Proceedings of EMNLP*, pages 254–263.
- Sprouse, J. (2011). A validation of Amazon Mechanical Turk for the collection of acceptability judgments in linguistic theory. *Behavior Research Methods*, 43:155–167.
- Stalnaker, R. (1970). Pragmatics. *Synthese*, 22:272–289.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25:701–721.
- Takahashi, S. and Fox, D. (2005). MaxElide and the rebinding problem. In Georgala, E. and Howell, J., editors, *Proceedings of SALT XV*, pages 223–240, Ithica, NY. Cornell University.
- Tanaka, H. (2011). Voice mismatch and syntactic identity. *Linguistic Inquiry*, 42:470–490.
- Tancredi, C. (1992). *Deletion, Deaccenting, and Presupposition*. PhD thesis, Massachusetts Institute of Technology.

Tanenhaus, M. K. and Carlson, G. N. (1990). Comprehension of deep and surface verbphrase anaphors. *Language and Cognitive Processes*, 5:257–280.

von Stechow, P. (2000). What is presupposition accommodation? Unpublished manuscript, Massachusetts Institute of Technology.

von Stechow, P. (2008). What is presupposition accommodation, again? *Philosophical Perspectives*, 22:137–170.

Williams, E. (1977). Discourse and logical form. *Linguistic Inquiry*, 8:103–139.