

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

Comprehending Ellipsis

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Linguistics

by

Margaret Kroll

June 2020

The Dissertation of Margaret Kroll
is approved:

Associate Professor Matt Wagers, chair

Professor Pranav Anand

Professor Adrian Brasoveanu

Assistant Professor Amanda Rysling

Quentin Williams
Acting Vice Provost and Dean of Graduate Studies

Copyright © by

Margaret Kroll

2020

Contents

Contents	iii
Acknowledgments	ix
1 Introduction	1
2 The pragmatics of sluicing	6
2.1 Introduction	6
2.1.1 Polarity reversals in sluicing	6
2.1.2 Methodological preliminaries	8
2.2 Sluicing in the literature	10
2.2.1 e-GIVENness	11
2.2.2 Hybrid Theories	13
2.2.3 Inquisitive entailment	13
2.2.4 Scopability	15
2.3 A modified account	16
2.3.1 The Well-Formedness Condition	16
2.3.2 Local Givenness	18
2.3.2.1 Context update and dynamic interpretation systems . .	18
2.3.2.2 Local Givenness	22
2.3.2.3 Application of Local Givenness	23
2.4 Deriving polarity reversal sluices	27
2.4.1 Polarity reversals are not semantic entailment: Neg-raising polar- ity reversals	28
2.4.2 Polarity reversals are not syntactic: <i>Remember</i> polarity reversals	33
2.4.3 Polarity reversals are not bidirectional pragmatic entailment: Disjunction polarity reversals	37
2.4.4 Polarity reversals are not entailment at LF: More disjunction . .	40
2.4.5 Non-factive negative attitude verbs	43
2.4.6 <i>Until</i>	45
2.4.7 Attitude holders	47
2.4.8 Failure to license	48
2.5 Structural constraints on sluicing	49

2.5.1	Data motivating syntactic constraints on sluicing	50
2.5.2	Local Givenness plus structural constraints	51
2.5.3	Structure plus pragmatics: A full picture?	54
2.5.4	Saliency	60
2.6	Non-isomorphic sluicing	66
2.6.1	Types of non-isomorphic sluices	67
2.6.2	Polarity reversals cannot be reduced to non-isomorphic sluices .	69
2.6.3	Summary	72
2.7	Conclusion	72
3	Violating Perspicuity	75
3.1	Processing anaphoric dependencies	76
3.1.1	Retrieval in anaphoric processing	77
3.1.2	Locality in anaphoric processing	83
3.2	Choosing an empirical domain	90
3.3	Experiment 1: Ellipsis interpretation shows a subjecthood preference . .	95
3.3.1	Method	96
3.3.1.1	Methodology	96
3.3.1.2	Participants	96
3.3.1.3	Materials	96
3.3.2	Procedure	101
3.3.3	Results	104
3.3.3.1	All Results	104
3.3.3.2	Subject vs. Object Responses	105
3.3.4	Discussion	106
3.3.4.1	Plural nouns make better antecedents	107
3.3.4.2	Subject nouns were preferred antecedents	109
3.3.4.3	Coordinated antecedents look like split antecedents . .	111
3.4	Experiment 2: Ellipsis interpretation shows a locality preference	112
3.4.1	Method	113
3.4.1.1	Methodology	113
3.4.1.2	Participants	113
3.4.1.3	Materials	113
3.4.2	Procedure	115
3.4.3	Results	116
3.4.3.1	First NP Analysis	119
3.4.3.2	Second NP Analysis	121
3.4.3.3	Kind & Conjoined NP Analysis	122
3.4.4	Discussion	125
3.4.4.1	Results show a proximity preference	125
3.4.4.2	Results show a null effect of structure	129
3.4.4.3	Results show overall preference for Kind interpretations	130
3.5	Experiment 2 items are moderately acceptable	133

3.5.1	Method	134
3.5.1.1	Participants	134
3.5.1.2	Materials	134
3.5.1.3	Procedure	135
3.5.1.4	Results	135
3.5.1.5	Discussion	137
4	Online ellipsis comprehension	140
4.1	Cataphoric ellipsis shows an online plausibility effect in the Maze task .	146
4.1.1	Method	146
4.1.1.1	Methodology	146
4.1.1.2	Participants	149
4.1.1.3	Materials	149
4.1.2	Procedure	152
4.1.3	Results	153
4.1.4	Discussion	158
4.1.4.1	Implausible nouns were read more slowly than plausible nouns	158
4.1.4.2	Indexing processing difficulty in the maze task	164
4.2	Self-paced reading does not show a plausibility effect	168
4.2.1	Method	169
4.2.1.1	Methodology	169
4.2.1.2	Materials	170
4.2.1.3	Participants	170
4.2.2	Procedure	170
4.2.3	Results	171
4.2.4	Discussion	173
4.3	Semantic relatedness drives reading times: Revisiting the Maze	175
4.3.1	Method	176
4.3.1.1	Methodology	176
4.3.1.2	Participants	176
4.3.1.3	Materials	176
4.3.2	Procedure	179
4.3.3	Results	179
4.3.4	Discussion	183
4.3.4.1	Noun plausibility modulates reading times to a greater degree under postcedent search	183
4.3.4.2	Indexing processing difficulty in the maze task	186
4.4	General Discussion	190
4.4.1	Theoretical implications	190
4.4.2	Methodological considerations: SPR vs. Maze	199
4.4.3	Further questions	200

5	Implicit Causality	201
5.1	NPE shows an implicit causality and object bias in the Maze task	204
5.1.1	Method	204
5.1.1.1	Methodology	204
5.1.1.2	Participants	204
5.1.1.3	Materials	204
5.1.2	Procedure	208
5.1.3	Results	209
5.2	NPE shows an implicit causality and object bias in self-paced reading .	212
5.2.1	Methods	212
5.2.1.1	Methodology	212
5.2.1.2	Materials	213
5.2.1.3	Participants	213
5.2.2	Procedure	213
5.2.3	Results	214
5.3	General Discussion	216
5.3.1	Implicit causality applies across anaphora and ellipsis	216
5.3.2	Implicit causality interacts with proximity preference	218
5.3.2.1	Question response time analysis	218
5.3.2.2	Origin of object response bias	224
6	Conclusion	229
6.1	Why we should look across the aisle	229
6.2	Truth and Beauty?	234
A	Experiment 1 items	241
B	Experiment 2 items	251
C	Maze/SPR experimental items	257
D	Implicit causality items	268
E	GloVe Cosine Measures, NPE Maze Experiments	271
	Bibliography	274

Abstract

Margaret Kroll

Comprehending Ellipsis

This dissertation is about ellipsis, a natural language construction in which a word or phrase is understood even though it is not pronounced. This project brings together two strands of research: formal theoretical research on the relationship between the interpretation of an ellipsis site and its surrounding linguistic context, and psycholinguistics research on the processes by which we build comprehensive meaning from the silence of an ellipsis input. I argue that only by considering these two research programs together can we fully understand ellipsis as a natural language phenomenon. I first present novel English sluicing data that challenge even the most successful existing theories of the relationship between antecedent and elided content in ellipsis constructions by showing that the elided content and antecedent content in a sluicing construction can mismatch to a greater degree than previously thought possible. I use this data to argue that an interpretation condition for ellipsis must be sensitive to pragmatic content, and motivate a proposal in which sluicing is treated as a pragmatics-sensitive phenomenon licensed by local contextual entailment. I then present the results of eight experiments on ellipsis comprehension that use offline comprehension measures and online measures of incremental processing. I show that anaphoric and cataphoric ellipsis comprehension is subject to a proximity bias, wherein comprehenders prefer to resolve an ellipsis site to the candidate antecedent in greatest proximity to the ellipsis site, regardless of whether the site precedes or follows the antecedent. I also show that cataphoric ellipsis comprehension is subject to an active, forward search strategy. I argue that ellipsis processing ultimately shares features both with anaphoric processing and with the processing of long-distance dependencies, such as filler-gap processes.

For Julia Bencie

Acknowledgments

I can never hope to put on paper the gratitude I feel for the support I have received throughout graduate school. It is with a profound state of unease that I can say only that I hope I have not disappointed those who have been invested in my endeavors, failures, and successes, and that I hope I have positively impacted people in return.

My road to linguistics and to finishing this degree was hardly straight or inevitable, and I am indebted to many people who anchored me at certain points along the way. Jerry Sadock, who agreed to advise a master's thesis of a student who was not only not a linguist, but knew very little about what linguistics was. Without Jerry I would never have taken semantics, and would certainly never have continued in linguistics. Chris Kennedy for supporting me despite the fact that I was not a linguistics student and despite my less than stellar performance early on in his semantics class. Chris provided guidance very early in my career, and I am grateful for his time and patience. Thank you also to Anna Szabolci and Lucas Chompollion for their early encouragement and advice.

At a different point in time but for similar reasons, I am grateful to Bruce Hayes. Not many semanticists have an MA thesis on phonology, and I could not have asked for a better advisor for that. I am thankful for Bruce's support, and aspire to one day dispense advice with such sagacious wit.

I am fortunate to have had two wonderful cohorts to share the grad school experience with. At UCLA: John Gluckman, Margit Bowler, Jesse Zymet, and Phil Côté-Boucher; I will always look upon partially ordered sets fondly. At UCSC: Chelsea Miller, Nate

Clair, Ben Mericli, Jeff Adler, Hitomi Hirayama, and sticking it out to the bitter end, Steven Foley and Jed Pizarro-Guevara. I couldn't have asked for a better group of people to go through grad school with, and I am so proud of what we have all accomplished.

Thank you to Tuomo Tiisala, for believing in me and supporting me in innumerable ways early on. Thanks to Craig Sailor for first encouraging me to look at ellipsis and anaphora, and for his unceasing quest to convince me of the importance of linguistic contexts. Jason Bishop, for always providing a voice of reason, and for teaching me to block out the haters and the takers. Kelsey Kraus, for being a supportive office-mate and roommate, being a seemingly endless source of patience and advice when listening to my complaints and struggles, and for baking the best dissertation cake I could ever have asked for. Thank you to Deniz Rudin for always believing in the value of my research even when I didn't, and for being a vocal advocate of it. I am grateful for his friendship, support, and generosity.

I am equally grateful for the friendship and encouragement of those not directly involved in my academic life. Rusty Kingon for teaching me how to sail, allowing me to fulfill a lifelong dream and giving me the hobby that was instrumental in getting me through graduate school. Tom Faraola and the crew of Rio, for being wonderful people to sail with and spend time with, and John Armstrong for being my reliable surfing buddy even though I wasn't always.

My fellow graduate students form an incredible support network, and I could not have accomplished what I have without their emotional, technical, and practical support. Thank you to Kelsey Sasaki for fighting the good fight in discourse processing; Tom Roberts for the puns and being a supportive co-author; Karl DeVries for being an enthusiastic stat discussant; Netta Ben-Meir for being an old-timer with me; Alex Goebel for the shared experience of being an unrepentant ideological outcast; Paul Willis, from whom I am continually learning new esoteric facts; Stephanie Rich and Jack Duff for running

experimental items for me, and all the members of s/lab and SPLAP for their years of helpful comments and feedback. Thank you to my RAs Melanie Esver, Lenny Kirshbaum, and Joyce Hong for their help running participants in the UCSC lab; I would not have been able to complete the experiments I did without their help.

The experimental work in this dissertation specifically benefited from discussions with my committee, Jesse Harris, Lyn Frazier, Alex Goebel, and the UCLA Psycholinguistics Lab. I am grateful to Stephanie Rich and Steven Foley for lending me some pieces of R code, and to Jed Pizarro-Guevara for helping me understand Helmert coding.

I am also incredibly thankful for the amazing staff of our department: Ashley Hardisty, Maria Zimmer, Gwyn Vandevere, and Logan Roberts. Thank you for all of your help with the shocking amount of bureaucracy at UCSC, and your support and encouragement over the past six years.

I have been fortunate to work with incredibly smart, interesting faculty while in grad school, either as a student or as a TA: Kie Zuraw, Russ Schuh, Emma Borg, Jessica Rett, Ed Stabler, Jaye Padgett, Donka Farkas, Sandy Chung, Bill Ladusaw, Dan Hardt, Ivy Sichel, and Maziar Toosarvandani. I have learned much more than linguistic theory from each of them.

I am especially grateful to Jim McCloskey. Jim was not my advisor in an official capacity, which makes his years of support and encouragement even more meaningful. Because of Jim's well-attested aversion to public praise, I will not go on at length about my feelings of respect for him, his kindness, and his dedication. Let there be no doubt, however, that I could.

I am grateful to Adrian Brasoveanu for his general enthusiasm, and for his support of decisions he feels are beneficial regardless of whether they are widely popular. Adrian has taught me to constantly look for new ways to improve upon whatever methods I am currently using, and I will continue to benefit from this lesson long after grad school.

Matt Wagers is my chair and advisor, and I would not have written this dissertation without his support, guidance, and advice. Whatever claim I have to the label of psycholinguist is because Matt was willing to take on a semanticist who self-proclaimed to have little interest in the traditional questions studied within psycholinguistics. I am incredibly thankful for his patience and support throughout the past four years, and for his willingness to let me follow my own path.

Pranav Anand has probably been the single greatest source of support I have had at UCSC. I have benefited in innumerable ways from knowing Pranav over the past six years that extend beyond my development as a scholar and researcher, though he was unquestionably influential in those ways as well. Pranav was instrumental in my entering experimental research, and I am thankful for all of our conversations, his thoughtful advice, the earnestness with which he tackles every question or concern addressed to him, and for his role in shaping the path I have taken.

I cannot imagine the last few years without Amanda Rysling in the department. Although Amanda has known me the shortest time of anyone on my committee, she has had an outsized influence on my academic life. Amanda, more than anyone, is constantly pushing me to be better in whatever I am doing, and, importantly, she never lets me get away with making excuses. I am grateful for her support, encouragement, and incredible dedication to her students. Thank you for helping me find the intellectual hill that I am steadfastly willing to die on.

Finally, thank you to Will Rush and to my family, Diane Kroll, Mike Kroll, Willa Levin, and Matt Kroll. I have accomplished what I have due to your unshakeable belief in me and in the innate value of earning a PhD. I am so lucky to have you all in my life.

Chapter 1

Introduction

Heard melodies are sweet, but those unheard
Are sweeter
(*Ode on a Grecian Urn*)

An idiosyncrasy of working on language is that you can never escape it. Since I began researching ellipsis, I have been the grateful recipient of many encountered-in-the-wild elliptical phrases that have been uttered, heard, or otherwise noted by friends and colleagues. Those phrases recorded by linguists also inevitably arrive with documented disagreements about what the encountered elliptical phrase means, could possibly mean, and a conclusion about whether the phrase in question was ultimately ‘possible’. The fact that we are still debating and discovering the possibilities of a construction that is, once you look, truly ubiquitous, is to me what makes ellipsis such a compelling area of study. I hope that a reader of this dissertation will feel the same.

Ellipsis comprehension is at its core the ability to recover meaning from silence. Consider the sentences in (1). All contain greater meaning than can be recovered from the meaning of the words in the sentence alone. Furthermore, each sentence communicates only a single meaning. Although the meaning of each sentence exceeds the words in the

sentence, all comprehenders converge on the same interpretation of what is not explicitly said.

- (1) a. Mary sailed a boat, but I don't know which boat.
- b. Mary sailed a boat, and Janet did, too.
- c. Mary sailed Sue's boat, and Janet sailed Amy's.

Conventional wisdom on ellipsis asserts that the ellipsis site is licensed by a preceding bit of linguistic context called the *antecedent*. In (2), the elided material in the ellipsis site – indicated by the struck-through material – is licensed by its relationship with the bolded antecedent material earlier in the sentence.

- (2) a. **Mary sailed** a boat, but I don't know which boat [~~Mary sailed~~].
- b. Mary **sailed a boat**, and Janet did [~~sail a boat~~] too.
- c. Mary sailed Sue's **boat**, and Janet sailed Amy's [~~boat~~].

In (2), the form of the antecedent matches the form of the elided material. However, the relationship between an antecedent and the elided content is not always so straightforward. It has been known for some time that the interpretation of an ellipsis site does not involve simply copying the form of an antecedent. For example, (3)-(8) represent a range of examples in which a presumptive antecedent is available in the preceding discourse, but the examples are nevertheless notable in some way.

- (3) I should probably get going, and you do, too.
(reported by Amanda Rysling as said by John Kingston)

- (4) In English you can extract over these clauses, but in Icelandic you're not supposed to be.
(reported by Stephanie Rich as said by Matt Wagers)

- (5) What is fluffy coffee, and why?

(thecut.com/2020/04/fluffy-coffee-is-what-exactly-investigating-the-viral-drink.html)

- (6) How to Behave and Why.

(Munro Leaf's 'How To Behave and Why', 1946)

- (7) Bring the broth to a boil – but make sure the polenta doesn't.

(foodandwine.com/cooking-techniques/how-to-make-no-cook-easy-polenta-tips; hat tip Amanda Rysling)

- (8) Life without pasta is not.

(poster on Whole Foods in L.A.)

More established examples include strict-sloppy readings as in (9), which allows both the interpretation that Jane sailed her own boat, and that Jane sailed Mary's boat (Fiengo and May, 1994). Example (10) showcases a split antecedent reading, in which the interpretation of the ellipsis site seems not to copy any preceding structure directly, but is instead interpreted to be along the lines of '...neither of them can do the adventure they want' (Webber, 1978; Hardt, 1993; Asher, 1993).

- (9) Mary sailed her boat, and Jane did [] too.

- (10) Bob wants to sail round the world and Alice wants to climb Kilimanjaro, but neither of them can [], because money is too tight.

Chapter 2 presents and discusses a new type of mismatch between elided material and the preceding context, called *polarity reversals*. An example of a polarity reversal sluice is given in (11), which is taken from a fight song of OSU. Notably, the interpretation of the ellipsis site contains negation that is not present in the preceding context.

(11) We're here to do or die. Ohio, Ohio. **We'll win the game** or know the reason why [~~we did not win the game~~].

Another way in which forming the interpretation of an ellipsis site is not merely copying the form of an antecedent is found in examples in which more than one candidate antecedent is available. For example, (12) has two possible interpretations, one in which the speaker does not know which tribbles hate Enterprise crew members, and one in which the speaker does not know which Enterprise crew members the tribbles hate. Similarly, (13) has a possible interpretation in which the speaker doesn't know if she has any apples, and one in which she doesn't know if she has any oranges. Given the possibility of more than one interpretation of the ellipsis site, we as comprehenders must choose a particular interpretation. Chapters 3-5 probe how we build interpretations of the ellipsis site when more than one interpretation is possible.

(12) I've heard that some tribbles hate some Enterprise crew members, but I don't know which ones [].

(13) Apples taste great with oranges, but I don't know if we have any [].

I argue in this dissertation that to understand ellipsis requires understanding how it shares properties with both anaphora resolution processing and with long-distance dependency processing. To this end, there are two main threads running throughout the dissertation. The first involves the nature of the relationship that must hold between an ellipsis site and its surrounding linguistic context. This question is covered mostly in Chapter 2, where I argue that bidirectional entailment accounts are too restrictive to account for the full range of possible interpretations of an ellipsis site. I show that the prevailing view in the literature – that the interpretation of an ellipsis site must be restricted to a strict matching relationship with a piece of previous linguistic structure – is too restrictive to account for the full range of attested data.

The remaining chapters present experiment findings suggesting that ellipsis interpretation is subject to a backward search for anaphoric ellipsis sites, and a forward search for cataphoric ellipsis sites. Chapter 3 presents offline interpretations showing both a subject and a proximity bias in ellipsis interpretations. The experimental findings provide evidence that comprehenders attempt to resolve ellipsis sites to a salient representation in the discourse, even if that representation doesn't have a structural antecedent. Chapter 4 presents the results of three online experiments probing the incremental processing of cataphoric ellipsis sites. I argue that cataphoric ellipsis comprehension – a construction in which the ellipsis site precedes the antecedent – involves an active, forward looking search akin to the more widely studied phenomenon of filler-gap processing. Chapter 5 shows that the interpretations of noun phrase ellipsis sites are subject to implicit causality biases at the same rate as pronominal anaphora. Finally, in Chapter 6 I lay out a theory of ellipsis comprehension in which I put theoretical licensing constraints on ellipsis into conversation with processing accounts of ellipsis comprehension. I argue that only by combining these two historically independent strands of research can we form a truly comprehensive understanding of elliptical phenomena.

Chapter 2

The pragmatics of sluicing

2.1 Introduction

2.1.1 Polarity reversals in sluicing

Sluicing, first noted by Ross (1969), is an ellipsis phenomenon in which the TP of an interrogative is elided under some identity condition, stranding an overt wh-phrase in the CP domain. An example is given in (14).

(14) Bernie knows that someone in Iowa voted, but he doesn't know who.

Research on sluicing mainly focuses on (at least) two major questions: whether or not there is syntactic material present in the ellipsis site, and how we should characterize the nature of the identity condition that licenses elision. This paper does not contribute anything new to the former question, and I assume an account of underlying syntactic structure elided at phonological form as presented in Merchant 2001 (cf. Hardt 1993; Chung et al. 1995; Ginzburg and Sag 2001; Barker 2013). The focus of this paper is on the latter question, the nature of the licensing condition for sluicing. The empirical contribution of the paper is to introduce a sluicing phenomenon that has previously gone

unnoticed; I'll refer to this as *polarity reversals under sluicing*.¹ Polarity reversals are sluices in which the presumed antecedent content (A) and the ellipsis site (E) differ in polarity. For example, the presumed antecedent in (15), *California will comply*, has positive polarity while the interpretation of the ellipsis site, *California won't comply*, has negative polarity.²

(15) I don't think that [California will comply]_A, but I don't know why [_{TP} ~~California won't comply~~]_E.

Similarly, the presumed antecedent in (16), *John didn't do an extra credit problem*, has negative polarity, while the interpretation of the ellipsis site, *he did do*, has positive polarity.

(16) **Context:** Students were given the option to do an extra credit problem, but were required to mark which problem they did next to their name on a spreadsheet. There is no mark next to John's name. The TA says:
Sluice: Either [John_j didn't do an extra credit problem]_A, or he_j didn't mark which one_i [~~he_j did do t_i~~]_E.

The polarity reversal data show a greater mismatch between antecedent and elided content than has been previously thought possible. Because such mismatches were not believed possible, data like (15) and (16) challenge even the most successful existing theories of the licensing condition for sluicing. Specifically, I show that the polarity reversal data are unable to be accounted for under theories that require any type of strict identity between the elided content and an antecedent in the discourse. The account

¹I use this label pre-theoretically and for convenience. As we will see, no actual "reversal" of polarity takes place.

²Note that there is a reading of the ellipsis site in (15) in which the antecedent and ellipsis sites include the matrix clause, but this reading is pragmatically odd.

presented here builds on the many successes of previous licensing conditions, but allows greater flexibility in the relationship between material preceding the sluice and the elision site by proposing that sluices rely on contextual entailment to license their elision sites. The account draws on theories from dynamic semantics and discourse coherence, which are traditionally largely unconnected with ellipsis.³ However, the tools used here have been independently established and well-motivated in their individual domains. Overall, this paper demonstrates that new and initially challenging data can be accounted for by combining traditional theories of ellipsis with insights from other areas in the literature.

2.1.2 Methodological preliminaries

A methodological aside on the data used throughout: The corpus examples given here were reviewed by at least eight members of the Santa Cruz Ellipsis Project.⁴ The initial round of annotation was conducted by either two to three independent undergraduate research assistants and the results were reviewed by a faculty supervisor. In the second phase of the annotation procedure, the initial annotations were reviewed by three undergraduate research assistants working with a faculty supervisor. The specific examples and judgments used here were additionally discussed and verified by a separate research group of two faculty members and four graduate students, in consultation with naïve speakers. Many of the examples presented here have more than one possible interpretation for the pre-sluice (that is, the unelided form of the sentence). The claim here is not that the pre-sluices provided for these examples are the *only* interpretation available, but merely that they are a felicitous, freely available interpretation in the context in which the sluice was found or constructed.

³Though see Elliott and Sudo 2016 for an exception for dynamic semantics.

⁴<https://babel.ucsc.edu/SCEP/>

Additionally, I exclude here two types of sluices that, to my knowledge, have not been explicitly excluded from previous analyses. The first is root sluices, like those in (17) and (18):⁵

(17) A: John bought a new sailboat.

Q: How long?

(18) Context: Taken from a blog post about a particular kind of concert Miley Cyrus has announced she will be holding:

“Miley has yet to confirm the news or provide further details – including, for example, ‘why.’”⁶

I exclude these sluices here because their range of use is clearly wider than that of embedded sluices. For example, there is no obvious linguistic antecedent for the *why* sluice in (18), though the meaning of the sluice is recoverable in context. While root sluices deserve to be studied in detail, their analysis should be separate from that given here.

The second group of sluices I exclude is semi-idiomatic uses like that in (19):

(19) Mary got a new climbing partner – guess who!

Although these examples go as far back as Ross 1969 and are more recently discussed in Barros 2014, they should be considered separately from the examples discussed here. Most obviously, these examples are special in requiring no linguistic antecedent, i.e. they are acceptable in out-of-the-blue contexts. For example, “Guess who!” can accompany

⁵Root sluices are interrogative forms that do not occur as complements of question-embedding verbs. While it is debated whether or not root sluices are derived via TP deletion (Bechhofer, 1976; Hankamer, 1977; Merchant, 2001), they are distinct from the examples under consideration here in occurring only in root (non-embedded) contexts.

⁶<http://jezebel.com/a-miley-cyrus-nude-concert-is-the-most-miley-idea-ever-1736444064>

an unexpected knock on the door, and “Guess what?” can be uttered at the start of a discourse with no antecedent – linguistic or otherwise – at all. Because these examples appear to have different licensing requirements, I put them aside here.

2.2 Sluicing in the literature

Numerous theories of sluicing have been proposed since the original syntactic isomorphy approach given in Ross 1969. A large part of the debate in the literature has been oriented around the question of licensing: what is the relationship between the content of an ellipsis site and the preceding discourse that licenses the elision of the site’s material.⁷ Traditionally, this licensing has been approached as a relationship between some salient antecedent and the interpretation of the ellipsis site. The identity condition underlying many approaches in the recent literature is that of semantic entailment. Originally proposed in Merchant 2001, the semantic entailment identity condition has held much weight as it is flexible enough to allow for attested structural mismatches, such as the finiteness mismatch in (20), but restrictive enough to rule out most impossible interpretations.

(20) [Sally rock climbs]_A. She learned how [~~to rock climb~~]_E from her mother.

I show here that a semantic entailment identity condition is too restrictive to account for the polarity reversal data, and therefore cannot be the identity condition we need. This section discusses the predictions of Merchant (2001)’s semantic entailment condition as well as the predictions of theories that rely in part on such an identity condition.

⁷Even anaphoric accounts rely on some specified relationship between the interpretation of the ellipsis site and the preceding discourse.

2.2.1 e-GIVENness

Merchant's theory of ellipsis, called e-GIVENness, is a theory of both the syntax and the licensing conditions for ellipsis constructions. e-GIVENness proposes that sluiced clauses contain syntactic structure within the ellipsis site that goes unpronounced, that is, is deleted at phonological form. The unpronounced structure within the ellipsis site consists of a TP missing a *wh*-constituent, called a *remnant* or *wh-remnant*, that has moved up and out of the TP prior to the TP's deletion.⁸ The licensing condition of e-GIVENness is indebted to Schwarzschild (1999)'s theory of GIVENness, which is not itself a theory of ellipsis, but one of focus and deaccenting. GIVENness proposes that an expression can be deaccented if its existential focus closure is contextually entailed by the existential closure of an antecedent.⁹ e-GIVENness modifies the GIVENness theory of deaccenting into a theory of ellipsis (Rooth, 1985, 1992; Romero, 1997) by modifying the GIVENness entailment condition from a contextual, unidirectional entailment condition to a bidirectional semantic entailment condition. Specifically, e-GIVENness proposes that in order for a TP to be elided it must stand in a bidirectional semantic entailment relationship with some salient antecedent. The account is given formally as follows:

Focus condition on TP-ellipsis: A TP α can be deleted only if α is e-GIVEN.

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

⁸The theory presented here adopts this approach and discusses only the licensing conditions for ellipsis, but I refer the reader to Merchant 2001 and much subsequent work for a thorough defense of the syntactic proposal.

⁹Existential closure is a type-shifting operation that raises expressions to type *t* by existentially binding unfilled arguments. Existential-F-Closure (F-clo) is the result of replacing F(ocus)-marked phrases in an expression with variables and existentially closing the result (Schwarzschild, 1999).

Note that the entailment requirement here is that of semantic entailment and, unlike GIVENness, does not leave room for contextual entailment.

The semantic identity condition of e-GIVENness is permissive enough to allow for certain observed syntactic mismatches between ellipsis sites and their antecedents, such as tense (Merchant, 2001); however, the bidirectional entailment requirement is too restrictive to allow for polarity mismatches. Let's look again at (15), repeated here as (21).

- (21) I don't think that [TP California will comply]_A, but I don't know why [TP California won't comply]_E.

Applying e-GIVENness to (A) and (E) yields the following:

A entails F-clo(E): No.

$A = \lambda w. \mathbf{comply}(c)(w)$

$F\text{-Clo}(E) = \lambda w. \neg \mathbf{comply}(c)(w)$

E entails F-clo(A): No.

$E = \lambda w. \neg \mathbf{comply}(c)(w)$

$F\text{-Clo}(A) = \lambda w. \mathbf{comply}(c)(w)$

Neither the antecedent expression nor the elided expression in (21) entails the other, and so e-GIVENness incorrectly predicts that we should not be able to elide the TP.

A skeptical reader might propose that, instead of the antecedent given in (21), we should consider the entire first conjunct to be the antecedent, as doing so will capture the negation in the antecedent expression. However, as (22) shows, expanding the antecedent to include the matrix negation still does not yield semantic entailment in either direction.

- (22) [TP I don't think that California will comply]_A, but I don't know why [TP California won't comply]_E.

A entails F-clo(E): No.

$A = \lambda w'. \neg \forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)]$

$\text{F-Clo}(E) = \lambda w. \neg \mathbf{comply}(c)(w)$

E entails F-clo(A): No.

$E = \lambda w. \neg \mathbf{comply}(c)(w)$

$\text{F-Clo}(A) = \lambda w'. \neg \forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)]$

2.2.2 Hybrid Theories

Merchant (2013b) and Chung (2013) argue that e-GIVENness alone is too weak an identity condition on sluicing, as it fails to rule out impossible sluices such as the active/passive mismatch in (23).

(23) # $[\text{John was kicked}]_A$, but I don't know who_i [~~t_i kicked John~~]_E.

Merchant and Chung propose to rule out sluices such as (23) by including substantive syntactic restrictions on sluicing in addition to the bidirectional semantic entailment condition of e-GIVENness. As the accounts presented in these works are by their very purpose more restrictive than e-GIVENness, the objections in the last section regarding the too-restrictive nature of e-GIVENness apply equally to these accounts, as well.

2.2.3 Inquisitive entailment

AnderBois (2014) criticizes e-GIVENness for failing to predict the impossibility of sluicing out of doubly-negated indefinites and appositives. AnderBois argues that while examples (24) and (25) satisfy bidirectional entailment, they are not well-formed sluices (AnderBois 2014, pp. 19 & 23), respectively, brackets added):

(24) # $[\text{It's not the case that no one left}]_A$, but I don't know who [~~left~~]_E.

(25) #Joe, [who once killed a man in cold blood]_A, doesn't even remember who [~~he once killed in cold blood~~]_E.

AnderBois proposes a modified account of sluicing based on inquisitive semantic entailment over CPs (Groenendijk and Roelofsen, 2009). Inquisitive semantic entailment is a more restrictive identity requirement than e-GIVENness and rules out examples (24) and (25): neither antecedent contains inquisitive content and therefore cannot entail the inquisitive content of the question CP (see AnderBois §3.4 and §4 for discussion).

Recent work has challenged the inquisitive semantic entailment account on empirical grounds. Collins et al. (2014) provide experimental evidence that sluices out of doubly-negated constructions and appositives are indeed possible,¹⁰ and Barros (2014) observes that (25) above, modified in (26), is acceptable when the remnant *who* is replaced with the d-linked wh-expression *which man*.

(26) Joe, [who once killed a man in cold blood]_A, doesn't even remember which man [~~he once killed in cold blood~~]_E.

Additionally, Inquisitive Entailment was created to be a more restrictive entailment identity account than e-GIVENness. As I have shown that e-GIVENness is too restrictive to permit polarity reversal sluices, it follows that an account that was created to predict a subset of those constructions predicted by e-GIVENness is also too restrictive to permit the polarity reversals. Thus, while contributing many insights into the semantics and pragmatics of sluicing, particularly regarding *which* sluices out of disjunctive clauses, Inquisitive Entailment is too strict an identity condition to correctly predict the full range of sluicing possibilities.

¹⁰The results show that while speakers in the experiment did not judge the examples highly, they judged them no worse than the equivalent pre-sluice examples.

2.2.4 Scopability

Barker (2013) analyzes sluicing as anaphora to the semantic remnant of a clause from which a subconstituent has been removed (a *continuation*). Under this proposal, sluicing constructions contain gaps and silent proforms but no internal syntactic structure in the ellipsis site. This analysis, called *scopability*, uses a type logical (categorical) grammar that allows a tight connection between syntactic and semantic content (Barker, 2007). Certain facts about scope and case matching behavior in sluicing therefore fall out naturally from the system.

Scopability focuses on the scope facts of sluicing originally observed in Chung et al. 1995. Chung et al. notice that an example like (27) only allows a reading in which the indefinite *a book* takes wide scope over the quantifier *everyone* (Barker, 2013, p. 191).

(27) Everyone selected a book, but I don't know which book.

Chung et al. propose that in sluicing constructions, the correlate (their *inner antecedent*: the optionally-present constituent in the antecedent of a sluicing construction that corresponds to the wh-remnant) must take scope over the antecedent clause. This scope fact follows in the system of scopability because a sluicing antecedent is created by allowing the correlate to take scope over the rest of the antecedent clause. In 27, this means that the indefinite correlate [a book] must take wide-scope over the antecedent [everyone selected ___]; the scope facts then follow naturally.

I argue that scopability, too, is too restrictive to capture the polarity reversal data. While scopability rejects the semantic entailment condition, Barker argues that scopability avoids overgeneration by imposing effectively the same restriction through different means: "... the net effect of the mutual entailment requirement [of e-GIVENness] is that once we subtract the inner antecedent from the antecedent clause, and once we subtract the wh-phrase from the sluice, the remainders must be semantically equivalent exactly

what is guaranteed by the anaphoric [scopability] analysis” (213). Indeed, it is clear that Barker intends scopability to deliver semantic equivalency between a sluice and its antecedent. However, we have seen that a condition which imposes semantic equivalency between an ellipsis site and its antecedent is too restrictive to predict the novel data contributed by polarity reversal sluices.

In summary, bidirectional semantic entailment accounts such as e-GIVENness, and theories that impose semantic identity between an ellipsis site and its antecedent, are too restrictive and fail to predict the existence of polarity reversal data.¹¹ The next section proposes an alternative account that builds off the insights provided by the accounts discussed here.

2.3 A modified account

This section proceeds in two parts. The first subsection discusses a constraint proposed in Dayal and Schwarzschild 2010 that explains certain data that have been used to motivate a bidirectional entailment account. I propose that since these data can be ruled out on independent grounds, they should not be ruled out by a sluicing theory specifically. I then propose a new identity condition on sluicing called Local Givenness, which argues that sluices are licensed by local contextual entailment.

2.3.1 The Well-Formedness Condition

This subsection shows that we can rule out certain forms of unacceptable sluices independently from our sluicing licensing condition. Indeed, there are many reasons why a particular sluice may be unacceptable. Importantly, the reason a given sluice may be

¹¹Note that Ginzburg and Sag (2001) and Barros (2014) take a slightly different approach, combining syntactic and pragmatic constraints.

unacceptable is not always because of its failure to satisfy an ellipsis licensing condition. Dayal and Schwarzschild (2010) observed in their investigation of sluices with definite correlates that one reason a sluice may be unacceptable is because its pre-sluice or, the unelided form of the sluice is also unacceptable (see also Romero 1998). They therefore propose that infelicitous pre-sluices will yield infelicitous sluices. For example, (28a) is infelicitous, or pragmatically anomalous (in this case, because it contains two generally contradictory clauses). It is not surprising, then, that the corresponding sluice in (28b) is also infelicitous.

- (28) (a) #Hillary_i knows that the President lives in the White House, but she_i doesn't know where the President lives.
- (b) Hillary_i knows that the President lives in the White House, but she_i doesn't know where_j #~~[the President lives t_j]~~.

For ease of reference, I call this observation the Well-Formedness Condition (see discussions in Romero 1998, Dayal and Schwarzschild 2010, Tancredi 1992, and Merchant 2013a):

- (29) *The Well-Formedness Condition:*
If a pre-sluice is infelicitous, then the corresponding sluice will not be well-formed.

The Well-Formedness Condition is, to the extent of my knowledge, both empirically verifiable and intuitively satisfying. It seems desirable that a question that is infelicitous when uttered overtly will remain infelicitous when partially elided. The importance of this observation is that it rules out certain examples that we do not want to force our sluicing theory to rule out. The observation is not specific to polarity reversal sluices, but is helpful to keep in mind when considering the scope of the data that any theory of sluicing needs to account for. By adopting the Well-Formedness Condition, we are

able to rule out examples like (b) independently, obviating the need to account for such sluices in our theory of ellipsis.¹²

2.3.2 Local Givenness

This section proceeds in three parts. The first outlines the basic formal assumptions used in the current account. The second presents a first pass at the sluicing theory that is developed and argued for in this paper, called Local Givenness. The third independently motivates the theory by applying it to several non-polarity reversal sluices. The reader who is familiar with dynamic theories of context update can safely skip to the second section on Local Givenness.

2.3.2.1 Context update and dynamic interpretation systems

I use here a dynamic interpretation system outlined in Kadmon 2001 and based on File Card Semantics (Heim, 1983a) and the Context Change Potential (CCP) system of Heim 1983b. Dynamic interpretation systems such as CCP update the conversational context incrementally. They aim to capture the observation that clausal interpretation may crucially rely on the interpretation of previous intrasentential clauses. For example, the context need not be updated only at the end of a sentence as in a static system; instead,

¹²We should ask why examples like (a) are infelicitous. I follow previous researchers who propose it is because it is infelicitous to ask a question that already has (at least) a partial answer available in the discourse (Romero, 1997; Fitzpatrick, 2005; Ginzburg, 2012; Barros, 2014). For example, B's question in (i) is infelicitous without the inclusion of *other*, because A has already asserted that she has seen some tigers that day at the zoo, which is a partial answer to the question 'What animals did you see today at the zoo?'.
(i) A: I saw some tigers today at the zoo.
B: What/which #(OTHER) animals did you see today at the zoo?

semantic content can incrementally update the context at the completion of a proposition or earlier. Because of this incremental updating, the second clause of a sentence such as in a sentence containing conjoined clauses can be interpreted in a different context from the first clause of the same sentence, and in a different context from the global conversational context. Theories of incremental updating have been of particular interest in studies of presupposition projection (Stalnaker 1973, 1974; Karttunen 1974a; Gazdar 1979; Karttunen and Peters 1979; Heim 1983b; Soames 1989; Beaver 2001, a.o.) and anaphora (Geach 1962; Evans 1977; Roberts 1989; Heim 1990; Kamp and Reyle 1993; van Rooij 1997/2006; Nouwen 2007; Murray 2014, a.m.o.).

In order to be able to clearly refer to the context of an entire sentence versus the context in which a particular clause is interpreted, Karttunen (1973) first distinguished *global* contexts, contexts at which the current sentence is interpreted, from *local* contexts, contexts at which the current clause (or possibly some smaller constituent) is interpreted. This terminology is adopted in Heim 1983b and Kadmon 2001 under a similar meaning. It is further developed under the motivation of processing parsimony in Schlenker 2009, 2010, 2011a. As mentioned above, the local context in which a clause is interpreted is not necessarily identical to the global context of its containing sentence. One effect of distinguishing between the two types of context is that propositions can be entered into local contexts without being entered into the discourse common ground or context set (Stalnaker, 2002), i.e. without being accepted as true by the speakers of the discourse for purposes of the discourse. Note that this means that, throughout a discourse, a local context is not necessarily a continually narrowing set of worlds. I notate local contexts throughout as c_L to distinguish them from global contexts, though the reader is asked to keep in mind that this serves merely as a reminder that we are concerned with updating our derivations incrementally.

In Heim (1983b)'s CCP system, the context is defined as a set of worlds (or, alter-

natively, the context can be defined as a set of world-assignment pairs). Instead of a standard truth-conditional semantics, operators contribute a context change potential, which is a partial function from contexts to contexts. Context change potentials express partial functions because a new context is defined only when the presuppositions of the entering expression are defined (entailed by the context), or are accommodated. The basic assumptions I make here are given as follows:

A context c and a proposition p are defined as sets of worlds. Because c is a set of worlds, entailment is defined by the subset relation, such that if a context c entails a proposition p , then $c \subseteq p$. A context is updated with a new proposition p by conjoining, or intersecting, with p . Context updating is defined as follows:

Context update:

- a. If c entails the presuppositions of p , then $c + p = c \cap p$
- b. If c does not entail the presuppositions of p , then c is undefined¹³

Some Heimian rules for basic English propositional operators are given here (see e.g. Karttunen 1974a; Heim 1983b; Kadmon 2001):

Negation: $c + \neg p = c \setminus (c + p)$ ¹⁴

Conjunction: $c + (p \wedge q) = (c + p) + q$

Disjunction: $c + (p \vee q) = (c + p) + (c + (c \setminus p) + q)$

Conditional: $c + (\text{If } p, q) = c \setminus (c + p \setminus ((c + p) + q))$

¹³Heim (1988), following Lewis (1979), also allows for the possibility of accommodation of the presuppositions of p if c is undefined. Accommodation is defined as follows: $c + p = (c \cap \text{ps}(p)) \cap p$. As Heim notes, accommodation is not a mechanism of the formal system, but instead an act that a speaker chooses to do to avoid a breakdown in communication (pg. 401). Because accommodation is not required in this paper, I refer the interested reader to the works cited here for further discussion.

¹⁴Where $c \setminus p$ indicates the intersection of c with the complement of p . This is notated in Kadmon (2001) as $-$, though the traditional notation is kept here.

Note that the final results of applying these rules to the context is an updating of the global sentence context, or the context resulting from the interpretation of the entire expression. Importantly, the context change potentials of these complex expressions are derived compositionally from the context change potentials of their constituents.

Let's look at the conditional example in (30) to illustrate.

(30) If [the tax bill passes]_p, [the University of California will be concerned]_q.

Let's assume for expositional simplicity that the context at the beginning of the sentence in (30) is the set of all worlds, W :

- (i) $c = W$, or the set of all worlds and the local context of the antecedent p .
- (ii) $c + p$: Intersects p with W , which yields the set of worlds in which the tax bill passes. This is the local context of q .
- (iii) $c + p + q$: Intersects W with the propositions expressed by p and q , which yields the set of worlds in which the tax bill passes and UC is concerned.
- (iv) We intersect the results of (ii) with the complement of the results of (iii), which gives us the worlds in which the tax bill passes and UC is not concerned.
- (v) Finally, we intersect (i) with the complement of the results of (iv), which gives us the set of worlds in which the tax bill passes and UC is concerned.

While there is debate in the literature over the correct formalization of some of these rules, the differing implementations of the rules does not impact in any crucial way the ellipsis derivations that concern us here. The interested reader is encouraged to reference Karttunen 1974a; Heim 1983b; Kadmon 2001; Schlenker 2010, 2011b for additional discussion of these operator rules, as well as their respective benefits and drawbacks.

In summary, I use a basic Heimian account here because the CCP account is widely known and is sufficient to capture the data we are concerned with. However, the Local Givenness account that is developed here crucially does not rely on the use of this particular system. Other developed dynamic theories, such as the closely related Discourse Representation Theory (Kamp, 1981) and its expansions and Dynamic Predicate Logic (Groenendijk and Stokhof, 1990), could alternatively be used. Also available as alternatives are non-dynamic accounts that use local contexts to compute incremental processing, such as Schlenker 2009, 2010, 2011a.

2.3.2.2 Local Givenness

The theory of sluicing presented here eschews semantic identity in favor of pragmatics-based entailment. The spirit of the proposal is indebted to those accounts already discussed and to the contextual entailment allowance that was included, though not given an exposition, in Schwarzschild (1999)'s GIVENness theory. Informally, I propose that the TP of an interrogative can be elided if and only if the proposition expressed by the TP, modulo existential closure, is entailed by the context in which the sentence expressing the proposition would be uttered. Recall that the notation c_L is used to indicate the local context of p ; that is, the context into which p is entered. Formally, Local Givenness is expressed as follows:

Local Givenness (preliminary): A TP α can be deleted *iff* $ExClo([\alpha]^g)$ expresses a proposition p such that $c_L \sqsubseteq p$.

Because the theory uses contextual entailment as its licensing requirement, there is no reliance on antecedents built into the theory. For expositional clarity and in deference to the historical importance antecedents hold in accounts of ellipsis, I use antecedent labels throughout in the derivations of sluices in this paper. I ask the reader to please keep in mind, though, that this is a notational convenience and not a requirement of the

theory. The propositions labelled as antecedents should be thought of more accurately as licensers, in that they provide the main propositional content constraining the local context of the sluiced proposition. There are no actual antecedent requirements built into the account, however, other than what content impacts the local context of the elided proposition.

2.3.2.3 Application of Local Givenness

This subsection applies the theory of Local Givenness to non-polarity reversal examples. The examples were chosen to display the generalizability of the theory and to show that it correctly accounts for examples that are not polarity reversals. The first example is one in which the elided content is clearly not entailed by the global discourse context.

- (31) Roy says that he received a PhD, but from whom? Probably no-one. It's debatable whether he even graduated from high school.

There are two possible interpretations of the sluicing site:

- a. Roy_i says that he_i received a PhD, but from whom_i [~~does Roy_i say that he_i received a PhD~~ t_i]_E?
- b. Roy_i says that he_i received a PhD, but from whom_i [~~did he_i receive a PhD~~ t_i]_E?

I take the reading in (a) to be straightforwardly calculable as entailed in its context under a standard treatment of conjunction in which p is the proposition expressed by 'Roy says that he received a PhD.'¹⁵ I focus then on the embedded reading given in (b). Note that, as follows from the rules given above, a local context of an expression p is not influenced by information following p in the discourse (see Schlenker 2010). Therefore, the fact that the sluice is followed by information that suggests that the speaker

¹⁵For a contemporary analysis of the denotation of *but* as a conjunction, see Toosarvandani 2014.

disbelieves Roy's claim is irrelevant to our derivation. Recall that Local Givenness does not require the elided proposition to be entailed by the global context or the context set of the conversation, but only that it is entailed by the local context (see discussion earlier in this section). I argue that in (b), we interpret the elided proposition in a context which entails the proposition that Roy received a PhD; that is, the local context for E is the set of worlds in which Roy received a PhD. I use here the following (simplified) denotation of *say*:

$$\llbracket \text{say} \rrbracket^g = \lambda p. \lambda x. \lambda w. \text{say}(p)(x)(w)$$

I also use throughout a function C that takes a given function f and returns a set such that for any $\langle s, t \rangle$ function f , $C(f) = \{w \mid f(w) = 1\}$. This is merely a notational convenience, which allows us to move between function and set notation.

Let's assume for maximal generality that our starting context is W , the set of all possible worlds:

- (i) $c = W$
- (ii) $\lambda w. \text{say}(\lambda w'. \exists x[\mathbf{phd}(x)(w') \wedge \mathbf{received}(x)(r)(w')])(r)(w)$
- (iii) $W \cap C(\lambda w. \text{say}([\lambda w'. \exists x[\mathbf{phd}(x)(w') \wedge \mathbf{received}(x)(r)(w')])(r)(w))] \approx$
(via assertive content)
 $W \cap C(\lambda w. \exists x[\mathbf{phd}(x)(w) \wedge \mathbf{received}(x)(r)(w)])$
 $\{w \mid \exists x[\mathbf{phd}(x)(w) \wedge \mathbf{received}(x)(r)(w)]\}$
- (iv) $\{w \mid \exists x[\mathbf{phd}(x)(w) \wedge \mathbf{received}(x)(r)(w)]\} \subseteq \text{ExClo}(\llbracket E \rrbracket^g)$

Step (ii) shows the proposition expressed by the first clause of the conjunction in (b). We can then ask what effect on the local context this clause has. The matrix verb *say*, in particular, is in the set of communicative assertive verbs. These verbs have long been observed to be able to assert their complement as true in a local context, although

the truth is not necessarily projected up to a global context (for recent discussions see Schlenker 2010 and Anand and Hacquard 2014). In this example, the proposition expressed by the embedded clause, that Roy received a PhD (i.e., $\lambda w'. \exists x[\mathbf{phd}(x)(w') \wedge \mathbf{received}(x)(r)(w')]$), is entered into the local context produced by the first conjunct. Step (iii) shows that the proposition in (ii) restricts the local context to worlds in which Roy received a PhD; this is achieved by intersecting the context W with the proposition expressed by the sentence ‘Roy received a PhD’. The local context for (E) is then the set of worlds in which Roy received a PhD. This entails the proposition expressed by (E), and we predict felicitous elision of the proposition in step (iv).

As already mentioned, the derivation here is an example of a communicative assertive verb asserting its complement clause into the local context. This is one example of a larger, more general observation that embedding verbs are able to assert their clausal complement as either true or as the main point of the utterance. This observation has been discussed in detail by Higginbotham (1975) and Simons (2007, 2013). The observation more specifically is that clausal complements of embedding verbs such as *see*, *think*, and *believe* can behave in discourse as independent propositions that can be asserted, responded to, and questioned independently of the matrix clause. Because this issue is discussed at length elsewhere, and because the current theory has no new contribution to this topic, I assume here existing proposals and treat the complements of such verbs as entering into the local context.¹⁶

The next example shows that the account correctly predicts sluicing out of appositives, as is shown to be empirically available in (26). While appositives are canonically considered to be not-at-issue content, it has been widely observed that they participate in discourse level activities including ellipsis in the same manner as matrix or at-issue

¹⁶This can be achieved formally via a form of local accommodation, see discussions in Roberts 1989, 2015 and Kadmon 2001, Ch. 9.

content. Appositives can also in the right circumstances behave like at-issue content in their projection and truth conditional behavior (Potts 2005; Amaral et al. 2007; Syrett and Koev 2014; AnderBois et al. 2015; Kroll and Rysling 2019, a.o.). I assume here that appositive relative clauses are propositional (Potts, 2005), stand in an anaphoric referential relationship with their main clause anchor (notated here by the superscript x) (Nouwen, 2007), and differ from main clause content in that appositive content is automatically added to the global context, while main clause content is added in the local context and introduced as a proposal to update the global context (Murray, 2014; AnderBois et al., 2015).¹⁷ A slightly simplified version of example (26) is repeated in (32).

(32) Joe_j^x, [who^x killed a man in cold blood]_A, doesn't know which man_i [~~he_j killed _{t_i} in cold blood~~]_E.

(i) The appositive content of (32) updates the global context c :

$$c = W \cap \{w \exists x [\mathbf{man}(x)(w) \wedge \mathbf{kill\ in\ cold\ blood}(x)(j)(w)]\} = \\ \{w : \exists x [\mathbf{man}(x)(w) \wedge \mathbf{kill\ in\ cold\ blood}(x)(j)(w)]\}$$

(ii) The existential closure of E is as follows:

$$ExClo([\mathbf{E}]^g) = \{w \exists x [\mathbf{man}(x)(w) \wedge \mathbf{kill\ in\ cold\ blood}(x)(j)(w)]\}$$

(iii) $\{w \exists x [\mathbf{man}(x)(w) \wedge \mathbf{kill\ in\ cold\ blood}(x)(j)(w)]\} \subseteq ExClo([\mathbf{E}]^g)$

As there are no intervening updates (no intervening operators or propositions), the local context for the expression following the appositive is identical to the global context, and we correctly predict entailment and elision of the sluiced clause in step (iii).

¹⁷See AnderBois et al. 2015 for a detailed explanation of the theory which is simplified here for expositional purposes. Note that while their account is created to allow for the free occurrence of ellipsis over borders, they rely on the account of AnderBois 2014 to rule out sluicing specifically. See discussion above on why we want an ellipsis theory to derive this example.

Entering the appositive content immediately into the global context in this way correctly predicts the projection behavior of appositives and the existence of anaphoric and ellipsis possibilities across appositive and main clause content. The interested reader is referred to Potts 2005, Del Gobbo 2007, Nouwen 2007, 2014, Schlenker 2013, and AnderBois et al. 2015, a.o. for additional details.

This subsection has independently motivated the Local Givenness theory using examples that are also predicted by Merchant's e-GIVENness theory. The contextual entailment condition of Local Givenness is a less restrictive condition than the bidirectional semantic entailment condition of e-GIVENness, and therefore predicts not only those sluices predicted by e-GIVENness, but also sluices that the bidirectional entailment condition is too restrictive to capture. The following section turns to deriving polarity reversal sluices.

2.4 Deriving polarity reversal sluices

The following sub-sections apply Local Givenness to five main categories of polarity reversal sluices. The categories are formed by grouping the data based on a salient similarity. For example, polarity reversal sluices are easily constructed with neg-raising verbs, so one category is Neg-Raising Polarity Reversals. While I believe that these categorizations hold theoretical significance, I remain uncommitted to them as holding final explanatory power or as comprising an exhaustive subcategorization.¹⁸

¹⁸Many polarity reversal examples given here contain *why* and *which NP* correlates. The paucity of correlate types may raise concerns that polarity reversals comprise a restricted set of data and are not generalizable to a wider theory of sluicing. Polarity reversal examples are, however, available with a greater class of remnants. For example, manner *how* polarity reversal sluices are possible and have been found in corpus data:

2.4.1 Polarity reversals are not semantic entailment: Neg-raising polarity reversals

One class of polarity reversal sluices contains neg-raising verbs. For example, (15) is repeated here as (33):

- (33) I don't think that [California will comply]_A, but I don't know why [_{TP} California won't comply]_E.

That neg-raising is the relevant property in (33) can be seen by the minimal pair comparison in (34) and (35). Example (35) swaps the neg-raising verb *think* with the non-neg raising verb *hope*. While the neg-raising and non-neg-raising interpretations are both available for (34), (35) cannot receive the polarity reversal interpretation in (35a). Instead, the only available interpretation is the matrix clause reading, given in (35b).

- (34) (a) Mary doesn't think that California will comply, but she can't explain why [~~California won't comply~~]_E.
 (b) Mary doesn't think that California will comply, but she can't explain why [~~she_i doesn't think California will comply~~]_E.
- (35) (a) Mary doesn't hope that California will comply, but she can't explain why #[~~California won't comply~~]_E.

(i) Corpus example, Santa Cruz Ellipsis Project

Context: No one expects Apple to stand still, however, and that fuels the speculation. "I don't think Steve Jobs will let it be a boring MacWorld,"

Sluice: "we just don't know how [~~Jobs won't let it be a boring MacWorld~~]."

Degree *how* polarity reversal sluices are also possible, as given in 78. Examples containing argument remnants are also possible, as shown below and in 82:

- (ii) I don't think that NO one came to the party, I'm just not sure WHO_i [~~t_i came to the party~~].

- (b) Mary doesn't hope that California will comply, but she can't explain why
[~~she_i doesn't hope California will comply.~~]_E.

Neg-raising verbs are clause-embedding verbs that when negated allow a reading in which matrix negation takes scope in an embedded clause. As it is one of the dominant approaches in the literature, I use here the account of neg-raising given in Gajewski 2007. Gajewski's account draws importantly on an idea from Bartsch (1973) that the inference from the literal interpretation of a neg-raising sentence like the antecedent in (33), where negation takes matrix scope, to the neg-raised interpretation, where negation takes embedded scope, is a pragmatic inference. Specifically, Bartsch argues that neg-raising verbs license an excluded middle presupposition as a pragmatic inference. For a sentence like the antecedent in (33) that contains the neg-raising verb *think*, the presupposition is that the subject either believes that the proposition expressed by the complement of the verb is true, or believes that it is false. The assertion of the antecedent of (33) combined with this presupposition then pragmatically entails that the speaker in (33) has a belief that California will not comply. The pragmatic nature of the reasoning involved explains how negation comes to be interpreted low and also explains why the neg-raised reading is cancelable in context. The criticism leveled against Bartsch's original account is that no principled reason is given for why some verbs are neg-raising verbs and others are not (Horn, 1978). For example, no explanation is given for why the verb *think* can neg-raise while the epistemically stronger verb *know* cannot, or why neg-raising verbs are idiosyncratically distributed across different languages.

Gajewski proposes to alleviate this objection by categorizing the excluded middle presupposition of neg-raising verbs as a soft-trigger presupposition in the sense of Abusch 2009. Abusch's soft-trigger presuppositions are presuppositions that are easily cancellable in context and as such are distinct from hard-trigger presuppositions, which cannot be cancelled. Soft trigger presuppositions are carried by predicates that invoke lexically-

stipulated alternatives as a matter of convention. The invocation of these alternatives triggers a pragmatic presupposition that one of the alternatives is true. In the case of neg-raising verbs, the alternatives invoked are the literal interpretation of the sentence and the neg-raised interpretation of the sentence. In summary, Gajewski proposes to treat neg-raising predicates as soft triggers that invoke a pragmatic excluded-middle presupposition. This proposal intends to capture the behavior described in Bartsch’s account while providing a more principled explanation for why some verbs allow neg-raising and others do not.

With this theoretical background in place we can return to example (33). I have proposed that the assertion of the antecedent in (33) combined with the excluded-middle presupposition invoked by the verb *think* entails that the speaker in (33) has the belief that California will not comply. Formally, this is expressed as follows:

(36) [I don’t think that California will comply]_A, but I don’t know why [_{TP} California won’t comply]_E.

$$\begin{aligned} & \llbracket \text{I don't think that California will comply} \rrbracket_A^g = \\ & \lambda w'. \neg \forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)] \end{aligned}$$

DOX(*s*)(*w*) indicates the set of worlds compatible with the doxastic state of the speaker. Via the excluded middle presupposition conventionally associated with the verb *think*, (A) presupposes that the world of evaluation of *think* meets the following restriction:

Excluded middle presupposition of (33)’s antecedent:

$$\lambda w'. [\forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)] \vee \forall w [w \in \text{DOX}(s)(w') \rightarrow \neg \mathbf{comply}(c)(w)]]$$

The denotation of *think* assumed here is given as follows (following Uegaki 2015):

$$\llbracket \text{think} \rrbracket^g = \lambda p. \lambda s. \lambda w'. [\forall w [w \in \text{DOX}(s)(w') \rightarrow p(w)] \vee \forall w [w \in \text{DOX}(s)(w') \rightarrow \neg p(w)]] . [\forall w [w \in \text{DOX}(s)(w') \rightarrow p(w)]]$$

Because (A) expresses that the first disjunct of the excluded middle presupposition is false, the presupposition of (A) and the assertion of (A) together entail the second disjunct of the presupposition. This entailment produces the stronger reading that the speaker uttering (33) has a belief that California will not comply. The following steps apply Local Givenness to (33).

(i) Starting Context:

$$c = W$$

(ii) Semantic Denotation of (A):

$$\llbracket \text{A} \rrbracket^g = \lambda w'. \neg \forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)]$$

(iii) Excluded Middle Presupposition of (A):

$$\lambda w'. [\forall w [w \in \text{DOX}(s)(w') \rightarrow \mathbf{comply}(c)(w)] \vee \forall w [w \in \text{DOX}(s)(w') \rightarrow \neg \mathbf{comply}(c)(w)]]$$

(iv) Strengthened Neg-Raised Interpretation of (A):

$$\lambda w'. \forall w [w \in \text{DOX}(s)(w') \rightarrow \neg \mathbf{comply}(c)(w)]$$

(v) Local Context for (E):

$$W \cap C(\lambda w. \neg \mathbf{comply}(c)(w)) = W \cap \{w : \neg \mathbf{comply}(c)(w)\} = c_{LE}$$

(vi) Denotation of (E):¹⁹

$$\text{ExClo}(\llbracket \text{E} \rrbracket^g) = \{w \neg \mathbf{comply}(c)(w)\}$$

¹⁹The observation that adjunct *wh*-traces do not participate in semantic parallelism relationships in ellipsis is part of an ongoing discussion in the literature (see Schwarzschild 1999, Merchant 2001, Hartman 2011, Barros 2014, Messick and Thoms 2016). This project has nothing to add beyond the current state of the literature on this question, and I refer the interested reader to the given citations. However, we might think that every event entails a reason (as well as a time, place, and manner) that it occurred even if that

(vii) Local Givenness:

$$c_{LE} \subseteq \{w \rightarrow \text{comply}(c)(w)\}$$

The semantic denotation of (A) in step (ii) asserts that it is not true that the speaker believes that California will comply. The pragmatic excluded middle presupposition in (iii) carried by (A) conventionally associated with the verb *think* requires that the speaker either believes that California will comply or believes that California will not comply.²⁰ Steps (ii) and (iii) together derive the strengthened neg-raised interpretation: Because (ii) asserts that it is not true that the speaker believes that California will comply, it follows from (iii) that the speaker believes that California will not comply. The utterance of (A) thus asserts the strengthened meaning given in (iv). Step (iv) pragmatically asserts that California will not comply. This assertion creates a local context c_L in which the worlds under consideration are only those in which California does not comply, given in step (v).²¹ Step (vi) shows the set of worlds in which California does not comply, and step (vii) shows that Local Givenness is satisfied because the local context given in (v) entails the elided proposition given by *California will not comply*. The theory therefore predicts the felicitous elision of the TP in (33).

reason is not known, such that existentially closing over that variable does not alter the entailment relations with other event denotations. Existentially closing over such a variable in 33 would yield the following: $\{w \exists e[-\text{comply}(c)(e)(w)] \wedge \exists r[\text{reason}(r)(e)(w)]\}$.

²⁰Note that the derivation does not fail here because the presupposition is "soft" in Abusch (2009)'s sense. That is, the presupposition is not a definedness requirement of the context, but is introduced by the lexical item, which invokes alternatives as a matter of convention. Thank you to a reviewer for requesting clarification on this step.

²¹See discussion of assertive verbs above.

2.4.2 Polarity reversals are not syntactic: *Remember* polarity reversals

The reader may at this point raise the objection that the previous example wrongly dismissed the possibility of a syntactic account of neg-raising as an explanation for the inference from $\neg\phi(p) \rightarrow \phi(\neg p)$. Indeed, the classic analysis of neg-raising originally advanced by, among others, Fillmore (1963) and Ross (1973), and revived recently by Collins and Postal (2014) argues for a syntactic explanation. This section shows that an appeal to a syntactic account of neg-raising will not save a semantic entailment account of sluicing. Instead, the inference $\neg\phi(p) \rightarrow \phi(\neg p)$ must, at least in some cases, be pragmatic in nature.

Example (37) is a corpus polarity reversal sluice containing *remember*.

(37) [Corpus example 91594, Santa Cruz Ellipsis Project]

Context: [O]n the day the Japanese invaded Pearl Harbor, Hummel was rounded up and locked in an internment camp along with about 2,000 other foreigners. . . So he and a British friend engineered an escape with the help of Nationalist guerrillas concealed nearby. He crawled over barbed-wire and walked most of the night and the next day. He was 20 and had no military training. But he was handed a small Belgian pistol, and he had little choice but to stay and help, harassing Japanese patrols by night and trying to defend a small patch of land against a communist takeover.

Sluice: “I don’t know why [I-wasn’t-scared], but I really can’t remember being scared.” [Hummel] said. “It all seemed like great fun.”

Example (37) is illustrative in that it appears to behave like the neg-raising examples: $\neg\textit{remember } p$ is interpreted in context as entailing $\neg p$. However, *remember* is not classified as a neg-raising verb in the literature and, indeed, the inference is more contextually

dependent than that carried by neg-raising verbs. For example, A's utterance in (38) is perfectly acceptable, while A's utterance in (39) is grammatical but a bit unwieldy.

(38) I don't remember being scared, but apparently I was!

(39) ?I don't think that Jane went to the party last night, but that's because I don't know anything about her whereabouts last night.

Karttunen (1971) classifies *remember* as an *implicative* verb. As such, it has the following properties when taking an infinitival complement: *remember p* \rightarrow *p*, \neg *remember p* \rightarrow \neg *p*. For example, in (40) there is a strong intuition that the assertion of the sentence commits the speaker to believing that she did not shut the door.

(40) I didn't remember to shut the door.

Higginbotham (2003) proposes that *remember* (along with *imagine*) in its usage with a gerund complement carries an obligatory *de se* reading when the embedded subject is PRO. For example, while (41) has both a possible *de re* and a possible *de se* reading, (42) carries only the *de se* reading, under which John remembers he himself going to the movies.

(41) John remembered his going to the movies.

(42) John remembered going to the movies. [Higginbotham 2003 7&10]

Based on these discussions, I propose that the inference \neg *remember p* \rightarrow \neg *p* in (37) is licensed by two defeasible contextual assumptions. The first assumption is that the speaker has a memory about the particular event represented by *p*.²² That is, the speaker

²²I abstract away here from concerns about negative events, and assume that the event in question in (37) exists and that it was either an event of being scared or an event of being not scared. Another way to approach this is to say that the speaker either remembers the event *e* or remembers the maximal eventuality *S* of all eventualities *e'* in the relevant time period and $e \notin S$ (see also Krifka 1989 and de Swart 1996 in which the following definition of event negation is used: $\lambda P.\lambda s.[\text{MAX}(s) \wedge \neg \exists e[P(e) \wedge e \subseteq s]]$).

is informed about the event under discussion. This is analogous to the Competence Assumption that is used by researchers in computing scalar implicatures (van Rooij and Schulz 2004; Geurts 2009, a.o.) I argue that this assumption is stronger in cases in which the subject of *remember* is remembering their own experience of the particular event, as in Higginbotham's *de se* examples. The second assumption is based on the idea that insofar as our memory of eventualities track with our beliefs about those eventualities, a speaker's memory represents the speaker's beliefs about the way the actual world was in the past. An assertion of memory can therefore in context be taken as doxastic evidence for or against a description of a particular eventuality, and can license inferences from memory to belief. These assumptions are defeasible in that a speaker can have the reliability of their memory challenged. The following steps apply Local Givenness to (37).

(i) Starting Context:

$$c = W$$

(ii) Assumption of Speaker Memory:²³

$$\lambda w'. \forall w [w \in \text{MEM}(s)(w') \rightarrow \exists e \exists t [\neg \text{scared}(s)(e)(w) \wedge \text{AT}(t)(e)(w)]] \vee \forall w [w \in \text{MEM}(s)(w') \rightarrow \exists e \exists t [\text{scared}(s)(e)(w) \wedge \text{AT}(t)(e)(w)]]$$

(iii) Semantic Denotation of (A):

$$[[A]]^g = \lambda w'. \neg \forall w [w \in \text{MEM}(s)(w') \rightarrow \exists e \exists t [\text{scared}(s)(e)(w) \wedge \text{AT}(t)(e)(w)]]$$

(iv) Enriched Denotation of (A):

$$\lambda w'. \forall w [w \in \text{MEM}(s)(w') \rightarrow \exists e \exists t [\neg \text{scared}(s)(e)(w) \wedge \text{AT}(t)(e)(w)]]$$

(v) Assumption of Speaker Consistency:

$$\lambda w'. \forall w [w \in \text{DOX}(s)(w') \rightarrow \exists e \exists t [\neg \text{scared}(s)(e)(w) \wedge \text{AT}(t)(e)(w)]]$$

²³Contextual domain restriction assumed throughout.

(vi) Context Update:

$$\begin{aligned} W \cap C(\lambda w. \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]) &= \\ W \cap \{w \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]\} &= \\ \{w \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]\} &= c_{LE} \end{aligned}$$

(vii) Existential Closure of (E):

$$ExClo(\llbracket \mathbf{E} \rrbracket^g) = \{w \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]\}$$

(viii) Local Givenness:

$$\begin{aligned} c_L \subseteq ExClo(\llbracket \mathbf{E} \rrbracket^g) \\ \{w \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]\} \subseteq \\ \{w \exists e \exists t [\neg \mathbf{scared}(s)(e)(w) \wedge \mathbf{AT}(t)(e)(w)]\} \end{aligned}$$

The pragmatic assumption associated with (A) is that the speaker has a memory of the particular event being discussed, namely an event of being scared or being not scared. This is given in the Assumption of Speaker Memory in step (ii); $\mathbf{MEM}(s)(w)$ acts as an information state of the speaker containing all those worlds compatible with the memory of the speaker. The semantics of (A) given in (iii) expresses that the speaker does not remember an event of being scared: in all the worlds compatible with the memory of the speaker there was no event (in the relevant time period) in which the speaker was scared. Steps (ii) and (iii) together entail the proposition that the speaker remembers an event of his being not scared. Therefore, an assertion of (A) expresses the proposition given in (iv). Under the assumption that the speaker's memories of the past represent the speaker's beliefs about the history of the actual world, we infer the proposition in (v) from (iv). Step (v) pragmatically asserts that the speaker was not scared; the context is then updated with this proposition in (vi) in the same manner as in example (33). Step (vii) provides the existential closure of (E), and step (viii) shows that the existential closure of (E) is

entailed by its local context, and we correctly predict felicitous elision of (E).²⁴

2.4.3 Polarity reversals are not bidirectional pragmatic entailment: Disjunction polarity reversals

The resourceful reader might at this point object that, instead of jettisoning our familiar bidirectional entailment account, a simpler path is to simply enrich the bidirectional entailment condition to include pragmatic and not merely semantic content. This section shows that a pragmatically enriched bidirectional entailment account still fails to generate the full range of polarity reversal data. Consider example (43) containing disjunction.

- (43) **Context:** Students in a semantics class were given the option to do an extra credit problem, and were required to mark the number of the problem that they did on a spreadsheet accessible by the course's professor and TA. Both the professor and TA thought that John, a student in the class, would have chosen to do a problem. They look at the spreadsheet and see that nothing is marked down under John's name. The TA says to the professor:

Sluice: Either [John_j didn't do an extra credit problem]_A, or he_j didn't mark which one_i [_{TP} he_j ~~did~~ t_i]_E.

In (43), we see that negation is present in the antecedent but not in the ellipsis site.²⁵

²⁴A reviewer correctly observes that the derivation here is very similar to the derivation of the neg-raising example in (33), as both rely on the excluded middle. However, while I follow Gajewski in the proposal that neg-raising verbs induce an excluded middle presupposition, the pragmatic inferences utilized for *remember* are less automatic and more easily defeasible than the presuppositions for neg-raising verbs. We can easily find contexts in which both the Assumptions of Speaker Memory and Speaker Consistency fail to hold; as expected in such contexts, the inference $\neg \text{remember } p \rightarrow \neg p$ also fails to hold.

²⁵Thank you to Jason Merchant (p.c.) for pointing out that these data run counter to the claim made in Merchant 2013a, p. 15 that negation present in the antecedent of a sluicing construction requires a corresponding negation present in the ellipsis site.

An utterance of (43) asserts that either (A) John didn't do an e.c. problem or (E) John did an e.c. problem.

To my skeptical readers, I can provide assurance that such examples, while easily constructed, are also found in nature. For example, (44) is a snippet from a song of the Ohio State University; note the bolded final line.²⁶

(44) Ohio, Ohio
The hills send back the cry [O - H!]
We're here to do or die [I - O!]
Ohio, Ohio
We'll win the game or know the reason why!

The disjunctions in examples (43) and (44) contain contrary propositions: they cannot both be true at the same time. Although we are unable in principle to tell whether the disjunctions are inclusive or exclusive, as the two cases collapse in this instance, an inclusive rule of disjunction would fail to make the correct predictions for the update of the local contexts.²⁷ The dynamic literature provides an existing rule for exclusive disjunction, however, that we can use (Karttunen, 1974a; Heim, 1983b; Kadmon, 2001).

Disjunction for Propositions:

For propositions p, q such that p or q is uttered in a context c :

c_L for $p = c$,

c_L for $q = c + (c \setminus p)$

The proposal says that the local context for the first disjunct of an exclusive disjunction construction is the context c of the conversation at the time at which p is uttered. The local context for the second disjunct is c intersected with the complement of the first disjunct. The intuition for this proposal is that the context for the second disjunct must be

²⁶Thanks to Deniz Rudin for this gem of an example, the best thing OSU has done for me since November, 2003. Go Blue. Full lyrics and history, including, oddly, the song's reliance on an older, University of Michigan song, can be found at http://www.sgsosu.net/osu/songs/i_wanna_go_back.html

²⁷Note that this fact holds regardless of whether the example contains ellipsis.

allowed to include worlds incompatible with the first disjunct, in order to correctly predict that the second disjunct is defined. The following steps show how Local Givenness correctly predicts the availability of the sluice in (43).

(i) Starting Context:

$$c = W$$

(ii) Denotation of (A):

$$C(\llbracket A \rrbracket^g) = \{w \rightarrow \exists x[\mathbf{extra\ credit\ problem}(x)(w) \wedge \mathbf{do}(x)(j)(w)]\}$$

(iii) Denotation and Existential Closure of (E):²⁸

$$ExClo(\llbracket E \rrbracket^g) = \{w \exists x[\mathbf{extra\ credit\ problem}(x)(w) \wedge \mathbf{do}(x)(j)(w)]\}$$

(iv) Local Contexts for (A) and (E):

$$c_{LA} = c = W$$

$$c_{LE} = W \cap \{w \rightarrow \exists x[\mathbf{extra\ credit\ problem}(x)(w) \wedge \mathbf{do}(x)(j)(w)]\}$$

(v) Local Givenness:

$$c_{LE} \subseteq ExClo(\llbracket E \rrbracket^g) =$$

$$\{w \rightarrow \exists x[\mathbf{extra\ credit\ problem}(x)(w) \wedge \mathbf{do}(x)(j)(w)]\} \subseteq$$

$$\{w \exists x[\mathbf{extra\ credit\ problem}(x)(w) \wedge \mathbf{do}(x)(j)(w)]\}$$

Local Givenness is satisfied in step (v) because the local context for E entails the proposition expressed by (E), assuming a classical logic in which a doubly negated

²⁸Note that the *wh*-phrase *which one* is d-linked in the sense of Pesetsky 1987, meaning that it ranges over a salient set in the discourse. One could assume here, following ?, that d-linked *wh*-phrases are referential and therefore leave behind a referentially indexed trace. Existentially closing over this trace would then restrict the possible identity of the thing to which the existentially bound variable can refer to a member of a particular set present in the discourse. However, as the d-linking is orthogonal to the example here, I suppress this issue for the sake of expositional clarity.

proposition equals its unnegated equivalent. We therefore correctly predict felicitous elision of (E).

The possibility of polarity reversal sluices in disjunction constructions illustrates the necessity of local contextual entailment in the current account. The global context of (43) does not entail the proposition that John did any extra credit problems, as both possibilities – of John having done extra credit problems and of him not having done any – are being entertained as possibilities. It is only in the local context of the second disjunct that the proposition that John did extra credit problems is entailed, as the local context excludes those worlds in which John didn't do any extra credit problems. Furthermore, examples such as (43) show that a pragmatically-enriched bidirectional entailment account is insufficient to explain the polarity reversal data, as no pragmatic enrichment of the semantic content of (A) and (E) in (43) will yield bidirectional entailment of the propositions. Instead, the crucial licensing factor in this example is the disjunctive operator – which contributes its heritage properties²⁹ to (A) and (E) – and not the propositional content of (A) and (E) themselves.

2.4.4 Polarity reversals are not entailment at LF: More disjunction

I have so far addressed concerns that the polarity reversal examples might be explained by appeals to syntactic accounts of neg-raising or to a pragmatically enriched bi-directional entailment account, and I have shown that both possibilities fail to account for the entirety of the polarity reversal data. A third possibility is that negation is scoped out of the relevant antecedent at logical form. Proponents of this view might argue that the example above is derived using the following LF form:

²⁹See Karttunen and Peters 1979, Heim 1983b, Kadmon 2001, a.o.

- (45) Either [not [John_j did an extra credit problem]_A], or he_j didn't mark which one;
[TP he_j did t_i]_E.

In such a construction, a standard bi-directional entailment account would predict the possibility of the sluice.

There are two reasons why this possibility does not save a bi-directional entailment account. First, it has for some time been observed that the scope of negation is fixed by its surface position, unlike quantifiers and modals (see an early discussion in Ladusaw 1988 and more recent discussions in Iatridou and Sichel 2011 and Potsdam 2013). That is, negation does not raise or lower at LF. Second, polarity reversal examples are also found in which negation is added into the ellipsis site. For example, the modified corpus example in (46):

- (46) [modified corpus example]

Context: On Dec. 10, Senator McCain sent a letter to the FCC urging the five-member board to end two years of deliberations and decide whether Paxson Communications should be given a license for a Pittsburgh station. Angela J. Campbell, an attorney for opponents to the deal, told the Globe that McCain's letter likely 'tipped' the scales in favor of the decision.

Sluice: "Senator McCain said, 'Either the Board grants the license by December 15 or it explains why [~~the Board didn't grant the license by December 15~~]' and the commission jumped to it and did it that very day," Campbell told the Globe.

This example contains a positive antecedent and a negative elided phrase, the reverse of (43). There is no possibility of scoping the negation out of the elided phrase; because it is unpronounced, the negation is uncontroversially inside of the ellipsis site. The following steps apply Local Givenness to the disjunction in (46).³⁰

³⁰Note that the disjunction does not consist of imperatives, as may appear at quick glance. Imperative

(i) Starting Context:

$$c = W$$

(ii) Denotation of (A):

$$C(\llbracket A \rrbracket^g) = \{w : \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\}$$

(iii) Denotation and Existential Closure of (E):

$$ExClo(\llbracket E \rrbracket^g) = \{w : \neg \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\}$$

(iv) Local Context for (A) and (E):

$$c_{LA} = c = W$$

$$c_{LE} = W \cap \{w : \neg \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\} = \\ \{w : \neg \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\}$$

(v) Local Givenness:

$$c_{LE} \subseteq ExClo(\llbracket E \rrbracket^g) = \{w : \neg \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\} \subseteq \\ \{w : \neg \mathbf{grant\ the\ license\ by\ December\ 15(b)(w)}\}$$

Local Givenness is satisfied in step (v) because the local context for (E) entails the proposition expressed by (E), as they express identical propositions. We therefore correctly predict felicitous elision of (E).

subjects must refer to the addressee (Portner 2004, Kaufmann 2012, a.o.). Additionally, imperatives are argued to express speaker preference (Condoravdi and Lauer 2011, 2012, a.o.). However, this construction does not require the speaker to have a preference for either disjunction. For example: Either Ohio State wins the football game or their fans riot (but both are terrible outcomes). The propositions instead express a simple present futurate (Prince 1971, Dowty 1979, Vetter 1973, Copley 2014, a.o.). For expositional clarity, and because concerns of tense are orthogonal to the current project, I set aside the interesting nature of this use of tense and focus on the core propositional content.

2.4.5 Non-factive negative attitude verbs

Another class of polarity reversal sluices contains non-factive negative attitude verbs. These verbs form a natural class grouped by Asher (1987) as Negative (Indefinite and Definite) Non-factives. I assume here a standard definition of negative verbs as generally those that admit downward entailing inferences, in the sense of Ladusaw 1979. Example (47) illustrates a naturally occurring example.

(47) [modified corpus example 99105, Santa Cruz Ellipsis Project]

We doubt that [Iraq will comply with the mandate]_p]_A, but we don't know why
~~[Iraq won't comply with the mandate]_{E/¬p}.~~

Following Anand and Hacquard 2013, we can think of an assertion of *doubt that p* as having three meaning components. The first is a felicity condition that *s* is uncertain about the truth of *p*. The second is that *doubt* semantically encodes a weak possibility assertion about *p*; that is, *s* doubts that *p* commits *s* to entertaining the doxastic possibility of *p*. Last, *doubt* expresses a preference assertion that $\neg p$ is more likely to the speaker than *p*. Formally, this is expressed as a probability ordering on propositions via direct comparison of the worlds contained within those propositions (see Kratzer 1991's better possibility ordering). The preference assertion that $\neg p$ is more likely to the speaker than *p* is what does the work for our purpose here.

Application of Local Givenness to (47):

(i) Starting Context

$$c = W$$

(ii) Felicity Condition on Assertion of (A):

$$\llbracket A \rrbracket = \llbracket \text{doubt}(p)(s)(w) \rrbracket^{wSg} \text{ is defined iff } [\exists w' [w' \in S' \wedge w' \in p] \wedge \exists w'' [w'' \in S' \wedge w'' \notin p]], \text{ such that } S' = \text{DOX}(s)(w)$$

Doxastic Requirement on (A):

$$\exists w' [w' \in S' \wedge w' \in p]$$

- (iii) Assume a QuD: *Will Iraq comply with the mandate?* This QuD splits the future of the discourse into two spaces. The assertion of (A) expresses a preference of the speaker toward future discourse space II, licensing the stronger inference of belief that $\forall w' [w' \in S' \rightarrow \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w')]$.

$$\text{I. } \{w : \mathbf{comply}(\text{mandate})(\text{iraq})(w)\}$$

$$\text{II. } \{w : \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w)\}$$

- (iv) Local Context Updated with Assertion of (A):

$$W \cap \{w' : w' \in S'\} = \{w' : \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w')\} = c_{LE}$$

- (v) Semantic Denotation of (E):

$$ExClo(\llbracket E \rrbracket^g) = \{w' : \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w')\}$$

- (vi) Local Givenness:

$$c_L \subseteq ExClo(\llbracket E \rrbracket^g) = \{w' : \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w')\} \subseteq \{w' : \neg \mathbf{comply}(\text{mandate})(\text{iraq})(w')\}$$

In step (ii), the speaker must be uncertain about whether p or $\neg p$ in order to felicitously assert (A). If felicitous, the utterance of A semantically asserts that the speaker believes that p is possibly true. The utterance of (A) also asserts a speaker preference for $\neg p$ over p . In context, the assertion of this preference licenses the strengthened inference that the speaker believes that $\neg p$, as shown in step (iii). The assertion of (A) pragmatically asserts that Iraq will not comply, shown in step (iv). The semantic denotation of (E) is given in step (v), and step (vi) shows that Local Givenness is satisfied and that we correctly predict felicitous elision of (E) in (i).

2.4.6 *Until*

Finally, the last class of polarity reversal sluices involves those licensed by punctual *until*.³¹ Example (48) illustrates a naturally occurring example.

(48) [corpus example 94827, Santa Cruz Ellipsis Project]

Context: Some directors said they viewed Clinton’s proposals [to regulate teenagers’ access to tobacco]_A as part of a larger campaign to increase government restrictions on the tobacco industry, and several criticized the administration for emphasizing to teenagers the adverse effects of smoking rather than of alcohol and drugs. “Tobacco_i has been in this country 200 years, but [the administration_j has never talked about it_i until now]_A,” said B. Frank Strickland of Lakeland, Ga.

Sluice: “I don’t know why [~~they_j are talking about it_i now~~]_E. But I do know tobacco does not do what alcohol and dope do to people. Yet they jump on tobacco. Why don’t they jump on the dope crowd?”

The division of pragmatic and semantic labor with *until* is debated in the literature (Karttunen, 1974b; de Swart, 1996; Giannakidou, 2002). I adopt a version of de Swart’s truth conditions containing the actualization of the event, though nothing crucial in the account here hinges on this choice. The truth conditions for an utterance with punctual *until* can therefore be expressed as follows:

$$\begin{aligned} \llbracket \text{until} \rrbracket^g = & \lambda Q. \lambda P. \lambda e. \exists t \exists t' \exists t'' [Q(t') \wedge P(e) \wedge \text{AT}(e, t'') \wedge \\ & \neg \exists e' \exists t'' [P(e') \wedge \text{AT}(e, t'') \wedge t \leq t'' < t']] \end{aligned}$$

In which Q expresses the clock expression in the *until* phrase, P is a property expressing some eventuality *e*, the variables *t* are times with the domain T of times (points or intervals on the time axis and a precedence relation providing a total order on

³¹See Karttunen 1974b and de Swart 1996, cf. Smith 1974; Mittwoch 1977; Declerck 1995, for a discussion of punctual vs. durative *until*.

T) and AT maps eventualities to their location on the time axis. AT(e, t) represents that eventuality e holds at t .

Application of Local Givenness to (48):

(i) Denotation for (A):

$$\begin{aligned} & \exists e \exists t \exists t' \exists t'' [\mathbf{n}(t') \wedge \mathbf{admin\ talking\ about\ tobacco}(e) \wedge \mathbf{AT}(e, t'') \wedge \\ & \neg \exists e' \exists t'' [\mathbf{admin\ talking\ about\ tobacco}(e') \wedge \mathbf{AT}(e, t'') \wedge t \leq t'' < t']] \end{aligned}$$

(A) asserts the following:

- a) that there exists an event e of the administration talking about tobacco and e occurs at time t'' ;
- b) there is a time t' which is the lower bound of time at which e can occur;
- c) there is no event of the administration talking about tobacco that occurs before t' ;
- d) there is a contextually determined time interval $t < t'$ within which e is expected to occur (accounting for the feeling of ‘lateness’ of e (Karttunen, 1974b)).

(ii) The context of (48) implicates that the event did not occur later than the time denoted by Q, or the utterance time; therefore $t'' = t'$.

Temporally Enriched Denotation of (A):

$$\begin{aligned} & \exists e \exists t \exists t' [\mathbf{n}(t') \wedge \mathbf{admin\ talking\ about\ tobacco}(e) \wedge \mathbf{AT}(e, t') \wedge \\ & \neg \exists e' \exists t'' [\mathbf{admin\ talking\ about\ tobacco}(e') \wedge \mathbf{AT}(e, t'') \wedge t \leq t'' < t']] \end{aligned}$$

(iii) Context Update with (A):

$$\begin{aligned} c_{LE} = W \cap \{w : & \exists e \exists t \exists t' [\mathbf{n}(t')(w) \wedge \mathbf{admin\ talking\ about\ tobacco}(e)(w) \wedge \\ & \mathbf{AT}(e, t')(w) \wedge \neg \exists e' \exists t'' [\mathbf{admin\ talking\ about\ tobacco}(e')(w) \wedge \\ & \mathbf{AT}(e, t'')(w) \wedge t \leq t'' < t']]\} \end{aligned}$$

(iv) Denotation of (E):

$$\{w : \exists e \exists t' [\mathbf{admin\ talking\ about\ tobacco}(e)(w) \wedge \mathbf{n}(t')(w) \wedge \mathbf{AT}(e, t')(w)]\}$$

(v) Local Givenness:

$$\begin{aligned} c_{LE} \subseteq \text{ExClo}(\llbracket \mathbf{E} \rrbracket^g) &= \{w : \exists e \exists t \exists t' [\mathbf{n}(t')(w) \wedge \\ &\mathbf{admin\ talking\ about\ tobacco}(e)(w) \wedge \mathbf{AT}(e, t')(w) \wedge \\ &\neg \exists e' \exists t'' [\mathbf{admin\ talking\ about\ tobacco}(e')(w) \wedge \mathbf{AT}(e, t'')(w) \wedge t \leq t'' < t']]\} \subseteq \\ &\{w : \exists e \exists t' [\mathbf{admin\ talking\ about\ tobacco}(e)(w) \wedge \mathbf{n}(t')(w) \wedge \mathbf{AT}(e, t')(w)]\} \end{aligned}$$

Step (v) shows that the local context of (E) entails the proposition expressed by (E) and Local Givenness is satisfied.³²

2.4.7 Attitude holders

As a last note, Dan Hardt (p.c.) and a reviewer for *Semantics & Pragmatics* pointed out to me that polarity reversal examples appear to contain a consistent attitude holder through the context and the embedding clause of the elided phrase. This, however, does not necessarily need to be the case, though context is generally needed to get the desired reading:

- (49) **Context:** Mary and John have been discussing the administration's position on immigration reform. Mary believes that the administration officials regret having their current hardline position, and may enact different policies. John has been trying to convince her that they will maintain their current position, though he

³²Karttunen 1974b observes that punctual *until* focuses on the onset of the event denoted by the main clause as opposed to the absence of that event in the time period leading up to Q. A proposition expressing the realization of the event will therefore also be more salient in the context than the proposition expressing the negation of that event. I believe that speaker variation on this example is due to individual variation on the relevant level of salience needed for elision. See below for a discussion of the salience requirement.

admits that he doesn't understand why the officials won't change their minds given the political unpopularity of the current policy. After debating both issues for some time, John thinks he has convinced Mary to see his side, though he still hasn't come up with an explanation for why officials are being so inflexible. So at the end of the argument we could say:

Sluice: Even though Mary finally doubts that the administration will change its position on immigration reform, John still can't tell her why [~~the administration won't change its position on immigration reform~~].

While this generalization does generally hold for the corpus and constructed polarity reversal examples is an interesting question, and unfortunately one that I must leave to future research.

2.4.8 Failure to license

This section provides an example of the account's ability to correctly rule out an impossible ellipsis site. Examples (43) and (50) contain minimally different contexts and sluices; however, while the sluice in (43) is acceptable, the sluice in (50) is not.

(50) **Context:** Students in a semantics class were given a set of extra credit problems, which they could choose to do up to half of. All students were required to put a mark on a preadsheet next to each question, indicating whether they did or didn't do it. The professor and TA look at the spreadsheet and see that John has not put a mark next to all of the questions. The TA says to the professor:

Impossible Sluice: [John_j marked which problems he did]_A, but he_j didn't mark which problems_i #~~[he_j didn't do t_i]~~_E.

While the example is acceptable in its un-elided form, the sluiced interpretation given in (50) is impossible. We predict this result because the conjunction in (50) does

not yield the same local context for the second conjunct as the disjunction does in (43). Recall our dynamic rule of conjunction stated above, which is motivated in the literature on presupposition projection in conjoined clauses: $c + (p \wedge q) = (c + p) + q$. Because the local context for the elided phrase contains the global context plus the proposition expressed by (A), $\{w : \exists x[\text{PL}(x)(w) \wedge \mathbf{problem}(x)(w) \wedge \mathbf{mark}(x)(j)(w) \wedge \mathbf{do}(x)(j)(w)]\}$, the local context clearly does not entail the proposition expressed by (E), $\{w : \exists x[\text{PL}(x)(w) \wedge \mathbf{problem}(x)(w) \wedge \neg \mathbf{do}(x)(j)(w)]\}$. Local Givenness is therefore not satisfied and we correctly predict the infelicity of the sluice in (50).

This section has discussed the possibilities of accounting for the polarity reversal sluices with a bidirectional semantic entailment account, a bidirectional pragmatic entailment account, and accounts that scope negation outside of the relevant antecedents either syntactically or at LF. I have shown that none of these possibilities can capture the full scope of the polarity reversal examples. I have additionally shown that a contextual entailment condition like Local Givenness successfully accounts for the full range of examples. The next section addresses concerns that have historically been levied against non-structural ellipsis licensing conditions.

2.5 Structural constraints on sluicing

The analysis that has been proposed so far is necessarily more permissive than existing syntactic or semantic entailment accounts of sluicing. While this additional permissiveness is required in order to capture the structural and semantic differences between the preceding linguistic context and the elided phrases in polarity reversal sluices, it also invites concerns of overgeneration. There are two ways Local Givenness can avoid these concerns. The first is to follow existing accounts in combining the contextual constraint of Local Givenness with an independent syntactic constraint. Section 2.5.1 and 2.5.2

outline what such a syntactic constraint would need to look like. While this approach is promising, §2.5.3 shows that even combining Local Givenness with a syntactic constraint does not yield a theory sufficient to capture the data. Instead, some notion of salience must also be integrated into the account. Because a salience constraint is independently necessary, §2.5.4 proposes a way in which we might define a notion of propositional salience for ellipsis.

2.5.1 Data motivating syntactic constraints on sluicing

A concern levied at non-structural analyses of sluicing is that they fail to capture certain identity requirements of sluicing constructions that appear to have a syntactic basis. One of these requirements is case matching, which was first noted in Ross 1969 and is discussed at length in Merchant 2001. Case matching is the generalization that the remnant in a sluiced construction must match in case with its correlate. Merchant (2001) accounts for this behavior by proposing that the correlate is assigned case in the elided structure, before movement. The case matching facts are then an argument for the existence of syntactic structure in the elision site. Because I adopt the position that the elision site contains syntactic material, Merchant's account of case matching applies equally to the pragmatic theory outlined in this paper as it does to his own account of bidirectional semantic entailment.

A second syntactic identity condition, proposed in Chung 2013, is a generalization that has since been named Chung's Generalization. The generalization states that a preposition can be stranded by a remnant in the ellipsis site only when the remnant corresponds to syntactic material in the antecedent clause. The Generalization captures the paradigm given in (51)-(53).

- (51) [John is flirting]_A, but I don't know with who_i [~~John is flirting~~ _{t_i}]_E.

(52) [John is flirting]_A, but I don't know #who_i [~~John is flirting with t_i~~]_E.

(53) [John is flirting with someone]_A, but I don't know who_i [~~John is flirting with t_i~~]_E.

These data present a difficulty for meaning-based accounts relying only on entailment licensing conditions. If we assume that the act of flirting entails the act of flirting with someone and vice versa, then bidirectional semantic entailment holds between the antecedent and elided clauses in (51)-(53). A meaning-based entailment condition would therefore fail to rule out the impossible sluice in (52). Since these observations, researchers advocating for meaning-based licensing conditions have opted to include a structural constraint in their theory in order to account for these facts. The next section discusses some ways in which this has been undertaken, and proposes a path forward that unites structural constraints of sluicing with Local Givenness.

2.5.2 Local Givenness plus structural constraints

Since Chung's observations of the facts above, meaning-based accounts have incorporated an independent structural constraint to account for the data. For example, AnderBois (2014) follows Merchant (2007) in using the constraint No New Morphemes, which states that the morphemes contained within the elided expression of a sluice must be a subset of the morphemes contained within the antecedent expression. The polarity reversal data show us, however, that No New Morphemes cannot be the generalization needed. For instance, example (46), repeated here as (54), contains a negation morpheme (bolded) in the ellipsis site that is not present in the antecedent site.

(54) Senator McCain said, 'Either the Board grants the license by December 15 or it explains why [~~the Board did **not** grant the license by December 15~~]' and the commission jumped to it and did it that very day,' Campbell told the Globe.

The morphemes contained within the ellipsis site of 54 are not a subset of the morphemes contained within the antecedent, in opposition to the predictions of No New Morphemes.

While No New Morphemes is not the formulation we need, it is still possible to adopt a limited syntactic identity constraint that works in conjunction with Local Givenness to account for the facts. Such a path follows Chung (2013), Merchant (2013b), Barros (2014), and AnderBois (2014) in proposing a theory that relies on limited structural constraints in conjunction with a contentful meaning-based licensing constraint. One obvious way to proceed is to adopt into the current account Chung's more finely articulated constraints on which No New Morphemes is based. Chung proposes two constraints, given below:

(55) Limited Syntactic Identity in Sluicing:

- (i) Argument structure condition: If the interrogative phrase is the argument of a predicate in the ellipsis site, that predicate must have an argument structure identical to the corresponding predicate in the antecedent clause.
- (ii) Case condition: If the interrogative phrase is a DP, it must be Case-licensed in the ellipsis site by a head identical to the corresponding head in the antecedent clause.

Merchant (2013b) operationalizes these conditions to account for both verb phrase ellipsis (VPE) and sluicing facts by proposing a limited syntactic identity condition that holds between syntactic heads within an ellipsis site and syntactic heads within some preceding structure. While not aware of the polarity reversal data, Merchant's account proposes that the syntactic identity condition for sluicing holds over a greater swath of syntactic structure than it does for VPE. Specifically, that sluicing requires syntactic identity from the bottom of the tree to a head above Voice, while VPE only requires matching from the bottom of the tree to vP.

While Merchant leaves open the exact formulation of the limited syntactic identity constraint needed, Rudin (2019) and Kroll and Rudin (2017) propose one possible implementation of Chung and Merchant’s ideas. The account proposes that syntactic identity holds only over heads within the eventive core of the elided clause, namely heads including and below the highest vP that is associated with an event-introducing verb. Crucially, because the syntactic identity condition holds only over the eventive core of the elided clause (bolded in (56)), any mismatch in structure or lexical material is allowed above the highest vP domain. This account correctly predicts the possibility of polarity mismatches, as polarity nodes are located above the highest vP domain (Laka, 1990; Ladusaw, 1992). For example, (15) is repeated as (56).

(56) I don’t think that [California will comply]_A, but I don’t know why [~~California won’t comply~~]_E.

A: [TP California [T will [P_oIP [vP **[vP comply]**]]]]

E: [TP California [T will [P_oIP not [vP **[vP comply t]**]]]]

The account also correctly rules out the data violating Chung’s generalization. It is clear in (52), repeated as (57), that the elided PP has no syntactically matching antecedent in the preceding discourse, and therefore violates the limited syntactic identity condition.³³

(57) [John is flirting]_A, but I don’t know #who_i [~~John is flirting with t_i~~]_E.

A: [TP John [T [P_oIP [vP **is [vP flirting]**]]]]

E: [TP John [T [P_oIP [vP **is [vP flirting [PP with t]**]]]]]]

³³Note that traces are not included in the structure-matching requirements of the theory. See Rudin 2019 for an extensive defense of these claims. For space concerns and because the current account is not dependent upon the specific implementations of this theory, I omit the details here.

In summary, one way to deal with Chung’s facts is to follow previous accounts and integrate into Local Givenness an independent limited syntactic identity constraint, as proposed in Chung 2013 and Merchant 2013b. Rudin (2019) and Kroll and Rudin (2017) argue that their account presents one possible way of doing this. However, Local Givenness need not rely on any particular formalization of the syntactic facts discussed here. The data have presented tenacious challenges for purely meaning-based approaches to sluicing ever since the original observations. As Merchant (2013b) states, the exact form of a syntactic identity formulation for ellipsis remains to be fully explicated, and the proposal here is not intended to advance this particular line of research. I hope merely to point out the arguments that have been presented in favor of combining meaning-based sluicing accounts with a limited syntactic identity condition. Ultimately, any limited syntactic identity constraint that captures the facts is in principle available to us. The contribution of the current proposal is to observe that such a constraint, whatever its exact formulation, must be permissive enough to allow for high syntactic mismatches such as polarity.

2.5.3 Structure plus pragmatics: A full picture?

In the previous section, we saw the arguments in favor of integrating an independent limited syntactic constraint into a meaning-based licensing account. However, even combining a meaning-based account with a syntactic constraint is not sufficient to alleviate over-generation concerns. Consider (a)-(b) (Cantor, 2013). In the desired interpretations, the remnants *who* are linked to a discourse referent that was introduced, via *some pedestrian*, within a discourse subordinated relative clause:

- (58) (a) #That John rented a car that hit some pedestrian surprised everyone, but the report didn’t say who_i [~~that John rented a car that hit t_i surprised everyone~~].

- (b) #A car that hit some pedestrian crashed into the wall last night, but the report didn't say who_i [~~a car that hit t_i crashed into the wall last night~~].

The sluices are easily judged to be infelicitous. Similarly, as we saw in (25), repeated as (59), sluicing with a simple *wh*-remnant out of a discourse subordinated appositive clause is also infelicitous.

- (59) #Joe, who once killed a man in cold blood, doesn't even remember who [~~he once killed in cold blood~~].

Because (a)-(b) and (59) all have matching linguistic antecedents, any theory that licenses sluiced clauses based only on the existence of an appropriately matched antecedent will fail to rule the examples out. Additionally, the infelicity of the examples is not due to a contextual licensing failure or a structural constraint failure: the examples have matching linguistic antecedents, satisfy Local Givenness, and do not violate any of the discussed syntactic constraints. We can confirm that there is no structural or licensing factor driving the infelicity of examples (a)-(b) by noticing that they become acceptable with the d-linked remnants *which pedestrian* (Cantor, 2013, p. 27-28):

- (60) (a) That John rented a car that hit some pedestrian surprised everyone, but the report didn't say which pedestrian_i [~~that John rented a car that hit t_i surprised everyone~~].
- (b) A car that hit some pedestrian crashed into the wall last night, but the report didn't say which pedestrian_i [~~a car that hit t_i crashed into the wall last night~~].

Additionally, we already observed in (26) the improvement of (59) with the d-linked remnant *which man*.

D-linking the remnant in (60) increases the acceptability of the sluice by providing an overt cue to the proposition that must be recovered. That is, the d-linked referent

signals that the elided proposition contains an extracted argument with the semantic and syntactic properties given by the remnant, thereby making the most recent propositional content given about that argument salient at that point in the discourse (see Martin and McElree 2011 and Harris 2019, a.o.). For comparison, the most salient correlate for the remnant in 58 is, due to a combination of recency and discourse status (see e.g. Jarvella 1971, van Dijk and Kintsch 1983, Kintsch 1988, Anderson 2007, a.o.), *everyone*, which leads to an infelicitous interpretation of the question. Note the improvement if an existential quantifier is used as the correlate instead:

- (61) That John rented a car that hit some pedestrian surprised someone, but the report didn't say who_i [~~that John rented a car that hit some pedestrian surprised t_i~~].

The ability of a d-linked remnant to link to a discourse subordinated correlate more easily than a bare wh-remnant is perhaps not surprising given the many years of careful experimental and theoretical work on related phenomena. Though I remain agnostic here about the particular analysis of d-linked phrases one may wish to take, the data here fit naturally with studies of anaphors showing that the more semantic content an anaphor has, the greater the distance that may separate the anaphor from its antecedent and the less prominent the antecedent can be while still maintaining successful reference (Givón, 1983; Ariel, 1990; Gundel et al., 1990; O'Brien et al., 1997; Foraker and McElree, 2006). The related observation that the amount of semantic information in an extracted element appears to affect its ability to be felicitously extracted has also been made in the theoretical literature (Karttunen 1977, Rizzi 1990, Cinque 1990, Pesetsky 1987, Chung 1994, Hofmeister 2007, 2011, a.o.).³⁴ While these questions are still the subject of active research, I bring them up to demonstrate that the sluicing data fit within a much larger research project of the role of salience in anaphoric and syntactic dependencies in

³⁴Thank you to Shayne Sloggett (p.c.) for bringing these works to my attention.

language.

The data discussed so far in this section demonstrate that structural and contextual/antecedent licensing requirements by themselves are not enough to constrain a theory of sluicing sufficiently to avoid overgeneration. This observation is not a novel one, and has been made by many previous researchers working on ellipsis. Because of this, explicit salience requirements are built into most if not all existing sluicing accounts, including each sluicing theory discussed here. For example, e-GIVENness (Merchant, 2001) requires that the elided constituent have “a salient antecedent,” Inquisitive Entailment (AnderBois 2014) requires “some salient antecedent CP,” and Barker (2013)’s scopability analysis requires “a silent proform that is anaphoric to some salient discourse object” (pg. 193). The thread tying all these accounts together is the recognition that salience is intricately tied to our ability to elide and recover propositional or sub-propositional expressions in conversation.

One option here is to follow in the footsteps of these analyses and build a notion of salience into Local Givenness. This constraint is given in (62).

- (62) **Local Givenness (final):** A TP α can be deleted *iff* $ExClo([\alpha]^g)$ expresses a proposition p , such that $c_L \subseteq p$ and p is maximally salient.

The final version of Local Givenness proposes that a proposition can be elided if the existential closure of the proposition is entailed by the local context and is maximally salient.³⁵ The requirement that the elided proposition must be maximally salient is

³⁵An earlier version of this theory used uniquely salient; this was changed after a reviewer asked about ambiguous readings of sluices such as in (i):

- (i) John told me that Mary is going to Canada, but I don’t know why [].
a. why John told me Mary is going to Canada
b. why Mary is going to Canada

The reviewer asks whether Local Givenness’s salience constraint predicts the possibility of such readings,

motivated by the common sense principle that in order for a speaker to felicitously not pronounce some part of an utterance, the meaning of the unpronounced piece of the utterance must be recoverable in the discourse. Put another way, this principle encompasses the idea that the content of an elided phrase must be sufficiently salient in the discourse such that it is recoverable in the discourse even though it is unpronounced.³⁶

I have argued in this section that, regardless of appeals to limited syntactic constraints, any theory of ellipsis must include a notion of salience in order to account for facts that cannot be explained by contextual, antecedent-based, or structural licensing conditions. As an exercise in parsimony, we might wonder whether the predictions of the salience constraint can also capture other facts, such as those that have been proposed to be syntactic in nature. For the remainder of this section, I preliminarily sketch some possible ways in which a salience condition could be used to account for the facts discussed here.

One identity condition discussed above is the Case condition, which is repeated here which are widely observed in ellipsis and anaphora. There are two possibilities for capturing these facts. The first is that the ambiguity is one of the context and, as is the case with anaphoric dependencies, the salience of (a) and (b) are partially established by non-linguistic top-down information such as conversational goals and expectations (Brown-Schmidt et al., 2005). Under this view, either (a) or (b) is maximally salient in a given particular context. The second possibility is that in a given linguistic context both (a) and (b) are salient to the same degree, in which case propositional salience can be modeled formally as a partial order in which both (a) and (b) are maximally salient. I am in principle open to either of these possibilities, and leave the decision of which more accurately captures the data to future research.

³⁶A reviewer brings up Hartman (2009)'s asymmetrical sluicing examples:

(i) Someone won the chess game, but I don't remember who ~~#[lost the chess game]~~.

While the propositions expressed by *Someone won the chess game* and *Someone lost the chess game* are (in most standard circumstances) denotationally equivalent, the account here does not predict that any proposition denotationally equivalent to (A) can be elided. This type of interesting observation is what motivates salience requirements in ellipsis theories. Even if we accept that such examples are properly entailed, they are ruled out by a salience requirement, as I can think of no argumentation for such a proposition being salient over the preceding proposition expressed by *someone won the chess game*.

(Chung, 2013).

Case condition: If the interrogative phrase is a DP, it must be Case-licensed in the ellipsis site by a head identical to the corresponding head in the antecedent clause.

Chung 2013, 48 argues for the Case condition based in part on the data in (63):

- (63) #The message said [to show up in the square at midnight]_A, but it didn't say who_i [~~t_i should show up in the square at midnight~~]_E.

While the remnant here is a subject of a finite clause and will be assigned Case by the finite T in the ellipsis site, the corresponding T in the antecedent clause is non-finite and therefore will not assign case (antecedent bracketing and labelling my addition). However, naïve speakers judge (63) to be perfectly felicitous in the following constructed context:

- (64) **Scenario:** There are two gangs fighting for turf, the Sharks and the Jets. After a verbal spar at the gym, the two gangs decide that each gang is going to send a single member to fight at a particular location and at a particular time in two nights time. Each gang gets to pick the member of the other gang that will fight. Additionally, the Sharks get to pick the time of the fight and the Jets get to pick the location of the fight. That night, the Jets send a message to the Sharks telling them to send the Shark gang leader to the West Side Square for the fight. The next night, the Jets receive a message from the Sharks in return. A Jet member opens the message and reads it. He says to the other members in an annoyed voice:

Sluice: “the message says to show up at the Square at midnight, but it doesn't specify who.”

The judgment given for the continuation of the sluice is “who [should show up at the Square at midnight].” The discourse was constructed in such a way as to preclude the possibility that an antecedent for the sluice could be found not in the preceding sentence, but in the preceding discourse.³⁷ If the infelicity of (63) is truly rooted in syntactic ill-formedness, as the Case condition claims, it is unclear why the example becomes acceptable in context. A salience condition could possibly explain why additional context facilitates the sluiced structure.

Additional facts that remain difficult for meaning-based licensing conditions are examples like (52), repeated in (65).

(65) [John is flirting]_A, but I don't know #who_i [~~John is flirting with t_i~~]_E.

While a detailed analysis of these facts is outside the scope of the current paper, it is plausible that the *ExClo*([E]^g) in (65) satisfies the entailment condition of Local Givenness but fails to satisfy its salience condition, which is based not on entailment but on the prominence of the proposition expressed by (E) in the discourse. I leave this interesting question as an area of future research.

2.5.4 Salience

I have followed previous accounts by including a salience constraint in Local Givenness to account for certain facts that are not explained under purely meaning or structure-based accounts. I have also tentatively suggested that such a constraint can be expanded to encompass certain facts that have previously been attributed to constraints on syntactic structure. Because I have motivated the inclusion of a salience condition into Local

³⁷Pseudosluicing is a possible explanation for the lack of case connectivity effects observed here. However, such an argument would need to explain why pseudosluicing is not freely available in English and why such a strategy would only be available here in a specific context.

Givenness, it is worth discussing what such a condition would look like and whether we can constrain it enough to have predictive power. Determining the salience of a given proposition is, of course, not trivial. This paper does not attempt to provide an exhaustive theory of salience, which remains a challenging and important continuing line of research. Instead, I argue here for two points: that salience is indeed relevant to sluicing, and that a theory of propositional salience can be constrained in such a way as to offer concrete, testable predictions. The reader who is already familiar with these facts or is uninterested in accounting for salience can safely skip this section.

While there is currently no rigorous definition of what it means for a given proposition to be salient enough to license ellipsis, some existing theories encode salience in discourse organizational terms. For example, Ginzburg and Sag (2001) and Barros (2014) encode the notion of salience as relevance to a Question under Discussion (QuD) (Roberts, 1996; Ginzburg, 1996; Roberts, 2012), and require the interpretations of sluices to be resolved to a Maximal Question under Discussion.³⁸ Similarly, Frazier and Clifton (2005) encode salience in the notion of main assertion. They claim that, *ceteris paribus*, comprehenders prefer to resolve ellipsis and anaphoric expressions to the main assertion of the preceding sentence (see also Syrett and Koev 2014). In their proposal, this preference follows from the greater salience, or availability, of the main assertion in the discourse structure.

Another linguistic phenomenon known to be sensitive to discourse organization and salience is anaphora. For example, the anaphor *that* is known to specifically license propositional anaphora of salient propositions (Webber 1988, Gundel et al. 1990, Hwang 1992, Asher 1993, a.o.). As we see in (66), B's response is more easily interpreted as a reply to the matrix assertion about Diane winning the race than it is as a reply to the claim about her age.

³⁸Ginzburg & Sag also have an explicit salience requirement, SAL-UTT, that requires a salient utterance to serve as the correlate for sluicing constructions.

(66) A: Nancy, who's 79 years old, will lose the race.

B: That's not true!

We can tentatively note that this type of anaphoric reference to the sluiced content is also possible in the polarity reversal cases, as shown in (67).³⁹

(67) A: I don't know why [~~I wasn't scared~~]_i, but I really can not remember being scared.

B: That_i's impossible! You were just a child.

In (67), the deictic demonstrative *that* is anaphoric to the sluiced proposition expressed by *I wasn't scared*. That is, the meaning of the first sentence in B's utterance is "It's impossible that you weren't scared."⁴⁰ The ability to pick out the elided proposition with an anaphor that selects salient propositions provides independent evidence that, at least in the example above, the elided proposition is salient in the context.

³⁹The ability of elided content to contribute antecedents for discourse anaphora has been widely documented in the literature, notably in Hankamer and Sag 1976.

⁴⁰We do not predict or expect, of course, a one-to-one correspondence between propositions that can be anaphorically referred to with *that* and propositions that can be sluiced. Additional requirements exist on sluicing that do not exist on *that*, such as entailment. For example, in (i) the proposition expressed by the sentence *Roy would win the election* can be referred to with *that*, but it is clearly not entailed by the local or global context.

(i) Roy_i was mistaken that [he_i would win the election]_j; that_j was never a possibility.

There are also likely constraints on *that* that are not constraints on sluicing. For example, Murray (2014) argues that not-at-issue content, such as content contributed by the canonical use of appositives, does not necessarily automatically introduce a discourse referent. Therefore, although we predict that not-at-issue content can be sluiced, we would not necessarily predict that such a sluiced proposition could be referenced with *that*. The set of propositions that can be sluiced and the set of propositions that can be referred to with *that* are therefore likely overlapping but not identical sets.

Additional experimental and theoretical work on anaphora and ellipsis can also inform ways in which we can constrain our notion of salience. A reviewer points out the following contrast:

- (68) (a) Sue bought a car. It's blue. ?/I'm not sure where, though.
(b) Sue bought a car that is blue. I'm not sure where, though.

While the sluice is acceptable in (b), it is not in (a). It seems natural to attribute the difference to a matter of salience, as the reviewer suggests. And indeed, existing work on salience can help us understand why the first proposition is more difficult to target in (a) than in (b).

One way to explain such data is to appeal to a Frazier & Clifton-style account. Such an account predicts the availability of the sluice *I'm not sure where* [*she bought a car that is blue*] in (b), since the buying event is the main asserted content of the preceding sentence. That account would also rule out the sluice *I'm not sure where* [*she bought a car*] for (a), since the sentence *It's blue* intervenes between the sluice and its antecedent. However, many authors working on the relevance of discourse relations and structure to anaphora resolution and ellipsis have observed that, although recency is a strong tendency, it does not always hold that the most recently encountered proposition contains the most salient topic (Hobbs 1978, Polanyi 1986, Webber 1988, Asher 1993, Kehler 2002, Asher and Lascarides 2003, Hardt and Romero 2004, Asher 2008, Harris 2015, a.m.o.). For example, in (69) the elided content is about Mary kissing someone, not about John leaving (Asher, 2008, 9a):

- (69) Mary kissed someone because John left for some other party. You'll never guess who.

To account for such data, Hardt and Romero (2004) propose a Discourse Condition on Ellipsis Resolution, in which an antecedent clause must locally c-command the elided

clause in the discourse tree. While Local Givenness cannot directly integrate such an account, as the theory has no explicit antecedent clause, the force of the proposal is to enforce a kind of local Right Frontier Constraint⁴¹ on a sluiced clause, which in turn requires the elided proposition to be salient or “on topic.” Asher (2008) takes this approach to data such as (69), and proposes that the Right Frontier Constraint holds for sluiced and VPE clauses. Asher’s theory of sluicing resolution requires two things: that the elided material of a clause be recovered from the discourse constituent that the clause is attached to, and that the ellipsis clause be attached to the antecedent clause by at least the structural relations Parallel or Contrast.

There are (at least) two interesting consequences of this theory. The first is that the theory, similar to a Frazier & Clifton-style account, imposes a type of locality restriction on sluiced clauses; that is, the clause containing the elided proposition must attach by a certain discourse relation to a discourse constituent on the right frontier of the discourse. This requirement forces the elided clause to be “about”, in an atheoretical sense, a recent proposition in the discourse. Asher (2008) and Hardt and Romero (2004) argue convincingly that the locality constraint is not sensitive to mere discourse constituent adjacency, but to the richer representations given by theories of discourse relations. However, these constraints still result in a strong advantage given to the last discourse constituent introduced in the discourse, such that we generally try to resolve the ellipsis site to modifying the most recent constituent in the discourse. For example, in (a) the oddness of the sluice disappears once the remnant is one that can plausibly be about the same topic as the previous proposition:

⁴¹The Right Frontier Constraint (see Polanyi 1986, Webber 1988, Asher 1993, Asher and Lascardes 2003, a.m.o.) is a discourse structure constraint requiring new discourse constituents to attach into an existing discourse only by attaching to the last simple constituent in the discourse or by attaching to any constituent that dominates the last simple constituent in the discourse. This constraint, very roughly, captures some notion of topicality.

(70) Sue bought a car. It's blue. I'm not sure on what parts [~~it's blue~~], though.

The second consequence of Asher's theory is that the elided clause must bear a contrasting or parallel relationship to the recent proposition in the discourse. This forces a particular type of relationship to hold between the elided proposition and the proposition to which it is attached in the discourse, and again requires that the elided proposition is salient or on topic in the discourse. An Asher-style account, for example, could rule out the sluice *I'm not sure where* [~~she bought a car~~] in (a) above by arguing that it does not attach to its antecedent clause by either a Parallel or Contrast relation.

I am not advocating here that an appeal to discourse relations will completely capture the full array of facts. However, theories of discourse relations are one existing way of spelling out in some detail what a theory of propositional salience may look like. There are challenges, however, in determining when certain propositions are subordinating or coordinating, and there is disagreement in the literature about the status of various relations (Kehler, 2002; Asher and Lascarides, 2003; Asher and Vieu, 2005). The d-linking data also present a challenge to an account of salience that relies purely on subordinating and coordinating relationships in the discourse. For example, the remnant in (a) can be interpreted as modifying the subordinated constituent if the remnant is modified such that this is the most salient interpretation:

(71) Sue bought a car that is blue. I'm not sure on what parts [~~it's blue~~], though.

The possibility of this reading indicates that the subordinated constituent content is in principle available in the discourse; the interpretation of the sluice is merely dependent on the remnant signaling the most salient and available interpretation of the elided material.

I have outlined here some possible ways to constrain a theory of propositional salience in relation to sluicing. The discussion here is not an attempt to exhaustively resolve this issue, which is an area of interesting future research. I hope merely to have convinced

my reader of two things. The first is that salience is indeed relevant to sluicing, as has been argued by many researchers before. The second is that a theory of salience can be constrained in such a way as to offer concrete, testable predictions. The discussion here has presented some possible avenues for future research that operationalize propositional salience using two factors: the recency of discourse constituents, and the discourse status of propositions within an articulated framework of discourse organization and relations.

2.6 Non-isomorphic sluicing

A reviewer mentions the possibility that the sluices discussed here are derived not by the underlying syntactic structure presented in this paper, but by some alternative, smaller structure that is pragmatically licensed by Local Givenness. Since we cannot see what is elided in ellipsis constructions, we oftentimes cannot argue with certainty that the form proposed is indeed the actual underlying linguistic structure. However, previous research has laid out a landscape of possibilities for possible “short form” sluices, more generally called non-isomorphic sluices. This subsection discusses these possibilities and shows that, while it is possible that some polarity reversal sluices are of the non-isomorphic variety, it is implausible that all polarity reversal sluices are non-isomorphic. Therefore, at least some of the polarity reversal sluices presented here must have a richer underlying structure.⁴²

⁴²Alternatively, it could be possible that the polarity reversal cases embody some new form of ellipsis that is separate from those that are described as sluices in the literature. I am however unaware of any convincing argumentation to motivate this view.

2.6.1 Types of non-isomorphic sluices

The possibility of non-isomorphic sluicing was first noted by Pollmann (1975) and Erteschik-Shir (1977). Much of the discussion since centers on the debate over whether or not sluicing ameliorates islands. For example, (72) is judged as acceptable, but under an isomorphic reading contains extraction of the remnant from within a relative clause island (Merchant, 2001, p. 152):

- (72) They want to hire someone who speaks a Balkan language, but I don't remember which Balkan language_i [~~they want to hire someone who speaks t_i~~].

Theories of non-isomorphic sluicing propose that, instead of the elided content in 72, such putatively island-violating examples are instead derived from a non-island-containing elided structure, such as that given in (73):⁴³

- (73) They want to hire someone who speaks a Balkan language, but I don't remember which Balkan language_i [~~it was t_i~~].

Because the elided clause does not contain a relative clause, we avoid any concerns about extraction from within an island.

Non-isomorphic sluices can be grouped into three main types. The first group is true pseudosluices, which are derived by a null copula and a null subject instead of by TP deletion (Vicente, 2018). Pseudosluices have been argued to exist in Japanese (Kizu, 1997; Merchant, 1998) and in Uzbek (Gribanova, 2013). However, since pseudosluices have not been argued to exist in English (Vicente, 2018), I do not discuss them further here but refer the reader to the cited works.

⁴³For clarity, I refer to sluices such as (73) as non-isomorphic sluices, though the term *pseudosluicing* is often also used as a general term for non-isomorphic sluices (Merchant, 2001); I follow Vicente (2018) here in using *pseudosluicing* in the strictest sense in which it is used in the literature.

The second group of non-isomorphic sluices contains those derived by the deletion of a copula clause. Two types of copular clause deletion have been proposed: specificational copular clause deletion (e.g. see Rodrigues et al. 2009 for arguments about Spanish) and predicational copular clause deletion (Barros et al., 2013). Specificational copular clause deletion is often discussed to account for the apparent violation of Merchant (2001)'s P-stranding condition in languages such as Finnish, Indonesian, and Spanish (see Vicente 2018 and works cited within for extensive defenses of these claims). The proposal in summary is that the apparent violations of Merchant's generalization are not actual counter-examples because the underlying linguistic structure contains a specificational copular clause, as in (74):

(74) John is jealous of some woman_j, but I don't know who_i [~~she_j is t_i~~].

The concern with extending this analysis to the examples discussed here is that Merchant's generalization does appear to hold in English. Therefore, if one wants to attribute all the sluices here to specificational copular deletion, one would need to have a theory of when such a strategy is and is not available in English and, moreover, why it is not freely available.

Predicational copular clause deletion is proposed by Barros et al. 2013 to be available as an island-evasion strategy in English. For example, in (75) they give the following sentence in (a) in which the elided structure contains a predicational copular clause, instead of an isomorphic left branch extraction as in (b) (ibid., 28):

(75) (a) Mary married a tall man, but I'm not sure how tall_i [~~he was t_i~~].

(b) *Mary married a tall man, but I'm not sure how tall_i Mary married a t_i man.

Barros et al. use such data to expand the argumentation against true island repair in sluicing, suggesting that instead some apparent ameliorations are actually caused by evasion strategies such as copula deletion.

The last group of non-isomorphic sluices are those derived by the deletion of a cleft (Erteschik-Shir, 1977; Pollmann, 1975; Rosen, 1976; Merchant, 2001; van Craenenbroeck, 2010; Barros, 2014). For example, in (76) the sentential subject extraction violation of (b) is avoided in (a) by eliding a cleft instead of an isomorphic island structure (Vicente, 2014, p. 33-34):

- (76) (a) That they will hire someone is possible, but I don't know who [~~it will be~~].
(b) *That they will hire someone is possible, but I don't know who_i that they will hire t_i is possible.

Merchant (1998, 2001), importantly, observes that the clefting strategy is not available with wh-adjuncts or implicit arguments, as illustrated in 77 (Merchant, 1998, 53a).

- (77) He fixed the car, but I don't know why *[it was].

2.6.2 Polarity reversals cannot be reduced to non-isomorphic sluices

The last subsection enumerated three main types of non-isomorphic sluices. The following examples show that no combination of these three possibilities can extend to the full range of polarity reversal sluices. Because pseudosluicing has not been argued to exist in English, I focus on copula clauses and clefting. Let's take the polarity reversal example in (78). I have proposed that the elided content contains the structure as shown here.

- (78) **Context:** Students in a semantics class were given the option to do extra credit problems, and they were required to write down the number that they did on a spreadsheet accessible by the course's professor and TA. Both the professor and TA thought that John, a student in the class, would have done at least some extra credit problems. They look at the spreadsheet and see that nothing is marked down under John's name. The TA says to the professor:

Sluice: Either John_j didn't do any extra credit problems, or he_j didn't tell us how many (extra credit problems) [he_j did].

The first non-isomorphic alternative is to use a copula in the elision site, as in (79). However this strategy appears unavailable, as speakers judge the example to be degraded or unacceptable.⁴⁴

(79) ??/*John_j either didn't do any extra credit problems, or he_j didn't tell us how many [the number he did was].

The second non-isomorphic alternative is to use a cleft, as in (80). This is also unacceptable, and so clefting appears also to be unavailable as the underlying elided structure.

(80) *John_j either didn't do any extra credit problems, or he_j didn't tell us how many [it was].

Copula deletion and clefting are both unavailable for (78). However, Merchant (2001) does note that one possible way to improve clefting examples containing adjuncts is to retain the presuppositional portion of the cleft. This is shown in (81):

(81) John_j either didn't do any extra credit problems, or he_j didn't tell us how many ??[it was that he_j did].

Retaining the presupposition material improves the example slightly, but it is still judged by speakers to be significantly degraded. Merchant does not propose an explanation for why retaining the presuppositional material improves some adjunct clefts;

⁴⁴Note that we are not ruling these examples out based on the Well-Formedness Condition, which is concerned with pragmatic well-formedness. Non-isomorphic strategies are proposed in part to counter the claim that ellipsis ameliorates islands; therefore, non-isomorphic theories assume that an underlyingly ungrammatical structure will lead to an ungrammatical sluice. Thus the ungrammaticality of a pre-sluice is used to predict whether or not that structure is available as the structure in a sluiced clause.

however, we should be cautious when proposing an explanation that is based on the presence of overt or non-overt material within the ellipsis site.

Even if we accept (81), a reviewer brings up another way of probing the presence of non-isomorphic structure, which is to use contrast sluices. They note that contrast sluices with *else* create a semantic clash between the presuppositional properties of *else* and the exhaustivity requirement of a cleft construction (that is, the uniqueness requirement that clefts impose on the correlate). If we can find contrast polarity reversal sluices, then this presents strong evidence against a clefting non-isomorphic strategy. As it turns out, contrast sluices are easily constructed with the polarity reversal property, as is shown in (82).

(82) Q: Do any of the candidates practice for their debates?

A: I don't think that BERNIE practices for them, but I don't know who ELSE_i [_{t_i} doesn't practice for them].

As predicted, (82) is not acceptable with a clefting strategy, as shown in (83a). Note in (83b) that retaining the presuppositional material does not improve the example.

(83) (a) I don't think that BERNIE practices for them, but I don't know who ELSE
*[it was].

(b) I don't think that BERNIE practices for them, but I don't know who ELSE
*[it was that doesn't practice for them].

We might wonder whether this particular example can be attributed to a copular non-isomorphic strategy instead, since a clefting strategy is unavailable. However, a copula strategy is also unsuccessful, as shown in (84).

(84) I don't think that BERNIE practices for them, but I don't know who ELSE *[is someone that doesn't practice for them].

Not only is such an example unacceptable, it is also not a copula construction that has been proposed or shown to be available for non-isomorphic sluicing.

2.6.3 Summary

I have shown in this subsection two different polarity reversal examples that cannot be explained by a non-isomorphic sluicing strategy. While this demonstrates that the polarity reversal data as a whole cannot be subsumed under a non-isomorphic sluicing analysis, I do not want to make the strong claim that no polarity reversal sluices can be non-isomorphic. There are strong arguments for the existence of non-isomorphic sluices in the literature, and I have no basis from which to argue that polarity reversal sluices somehow form an exception to these arguments. The landscape may very well be as proposed by van Craenenbroeck (2010), in which cleft/copula sluices can be used as last resort strategies to avoid ungrammatical structures, but are not universal. Contrary to being problematic, this picture dovetails well with the theory presented here. Barros et al. (2013) point out that their observations challenge accounts of sluicing that require isomorphy or strict syntactic and/or semantic equivalence between an antecedent and a sluicing site. Far from being a challenge to the current theory, the non-isomorphic sluicing data provide further support for the argument presented in the current work, which is that a complete theory of sluicing must be more permissive than current theories allow.

2.7 Conclusion

This chapter has presented and discussed polarity reversal sluicing data that present a new challenge to the enterprise of determining the conditions under which linguistic content can be felicitously elided. I argue that, counter to its dominant treatment in the syntactic

literature, sluicing is a pragmatics-sensitive phenomenon subject to contextual licensing. I show that the ability to elide linguistic content fits naturally into general theories of constraints regulating coherent discourses, and have detailed one way to account for the pragmatic sensitivity of data that present serious challenges for non-pragmatic theories.

A natural topic of exploration is whether and to what extent a pragmatic account can be extended to different forms of ellipsis. For example, there is some evidence that the polarity reversal phenomenon can also be found in VPE (example from Jim McCloskey, p.c.):

- (85) **Context:** In an internet discussion of the controversial political book *Fire and Fury*, which some discussants argue did not properly document all its claims:
Commenter: “Much of the book is unsourced, but some is [sourced].”

The full scope of ellipsis possibilities is still being determined; for example Merchant (2013b) observed that VPE allows voice mismatches, a mismatch that had previously been thought impossible. It remains to be explored whether or not current analyses can account for cases of polarity reversals under VPE, if such examples are indeed robust. I leave this interesting question for future research and close by noting that, while the current project seeks both to challenge current assumptions on sluicing and to provide an answer to this challenge, much additional work on the road to developing a complete theory of sluicing and ellipsis, of course, remains.

I conclude with a final broad thought about the use of ellipsis in natural language. The question implicit throughout the current chapter is a question of use: why do we choose to utilize sluicing constructions? A Gricean story of Quantity is unlikely, as it is far from clear that the cognitive burden placed on listeners in reconstructing ellipsis sites overrides the moderate benefits received by the speaker in physically uttering fewer words. Additionally, our explanatory theory should unite our dislike of repetition, manifested in

MaxElide, deaccenting, and the general preference for elliptical constructions over their non-elliptical counterparts, with our dislike of stating the “already known,” manifested in Maximize Presupposition. I propose a simple but I hope not simplistic explanation.

If the (simplified, overarching) goal of conversation is to determine which world out of the set of possible worlds is the actual world (as proposed in e.g. Roberts 2012), then each move of a conversation is a step toward achieving this goal. Each utterance in a conversation then acts as a cue to where the participants are in achieving their goal of determining the characteristics of the actual world. It is therefore not so much as Lewis (1979) puts it, that an utterance that violates Maximize Presupposition has ‘no conversational point,’ but that such a violation gives the incorrect cue as to where we are on the road to achieving the conversational goal. Under this view of conversation, Maximize Presupposition is motivated by the desire for each utterance to correctly reflect the current progress of the conversation by presupposing all the information that has already been agreed upon. The same reasoning extends to elliptical constructions. Just as focus is used to cue participants that new information is being entered into the discourse, ellipsis is a cue to conversational participants that the information contained in the elided phrase is already available in the discourse (though has not necessarily been agreed upon). Ellipsis can therefore be completely separated from notions of speaker/listener efficiency; it is merely a road sign to listeners that the information contained in the elided clause has already been entered into the discourse.

Chapter 3

Violating Perspicuity

In the last chapter, I showed that our interpretation of ellipsis sites is not always dependent on strict syntactic, semantic, or pragmatic identity with an antecedent. I argued for a theoretical account under which our interpretation of ellipsis sites is instead constrained by contextual entailment, and is at least partially determined by the salience of particular representations that are available to the interpreter at the point of comprehension of the ellipsis site. In the chapter I discussed one way that a notion of propositional salience can be cashed out within a theory of discourse coherence. I also argued that much remains unknown about propositional salience, and that the field currently lacks a rigorous definition of salience.

This chapter focuses on one measure of salience that is well-established in the anaphora and ellipsis literature, namely locality. I present the results of three experiments probing offline interpretations of ellipsis sites. The experiments manipulated the location of candidate ellipsis antecedents relative to the ellipsis site to test how we select from competing representations when determining the meaning of an ellipsis site. The experiments in this chapter and in the following two chapters implicitly assume some knowledge of the existing research on memory architecture and anaphoric resolution processes. Therefore, before beginning to discuss the experimental findings, I provide a brief overview of memory retrieval processes and the role of locality in anaphoric

processing. Although I cannot provide a comprehensive overview of these broad topics here, §3.1 will be sufficient background for an unfamiliar reader to follow the remaining chapters. I have included citations throughout for readers who are interested in delving deeper into particular topics. The reader who is intimately familiar with the state of the art on cognitive architectures can safely skip to §3.1.2.

3.1 Processing anaphoric dependencies

There is a general (though not universal) consensus among researchers on how anaphoric processes work. As language users, we have a capacity-limited memory. We are able to hold about a clause clearly in memory, in a storage unit called the focus of attention. Any linguistic material within the focus of attention is immediately available for integration with any incoming material. Once material has been pushed out of the focus of attention, it is shuttled quickly into long-term memory in chunks and stored as a meaning representation. Each representation has an activation level given as a function of the chunk's baseline activation level and the recency of the chunk. Once a representation has been shuttled into long-term memory, it must be retrieved in order to be integrated into any new linguistic material. How easily a representation can be retrieved is a function of three things. First is its activation level: *ceteris paribus*, the higher the activation level of a representation, the more likely it will be retrieved. Second is the cues that tag a particular representation in memory, where possible cues include phi features, syntactic category, or semantic part of speech. The more specific the cues a representation has, the more likely it is to be retrieved. Third is the number of competing representations in memory. Generally speaking, the more cue overlap between competing representations in memory, the less likely a given representation will be successfully retrieved.

The following two subsections delve into greater detail on existing theories of retrieval

in anaphoric processes, and the role that locality plays in interpreting an ellipsis site when there are multiple competing antecedents in play.

3.1.1 Retrieval in anaphoric processing

A main topic of investigation in anaphoric processing is the nature of the mechanism used to locate antecedent representations in memory. This mechanism is most often described as a search procedure. Under this search procedure, a comprehender encounters an anaphor, such as a pronominal anaphor or an ellipsis site, and then engages in a search of their memory representations to find the appropriate antecedent to fill in the anaphor's meaning.

The memory literature offers two main types of memory search procedures. Early theories of pronoun resolution, inspired by the memory scanning retrieval models of the times, assumed an automatic (Neely, 1977; Posner and Snyder, 1975; Ratcliff and McKoon, 1981) and serial search process for pronoun resolution (Hobbs, 1978; Clark and Sengul, 1979; Corbett and Chang, 1983; van Dijk and Kintsch, 1983). Serial searches are characterized by relatively slow processes, and involve searching the available discourse space for an appropriate antecedent representation by cycling through the available representations one-by-one, generally in descending order from the most recent interpretation. Because such searches are serial, they are, for example, argued to be necessary for recovering temporal and spatial order information. Anaphoric processing theories which claim that anaphora resolution involves serial search processes propose that the entire discourse space is available to be searched. The search for an appropriate antecedent representation proceeds serially through memory until an appropriate antecedent is found or the entire space of the discourse is exhausted (Sternberg, 1966).

Recent theories of anaphoric retrieval have moved away from models of serial search. Greene et al. (1992) proposed that, instead of a serial search process, pronouns serve

as cues to the most likely entity in a discourse representation. Greene et al. propose that pronouns are resolved only if their intended referent is sufficiently more highly accessible in the comprehender's discourse model relative to all other discourse entities. Instead of a serial search, an anaphor's cues, such as gender and number, are "matched" automatically and in parallel against the semantic features of all entities in the current discourse model, making the matching process a competition between all entities in the discourse space.

This approach toward anaphora resolution is harmonious (though not identical) with current models of memory retrieval, which consist of cue-based content-addressable approaches. In content-addressable theories, a representation is retrieved from memory using a direct-access cue-matching procedure (McElree et al., 2003; Dyke and Lewis, 2003; Lewis and Vasishth, 2005; Dyke and McElree, 2006). Cues in these models consist of features such as number and gender. In content-addressable search theories, available memory representations are compared simultaneously against a set of search cues to find the best match (see Wagers 2014 for an overview). The search for an appropriate antecedent representation then does not proceed serially, but all at once.

An alternative view of anaphoric processing is that no search/retrieval or matching is involved, but that antecedents are held in a special cognitive space, called the focus of attention. Focus of attention theories propose that there is a significant limit on the amount or scope of information that can be maintained and made available to on-going processing (Broadbent, 1958; McElree, 2001; Wagers and McElree, 2012). The limited information that is directly accessible to cognitive processes is said to be in focal attention. Because information in focal attention is directly and concurrently accessible to cognitive processes, it does not need to be retrieved from memory and is therefore available for immediate resolution when encountering a pronoun (Grosz and Sidner, 1986; Grosz et al., 1995; Gundel et al., 1999; Foraker and McElree, 2006). If anaphoric resolution is

accomplished using information in focal attention, then no search procedure must be undertaken. Focus of attention theories are compatible with cue-based retrieval theories. For example, it could be that anaphors are resolved to a referent in the focus of attention if possible. If the material in focus of attention fails to resolve the anaphor, then a cue-based search can be initiated (O'Brien et al., 1997; Anderson, 2007), perhaps at loss of coherence (Grosz and Sidner, 1986; Grosz et al., 1995).¹

The three types of anaphoric retrieval processes make different predictions as to the speed and accuracy of anaphoric resolution. A serial search, due to its serial nature, predicts that retrieval speed will vary with the recency of a representation and with its serial position in a list of items, such that less recent items and items further back in a list will be retrieved more slowly, assuming a backward search through memory (McElree and Doshier, 1993; Neath, 1993; Neath and Knoedler, 1994). A content-addressable search, however, predicts no retrieval speed difference based on the recency or serial positioning of a possible antecedent. Because the search space is searched simultaneously in a content-addressable search, the recency of the antecedent representation is irrelevant to the speed of retrieval. However, such theories do predict variation in retrieval accuracy based on the specificity of the cues used to search for the antecedent and on the number of overlapping cues on competing representations in memory. The more competing representations in memory and the more cue-overlap between these representations and the target representation, the lower retrieval accuracy of the representation will be.² Therefore, under these accounts the primary source for unsuccessful retrieval (i.e. the retrieval of the incorrect representation) and forgetting is the overlap in memory

¹Note that the idea of a focus of attention is not unique to theories of language processing. Discussion of salience in visual attention also propose a visual focus of attention, see e.g. Itti et al. 1998.

²Note the relevance here with the d-linked sluicing examples in the last chapter, where we saw that the more specific cues on the d-linked remnants allowed sluicing out of constructions otherwise thought impossible. See also Harris 2019.

between representations sharing cues with the intended representation and the intended interpretation itself. This process of forgetting or incorrect retrieval is called *cue-based retrieval interference*, and is a function of cue-overload as diagnosed by the match between the search cues and the cues on the intended representation, divided by the match between the search cues and the cues on additional representations in memory (Keppel and Underwood, 1962; Waugh and Norman, 1965; Nairne, 2002; Martin, 2016, 2018).

Finally, focus of attention theories predict a distinction in the speed between the accessing of representations that are in focal attention and representations that are not in focal attention. Because representations not in focal attention must be retrieved from memory, they should be accessed more slowly. For anaphoric resolution in particular, Foraker and McElree (2006) propose that such theories predict that prominent or ‘focused’ antecedents will differ from less prominent ones in the speed in which the antecedent is accessed upon encountering a pronoun. Contrastingly, theories that do not assume a distinct focus of attention module predict that comprehenders will be more likely to access a prominent antecedent, resulting in a higher quality interpretation. Foraker and McElree test this hypothesis for pronominal antecedents that are in clefted clauses, and find that, at least for clefting, there is no evidence that the antecedent representation in the clefted constituent was actively maintained in focal attention. However, I know of no experimental investigations testing the predictions of focus of attention on ellipsis comprehension, specifically.

Recent experimental investigations of ellipsis provide evidence in favor of a simultaneous, cue-based retrieval mechanism operating for ellipsis comprehension. Both Martin and McElree (2008, 2009, 2011), using a speed-accuracy tradeoff methodology, and Harris (2015), using eye-tracking experiments, test the predictions of a cue-based retrieval system for verb-phrase ellipsis and sluicing and find results consistent with a

direct-access operation over content-addressable representations. While Harris (2015) remains agnostic on the exact model of sentence processing utilized by ellipsis resolution, Martin and McElree argue that ellipsis resolution involves a cost-free pointer mechanism to a representation in memory. This theory is similar to that proposed by Frazier and Clifton (2001, 2005), in which ellipsis resolution entails a cost-free copying of the structure of an antecedent clause into the ellipsis site. However, Martin and McElree's proposal posits no syntactic material within the ellipsis site, while Frazier and Clifton's theory does.

A related but separate question is whether the retrieval process can fail upon an initial parse or fail completely. For Greene et al., the matching process for anaphors can fail if no discourse representation match is sufficiently higher and better than all other discourse entities or if no referent is identified during the cue-matching search. In the case of failure, the selection of an antecedent may be delayed. The comprehender can either wait for additional disambiguating information from the discourse or can engage in a strategic problem solving process to resolve the anaphor (see also Yule 1982, Webber 1983). For example, Greene et al. present evidence that comprehenders do not resolve referents of pronouns unless they are given time (about 450ms as compared to 250ms), are motivated to do so by a task requiring the resolution of the referent, and have short enough materials in order to resolve the references.

Love and McKoon (2011) follow up on these experiments and show that participant engagement in the task is a factor in the probability of reference resolution for an anaphor. Love and McKoon used probe-recognition tasks to test whether the antecedent of an anaphor was re-activated upon comprehenders' interpretation of the anaphor. They found that the length of experimental materials determined whether pronouns were automatically resolved, with only longer stories of 4-8 sentences leading to automatic and correct resolution. They attribute this finding to readers' increased engagement with the

longer stories. Notably, the failure to resolve the anaphor's referent did not lead to parsing failure; participants were able to complete discourses without resolving the referents of pronouns contained in the discourses. The outcome of a failed or delayed search is, in this way, importantly different from a failure to resolve a filler-gap dependency in constructions like *wh*-movement, whose searches are automatic and obligatory (Frazier, 1987; Frazier and Clifton, 1989; Clifton and Frazier, 1989).

I know of no experimental studies that investigate the possibility of search or retrieval failure in ellipsis directly. Chapter 4, however, discusses a closely related question involving the search procedure for cataphoric ellipsis. The chapter investigates whether we actively resolve ellipsis sites when parsing a sentence, or whether the sites are resolved only after an initial parse is generated. If ellipsis patterns like anaphora, then we might expect that the resolution of ellipsis sites could be delayed or left unresolved completely, as found by Greene et al. (1992) and Love and McKoon (2011). Chapter 4 provides evidence consistent with a theory under which ellipsis sites are in fact resolved using an online, active search procedure.

The last big question about anaphoric processing involves the accessibility of possible antecedents. That is, do comprehenders search an entire discourse space for candidate antecedents? Or is the search space constrained by syntactic, semantic, or other grammatical factors. One theory that argues that the space of possible antecedents is constrained is the Right Frontier Constraint (Hobbs, 1978; Polanyi, 1986; Webber, 1988; Asher, 1993). The Right Frontier Constraint is a grammatical constraint that proposes that 1) discourse is organized in a hierarchical structure, and 2) that only the leading edge of the structure, or the "right frontier" of the discourse, is available for attaching new discourse units into the existing structure. The constraint therefore places boundaries on what discourse entities may be used as antecedents to resolve anaphoric processes, namely only representations that occur on the right edge of the discourse hierarchical structure.

The Constraint, notably, is claimed to be a grammatical constraint, and not a constraint on the scope of the cognitive architecture search mechanism.

Although the Right Frontier Constraint has received some notoriety, it is an unresolved question whether the entire discourse space is in principle available to serve up antecedents to anaphoric processes. This project unfortunately does not attempt to contribute to this interesting question. I encourage the interested reader to refer to van Dijk and Kintsch (1983); Kintsch (1988); Ratcliff and McKoon (1988) for arguments that the entire discourse space is in principle available, and to Polanyi (1986); Webber (1988); Asher (1993); Asher and Lascarides (2003) for arguments that certain portions of the discourse are in some sense “blocked” from anaphoric search.³

3.1.2 Locality in anaphoric processing

In the previous section, I discussed processing theories in which the likelihood of retrieving a particular meaning representation from memory is based on the representation’s activation level and the specificity of its cues: when I read or hear an anaphor, I search for the most highly-activated stored representation that has cues matching the anaphor (Lewis and Vasishth, 2005). In principle, there is no reason why I cannot retrieve a representation from far back in the discourse. Even if an antecedent has low activation, our cue-based retrieval system should be able to retrieve it if provided with sufficient cues (van Dijk and Kintsch, 1983). I also discussed that there is support within the general memory literature and current language processing theories that the main factor for forgetting and incorrect retrieval is not due to low activation via decay, but to cue-based interference. In cue-based interference, competing representations overlap with the cues

³There is also an expansive literature on whether the search space for pronouns is constrained by binding-theory. This question is slightly orthogonal to the current project, so will not be discussed. See e.g. Sloggett 2017 for discussion.

used to search for a representation in memory, thereby decreasing the diagnosticity of those cues (Wagers and McElree, 2012).

The current subsection zooms in slightly from the predictions of these specific theories to discuss an empirical phenomenon that is widely documented in the experimental ellipsis literature. Both online and offline experiments on ellipsis interpretation provide evidence that ellipsis interpretation is subject to what is called a *locality bias*. The locality bias is instantiated as a preference for choosing an ellipsis antecedent that is in the most local position to the ellipsis site, given cases in which two or more candidate antecedents are available as the interpretation of the ellipsis site (Frazier and Clifton, 1998, 2005; Carlson et al., 2009; Martin and McElree, 2011; Harris, 2015, 2019).

For example, Martin and McElree (2011) found that sluicing out of the first conjunct of two conjoined candidate antecedents was penalized in the likelihood of successful comprehension, though not in the time taken to retrieve and interpret the antecedent. Harris (2015) found a penalty in eye-tracking while reading when a sluice had to be resolved to a non-local correlate; that is, when a sluice had to jump over a candidate antecedent to find a felicitous antecedent that was further back in the sentence. Harris (2015) captures this generalization in what he calls the *Locality Bias* principle, given for sluicing specifically in (86). To give an example, the Locality Bias predicts that comprehenders will prefer to resolve the sluice in (87) to *which crew members* over *which tribbles*.

(86) Locality Bias:

Associate the remnant of clausal ellipsis with a correlate occupying the structurally most local position.

(87) I've heard that some tribbles hate some Enterprise crew members, but I don't know which ones.

The origin of the Locality Bias is not settled. The most basic explanation for this preference is that *crew members* is the most salient NP when the comprehender hits the ellipsis site, and is therefore the preferred interpretation. However, this notion of salience is exactly what we would like to explain. One possibility is that because *tribbles* and *crew members* share retrieval cues, the search process relies on the relative activation level of each representation when choosing an interpretation of the ellipsis site.

The theory of activation for language representations in memory is laid out in detail in Anderson (2007) and Lewis and Vasishth (2005), but for our purposes here we care mostly that the speed and success of retrieval of a given representation in memory is determined not only by the cues used to search for the representation in memory, but also by the representation's activation level. The general idea is that the activation of a given representation in memory is determined in part by a base activation level, which is a function of the total number of times a representation has been retrieved and the time since it was last retrieved. This base-level activation is subject to a decay rate over time. The total activation for a particular representation in a particular context is determined by the representation's base activation level and its strength of association with elements in that particular context. The retrieval latency and probability of retrieval of a given representation in a given context is then determined by the total activation level of the representation. That is, *ceteris paribus*, the lower the total activation level of the representation, the more slowly and less successfully the representation will be retrieved (Anderson and Schooler, 1991).

Given this definition of activation, it follows that for two competing representations that are equally matched on cues and contextual association measures, the more recent representation will be the most highly activated and therefore stands a greater chance of successful retrieval. Such a recency advantage is well-established in the general dependency-resolution literature (Frazier, 1978; Neath, 1993; Neath and Knoedler,

1994; Gibson et al., 1996; Pearlmutter and Gibson, 2001; Sturt et al., 2002). As already discussed, within the domain of ellipsis processing, Martin and McElree (2011) and Harris (2015) provide congruent evidence suggesting that the distance between a candidate antecedent and an ellipsis site and the intervention of competing representations in the context increased the processing difficulty of interpreting a sluicing site. Martin and McElree argue that recency confers the advantage not only of higher activation, but also of less cue overload, and therefore greater cue-specificity and less interference.⁴ While the results of Harris 2015 are compatible with this story of ellipsis retrieval, Harris points out, quite correctly I believe, that the mechanisms responsible for searching for a correlate in sluicing constructions may not be analogous to the mechanisms responsible for searching for the representation for the entire ellipsis site.

These studies provide important insights into the mechanisms used to search and retrieve ellipsis representations. However, the studies relied on comprehenders' evaluations of sluices that were manipulated to be either grammatical or ungrammatical under a particular resolution of the ellipsis site. They therefore limited the possible interpretations that comprehenders had available for the ellipsis site. Furthermore, the studies looked only at ellipsis constructions in which the antecedent followed the ellipsis site. Because of the focus on anaphoric ellipsis processes, structural locality, recency, and proximity were inexorably intertwined. We are therefore left with the potential to be conflating a recency bias of antecedent material with that of a structural locality preference. Such a possible conflation is discussed briefly by Harris (2015), who acknowledges that the

⁴Martin and McElree also gesture toward an account in which activation plays no role in recency/locality effects. My understanding of such a theory is that it requires the encoding cues of more recent representations to be more similar to the site at which the representations are retrieved than less recent representations. For example, if time is somehow encoded as a cue on representations, then the more recent a representation was encountered, the more cue similarity it will have to the location of retrieval. Thanks to Matt Wagers for discussion on this point.

existing literature on ellipsis comprehension has not yet determined whether the effect of candidate antecedent positioning is truly due to structural prominence, or is instead an effect of linear distance, linear ordering, or temporal precedence. As Harris states, “the nature of the mechanisms that underlie [the] putative advantage [of structural locality] will remain unsettled until an effect of structural privilege is replicated in a design that dissociates structure from other factors, like linear order” [p. 16].

Conveniently, ellipsis sites do not always have to follow their antecedents. We know that in some cases ellipsis sites can precede their antecedents, as seen in cataphoric ellipsis processes. Existing theories of ellipsis comprehension make no hypothesis about what will happen when an ellipsis site precedes its potential antecedents. Because existing studies looked solely at anaphoric ellipsis, the search mechanisms discussed were necessary backward-looking in nature. However, because the ellipsis site in cataphoric constructions precedes any potential antecedent representations, a backward-looking search cannot explain the resolution process. This raises the question, then, of what procedure is initiated when a comprehender encounters a cataphoric ellipsis site.

I propose two main hypotheses. The first is that a forward-looking search will be initiated. This search shares properties with the backward-looking searches described above, but will operate over incoming representations as parsing happens in real-time. Under this hypothesis, cataphoric ellipsis resolution maintains its similarities with other dependency searches, and relies fundamentally on similar search mechanisms as in anaphoric ellipsis. The main difference between anaphoric ellipsis resolution and cataphoric ellipsis resolution under this hypothesis is how the search unfolds in real-time. The second hypothesis is that no forward-looking search is launched. Instead, the ellipsis site is not resolved until an initial clausal or sentential parse is complete. Under this hypothesis, the search mechanism used for cataphoric ellipsis is identical to that used for anaphoric ellipsis resolution. The main difference is that the search process is initiated at

a later time relative to the ellipsis site in cataphoric ellipsis than in anaphoric ellipsis.

The two hypotheses make different predictions for offline interpretations and online processing of ellipsis sites. The online hypotheses are discussed and tested in Chapter 4. The current chapter focuses on the offline interpretation predictions. If a forward-search is initiated in cataphoric ellipsis comprehension, then we predict that ellipsis comprehension will show a preference for the most proximate candidate antecedent to the ellipsis site, regardless of whether that antecedent precedes or follows the ellipsis site. I call this hypothesis *Locals Only*.

If no forward search is initiated, then we predict that ellipsis comprehension will show a preference for the candidate antecedent that is most recent to the point at which the search is initiated. We have evidence that the search for an antecedent in anaphoric ellipsis is initiated at the ellipsis site (Martin and McElree, 2011; Harris, 2015). The search for the antecedent in cataphoric ellipsis would plausibly occur at a main clausal or sentential boundary following the ellipsis site (Jarvella, 1971; Just and Carpenter, 1980; Warren et al., 2009; Stowe et al., 2018); however, it may also pattern similarly to anaphor resolution as seen in Greene et al. (1992) and Love and McKoon (2011), in which comprehenders wait until the end of a discourse segment, or do not resolve the ellipsis at all until prompted or required. I call this hypothesis *Wait for It*.⁵

Alternatively, some combination of these two hypotheses could be found; I call this hybrid hypothesis *Liger*. The predictions of each theory are given in examples

⁵These hypotheses leave out a very interesting question that I unfortunately cannot answer here, which is whether there is a preference for searching for an antecedent prior to or following ellipsis sites, when both are potentially available. See Liversedge and van Gompel 2003 for experimental evidence that comprehenders attempt to resolve pronouns both cataphorically and anaphorically to a particular antecedent. That is, given a choice to establish pronominal reference cataphorically to a provided postcedent, or anaphorically to a provided antecedent, comprehenders attempt to resolve both the referent in both directions, seeking maximal coherence throughout a mini-discourse.

(88)-(90). Note that I have chosen to move away from the term *locality* because it evokes comparisons with theories of structural locality preferences in syntactic theory, which at this point is undesired.

(88) Locals Only:

Ellipsis interpretation will show a preference for resolution to the closest candidate antecedent in linear distance to the ellipsis site, measured by word count.

(89) Wait for It:

Ellipsis interpretation will show a preference for object resolution, replicating the findings for anaphoric ellipsis by Martin and McElree (2011); Harris (2015); Harris and Carlson (2016, 2015).

(90) Liger:

Some combination of the preferences in Locals Only and Wait For It will emerge. For example, perhaps there is a strong object preference overall, but the interpretations are pulled toward the subject in the cataphoric cases. Or perhaps we see no overall preference patterns, suggesting participants are utilizing a variety of processes.

(91) Anaphoric Ellipsis

Bowties look cool with **fezzes**, but I don't know anyone who wears any [].

Cataphoric Ellipsis

I don't know anyone who wears any [], but **bowties** look cool with **fezzes**.

Each hypothesis makes a different prediction for the preferred interpretations of an example like (91). Local Only predicts that the preferred interpretation in the anaphoric example will be *fezzes*, but the preferred interpretation in the cataphoric example will be *bowties*. Wait for It predicts that the preferred interpretation in both the anaphoric

and the cataphoric example will be *fezzes*.⁶ Liger predicts that some other pattern will emerge. Perhaps both the anaphoric and the cataphoric example will see a preferred object interpretation, but the cataphoric case will see a relative increase in a preference for subject interpretation, even if it is not the majority interpretation. Such a finding would confirm the object bias found in previous studies, but also indicate that proximity of candidate antecedents to the ellipsis site has some effect on interpretation preferences. Recall that previous experiments were unable to distinguish between these different strategies because all items were anaphoric, and therefore the object NP was always the most recent candidate antecedent seen by the reader.

3.2 Choosing an empirical domain

The empirical domain of ellipsis in English is broad and well-documented in the theoretical literature. This presents us with a question of where to begin when investigating how we as language users understand and process ellipsis constructions. While the syntactic and semantic literature going back to Ross (1967) has achieved substantial breadth in mapping the landscape of ellipsis constructions, experimental investigations of ellipsis are both more recent and less expansive in scope. Experimental investigations of ellipsis processes have generally focused on verb-phrase ellipsis (Shapiro et al. 2003; Frazier and Clifton 2005; Arregui et al. 2006; Martin and McElree 2008, 2009; Frazier and Clifton 2010; Miller and Hemforth 2014; Poppels and Kehler 2019) and sluicing (Frazier and Clifton 1998; Gullifer 2004; Frazier and Clifton 2005; Poirier et al. 2010; Martin and

⁶Note that this prediction is also accounted for under a theory of ellipsis resolution that allows interpretation revision for the cataphoric cases. Offline experiments cannot conclusively distinguish between a theory in which the interpretation is constructed only once the parsing of the example is complete, or a theory in which the interpretation is revised to the last NP after a complete parsing of the example. See Chapter 4.

McElree 2011; Yoshida et al. 2012; Harris 2015, 2019). There are notable exceptions to this trend, of course: Carlson (2002) looks at the role of prosody in processing gapping constructions, Miller (2016) looks at syntactic reactivation in noun-phrase ellipsis, and Harris and Carlson (2016, 2015) look at the role of focus in *let-alone* ellipsis. However, the experimental literature has yet to achieve the breadth or depth of the formal theoretical literature on ellipsis.

How do we determine, then, which type of ellipsis to focus our experimental investigations on? Furthermore, should we think that the choice of particular ellipsis construction even matters? I think that the choice of construction is not inconsequential, and that there are benefits and drawbacks to using each type of ellipsis construction. I also believe that these considerations have guided previous researchers in their choice of construction for their investigations, even if implicitly. For example, sluicing provides a *wh*-remnant that stands in a semantic relationship with an optional correlate. No other ellipsis construction has this particular configuration, namely that in which a *wh*-remnant has moved out of the ellipsis site and provides an overt link to ellipsis site antecedent. The existence of the remnant and (optional) correlate affords a great deal of flexibility in controlling possible interpretations of the ellipsis site, as it provides a direct cue to the semantic and syntactic properties of its antecedent (see e.g. Harris 2019). Gapping and *let-alone* ellipsis, on the other hand, are particularly good constructions for studying the interaction of focus and ellipsis. Verb-phrase ellipsis provides a minimal pair with event anaphora, and so is a sensible choice for direct comparisons between anaphors and ellipsis sites. The point I wish to make, here, is simply that the choice of ellipsis construction in an experimental investigation of ellipsis does matter, as the properties of the construction will limit and/or facilitate the types of questions an experiment is able to investigate. The construction chosen, then, should harmonize with the research question.

With this objective in mind, I chose to use noun phrase ellipsis (NPE) in the experi-

mental investigations of ellipsis comprehension undertaken here. NPE was chosen as the appropriate choice of construction because the piece of syntactic/semantic/phonological material that is elided is small. The small nature of the elided material affords a particularly valuable opportunity not offered by alternative ellipsis constructions, namely that candidate antecedents are also small. Because the goal of the current experiments is to isolate offline and online effects of ellipsis processing, it is arguably more likely that we will find consistent effects if the search space is constrained.

For example, Yoshida et al. (2012) tested examples such as (92) and (93) to investigate whether comprehenders actively search for an antecedent in the comprehension of cataphoric ellipsis. Yoshida et al. hypothesized that if comprehenders engage in such a search, called an *active search strategy*, then we would find a plausibility effect manifested as a reading-time slow-down on the verb *notified* in (92). That is, they hypothesized that an active search strategy for cataphoric sluicing would involve parsing the clause directly following the cataphoric ellipsis site, and then attempting to resolve that clause to the ellipsis site. Yoshida et al. propose that comprehenders will try to resolve the ellipsis site at the verb because *wh*-phrases receive their thematic role from the verb in both sluicing constructions and in non-elliptical *wh*-questions. Therefore, the processing of cataphoric ellipsis should involve a long-distance dependency between the *wh*-remnant and the verb of the elided phrase/antecedent phrase, just as there is a dependency between a *wh*-word and a clausal verb in non-elliptical *wh*-questions.

(92) I don't know which book, but the editor notified the publisher about a new book.

(93) I don't know which book [~~#the editor notified~~], but the editor notified the publisher about a new book.

This hypothesis is well-founded; however, it is merely a hypothesis as to how comprehenders would attempt to resolve the ellipsis site. I believe it is also reasonable to

believe that comprehenders would instead try to resolve the first subject position to the correlate, e.g. *I don't know which book, but some book is...* Under this procedure, we would expect a reading time slow-down to occur on the noun *editor*, or on the definite article. Because sluicing elides an entire TP, the possible resolution sites are simply not constrained.

The use of NPE in the current experiments is therefore motivated by concerns of practicality and precision, not by an interest in NPE as a construction per se. It is my hope that the results here can be extended to additional forms of ellipsis, such as verb-phrase ellipsis and sluicing. However, since investigations of ellipsis comprehension are still in the early stages, I believe breadth must be sacrificed in favor of precision. It remains an open question, of course, whether all ellipsis constructions can be united under a single grammatical licensing condition and under a single processing mechanism. I make the implicit assumption here that they can. The intended scope of the claims made here are therefore not about NPE in particular, but about the phenomenon of ellipsis itself. The question of whether this assumption turns out to be correct will be answered by future work. The remainder of this section provides a brief background on NPE as a way to set the stage for the subsequent experimental investigations.

Noun phrase ellipsis is a form of ellipsis in which an N' is elided, as in (94). NPE is licensed by the head of a DP, and is ungrammatical without such a licenser, as in (94d). NPE in English is licensed by quantifiers, numerals, and possessives.

- (94) (a) River loves every doctor, but Rose only loves a few [dɒktərs]
(b) Rose remembers every adventure and Donna doesn't remember any [ədventʃə].
(c) Rose loves the Doctor's face and Rory loves Amy's [fæʃ].
(d) Amy loves cats and Rory loves #[eɪts], too.

The licensing conditions of NPE are most notably spelled out by Lobeck (1995). Lobeck argues that NPE is licensed by a head containing a [+poss] or a [+partitive] feature. This feature requirement accounts for the licensing ability of possessives, singular *one* and *each*, and all quantifiers except singular *every*, which cannot license NPE. Lobeck accounts for the licensing ability of numerals by arguing that the head NUM can be specified for strong agreement, which allows a head to license complement deletion.

A notable property of NPE is that, like VPE and sluicing, it can occur in constructions in which the ellipsis site precedes the antecedent. I refer to ellipsis sites that precede their antecedent in this way as *cataphoric* ellipsis. Ross (1967) observed that NPE, like pronominal anaphora, appears in these cases to obey the Backwards Anaphora Constraint. The Backwards Anaphora Constraint says that pronouns can precede their antecedents, but only when in a discourse subordinated clause. For example, consider the minimal pairs in (95). While the ellipsis site in (a) is in a discourse subordinated clause, it is in a conjoined clause in (b). Although the exact licensing conditions for cataphoric ellipsis remain unknown, the current studies will use only uncontroversially acceptable examples.

- (95) (a) While Donna doesn't remember any #[adventure], Rose remembers every adventure.
- (b) Donna doesn't remember any #[adventure] and Rose remembers every adventure.

There is a rich cross-linguistic literature on the language-specific licensing facts of NPE, as well as analyses of NPE in various historical syntactic frameworks. However, such discussions would unfortunately take us too far afield for current purposes. I encourage the interested reader to consult Lobeck (1995) and Lobeck (2006), *inter alia*, for detailed discussion. The important points for our purposes here is to have established

that NPE is uncontroversially licensed by quantifiers and possessives, and that it can occur cataphorically in discourse-subordinated contexts.

3.3 Experiment 1: Ellipsis interpretation shows a subjecthood preference

The first experiment tests how the proximity of candidate antecedents affects preferences for those antecedents in NPE comprehension. The experiment uses both anaphoric and cataphoric ellipsis sites that have two structurally and semantically licit candidate antecedents. For example, in (96), both *clarinets* and *flutes* are candidate interpretations for the NPE site. The experiments asked for offline interpretations that allow us to disentangle two possible influences on the comprehension of the NPE site. If **strict proximity** is the dominant force driving interpretation preferences, then we predict that *flutes* will be the dominant interpretation in the anaphoric example, and *clarinets* will be the dominant interpretation in the cataphoric example. Such a result would support the Locals Only hypothesis. If **recency** is the strongest dominant force driving interpretation, then we predict that *flutes* will be the dominant interpretation in both examples. This result would support the Wait for It hypothesis.

(96) Anaphoric NPE

Clarinets would sound good with **flutes** during the reception, if we can find any [] by this evening.

Cataphoric NPE

If we can find any [] by this evening, **clarinets** would sound good with **flutes** during the reception.

3.3.1 Method

3.3.1.1 Methodology

The experiment collected offline, free-response interpretations of NPE sites. The experimental task was set up such that participants were instructed to read short exchanges in which a ‘confusion’ arose between two speakers. Participants were asked to resolve the confusion between the speakers in whatever way they felt was most natural. In order to encourage participants to fully read the exchanges and engage in the most natural comprehension process possible given the restrictions of an experimental setting, the experimental items were set up as a small text message discourse between two people (see Kroll and Wagers 2017 and Figure 3.1).

3.3.1.2 Participants

Forty-eight participants were recruited from the University of California, Santa Cruz. All participants were undergraduate native speakers of English; all received course credit for their participation.

3.3.1.3 Materials

Materials consisted of 42 critical items comprising six conditions. The design was a 2 x 3 with the factors *ellipsis* (Anaphoric ellipsis, Cataphoric ellipsis) and *antecedent* (Subject antecedent, Object antecedent, and Ambiguous antecedent). An example experimental item is given in Table 3.1.

Ellipsis	Antecedent	Sentence
Anaphoric	Subject	Honeycrisp apples would taste great with this navel orange for breakfast, but I don't know if we can find any right now.
Anaphoric	Object	This honeycrisp apple would taste great with navel oranges for breakfast, but I don't know if we can find any right now.
Anaphoric	Ambiguous	Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any right now.
Cataphoric	Subject	I don't know if we can find any right now, but honeycrisp apples would taste great with this navel orange for breakfast.
Cataphoric	Object	I don't know if we can find any right now, but this honeycrisp apple would taste great with navel oranges for breakfast.
Cataphoric	Ambiguous	I don't know if we can find any right now, but honeycrisp apples would taste great with navel oranges for breakfast.

Table 3.1: Experiment 1 example item.

Looking at the antecedent conditions first, the names of the levels – Subject, Object, and Ambiguous – refer to the grammatical position of the preceding NP that was the most favorable candidate antecedent for the NPE site. All antecedents were given in the preceding sentence, and favorability was manipulated using definiteness. I hypothesize that interpretations of the ellipsis site will show a preference for plural, non-specific antecedents over singular, specific antecedents. This preference is hypothesized to be due to the selectional restriction of the indefinite quantifiers *any* and *some*. See §3.3.1.3 for additional discussion.⁷ This manipulation serves as a baseline measure of comprehension

⁷Previous iterations of this manipulation included using bare indefinites, though informal polling showed that people were fairly easily able to coerce such readings, e.g. Clarinets would sound good with a flute, but I don't know if we can find any [flutes] right now. Thanks to Craig Sailor for discussion on this point, see also Saab (2018).

preferences. The definiteness preference should not be affected by the location of the ellipsis site relative to the candidate antecedents, and therefore provides a baseline measure of how participants' preferences can be pushed toward one candidate antecedent over another.

For example, both sentences in example (97) contain NPE licensed by the quantifier *any*. Because *any* selects for a plural or mass noun, the possible antecedent *this streamer of mine* in (a) and *this balloon of mine* in (b) are degraded. The possible antecedents *balloons* and *streamers* are therefore considered the most favorable antecedents, respectively, for the NPE in each sentence.⁸

- (97) (a) Balloons would look good next to this streamer of mine for the party, but I don't remember if we have any for tomorrow.
- (b) This balloon of mine would look good next to streamers for the party, but I don't remember if we have any for tomorrow.

In the Ambiguous condition, both NPs were plural indefinites and were therefore considered equally favorable candidate antecedents in terms of their definiteness. The question we are most interested in is which NP participants will choose in this case. Each antecedent condition was presented in both an anaphoric and a cataphoric form. A partial item example is given in (98). As discussed, anaphoric and cataphoric conditions were tested in order to see whether antecedent preferences flipped or altered based on the relative linear positioning of possible antecedents and the ellipsis site.

- (98) (a) Anaphoric Condition
Eggplants would look cool beside pineapples for the mantelpiece, but I forget if we need any this time.

⁸Note that in this particular example, the use of the possessive also adds a plausibility consideration in conjunction with the embedded question in the final clause.

(b) Cataphoric Condition

I forget if we need any this time, but eggplants would look cool beside pineapples for the mantelpiece.

Previous research has suggested that the interpretation of ellipsis sites is influenced by information structure and focus placement. Experimental results from Carlson et al. 2009 suggest that a preference for late focus placement in sluicing antecedent clauses drives findings of locality preferences in sluicing interpretation. Similarly, Harris and Carlson 2015 find that in *let alone* ellipsis there is a preference to pair the remnant with a correlate containing contrastive pitch accent. They find, however, that this preference is not strong enough to completely overcome the preference for pairing the remnant with the most local plausible correlate, regardless of accent placement.

To help ensure that focus placement does not interfere with the intended experimental manipulations in the current experiments, the experimental items were embedded within a mini-discourse in which two participants attempt to resolve some communication confusion. To control for focus placement, each item included an opening QuD, which encouraged narrow focus placement on both candidate antecedents in the critical sentences in the following manner. It has long been observed that in responses to constituent questions, focus falls on the constituent that corresponds to the *wh*-word in the interrogative form (Halliday, 1967; Rooth, 1985, 1992; Roberts, 1996). This pattern is called *Question-Answer Focus Congruence*. For example, in (99), whose *wh*-word is questioning an object, focus only felicitously falls on the object, and is infelicitous if expressed on the subject. However, in (100), whose *wh*-word is questioning a subject, we see the opposite pattern emerge, in which focus is felicitously realized only on the subject, and not on the object.

(99) What sport does Elizabeth like?

- (a) Elizabeth likes [sailing]_F.
- (b) #[Elizabeth]_F likes sailing.

(100) Who likes sailing?

- (a) #Elizabeth likes [sailing]_F.
- (b) [Elizabeth]_F likes sailing.

Now consider the example experimental item in (101). The *wh*-word is targeting nouns that satisfy the relational predicate of going well together. In the critical sentence, those nouns are given by ‘honeycrisp apples’ and ‘navel oranges’. The QuD therefore encourages readers to place focus equally on those two nouns. Note the difference with the hypothetical example given in (102), which encourages focus placement on ‘honeycrisp apples’. Without an explicit QuD, we have no way of knowing what context participants have reconstructed for the test sentences, and therefore no way of disentangling the possible effects of implicit focus placement. The inclusion of an explicit QuD allows us to control this factor (see also Kroll and Wagers 2017, 2018 and Kroll and Rysling 2019 for further discussion of the importance of explicit QuDs in semantic and pragmatic experimental investigations). Following the critical sentence, an additional question asking clarification about the content of the critical sentence was given. For the critical items, the clarification question always targeted the interpretation of the ellipsis site in the critical sentence.

(101) Opening QuD

What would go well together for breakfast tomorrow?

Critical sentence

Honeycrisp apples would taste great with navel oranges for breakfast, but I’m not sure if we can find any right now.

Clarification Question

Wait, what are you not sure we can find?

(102) Hypothetical Opening QuD

What would go well with navel oranges for breakfast tomorrow?

Hypothetical Critical sentence

Honeycrisp apples would taste great with navel oranges for breakfast, but I'm not sure if we can find any right now.

Hypothetical Clarification Question

Wait, what are you not sure we can find?

3.3.2 Procedure

The experiment used the IbexFarm experimental software (Drummond, 2014). Critical items were distributed into six Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, ensuring balanced responses across lists.⁹ Participants were trained on seven practice items and were given a cartoon break about half way through the experiment. Items were displayed as text message conversations that spanned two screens. Each item started with a question presented in a text bubble. Participants clicked the *Press to continue* button to display the additional text bubbles one at a time. Participants had as long as they liked to read each text bubble. The first screen showed the initial question and two additional text bubbles containing an answer to the question; each answer contained two clauses, and each clause was displayed in its own text bubble. An example of a complete first screen is given in Figure 3.1.

⁹See IbexFarm manual on how to generate separate links for each Latin Square list within an experiment (ibid.).

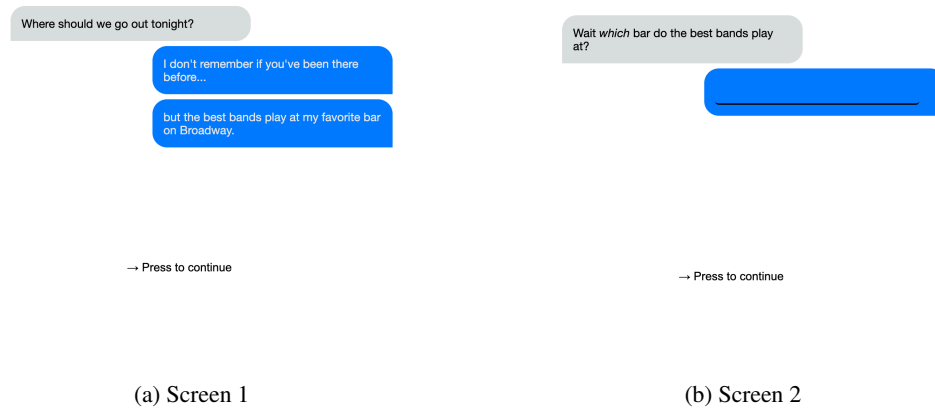


Figure 3.1: Experiment 1 example screens

The second screen displayed a text bubble with a question requesting clarification on the information given on the previous page, as discussed in Section 3.3.1.3. Participants were then instructed to write the answer that they felt was best in the blank provided in the last text bubble. The conversation was displayed over two screens purposefully with the intention of hiding the statement containing the possible ellipsis antecedents from the participants when they were responding to the final question, thereby preventing participants from re-reading the sentence after they were given the question to respond to. While writing their fill-in-the-blank response, participants instead had to consult their memory of their interpretation of the statement containing the antecedents. An example of the second screen is given in Figure 3.1.

I chose to use a free-response methodology instead of a forced-choice methodology for several reasons. First, the options for a forced-choice task would not have been consistent, as forcing participants to choose between the two provided possible antecedents would have required giving options that were not minimal pairs, such as the two choices *this banana of mine* and *apples*. Second, the free-response was easily integrated into the

text message display, which was intended to encourage participants to read the entire dialogue. Third, the free-response task obviated concerns that displaying certain NPs would itself influence participants' memory of the antecedent phrase or the interpretation of the ellipsis site. Last, and related to the previous point, the free-choice task allowed greater freedom to participants in reporting their interpretation of the ellipsis site. This led to unexpected responses, as discussed below in Section 3.3.3 and in Experiment 2.

A verbatim memory task was included in order to test whether participants were attending to the task and to encourage participants to read the full dialogues. Questions were given after 32 items throughout the experiment, and participants did not know after which items they would receive a question. In total, 11 critical items had a question, as well as 21 fillers. The questions themselves were either a verbatim repetition of the initial question on screen one of the immediately preceding item, or were an altered version of that question. Participants had to select whether or not they had read that exact question on the preceding item. Questions were balanced so that half were verbatim repetitions of the earlier question (correct answer *Yes*) and half were altered versions of the earlier question (correct answer *No*).

Finally, 84 fillers were included with the critical items. Forty-two fillers contained various types of ellipsis, and were split evenly between cataphoric ellipsis examples and anaphoric ellipsis examples. Twenty-eight fillers had an intended free-response of the pattern *this N*, which was intended to mask the definite antecedents in the critical items. By similar reasoning, 22 fillers had an intended free-response containing a possessive and a noun. Finally, fillers that contained anaphoric and cataphoric pronouns were also included.

3.3.3 Results

3.3.3.1 All Results

After the data were collected, I hand-coded all free responses. Each response was coded as belonging to exactly one of five bins: object, subject, conjoined, other, and blank. Example responses for each coding bin are given in (103)-(106). The total count for each bin is given in Table 3.2.

- (103) Example Response Coded for *Subject*
Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any [**apples**] right now.
- (104) Example Response Coded for *Object*
Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any [**oranges**] right now.
- (105) Example Response Coded for *Conjoined*
Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any [**apples and oranges**] right now.
- (106) Example Response Coded for *Other*
Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any [**pears**] right now.

Coding Bin	Total Response Count
Subject	1333
Object	510
Conjoined	74
Other	97
Blank	2

Table 3.2: Experiment 1 total count of ellipsis interpretations by coding bin.

Because there was no overarching pattern to the responses coded *other*, these responses were discarded. The conjoined cases are plotted by Antecedent condition in Figure 3.2. Because conjoined responses are less than 4% of the overall data, and because they were given predominantly in the Ambiguous condition, the responses were set aside




Figure 3.2: Experiment 1 proportion of conjoined antecedents chosen as the interpretation of the ellipsis site, by ellipsis condition and antecedent condition. Error bars indicate 95% confidence intervals.

for the general analysis. The conjoined responses are discussed in §3.3.4. Discarding the *conjoined* and *other* responses excluded 8.6% of the responses overall. The Subject and Object response bins were then analyzed as a binomial model. This analysis is given given in §3.3.3.2.

3.3.3.2 Subject vs. Object Responses

A model was created comparing Subject NP responses to Object NP responses. The coded results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the *brms* package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for

binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed-effects structure, including all interactions.

Table 3.3 provides the posterior estimates of the fixed effects as well as 95% credible intervals. The Antecedent conditions were coded with Helmert contrast coding. **Antecedent1** compares the *object* level to the mean of the *subject+ambiguous* levels. **Antecedent2** compares the *subject* and the *ambiguous* levels to each other. These groupings were chosen because the responses to the *subject* and *ambiguous* conditions patterned together to the exclusion of the responses to the *object* conditions. Table 3.3 therefore shows us a significant difference between the responses to the *object* conditions and responses to the other two conditions, as well as a significant difference between responses to the *subject* conditions and the *ambiguous* conditions. No significant effect of **Ellipsis** was found, and no significant interactions were found.

	estimate	est.error	2.5%	97.5%
Ellipsis	0.43	0.31	-0.20	1.03
Antecedent1	3.09	0.35	2.43	3.81
Antecedent2	1.36	0.46	0.48	2.31
Ellipsis:Antecedent1	0.25	0.32	-0.38	0.87
Ellipsis:Antecedent2	0.90	0.79	-0.61	2.52

Table 3.3: Experiment 1 Bayesian model estimates for subject and object NP antecedent choice.

3.3.4 Discussion

Let’s recall our two hypotheses. First, I hypothesized that we would find a proximity effect. This proximity effect could take one of three forms: *Locals Only*, under which cataphoric ellipsis will see a preference for resolution to the subject NP, while anaphoric




Figure 3.3: Experiment 1 proportion of subject antecedent chosen as the interpretation of the ellipsis site, by ellipsis condition and antecedent condition. Error bars indicate 95% confidence intervals.

resolution will see a preference for resolution to the object NP; *Wait for It*, under which both cataphoric and anaphoric ellipsis will see a preference for resolution to the object NP; and *Liger*, under which some combination of Locals Only and Wait for Is is observed. I also proposed that singular, specific antecedents would show a dispreference for serving as an antecedent compared to plural non-specific antecedents.

3.3.4.1 Plural nouns make better antecedents

Starting with the second hypothesis first, the findings confirmed a preference for plural antecedents. This preference is borne out in the main effects of **Antecedent1** and **Antecedent2** in Table 3.3, and can also be seen in Figure 3.4. Figure 3.4a shows the ratio




Figure 3.4: Experiment 1 proportion of antecedent NPs chosen as the interpretation of the ellipsis site.

of interpretation responses by antecedent condition. This plot collapses across ellipsis conditions. Figure 3.4b shows the ratio of interpretation responses by ellipsis condition, collapsing across antecedent conditions. Recall that in the Object condition, the object noun was plural, while the subject noun was specific singular. A majority of the responses in the Object condition yielded an object interpretation, showing a dispreference for choosing the singular subject noun as the antecedent for the ellipsis site. Recall that in the subject condition, however, the subject noun was a plural and the object noun was singular. Here we see that the vast majority of responses gave a subject interpretation for the ellipsis site, again showing a dispreference for the singular noun as antecedent.

It is notable that a non-trivial amount of singular specific antecedents were chosen as an antecedent. An example of this category of response is given in example (107). Even when the singular specific antecedent was chosen as the interpretation of the ellipsis site, participants wrote in the plural version of the singular antecedent noun, ensuring that the noun would fit grammatically into both the critical sentence and the prompt question.

This type of change is indicative of the repair strategy for ellipsis. The repair strategy proposes that the processor can repair a syntactic antecedent that closely matches the syntax of an ellipsis site if the repair involves few operations, there is evidence for those operations, and if the syntactic input corresponds to a natural speech error (Arregui et al., 2006; Frazier and Clifton, 2015). As all these criteria seemingly hold in these particular responses, it is reasonable to conclude that repair is happening in this constrained set of responses.

(107) Critical Sentence:

Honeycrisp apples would taste great with this navel orange for breakfast, but I don't know if we can find any [] right now.

Prompt Question:

Wait, you don't know if we can find any *what* right now?

Example Response:

navel oranges

3.3.4.2 Subject nouns were preferred antecedents

Turning to the hypothesis of proximity, the results failed to show any proximity preference. Instead, we found an overwhelming subject bias across all antecedent conditions and both ellipsis conditions. Looking at the antecedent condition, the Ambiguous condition – in which both subject and object noun were plural – shows a clear preference for a subject interpretation, with more than 75% of the total responses in the Ambiguous condition providing a subject interpretation. Counter to our hypothesis, no difference was found between the Anaphoric and Cataphoric conditions, suggesting that the location of the ellipsis site relative to the antecedent was irrelevant for the interpretation of the ellipsis site.

A subject bias is also seen in the Subject and Object conditions. Although a majority of responses in the Object condition gave an object interpretation, the percentage of object interpretations in the Object condition is less than that of the subject interpretations

in the Subject condition. That is, the singular manipulation had a harder time ‘pushing’ the interpretation of the ellipsis site to the object interpretation in the Object condition. Even when the subject noun was disadvantaged to serve as an antecedent, participants still chose it as the interpretation of the ellipsis site over a quarter of the time.

Why is subjecthood playing such a large role? Put another way, why do participants overwhelmingly choose the subject noun as the interpretation of the ellipsis site? This result is perhaps not surprising to those familiar with the privileged role of subjecthood within theories of discourse coherence and pronominal anaphora resolution. The preference of comprehenders to assign a pronoun the interpretation of the entity mentioned in the subject position of the previous clause is well-established in the anaphora literature (Frederiksen 1981; Crawley et al. 1990; Greene et al. 1992; Gordon et al. 1993; Garnham 2001; Kehler and Rohde 2013, a.o.). For example, centering theory uses the morpho-grammatical property of subjecthood as a strong predictor for the interpretations of pronouns (Grosz and Sidner, 1986; Grosz et al., 1995). Subjecthood was also integrated into early AI predictors of pronominal interpretation (Lappin and Leass, 1994). This preference has also been shown to track syntactic subjecthood, and not thematic roles (Gordon et al., 1993; Gordon and Chan, 1995). Given these existing findings, it is reasonable to conclude that the subject NPs in this experiment were conferred the same subjecthood benefit that has been found for subject antecedents of pronominal anaphors. This preference may have been especially effectual because the competing NPs were indirect objects, and were not occupying a primary role in the state or event described in the sentences.

The experimental results suggest that the interpretation of NPE sites may pattern more similarly to interpretations of pronominal anaphora than to sluicing, which has been shown to exhibit an object bias, as discussed in §3.1.2. The current results are compatible with multiple explanations for the subjecthood preference, and suggest several interesting

follow-up studies to further tease these options apart. For example, one could manipulate focus placement on the two candidate antecedents to see if focus on the object can pull the interpretation of the ellipsis site toward the indirect object (following the spirit of Carlson et al. 2009). For example, clefted structures or contexts placing narrow focus on the object could accomplish this. Passive structures could also be used to test whether the subjecthood preference found is indeed one for syntactic subjecthood, or whether it tracks semantic roles. I leave these interesting follow-ups for further research.

3.3.4.3 Coordinated antecedents look like split antecedents

An unexpected finding of the experiment was the responses that coordinated the two available antecedents. The coordinated responses are interesting in that they are interpretations of the ellipsis site that have no contiguous syntactic antecedent in the structure preceding the ellipsis site. These responses were found almost solely in the Ambiguous conditions, as shown in Figure 3.4. Recall that in the Ambiguous sentences, both candidate antecedents were plural nouns and were therefore equally felicitous as possible antecedents or postcedents. This is a reasonable finding, which can be explained by two preferences. One, it is pragmatically odd to coordinate two semantically non-parallel NPs, e.g. *I don't know if we have any apples or an orange*. Two, in the Object and Subject sentences, one NP was clearly a better candidate than the other, which in general pushed the interpretation of the ellipsis site toward the more pragmatically licit antecedent, as discussed above. It is therefore an unsurprising result that most of the coordinated responses occurred largely within the Ambiguous conditions.

The coordinated responses are reminiscent of the phenomenon of *split antecedents* for ellipsis. The exemplar of split antecedents is from Webber (1978):

- (108) Bob wants to sail round the world and Alice wants to climb Kilimanjaro, but neither of them can, because money is too tight.

Split antecedents are largely discussed in the literature on verb-phrase ellipsis, and are interesting because they do not contain a contiguous syntactic or semantic antecedent. Split antecedents have historically been a challenge for accounts that rely on matching syntactic or semantic antecedent material to an ellipsis site. As such, they have been used to argue in support of anaphoric pro-form accounts of ellipsis, which have no trouble accounting for their existence. For example, Hardt (1993, 1999) argues that VPE contains a null pro-form instead of elided syntactic content. Similarly, Asher (1993) argues that VPE is event anaphora, and that split antecedents are the referent of a sum event of the two antecedent VPs. Frazier and Duff (2019) argue for an account of syntactic material within an ellipsis site, but propose that split antecedents are a form of accommodation. Similar to the proportion of overall responses found here, Frazier and Duff also found that while coordinated VP responses to items such as (108) were given, they were not the most common response given. See §3.4 for additional discussion.

An in-depth exploration of the possibilities raised here is unfortunately beyond the scope of this dissertation. However, I close by noting that the findings here further establish the robust nature of ellipsis site interpretations that do not rely on a contiguous antecedent. The results here are also, to my knowledge, the first time that such readings have been observed with NPE, specifically.

3.4 Experiment 2: Ellipsis interpretation shows a locality preference

Counter to predictions, Experiment 1 found a strong subject interpretation preference for cataphoric and anaphoric NPE sites that had both subject and indirect object candidate antecedents. We also found no proximity preference in the interpretation of the ellipsis sites, counter to predictions. I proposed that one explanation for the failure to find a proximity effect is that subjecthood is such a privileged position over object position –

thereby lending the subject noun such greater prominence or salience than the object noun – that any proximity preferences were subjugated or washed out.

Experiment 2 follows up on the findings of Experiment 1 by more closely matching the salience of the two candidate antecedents. Experiment 2 follows the same general design as Experiment 1, but it places the two candidate antecedents in a conjoined structure in object position. This design obviates the overwhelming subject bias that we saw in Experiment 1. The hypotheses of this experiment are otherwise identical to those stated in the previous section.

3.4.1 Method

3.4.1.1 Methodology

As in Experiment 1, the current experiment collected offline, free-response interpretations of NPE sites. The experiment used the same basic instructions, methodology, and template as Experiment 1.

3.4.1.2 Participants

Forty-eight participants were recruited from the University of California, Santa Cruz. All participants were undergraduate native speakers of English; all received course credit for their participation.

3.4.1.3 Materials

Materials consisted of 36 items comprising four conditions. The design was a 2 x 2 with the factors *ellipsis* (Anaphoric ellipsis, Cataphoric Ellipsis) and *coordination* (NP coordination, TP coordination). All critical items contained two clauses. The Cataphoric ellipsis conditions contained NPE in the first clause with two possible postcedents in

the second clause, and the Anaphoric ellipsis conditions contained NPE in the second clause with two candidate antecedents in the first clause. The subject of both clauses was the first person singular pronoun *I*, which is not a felicitous antecedent for NPE and therefore was not competing with the intended antecedents for the interpretation of the ellipsis site.¹⁰ All ellipsis sites and antecedents were NPs headed by possessives containing a name.

Finally, in the NP coordination conditions the two candidate antecedents were coordinated directly (at the DP level), and in the TP coordination conditions the candidate antecedents were coordinated at the level of the TP. The TP coordinated structure added both additional structure and additional linear distance between the ellipsis site and one of the candidate antecedents compared to the NP coordinated structure: the first candidate antecedent in the anaphoric cases, and the second candidate antecedent in the cataphoric cases. I hypothesize that this added structure will encourage an interpretation of the ellipsis site of the structurally more local candidate antecedent. An example item is given in Table 3.4.

¹⁰This claim is intended informally; whether the ellipsis resolution process preliminarily searches the entire discourse space for a possible representation for the ellipsis site is not settled. However, based on the findings of structural search in Kazanina et al. (2007) and Patterson and Felser (2019), it is reasonable to hypothesize that if the pronoun in this situation is considered it is discarded early enough in the process to not interfere with the intended offline experimental manipulation in this particular experiment.

Ellipsis	Coord.	Sentence
Anaphoric	NP	I saw Mary's dog and Susan's cat yesterday, but I didn't see Jane's at the time.
Anaphoric	TP	I saw Mary's dog and I saw Susan's cat yesterday, but I didn't see Jane's at the time.
Cataphoric	NP	I didn't see Jane's at the time, but I saw Mary's dog and Susan's cat yesterday.
Cataphoric	TP	I didn't see Jane's at the time, but I saw Mary's dog and I saw Susan's cat yesterday.

Table 3.4: Example item from Experiment 2.

3.4.2 Procedure

The experiment used the IbxFarm experimental software (Drummond, 2014). Critical items were distributed into four Latin Square lists and were presented in pseudo-random order. The format and experimental procedure was then identical to that of Experiment 1. An example of a complete first and second screen for a critical item is given in Figure 3.5.

A verbatim memory task was again included in order to test whether participants were attending to the task and to encourage participants to read the full dialogues. The task was identical to that in Experiment 1. Questions were included after 27 items throughout the experiment, and participants did not know after which items they would receive a question. In total, eight critical items had a question, as well as 19 fillers. As in Experiment 1, questions were balanced so that half were verbatim repetitions of the earlier question (correct answer *Yes*) and half were altered versions of the earlier question (correct answer *No*).

Finally, 84 fillers were included with the experimental items. The fillers were modified slightly from Experiment 1 in order to mask the altered nature of the critical items. Fillers

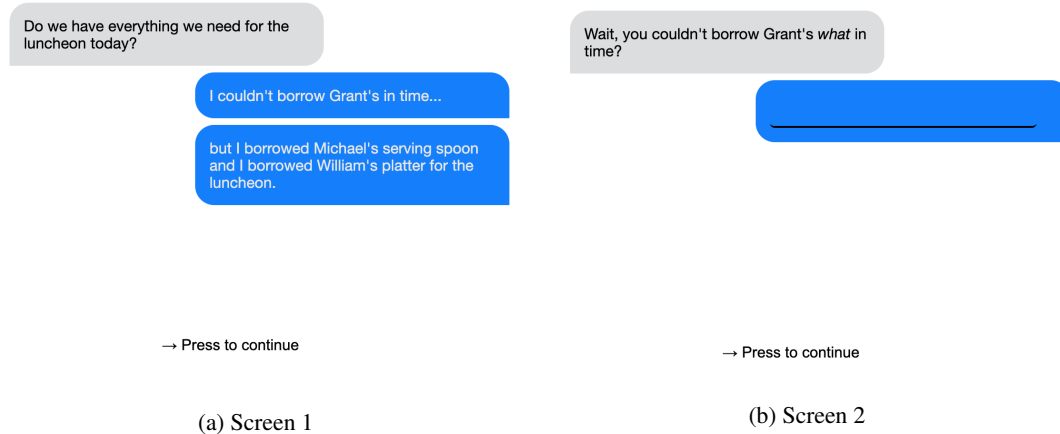


Figure 3.5: Experiment 2 example screens

included twenty-one items containing possessive NPs, and additional items containing ambiguities such as strict/sloppy VPE readings and ambiguous sluices.

3.4.3 Results

After all the data were collected, I again hand-coded all free responses. Each response was coded as belonging to exactly one of six bins: First NP, Second NP, Conjoined, Kind, Other, and Blank. Example responses for each coding bin are given in (109)-(113). The total count for each bin is given in Table 3.5.

- (109) Example Response Coded for *First NP*
 I saw Mary's dog and Susan's cat yesterday, but I didn't see Jane's [**dog**] at the time.
- (110) Example Response Coded for *Second NP*
 I saw Mary's dog and Susan's cat yesterday, but I didn't see Jane's [**cat**] at the time.
- (111) Example Response Coded for *Kind*
 I saw Mary's dog and Susan's cat yesterday, but I didn't see Jane's [**pet**] at the time.

- (112) Example Response Coded for *Conjoined*
 I saw Mary’s dog and Susan’s cat yesterday, but I didn’t see Jane’s [**dog {and/or} cat**] at the time.
- (113) Example Response Coded for *Other*
 I saw Mary’s dog and Susan’s cat yesterday, but I didn’t see Jane’s [**rabbit**] at the time.

Coding Bin	Total Response Count
First	438
Second	382
Kind	839
Conjoined	71
Other	366
Blank	28

Table 3.5: Experiment 2 total count of ellipsis interpretations by coding bin.

For the analysis, only responses from the First, Second, Kind, and Conjoined bins were used. This division excluded 19% of the responses overall. The 28 Blank responses that were given deserve no additional analysis, and are likely the result of accidental key presses. The Other responses are a grab bag of various responses existing on a spectrum of seriousness. However, a couple interesting patterns appear in the data that are worth commenting on. Fifty-three participants contributed at least one response to the category. Ninety-one of the 366 total responses consist of some version of the response ‘don’t know, didn’t specify, didn’t say, unclear,’ etc. These responses indicate a failure or unwillingness to resolve the ellipsis site at all. These participants simply failed to resolve the ellipsis site to either of the available antecedents. These interpretations were disproportionately contributed by so-called heavy-users: 29 of these responses came from participant 17, and 20 of them from participant 74. The next largest contributor was participant 7 at 9 responses, and participant 30 at 7 responses. These responses demonstrate the individual variation that occurs in strategies of ellipsis interpretation.

For these participants, the lack of a uniquely salient antecedent led to a failure to resolve the ellipsis site at all.

The other pattern seen in the Other responses is less frequent but still notable: fifty-two responses consisted of some version of the response ‘something, stuff, thing’. These responses were also disproportionately given by one participant, who gave 10 responses of this nature; no other participant contributed more than 3 responses. These responses show that the participants did resolve the ellipsis site, but similar to the Kind readings they failed to resolve it to either of the candidate antecedents. These responses are, in a sense, a less eloquent version of the Kind responses.

Figure 3.6 shows the total proportion of all types of noun interpretations across the anaphoric and cataphoric conditions. There are three important take-aways from the response pattern seen in the plot. The first is that the Kind and Conjoined responses showed no sensitivity to the ellipsis conditions. Whether the ellipsis site interpretation was anaphoric or cataphoric did not appear to influence whether or not a comprehender chose a Coordinated or Kind reading of the ellipsis site. In contrast, the relative number of First and Second noun interpretations flipped between the anaphoric and cataphoric ellipsis conditions: more Second NP interpretations were given in the anaphoric condition than in the cataphoric condition, and more First NP interpretations were given in the cataphoric condition than in the anaphoric condition. The following two subsections focus on the First and Second NP interpretations, and the final subsection will briefly discuss the Kind and Coordinated interpretations. For ease of interpretation, analyses were run for binomial models that collapsed across particular response bins. These analyses are given in sections 3.4.3.1 – 3.4.3.3.¹¹

¹¹Note that due to the free response methodology used no response times were captured.

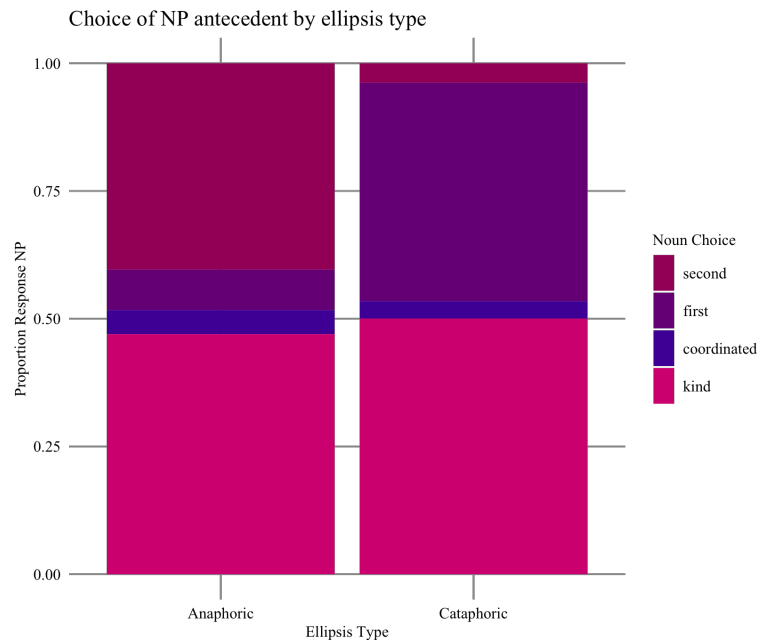


Figure 3.6: Experiment 2 proportion of antecedent NPs chosen as the interpretation of the ellipsis site, by Ellipsis condition.

3.4.3.1 First NP Analysis

This section compares First NP responses to non-First NP responses. The non-First NP responses collapse across the Second, Kind, and Conjoined response bins. The coded results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 3.6 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

Figure 3.7 shows the overall proportion of First NP responses, broken down by

	estimate	est.error	2.5%	97.5%
Intercept	-2.25	0.40	-3.05	-1.49
Ellipsis	-2.79	0.68	-4.17	-1.49
Coordination	0.31	0.26	-0.19	0.83
Ellipsis:Coordination	-0.10	0.58	-1.24	1.06

Table 3.6: Experiment 2 Bayesian model estimates for First NP antecedent choice.




Figure 3.7: Experiment 2 overall proportion of First NPs responses, by Ellipsis and Coordination conditions. Error bars indicate 95% confidence intervals.

Coordination conditions and Ellipsis conditions. The plot shows that while slightly less than 50% of total responses in the Cataphoric condition consisted of First NP responses, only about 10% of responses in the Anaphoric condition consisted of First NP responses. This finding is reflected in a main effect of **Ellipsis** in Table 3.6. No significant difference was found between the NP and TP Coordination levels, as demonstrated by the lack of a significant main effect of **Coordination** or a significant interaction of **Ellipsis:Coordination** in Table 3.6.

3.4.3.2 Second NP Analysis

This section compares Second NP responses to non-Second NP responses. The non-Second NP responses collapse across the First, Kind, and Conjoined response bins. The coded results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the *brms* package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 3.7 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

	estimate	est.error	2.5%	97.5%
Intercept	-2.71	0.37	-3.47	-2.03
Ellipsis	3.02	0.62	1.83	4.27
Coordination	-0.07	0.31	-0.68	0.54
Ellipsis:Coordination	0.54	0.63	-0.68	1.81

Table 3.7: Experiment 2 Bayesian model estimates for Second NP antecedent choice.

Figure 3.8 shows the overall proportion of Second NP responses, broken down by Coordination conditions and Ellipsis conditions. The plot shows that while about 40% of total responses in the Anaphoric condition consisted of Second NP responses, only about 4% of responses in the Cataphoric condition consisted of Second NP responses. This finding is reflected in a main effect of **Ellipsis** in Table 3.7. No significant difference was found between the NP and TP Coordination levels, as demonstrated by the lack of a significant main effect of **Coordination** or a significant interaction of **Ellipsis:Coordination** in Table 3.7.




Figure 3.8: Experiment 2 proportion of Second NPs chosen as the interpretation of the ellipsis site, by Ellipsis and Coordination conditions. Error bars indicate 95% confidence intervals.

3.4.3.3 Kind & Conjoined NP Analysis

This section looks at the Kind NP and Conjoined NP responses. I'll discuss the non-Kind NP responses first, which collapse across the First, Second, and Conjoined response bins. The coded results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 3.8 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

Figure 3.9 shows the overall proportion of Kind responses, broken down by Coordi-

	estimate	est.error	2.5%	97.5%
Intercept	-0.43	0.58	-1.59	0.69
Ellipsis	-0.29	0.27	-0.83	0.24
Coordination	-0.25	0.22	-0.68	0.17
Ellipsis:Coordination	-0.22	0.43	-1.08	0.63

Table 3.8: Experiment 2 Bayesian model estimates for Kind antecedent choice.

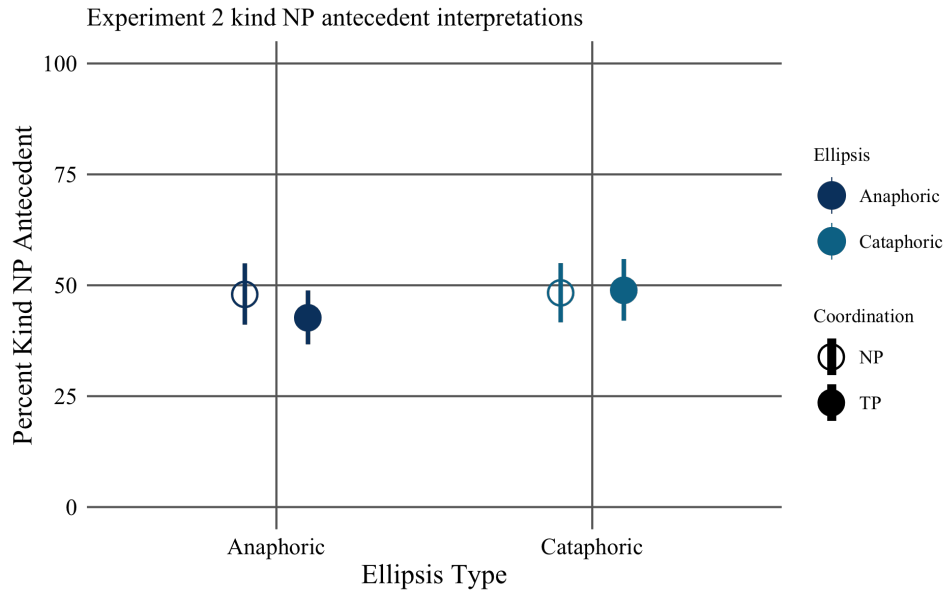


Figure 3.9: Experiment 2 proportion of Kind NPs chosen as the interpretation of the ellipsis site, by Ellipsis and Coordination conditions. Error bars indicate 95% confidence intervals.

nation conditions and Ellipsis conditions. The plot shows that slightly less than 50% of total responses in the Anaphoric condition and in the Cataphoric condition consisted of Kind responses. This finding is reflected in the failure to find a significant main effect of **Ellipsis** in Table 3.8. Again, no significant difference was found between the NP and TP Coordination levels, as demonstrated by the lack of a significant main effect of **Coordination** or a significant interaction of **Ellipsis:Coordination** in Table 3.8.

The Conjoined NP responses were compared to non-Conjoined NP responses, which

collapsed across the First, Second, and Kind response bins. The coded results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 3.9 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

	estimate	est.error	2.5%	97.5%
Intercept	-7.19	1.05	-9.57	-5.47
Ellipsis	0.92	0.93	-0.55	3.15
Coordination	-0.45	0.54	-1.58	0.52
Ellipsis:Coordination	1.74	1.03	-0.22	3.853

Table 3.9: Experiment 2 Bayesian model estimates for Conjoined antecedent choice.

Figure 3.10 shows the overall proportion of Conjoined responses, broken down by Coordination conditions and Ellipsis conditions. The plot shows that slightly less than 5% of total responses in the Anaphoric condition and in the Cataphoric condition consisted of Conjoined responses. This finding is reflected in the failure to find a significant main effect of **Ellipsis** in Table 3.9. And once again, no significant difference was found between the NP and TP Coordination levels, as demonstrated by the lack of a significant main effect of **Coordination** or a significant interaction of **Ellipsis:Coordination** in Table 3.9.




Figure 3.10: Experiment 2 proportion of Conjoined NPs chosen as the interpretation of the ellipsis site, by Ellipsis and Coordination conditions. Error bars indicate 95% confidence intervals.

3.4.4 Discussion

3.4.4.1 Results show a proximity preference

Recall our three hypotheses about the offline interpretation preferences of ellipsis, repeated in (114), (115), and (116).

(114) Locals Only:

Ellipsis interpretation will show a preference for resolution to the closest candidate antecedent in linear distance to the ellipsis site, measured by word count.

(115) Wait for It:

Ellipsis interpretation will show a preference for object resolution, replicating the findings for anaphoric ellipsis by Martin and McElree (2011); Harris (2015); Harris and Carlson (2016, 2015).




Figure 3.11: Experiment 2 proportion of antecedent NPs chosen as the interpretation of the ellipsis site, by ellipsis condition.

(116) Liger:

Some combination of the preferences in Locals Only and Wait For It will emerge.

The experimental results support the Locals Only hypothesis, which states that ellipsis interpretation would show a preference for resolution to the candidate antecedent in greatest proximity to the ellipsis site. This preference was borne out in a First NP preference for cataphoric ellipsis sentences, and a Second NP preference for anaphoric ellipsis sentences, as shown in Figure 3.11. This section discusses the implications of this finding. An unexpected finding is that neither the First NP nor the Second NP interpretation was the dominant interpretation chosen by participants; instead, almost 50% of responses gave a Kind interpretation. These responses are discussed in §3.4.4.3.

These results have two important takeaways. The first is that they are evidence that the result of Experiment 1 was in fact due to a subject bias overwhelming a proximity

bias. The results of Experiments 1 and 2 reveal that ellipsis resolution is sensitive to grammatical or structural effects that elevate certain representations in the discourse to a high level of salience. In Experiment 1, the greater prominence of the subject of the items overwhelmed the proximity bias that emerged in Experiment 2, when the candidate antecedents were more closely matched in structural prominence and semantic role. The finding that ellipsis resolution is subject to this type of prominence effect aligns a theory of ellipsis resolution with existing theories of anaphora resolution, in which the interaction of multiple factors of different weights or strengths determines the interpretation of the anaphor (Grosz and Sidner 1986; Greene et al. 1992; Lappin and Leass 1994; Grosz et al. 1995; O'Brien et al. 1997, a.m.o.).

The second takeaway of the Experiment 2 results is the finding of greatest interest: ellipsis resolution showed a preference to the NP in closest proximity to the ellipsis site, regardless of whether the NP was encountered prior to or after the ellipsis site. This preference was borne out through a greater proportion of Second NP interpretations in the Anaphoric condition than in the Cataphoric condition, and a greater proportion of First NP interpretations in the Cataphoric condition than in the Anaphoric condition. Experiment 2 is, to my knowledge, the first evidence that the interpretation preferences of an ellipsis site flip based on the location of the ellipsis site relative to candidate antecedents.

Recall that the motivation for testing the offline preferences of cataphoric ellipsis was twofold. First, existing theories of ellipsis processing have looked only at anaphoric ellipsis interpretation, and have therefore been unable to disentangle the effects of recency and proximity. The results here show that the preference for resolution of an anaphoric ellipsis site to an object antecedent seen in previous studies (Martin and McElree, 2011; Harris, 2015) is not a result of recency. When candidate antecedents are otherwise matched for salience, proximity of the candidate antecedent to the ellipsis site modulates

interpretation preferences. A theory of recency would fail to account for this behavior, as it has nothing to say about situations in which an ellipsis site precedes the candidate antecedents. Therefore, notions of recency should not be used when discussing ellipsis interpretation preferences.

Second, the search mechanisms discussed in previous studies were necessarily backward-looking in nature, and therefore failed to acknowledge the empirical fact that ellipsis sites can precede their antecedents. The fact that only backward-looking search processes had been investigated raised the question of whether cataphoric ellipsis relies on fundamentally different processing mechanisms than anaphoric ellipsis. The results of Experiment 2 provide evidence against a theory in which cataphoric ellipsis sites are not resolved until an initial clausal or sentential parse is complete. Recall that under such a hypothesis, the search process used for cataphoric ellipsis would be identical to that used for anaphoric ellipsis resolution. However, instead of initiating a search at the ellipsis site, in cataphoric ellipsis constructions the search process would be initiated at the end of a clause or boundary once an initial parse has been constructed. The results of Experiment 2 do not support such a theory. Instead, the results provide additional evidence for a backward-looking search for anaphoric ellipsis, and provide novel preliminary evidence for a forward-looking search process in cataphoric ellipsis.

The finding that comprehenders showed a preference for resolving a cataphoric ellipsis site to the first NP encountered after the ellipsis site suggests that comprehenders are resolving the ellipsis site in real-time as they are parsing the sentence. Under this model of cataphoric ellipsis processing, the difference between anaphoric and cataphoric processing is the nature of the search procedure that takes place. In both anaphoric and cataphoric ellipsis, a search is initiated once an ellipsis site is encountered. In anaphoric resolution, the search procedure simultaneously searches the available discourse space in memory for an appropriate antecedent using cue-matching procedures. In cataphoric

ellipsis, however, the search for a matching antecedent occurs *while* the parsing of the sentence is ongoing. Therefore, no search through memory takes place. Instead, representations are searched as they are parsed; if a representation matches the cues needed to resolve the ellipsis site, that representation is integrated into the representation of the ellipsis site directly from focal attention.¹² The results here cannot confirm that this search is taking place in real-time, given that the interpretations were collected offline. Chapter 4 investigates the real-time processing of cataphoric constructions and shows that the search strategy used is indeed conducted online as the sentence is being parsed.

3.4.4.2 Results show a null effect of structure

Another result of note is the failure to find an effect of structure: the results revealed no difference in preference between the NP and TP Coordination conditions, as demonstrated in Figure 3.6. The interaction of distance and interference has been investigated in previous online experiments of ellipsis comprehension. Martin and McElree (2011) found that increasing the distance between an ellipsis site and its antecedent and increasing the number of candidate antecedents both decrease the quality of the antecedent representations, in that these factors decrease comprehension accuracy. However, they found that these factors do not affect processing speed of the ellipsis interpretation. Martin and McElree also found that sluicing out of a conjoined clause presented difficulties, which they hypothesized might be attributed to the theory that coordinated clauses are ‘chunked’ in memory. Harris (2015) also presents converging evidence that interfering antecedents degrade the quality of antecedent representations in memory.

The results here interface with the results of these previous studies only indirectly, as both Martin and McElree and Harris were conducting online studies, while the results

¹²This explanation is vague on the status of syntactic material in the representation; see discussion of this in the Conclusion.

here provide only offline interpretations. Therefore, I can make claims only about offline comprehension. Keeping that restriction in mind, the results here suggest that interference plays a larger role than linear or structural distance in determining offline comprehension of the ellipsis site. The additional structure and linear distance between the ellipsis site and the candidate antecedents had no effect on pushing the interpretations of the ellipsis site toward or away from the more local candidate antecedent. For example, we may have thought that the NP Coordination examples would have been more likely to receive a Conjoined or Kind interpretation than the TP Coordination examples, in which the non-local candidate antecedent was structurally and linearly further away from the ellipsis site than in the NP Coordination examples. However, no such bias was found. The findings suggest that the degradation that occurs in retrieving a candidate antecedent across additional structure is mild enough to not survive in offline comprehension measures.

3.4.4.3 Results show overall preference for Kind interpretations

An unexpected finding of Experiment 2 is the overall preference for Kind interpretations of the ellipsis site. Participant feedback suggests that it was common for participants to feel unclarity (and even frustration) regarding which NP was supposed to serve as the antecedent/postcedent to the ellipsis site. In one sense this is an inevitable and intentional consequence of the design: I wanted to see what would happen when two candidate antecedents/postcedents are matched for salience. However, I think it is still an important consequence to note. The fact that this unclarity of interpretation arose is itself further proof that salience is an integral part of our interpretations of ellipsis sites; the mere presence of a felicitous antecedent is not sufficient for our understanding of ellipsis sites. Further, while it's clear the examples engender feelings of unease, they are judged by native speakers to be overall acceptable (see §3.5). I believe it would be incorrect to label these examples as ungrammatical, as it is unclear what rules of grammar they would be

violating. Instead, I think that the unease originates from our dislike of having to choose between two equally matched candidate antecedents/postcedents. Where exactly the root of this dislike should be located – in the grammatical rules of ellipsis, in the processing of ellipsis – is exactly the question I am probing. Furthermore, I believe that it is likely due to this unease that the Kind interpretations arose.

As noted in the discussion to Experiment 1, the Kind responses are notable in having no strict syntactic or semantic antecedent available in the linguistic context surrounding the ellipsis site. A plausible explanation for the Kind responses is that they are a consequence of participants accommodating a sub-QuD in the mini-experimental context.¹³ For example, a similar pattern of results was found in the first experiment of Frazier and Duff 2019. Frazier and Duff gathered free-response answers to split antecedent VPE examples such as (117). They found that 69% of responses to questions probing the interpretation of the elided clause did not contain material copied from the surface form of the antecedents. For example, most responses were variations of the form ‘neither of them has gone on their anticipated adventure.’ This pattern of response appears to be accommodating a sub-QuD in the context: e.g. ‘Wendy and Bruce both want to go on adventures.’ Indeed, Frazier and Duff argue that accommodation is exactly what is happening in these examples (although they do not propose or advocate for the position held here).

(117) Wendy is eager to sail around the world and Bruce is eager to climb Mt. Kilimanjaro, but neither of them has so far.

Evidence for this hypothesis is also found in the current experiment. For example, the number of Kind responses was not distributed evenly across items. In fact, certain items received a large number of identical Kind answers. These items also tended to

¹³I am grateful to Pranav Anand for raising this possibility.

be exactly those for which a salient sub-QuD could be inferred. For example, item (1) is given in (118). Out of the 59 responses for item (1), 32 participants responded ‘pet’, and no other Kind readings were given. Because *cat* and *dog* are both salient hyponyms of *pet*, it is reasonable to conclude that participants accommodated a sub-QuD in the context, such as ‘Whose pets did you see yesterday?’.

Similarly, out of the responses from item (17), given in (119), 13 participants said some version of ‘graham crackers’. This example is slightly different than example (118) because *graham crackers* is not a hyponym of the two candidate antecedents, *marshmallows* and *chocolate bars*. Instead, the interpretation and the two candidate antecedents are meronyms of a fourth noun, *s’mores*. The context set up by the Contextual QuD, a campfire, plus the more detailed context added by the Critical sentence (marshmallows + chocolate bars = s’mores) raises a sub-QuD like ‘Do we have what we need to make s’mores?’.

(118) Contextual QuD: What did you do yesterday?

Critical sentence: I saw Mary’s dog and Susan’s cat yesterday, but I didn’t see Jane’s at the time.

(119) Contextual QuD: What are we going to eat at the campfire tonight?

Critical sentence: I bought Zoe’s marshmallows and Olivia’s chocolate bars for the campfire, but I forgot to bring Stella’s with me.

However, items that had less of a clear semantic relationship between the two candidate antecedents and a third noun received fewer Kind interpretations. For example, the items with the two lowest Kind readings were items (2), which received two total Kind interpretations, and (7), which received three total Kind interpretations. These items are given in examples (120) and (121), respectively. Because there is no salient semantic relationship between *partner* and *colleague*, or between *pug puppy* and *niece*, there is no

salient sub-QuD raised by the Critical sentences in these items. If the Kind readings are indeed due to the raising of this sub-QuD, then the very low rate of Kind responses in these items is predicted.

(120) Contextual QuD: What did you do at the networking dinner?

Critical sentence: I talked to Chris's partner and Pete's colleague for a while, but I wasn't able to talk to Owen's at the dinner.

(121) Contextual QuD: What did you do at the beach today?

Critical sentence: I swam with Janice's pug puppy and Kelsey's niece for a while, but I wasn't able to swim with Ivy's for some reason.

3.5 Experiment 2 items are moderately acceptable

In the last section, I raised the question of at which level of representation we should locate the dislike of the items in Experiment 2. I argued that it would be incorrect to label these examples as ungrammatical, as it is unclear which rules of grammar they would be violating. For example, the items do not violate any syntactic or semantic rules, and they adhere to the licensing requirements of the theories of ellipsis discussed in §2, save the salience requirement. If the uneasiness around the items in Experiment 2 is indeed due to the violation of a salience requirement, as I believe to be the case, then the pertinent question is at what level of representation this salience requirement holds.

It is unlikely that a salience requirement holds in the grammar, although it is possible to claim this. I believe a more plausible state of affairs is that the requirement is at the processing level. One could argue that the requirement is a pragmatic requirement, because the items in Experiment 2 are certainly uncooperative to utter. However, it is unclear that speakers and listeners align their production and comprehension of language in a way that is optimized for both parties. Additionally, we should ask why the sentences

are uncooperative. Simple ambiguity does not necessarily lead to processing difficulty or decreased acceptability; in fact, some ambiguous pronouns have been shown to be read faster than unambiguous pronouns, indicating that the ambiguity facilitates processing in some way (see e.g. van Gompel et al. 2005; Clifton and Staub 2008 for discussion of whether or not there are processing advantages for structural ambiguities; see Gernsbacher 1989; MacWhinney and MacDonald 1990; Arnold J.E. and J.C. 2000; Badecker and Staub 2002; Stewart et al. 2007; Kaiser 2011 for discussions of the ambiguity advantage with pronouns.). Instead, it may be that the items in Experiment 2 are uncooperative utterances because they are difficult for a listener to comprehend. But if difficulty of comprehension leads to feelings of unease (e.g. Chomsky and Miller 1963), then we must ask why these examples are difficult to comprehend, which leads us back to processing difficulty.

The current experiment seeks to partially address these questions by soliciting acceptability judgments of the items in Experiment 2. I show that while the items are not ranked as perfectly acceptable, they are ranked overall as moderately acceptable.¹⁴

3.5.1 Method

3.5.1.1 Participants

Forty-eight participants from Prolific participated in the experiment; all self-identified as native speakers of English.

3.5.1.2 Materials

Materials are the same as those used in Experiment 2; four additional items of the same pattern were also included, for a total of 40 items. Items were run as fillers within another

¹⁴I am grateful to Jack Duff and Amanda Rysling for including these items in their experiment.

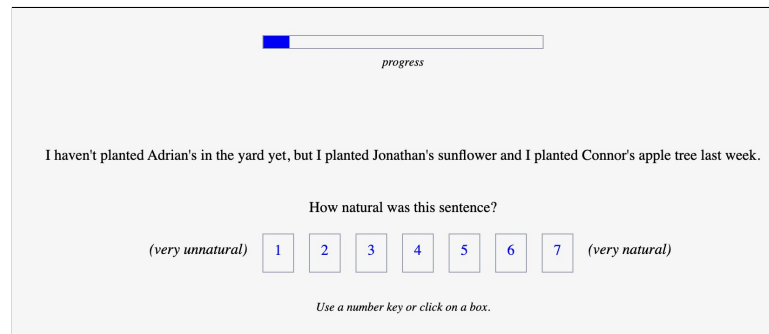


Figure 3.12: Acceptability experiment of items from Experiment 2.

experiment. Non-critical items contained an assortment of constructions, including ungrammatical items characterized by agreement mismatches, e.g. ‘The tall nurse who cared for the elderly widow definitely were reluctant to work long shifts’.

3.5.1.3 Procedure

The experiment used the IbxFarm experimental software (Drummond, 2014). Critical items were distributed into four Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, ensuring balanced responses across lists. Unlike Experiment 2, experimental items were presented as simple sentences centered on the screen. Sentences were presented with a Likert acceptability scale on which participants were asked to rate each sentence from 1 (very unnatural) to 7 (very natural). An example screen is given in Figure 3.12.

3.5.1.4 Results

Results were analyzed using a Bayesian mixed-effects ordinal regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. The cumulative family with logit link was used, following the recommendation given for ordinal data in Bürkner 2017. The model included the full random effects

structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 3.10 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

	estimate	est.error	2.5%	97.5%
Ellipsis	0.20	0.11	-0.01	0.42
Coordination	-0.37	0.10	-0.57	-0.18
Ellipsis:Coordination	0.04	0.20	-0.35	0.43

Table 3.10: Bayesian model estimates for the acceptability of Experiment 2 items.

The model shows a small but significant main effect of **Coordination**, such that the NP Coordination condition was rated significantly higher than the CP coordination. No other effects reached significance.

The mean ratings are given in Tables 3.11 - 3.13, with means computed over raw scores, over items, and over participants, respectively. Standard errors reveal the greatest standard error when means are computed over participants, suggesting more substantial participant variation than item variation in ratings. Because participant variation is the greatest standard error and therefore the most conservative, means computed over participants are used throughout the rest of the presentation and discussion of the results. Mean ratings over participant are given in Figure 3.13.

Ellipsis	Coord.	Rating Mean	St.dev.	Num. Items	St.error
Anaphoric	NP	4.220833	1.717278	480	0.0784
Anaphoric	TP	3.989583	1.725982	480	0.0788
Cataphoric	NP	4.108333	1.703715	480	0.0778
Cataphoric	TP	3.856250	1.668248	480	0.0761

Table 3.11: Experiment 2 acceptability ratings means, raw ratings.

Ellipsis	Coord.	Rating Mean	St.dev.	Num. Items	St.error
Anaphoric	NP	4.220833	0.4768688	40	0.0754
Anaphoric	TP	3.989583	0.4332953	40	0.0685
Cataphoric	NP	4.108333	0.4377975	40	0.0692
Cataphoric	TP	3.856250	0.4575995	40	0.0724

Table 3.12: Experiment 2 acceptability ratings means over items.

Ellipsis	Coord.	Rating Mean	St.dev.	Num. Items	St.error
Anaphoric	NP	4.220833	1.223347	48	0.1766
Anaphoric	TP	3.989583	1.312824	48	0.1895
Cataphoric	NP	4.108333	1.279932	48	0.1847
Cataphoric	TP	3.856250	1.251536	48	0.1806

Table 3.13: Experiment 2 acceptability ratings means over participants.

3.5.1.5 Discussion

The results show that the critical items from Experiment 2 were rated as moderately acceptable, averaging about a rating of 4 on a Likert scale of 1 (unacceptable) to 7 (perfectly acceptable). Figure 3.14 shows a box plot comparison of the critical items from Experiment 2 and the rest of the items included in the ratings experiment. The plots show that participants used the entire scale to rate both the fillers and the critical items. The median rating for the filler items was one step higher on the Likert scale than the critical items. The lower 25% of the inter-quartile range for the fillers overlaps with the upper 25% of the inter-quartile range for the critical items, showing that the critical items were generally rated lower than the fillers. The fillers showed greater variation in rating, with their inter-quartile range spanning the Likert range 4-7. The inter-quartile range for the critical items, however, spanned the Likert range 3-5. We can conclude that the critical items were rated lower overall than the filler items; however, they are not overall rated substantially lower than the fillers, or very low overall.

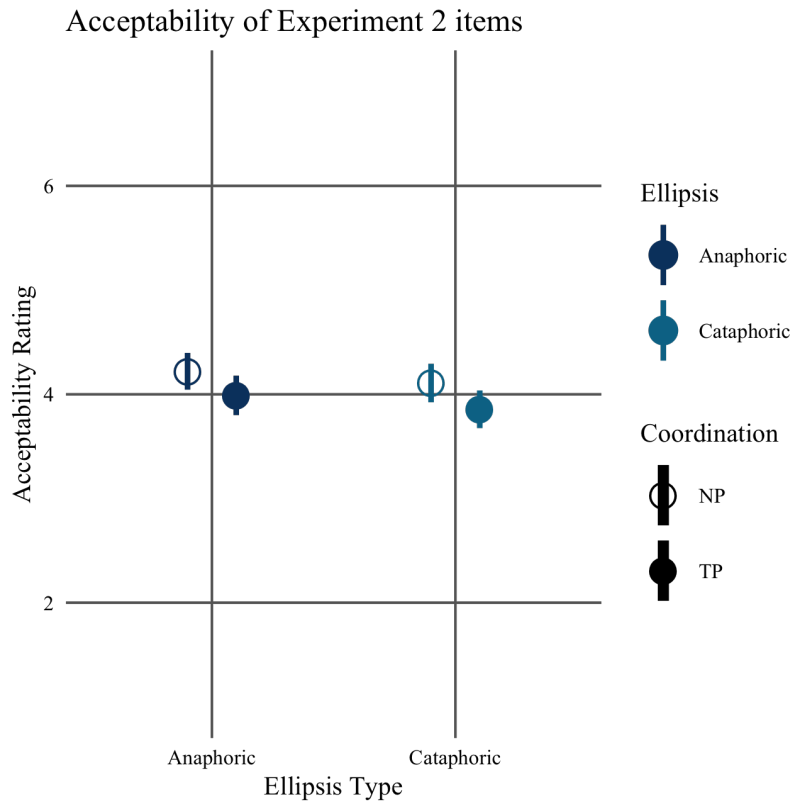


Figure 3.13: Experiment 2 item acceptability. Error bars indicate standard error over participants.

The analysis showed that the NP Coordination conditions were rated slightly but significantly higher than the TP Coordination conditions. I believe this effect is due to a slight repetition penalty within the TP Coordination condition. Because the TP Coordination conditions join two TPs which differ only in their DP, there is a repeated portion of the phrase that is avoided in the NP Coordination conditions. Due to the very small nature of the effect, I won't expound upon this any further.

The lack of an effect in the Ellipsis conditions is more interesting. I would a priori believe that cataphoric ellipsis sentences would be ranked lower than anaphoric ellipsis sentences for two reasons. First, comprehenders must recognize that cataphoric ellipsis




Figure 3.14: Experiment 2 item and filler acceptability.

has an ellipsis 'gap' that must be filled. Second, cataphoric ellipsis is generally a more marked structure than anaphoric ellipsis. Counter to these prior beliefs, the Cataphoric and Anaphoric conditions were not rated significantly different, showing that participants had no preference for one construction over the other.

Chapter 4

Online ellipsis comprehension

In the previous chapter, I discussed converging experimental evidence for a theory of ellipsis processing in which the resolution of ellipsis sites involves a retrieval search process conducted over content-addressable representations in memory. I also discussed that no experiments to date have successfully addressed the nature of the search mechanism in which no antecedent is available for retrieval in memory. That is, the current evidence for ellipsis comprehension has investigated only the ‘retrospective’ search that occurs in anaphoric ellipsis. I argued from offline experimental judgments that cataphoric ellipsis interpretation is subject to a search procedure, just like anaphoric ellipsis. However, I proposed that the properties of this search are slightly different than for anaphoric ellipsis. I hypothesized that both forms of ellipsis see a search procedure initiated at the ellipsis site; however, while the search in anaphoric ellipsis takes place over representations in memory, the search in cataphoric ellipsis operates over incoming representations as parsing of a sentence happens in real-time.

The experiments in this chapter test the processing of elliptical dependencies as a forward-looking process by looking at the real-time processing of cataphoric NPE. I used online measures of incremental processing comprehension to investigate whether comprehenders actively resolve ellipsis sites when parsing a sentence, or whether ellipsis sites are resolved only after an initial parse is generated. Note that in order to keep in

mind that we are dealing with cataphoric ellipsis in particular, I refer to the material that provides the interpretation of the ellipsis site as a *postcedent*.

Experimental evidence for the active processing of long-distance dependencies in language has been available in the processing literature for some time (Fodor, 1978; Crain and Fodor, 1985; Stowe, 1986; Frazier, 1987; Frazier and Clifton, 1989). Long-distance dependencies encompass a broad class of constructions that include antecedent-pronoun relationships and verb-argument relationships, but the arguably most widely discussed long-distance dependency relationship is that between a filler and its gap. Filler-gap dependencies are found in constructions such as relative clauses, topicalization, and *wh*-movement. Experimental investigations of these constructions provided early evidence that the parser does not wait to receive unequivocal evidence of a gap before positing one at the first structurally-licit position as a sentence is processed in real-time. For example, in so-called *filled-gap effect* constructions a reading-time slow-down is observed when a reader encounters an NP where a gap has been predicted. Stowe (1986) gives the following example, in which a reading time slowdown occurs at *us* in (122a) relative to (122b). This suggests that the comprehenders had already posited a gap following the verb *bring* in (122a) before they received unambiguous evidence on whether the direct object slot was filled or contained a gap.

- (122) a. My brother wanted to know who Ruth will bring us home to ___ at Christmas.
b. My brother wanted to know if Ruth will bring us home to Mom at Christmas.

The idea of an active filler strategy taking place for *wh*-words in particular was codified in Frazier 1987. Frazier proposed a parser that, upon encountering a *wh*-word, will actively look for a gap position to assign the filler to. Clifton and Frazier (1989) formulated the idea as the Active Filler Hypothesis given in (123).

(123) Active Filler Hypothesis:

When a filler of category XP has been identified in a non-argument position, such as COMP, rank the option of assigning its corresponding gap to the sentence over the option of identifying a lexical phrase of category XP. [pg. 292]

The dependencies under consideration here are, of course, not completely analogous to the phenomenon discussed by Frazier and Clifton, and cannot be directly subsumed under the Active Filler Hypothesis. Most notably, the phenomenon we are interested in here is one in which a gap is encountered *first*. The search that is initiated is for a filler, prompted by the gap itself. Nevertheless, I believe that hypotheses like the Active Filler Hypothesis are a natural starting place when considering the nature of the search strategy employed by cataphoric ellipsis (see Ng 2008 for a proposal of an active gap-filler strategy, suitably called an Active Gap Strategy, in Chinese). The current hypothesis, then, is a proposal that a search strategy can also occur in the gap-filler direction: upon encountering an ellipsis gap, a comprehender launches an active search for an antecedent to fill the interpretive gap. We might thus state this generalization as the Active Gap Hypothesis as in (124).

(124) Active Gap Hypothesis:

When an ellipsis gap of category α has been identified, rank the option of assigning its corresponding postcedent to the sentence over the option of identifying a non-postcedent phrase of category α .

This framing situates ellipsis resolution in a position more closely aligned with obligatory dependency resolutions such as *wh*-word processing, than with non-obligatory anaphoric dependencies such as pronominal anaphora resolution. In the Active Gap Strategy world, elliptical dependencies are not fundamentally different from other types of long-distance, or unbounded, dependencies in natural language. The Active Gap Strategy

hypothesis makes a prediction: if cataphoric ellipsis sites spur an active search process, then we predict that candidate postcedents will be assessed on their goodness-of-fit as a postcedent in real-time as the representations are encountered in parsing. Operationalizing this hypothesis, we predict that the plausibility of a candidate postcedent will modulate the processing time of that postcedent, such that an implausible postcedent will be read more slowly than a plausible postcedent.

The current study is inspired by the design of Yoshida et al. (2012), who found preliminary evidence that comprehenders actively searched for antecedent phrases in cataphoric sluicing. Yoshida et al. found a plausibility effect in cataphoric sluicing constructions located on the first clausal verb encountered after the *wh*-remnant. This plausibility effect was similar to that seen in (126), in which a dependency is formed between an argument noun and the verb that assigns the argument its thematic role. Yoshida et al. hypothesized that an active search process would manifest as a plausibility effect on the verb of a potential antecedent clause because the parser would actively build the clausal sluicing antecedent upon encountering preliminary evidence of the beginning of the clause. However, their results were ultimately not replicable.

For the current design I chose to investigate NPE instead of sluicing because any online processing effects of postcedent search would be highly local. An NPE site is relatively small and consists only of a missing noun; therefore, I hypothesize we are more likely to isolate any existing incremental processing effects with NPE than with sluicing, which is missing an entire TP. An additional concern with investigating sluicing is discussed by Harris (2015), who notes that the process of identifying a sluicing correlate in memory could be a distinct process from recovering the entire antecedent itself. The correlate is an optional word that corresponds to the remnant, which itself is an overt word that has moved out of the ellipsis site. The remnant-correlate relationship therefore serves as a cue to the interpretation of the ellipsis site. Harris makes the point, however,

that the search mechanisms responsible for identifying the correlate in memory are not a priori identical to the search mechanisms responsible for recovering the missing material in the ellipsis site itself. Thus, we should be cautious when generalizing from processing data on identifying sluicing correlates to ellipsis comprehension overall.

Given this discussion, where do we predict to find an online effect in NPE? I hypothesize that comprehenders engage in an active search strategy after encountering a cataphoric ellipsis site. That is, comprehenders will attempt to resolve the open dependency between the ellipsis site and its postcedent as soon as possible, and prior to unambiguous evidence that the piece of structure encountered is indeed a postcedent. Unlike in sluicing, I propose that the dependency formation in a cataphoric NPE construction holds directly between the N' gap of the ellipsis site and the noun that serves as a postcedent. It follows then that comprehenders will launch a search under which they will posit a postcedent for the ellipsis site at the first structurally licit N' location. My prediction therefore is that a plausibility effect will be observed on the first licit NP position following the ellipsis site.

- (125) a. While many [] are occasionally excited and energetic, most **students** are generally depressed.
- b. While many [] are occasionally excited and energetic, most **reports** are generally saying that students are depressed.

To illustrate the logic of the experimental design, consider example (125). Both (a) and (b) contain an elided NP in the subject position of the first clause. I have predicted that comprehenders will posit a postcedent at the first licit NP site following the ellipsis site. Assuming a linking hypothesis under which reading times positively correlate with processing difficulty, we predict that comprehenders will read more slowly if their prediction of encountering a postcedent is violated (Rayner et al., 2012). More specifically,

we predict reading times to vary with the plausibility of a predicted postcedent, such that an implausible postcedent candidate will be read more slowly than a plausible postcedent candidate.

The plausibility manipulation used here is methodologically akin to plausibility manipulations such as those investigating the active resolution strategy in argument-verb dependency formulation, as in Traxler and Pickering (1996). Phillips and Wagers (2007) point out that, although active dependency resolution processes are formulated as syntactic in nature, the literature uses semantic information in order to infer the syntactic processes taking place. For example, Phillips and Wagers provide the example in (126) of active resolution processes from Traxler and Pickering 1996.

- (126) (a) That's the pistol with which the heartless killer *shot* the hapless man yesterday afternoon __.
- (b) That's the garage with which the heartless killer *shot* the hapless man yesterday afternoon __.

Traxler and Pickering found that in topicalized sentences such as (126), eye-tracking measurements show a reading disruption at the verb *shot* when the verb is semantically anomalous with the fronted argument to which it would assign a thematic role, as in (b), compared to when the verb is not anomalous, as in (a). Comprehenders detect the anomalous verb at the point of the verb, suggesting that the dependency between the nouns *pistol* or *garage* and the verb *shot* was formed sometime before comprehenders encountered the verb itself. This paradigm also suggests that semantic information can be used to infer syntactic processes such as dependency formation.

Expanding upon this paradigm, we therefore expect that if comprehenders are actively positing a postcedent noun in the first structurally licit NP position in (125), they will posit the postcedent representation in the position directly following the quantifier in the

second clause, as this is the first structurally licit position. Note here that the quantifier sets up a clear structural position for the noun. If comprehenders posit the postcedent following the quantifier, we predict that they will read the NP *reports* in (b) more slowly than they will read *students* in (a), as the predicates *excited* and *energetic* more plausibly describe a student than a report.¹

4.1 Cataphoric ellipsis shows an online plausibility effect in the Maze task

The first experiment in this section probes the online processing of noun phrase ellipsis using the Maze task. As discussed in the introduction to this chapter, the experimental design relies on a plausibility manipulation on nouns that could serve as a possible postcedent for a cataphoric NPE site. My hypothesis is that, if the comprehension of ellipsis is performed online, meaning that people actively search for a piece of grammatical structure that could serve as the missing piece of the ellipsis site, then the plausibility of a candidate antecedent will modulate the reading time of that candidate antecedent.

4.1.1 Method

4.1.1.1 Methodology

This experiment used the *maze* task methodology (Freedman and Forster, 1985; Forster et al., 2009). The maze methodology is similar to its more widely-known kin, self-paced reading, in that it relies on reading times to measure the incremental processing of sentences. However, instead of each word in a sentence being presented in succession by button-presses, each word of the sentence is presented alongside a distractor word, or *foil*,

¹Note that this plausibility manipulation is not exactly like that tested by Yoshida et al. or Traxler and Pickering; please see the conclusion of this chapter for discussion.

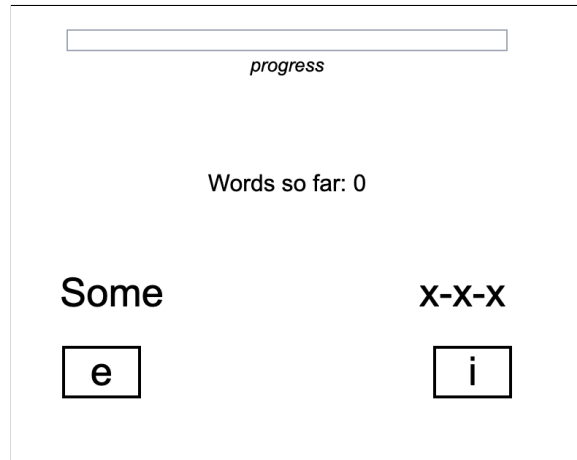


Figure 4.1: Example opening trial of Maze task.



Figure 4.2: Example mid-sentence trial of Maze task.

and the participant must choose which word will grammatically/felicitously continue the sentence. Participants are instructed to choose a word by pressing either the *e* key, for the left word choice, or the *i* key, for the right word choice. For the first trial of each sentence, the first word of the sentence is presented alongside the option ‘x-x-x’. Trials are randomized so that the correct word-choice alternates between sides, and the presentation is randomized across experimental presentations. Examples of an opening trial and a mid-sentence trial are given in Figure 4.1 and Figure 4.2, respectively.

One might wonder why the maze task is preferable over alternative, historically more prominent online methodologies. As discussed in detail by Forster et al. and Witzel et al. (2012), the maze task intends to remedy several well-known drawbacks to the self-paced reading (SPR) methodology. One concern with self-paced reading is that the nature of the task leads participants to select a constant rate of keypress. Another concern is that integration of a given word with the previous linguistic context often carries forward to a following word or words, leading to inconsistent hold-over or spill-over effects. Finally, self-paced reading relies on comprehension questions to ensure that participants are actually reading and processing the experimental sentences.

The maze task alleviates these concerns in the following ways. First, the reading times for the maze task are slower than those for self-paced reading. Therefore, participants are perhaps less likely to choose a constant rate of keypress, and therefore the assumption underlying both self-paced reading and maze – that the time taken to read a word reflects the time it takes to integrate the word into the previous linguistic context – is more likely to hold. Second, the maze task requires readers to commit to a structural parse at each trial; because readers must choose a grammatical continuation at each keypress, they are required to integrate each word into their existing structural representation of the sentence before making a choice to move on to the next trial. If this model of the task is correct, processing-time differences are highly localized, and there are no hold-over effects as in SPR. These localized effects therefore give maze data an interpretive clarity that is lacking in self-paced reading data.

Last, and related to the previous point, concerns that participants are not sufficiently attending to the task or fully comprehending the sentences they are reading are obviated because the nature of the task requires comprehension at each word, or trial, within the sentence. Forster et al. report that participants who prematurely make a choice that nonetheless ends up being correct feel unable to make an informed decision on the next

trial and are forced to guess. The task therefore forces comprehension at each trial of a sentence, and participants are unlikely to guess or luck their way through many trials (give there is a 50/50 chance at each trial).

One criterion I have not discussed so far is the naturalness of the experimental task. It is reasonably argued that the more natural the experimental task, the better, as natural tasks do not interfere – or interfere less – with normal comprehension processes (Mitchell, 2004). As pointed out in Forster et al. 2009 and Witzel et al. 2012, the maze task is highly artificial. The use of the maze task is therefore a methodological choice to prioritize the benefits of the maze task as outlined here over its artificiality.

4.1.1.2 Participants

Seventy-nine workers from Amazon’s Mechanical Turk participated in the experiment. Seven workers were excluded because they were unable to successfully complete any sentence. The data from the remaining seventy-two workers was included in the experimental analysis. Participants were paid \$5, and were awarded a \$2 bonus if they successfully completed 80% of the experimental sentences.

4.1.1.3 Materials

Materials consisted of 48 critical items comprising four conditions. The design was a 2 x 2 with the factors *ellipsis* (Ellipsis, No Ellipsis) and *plausibility* (Plausible, Implausible). All critical items consisted of two clauses. The two Ellipsis conditions contained cataphoric NPE in the first clause. Each ellipsis site was preceded by one of the following quantifiers: *while*, *most*, *many*, *(a) few*, *almost all*, *(nearly) all*, *countless*, or *much*. Discourse subordinating contexts were created for each initial clause in order to license the cataphoric ellipsis site; the connectives used were *while*, *although*, *even though*, *despite*, and *because*. The two No Ellipsis conditions contained a noun phrase – instead of an

ellipsis site – following the quantifier in the initial clause. To keep items as consistent as possible, Ellipsis sentences in each item contained 7 words in their first clause, and No Ellipsis sentences contained 8 words in their first clause.

The second clause in all critical items followed the pattern *quantifier + NP + to be + adverb*. All items and conditions then contained at least one additional spill-over region. In the Plausible conditions, the subject NP of the second clause was a plausible postcedent for the cataphoric ellipsis in the first clause. Implausible conditions were created by inserting an implausible postcedent in the subject position, followed by an embedding verb that embedded the noun and predicate clause from the Plausible conditions. Note that this pattern satisfied two conditions: the critical region of the second clause within each item differed only on the plausible/implausible NP, and all experimental items consisted of grammatical, felicitous sentences. The animacy of the critical NP in the second clause was balanced such that half the experimental items contained an animate NP and half contained an inanimate NP. Critical NPs were chosen so that a range of frequencies and word lengths were represented. An example item is given in Table 4.1.

Ellipsis	Plausibility	Sentence
No Ellipsis	Plausible	While many pre-meds are occasionally excited and energetic, most undergraduates are generally depressed.
No Ellipsis	Implausible	While many pre-meds are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.
Ellipsis	Plausible	While many are occasionally excited and energetic, most undergraduates are generally depressed.
Ellipsis	Implausible	While many are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.

Table 4.1: Example item from Experiment 3.

Distractor words, or foils, for all items were constructed using Boyce et al. (2020)'s

automatic foil generator. This program is directly implementable in IbeXFarm, and includes a Maze controller that is built off the Dashed Sentence controller used for self-paced reading tasks. Boyce et al.'s generator requires all conditions within an item to be the same length; therefore, for the critical items the Ellipsis-Implausible sentence was used to generate foils. The foils generated for this condition were then altered by hand for the additional three sentences/conditions within each experimental item. Foils were also individually checked and altered when the grammatical/felicitous continuation was insufficiently clear. All foils were altered such that the foils for the critical regions were identical across conditions within each item. An example item with foils for each condition is given in (127) - (130).

(127) **No Ellipsis, Local/Plausible**

Item:

While many pre-meds are occasionally excited and energetic, most undergraduates are generally depressed.

Foils:

x-x-x hear mixture thou consequences belongs go textbooks, ends ago net extension assertion.

(128) **Ellipsis, Local/Plausible**

Item:

While many are occasionally excited and energetic, most undergraduates are generally depressed.

Foils:

x-x-x hear thou consequences belongs go textbooks, ends ago net extension assertion.

(129) **No Ellipsis, Distant/Implausible**

Item:

While many pre-meds are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.

Foils:

x-x-x hear mixture thou consequences belongs go textbooks, ends ago net extension effort ring sentimental sin assertion.

(130) **Ellipsis, Distant/Implausible**

Item:

While many are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.

Foils:

x-x-x hear thou consequences belongs go textbooks, ends ago net extension effort ring sentimental sin assertion.

Each item also included a forced-choice question about the content of the item's sentence. Questions were balanced across critical items such that half asked about a property of the first clause of the sentence, and half asked about a property of the second clause.

4.1.2 Procedure

The experiment was administered on Amazon's TurkPrime and used the IbexFarm experimental software (Drummond, 2014). Informal piloting of the Maze task revealed that the task took significantly longer than an equivalent experiment using self-paced reading methodology. Because of concerns of length, and the accompanying concerns of fatigue and attention-deficiencies, the 48 critical items were split into two experimental lists of 24 items each. Each list was then run as a concurrent experiment and the results were combined upon their completion.

Within each experimental list, the critical items were distributed into four Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, ensuring balanced responses across lists. Participants were trained on six practice items and were given two cartoon breaks half and 2/3 of the way through the experiment. Participants could complete each Maze trial at their own pace. If they successfully completed a sentence, they received a screen informing them they were correct and prompting them to press any key to continue to the forced-choice question.

Any keypress would then bring up the question screen, which would time out after no activity for eight seconds. If a participant got a Maze trial wrong – if they chose a foil instead of the grammatical/felicitous continuation – the trial was aborted. After an incorrect choice, participants were given a screen telling them that their choice was incorrect and prompting them to press any key to continue to the question. Note that participants were required to answer the questions for each trial, regardless of whether they successfully completed the Maze task for that trial.

In addition to the 48 critical items, 24 items from an additional experiment were also included, as well as 24 additional filler items. Filler items included sentences with various types of elliptical and pronominal anaphoric/cataphoric dependencies, including verb phrase ellipsis, sluicing, and pronominal anaphora. All critical items and fillers were identical between the two experimental lists. At the completion of the experiment a debriefing form was given in which participants were asked to record any difficulties they had encountered.

4.1.3 Results

Maze trials were filtered so that only correct trials were included in the analysis. Following the suggestion in Forster et al. 2009, the first trial of each sentence was removed before any analysis; these trials were uninformative because participants saw the first word of the sentence and the non-word foil ‘x-x-x’. Reading times for the first trial also suggested that participants often used this first word as a break; RTs were therefore not representative of the RTs overall. After removing the first trials, the longest .3% of remaining RTs was removed, as well as any RT below 200ms. Eight regions were labelled; within each region, means were calculated for each condition over items, and then the mean over items within each condition. Means for each region are given in Figure 4.3.

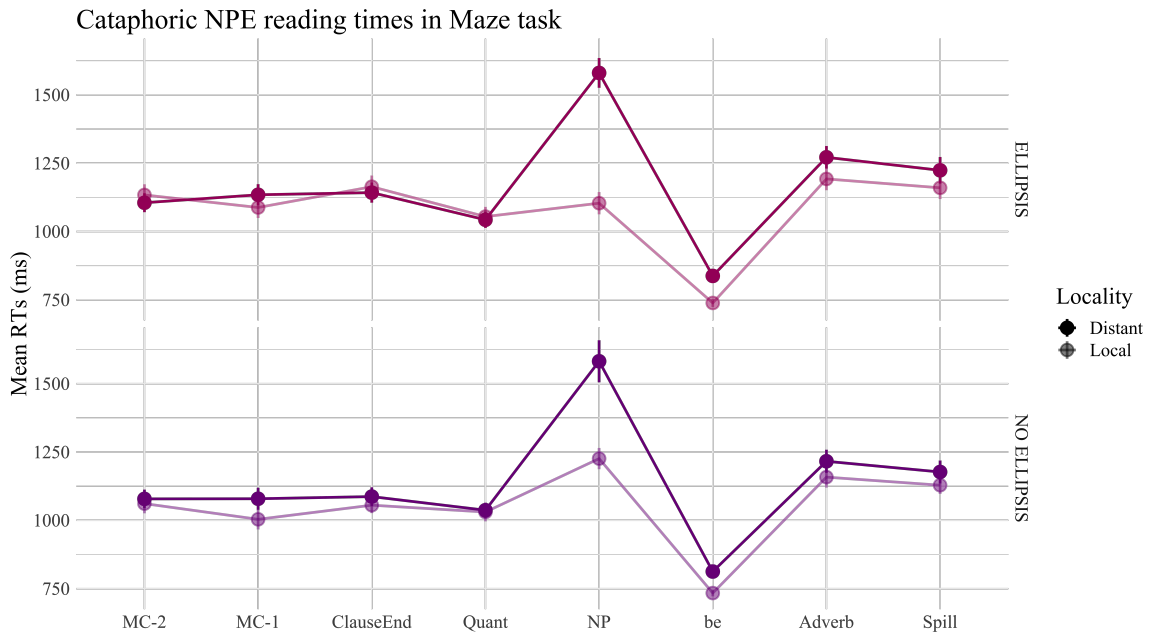


Figure 4.3: Maze task reading times by region. Error bars indicate standard error of the mean over items. N items = 48, N participants = 72

The critical region, NP, was analyzed using a Bayesian mixed-effects linear regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. An exgaussian (exponentially modified Gaussian) family was used, following the recommendation given for response times in Bürkner 2020. Exgaussian models are used to fit responses that are believed to be caused by two independent processes: a Gaussian distribution, and a decaying exponential. The models are therefore well-suited to fit the skewed distribution of response times. The exgaussian is a three parameter model characterized by μ , the posterior mean of the Gaussian, σ , the standard deviation of the Gaussian (models the variation around μ), and λ , the decay rate of the exponential, which models the tail of the long response times (note the exponential decay is sometimes modelled as $\tau = 1 / \lambda$). In brms, β is the inverse rate of the exponential component, σ is the standard deviation of the Gaussian component, and ξ is the mean of

the Gaussian component. In the model $\mu = \xi + \beta$ so that the main predictor term equals the mean of the distribution.

Because an exgaussian family was used, the model was estimated using raw RTs (as opposed to log-transformed RTs), which were trimmed as outlined in the previous paragraph.² The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions.

Table 4.2 provides the posterior estimates of the fixed effects as well as 95% credible intervals. I predicted that the critical region would show a slow-down for the Distant (implausible) condition compared to the Local (plausible condition). This prediction was borne out, with reading times showing a significant slow-down on the NP region in the Distant condition, but not in the Local condition. This finding is reflected in the main effect of **Locality** shown in Table 4.2. I also predicted that the RT slow-down would be significantly greater in the Ellipsis condition than in the No Ellipsis condition. This effect did not reach significance, as shown by the critical intervals in Table 4.2 spanning zero for the interaction **Locality:Ellipsis**.

	estimate	est.error	2.5%	97.5%
Intercept	1525.92	46.37	1434.66	1617.40
Locality	218.45	36.12	149.53	291.63
Ellipsis	-41.40	23.75	-89.79	4.20
Locality:Ellipsis	47.21	44.26	-41.41	133.34
Frequency	-73.36	12.09	-97.11	-49.85
sigma	79.66	15.73	45.81	108.24
beta	424.47	16.50	392.76	457.29

Table 4.2: Experiment 3 NPE Maze: Bayesian model estimates for Region *NP*.

²Log transformed RTs were also modeled using a gaussian family model; equivalent population-level effects were found.

Recall each item had the identical word for the Quant, be, and Adverb regions. The nature of the design, however, required the noun to differ in the critical NP region. To account for possible RT differences due to word length or frequency, both were estimated and included in the model. Word length did not reach significance in any model, and therefore was not included in the final model reported here. Word frequency was estimated using the SUBTLEXus corpus.³ Frequency is reported on the Zipf scale (van Heuven et al., 2014), which is simply the the $\log_{10}(\text{frequency per billion words})$ of the word in the SUBTLEXus corpus. Frequency was a significant main effect in each model, and was therefore included in the final model provided here.

As mentioned, the experiment was presented in two lists, with 36 participants each. To test if any significant difference was present between the two lists, a two-level factor List, with levels List1 and List2, was created and was included in the brm model as a fixed effect. The effect failed to reach significance, and therefore was not included in the final model reported above.

To check model fit, I ran a posterior predictive check. The posterior predictive check plots the observed data of the experiment against simulated data generated from the *posterior predictive distribution*, that is, the distribution of the outcome variable provided by the model after using our observed, or experimental, data to update our beliefs about the unknown model parameter θ (Gabry et al., 2019). If the model is a good fit to the observed data, then the simulated data from the posterior distribution, y_{rep} , should closely match the data we observed, y , which plots the vector of dependent variable observations. As stated in Gelman et al. 2013, “If the model fits, then replicated data generated under the model should look similar to observed data. To put it another way, the observed data should look plausible under the posterior predictive distribution.” pg. 143.⁴ The graphical

³<https://www.ugent.be/pp/experimentele-psychologie/en/research/documents/subtlexus>.

⁴While posterior predictive p values for Bayesian models do exist, they are not standardly used. See

posterior predictive check revealed good model fit, as shown in Figure 4.4.




Figure 4.4: NPE experiment 3 posterior predictive check generated by Bayesian model.

An additional analysis was run on region *be*, to test for any existing spill-over effects from the critical NP region. The analysis was identical to that of the NP region. Note that because all words in this region contained either *are*, *is*, *were*, or *will*, frequency was not included in the model.

Table 4.3 provides the posterior estimates of the fixed effects as well as 95% credible intervals for the region *be*. Although the reading times for the region drop off substantially from the reading times at the critical region,⁵ the Distant conditions were read significantly slower than the Local Conditions. As in the NP

region, these findings are reflected in a significant effect of Locality. No other effects reached significance. These results suggest a spill-over effect from the slow-down observed in the critical region.

discussion at e.g. p-values blog and p-values STAN.

⁵Note that the NP region contains nouns, while the ‘be’ region contains variations on the verb *to be*. The RT difference between the two regions is therefore expected.

	estimate	est.error	2.5%	97.5%
Intercept	770.86	12.44	746.75	795.34
Locality	29.12	7.24	15.10	43.28
Ellipsis	11.28	6.96	-2.50	24.71
Locality:Ellipsis	22.09	14.77	-7.24	50.98
sigma	46.17	4.78	36.63	55.45
beta	172.07	6.60	159.39	185.38

Table 4.3: Experiment 3 NPE Maze: Bayesian model estimates for Region *be*.

4.1.4 Discussion

4.1.4.1 Implausible nouns were read more slowly than plausible nouns

This experiment sought evidence of whether ellipsis resolution is more similar to syntactic dependencies that are actively resolved by the parser, such as those created in *wh*-questions, or whether ellipsis resolution is more similar to pronoun resolution, which is often resolved using general reasoning strategies after sentence parsing is complete. Recall our main hypothesis: if the cataphoric ellipsis site catalyzed an active search process for a postcedent, then the plausibility of a candidate postcedent would modulate the reading time on the postcedent such that an implausible postcedent would be read more slowly than a plausible postcedent. The results support this hypothesis, showing that participants read the critical noun at a significantly slower rate when it was an implausible postcedent (the Distant condition) than when it was a plausible postcedent (the Local condition). The findings therefore overall support a view under which comprehenders were engaging in an active search process to resolve the cataphoric ellipsis site.

An unexpected finding was that participants were also significantly slower to read the implausible NP than the plausible NP in the No Ellipsis condition. This appears at first blush to present a concern for my conclusion that the RT slow-down in the Ellipsis

condition was driven by the (im)plausibility of the NP as a postcedent. There was no ellipsis in the No Ellipsis conditions, and therefore there was no possibility that the critical NP was being considered as a candidate postcedent. The means in Figure 4.3 show that numerically the difference between the Distant and Local conditions in the No Ellipsis condition was smaller than the difference between those conditions in the Ellipsis condition. However, the difference did not reach significance, as demonstrated in the lack of a significant interaction of **Locality:Ellipsis** in Table 4.2.

I believe however that the RT slow-downs in the No Ellipsis condition do not invalidate the hypothesis I have proposed. It is known that reaction times in general are sensitive to the predictability of words in behavioral tasks (Fischler and Bloom, 1979; Forster, 1981; Schwanenflugel and Schoben, 1985; Schwanenflugel and LaCount, 1988; Arnon and Snider, 2010; Kuperberg and Jaeger, 2016). Furthermore, Brasoveanu and Dotlacil (2020) show evidence for a *semantic*, predictive parser that operates incrementally during real-time parsing. If the critical nouns violate the predictions of this parser, then the observed RT slow-down is expected. However, the question of interest then is *why* these particular nouns were unpredictable.

The main effect of **Locality** tells us that participants were sensitive to the plausibility manipulation both when the sentence had ellipsis and when it did not. Why then would comprehenders be sensitive to the plausibility of the critical noun even when it was not under consideration as a postcedent? The most plausible reason is that the comprehender has built up a structural and semantic representation based on the linguistic and non-linguist input up to the point of integrating the critical noun such that the critical noun was unexpected. Specifically, I believe that the critical noun was placed in a comparative relation with the corresponding noun in the first clause. Consider the examples in (131) and (132), which have the critical nouns in bold.

(131) Distant/Implausible Condition

While most **reptiles** appear frightening and potentially dangerous, almost all **toddlers** are actually positive that snakes are harmless to people.

(132) Local/Plausible Condition

While most **reptiles** appear frightening and potentially dangerous, almost all **snakes** are actually harmless to people.

The information structure of the items encourages a comparative relation between the noun in the first clause and the critical noun. First, the initial clauses in all items are discourse subordinated to the second, and begin with a conjunct such as *while*, *although*, or *despite*. These connectives engender the particular discourse relations Violated Expectation, Exception, and Denial of Preventer (Kehler, 2002). There is evidence that language comprehenders use inferences of causal and discourse relations during early stages of semantic processing and to help facilitate the processing of new words (Kuperberg et al., 2011; Rohde and Horton, 2014). Because the nouns are each the subject of their respective clauses, it is very plausible that comprehenders held expectations that these nouns will instantiate these discourse relations, as they in fact do. Additionally, the use of the two quantifiers further suggests a contrastive relationship between the two nouns, as quantifiers generally denote subsets of some salient set of objects. It is therefore extremely plausible that participants parsing the critical sentences up to the second quantifier had an expectation of seeing a noun in a contrastive relation that was contained in some salient subset with the noun in the first clause. This expectation was largely realized in the Local conditions. Thirty-one of the items had a hypernym relationship between the first NP and the Local noun, 12 items had the two nouns as members of a single kind class, and three items had the two nouns as synonyms of one another. However, this expectation was violated in the Distant condition. As a result of the plausibility design, the Distant nouns had no salient relationship with

the noun in the first clause.

To take a concrete example, the discourse relation in (131) is plausibly Violated Expectation (Asher, 1993; Kehler, 2002; Asher and Lascarides, 2003). The subject of the first clause, *most reptiles* evokes a salient set of reptiles. Cognitively, this representation activates representations of concepts closely related to reptiles (Anderson, 2007). The quantifier in the second clause further sets up the expectation of a contrastive relation between the set of reptiles and the set of entities denoted by the upcoming noun. If a comprehender then reads *snake*, this representation is fairly easy to integrate into the ongoing discourse representation because it is conceptually a neighbor to *reptiles*. However, if the comprehender reads *toddlers*, this representation is more challenging to integrate into their current discourse representation because it is cognitively distant to *reptiles*.⁶ This discrepancy in integration difficulty was reflected in reading time, just as the plausibility of the postcedent was.

While I do not have access to individual participants' cognitive representations of the nouns used in this study, a standard if imperfect measure can be obtained using semantic relatedness scores. I measure semantic relatedness here using GloVe pre-trained word vector representations (Pennington et al., 2014). The vectors contained 840B tokens, 2.2M vocab, and consisted of 300 dimensions. Semantic similarity was computed using the Cosine Similarity measure, which measures the similarity between two vectors of n dimension by finding the cosine of the angle between the vectors (see <https://github.com/margaretkroll/cosineSimilarity>). Measures are given in Appendix E.1. Cosine measures range from -1, reflecting opposite semantic similarity, 0, reflecting no semantic relationship, and 1, reflecting high semantic similarity.⁷ Table 4.4 shows summary statistics for the cosine measures of the items in the Local and Distant condi-

⁶For most of us.

⁷Note that GloVe is not a co-occurrence matrix; therefore negative values are possible

tions. The measures confirm that the nouns in the Local condition were more closely semantically related to each other than the nouns in the Distant condition were.

Condition	Cosine mean	Cosine s.d.	Cosine median
No Ellipsis, Distant	0.236	0.172	0.213
No Ellipsis, Local	0.553	0.170	0.570

Table 4.4: Experiment 3 NPE Maze: Cosine semantic similarity measures, No Ellipsis conditions.

To probe the influence of the semantic similarity of the nouns on reading times, I ran a Bayesian linear regression predicting the RTs of the No Ellipsis condition from Locality, word frequency (Zipf value), and semantic similarity (cosine). Word frequency failed to reach significance, and therefore was not included in the final model. Table 4.5 provides the posterior estimates of the fixed effects of Locality and Cosine as well as 95% credible intervals. The model used the *shifted_lognormal* family, as recommended for response times in Bürkner 2020. The *shifted_lognormal* family is a three parameter model. μ is the mean of the log-normal distribution; the mean of μ represents the median RT. σ is the standard deviation log-normal distribution. The *shift* parameter is the time of the earliest possible response; if *shift* = 0, then the model is a log-normal distribution with two parameters. In brms, the *shift* parameter is given by a positive parameter called *ndt*.

The model included random intercepts for Item and Participant, and random slopes for Item and Participant for each main effect; interactions between fixed effects were not included.

I predicted that semantic similarity between the two nouns would impact the reading time of the second noun such that the greater the semantic distance between the two nouns, the greater the reading time on the second noun. This prediction was borne out. Table 4.5 shows a significant main effect of Cosine, reflected in the 95% credible

	estimate	est.error	2.5%	97.5%
Intercept	7.06	0.09	6.89	7.24
Locality	-0.04	0.09	-0.21	0.13
Cosine	-1.00	0.20	-1.40	-0.62
sigma	0.47	0.02	0.42	0.51
ndt	437.71	25.57	381.15	480.96

Table 4.5: Experiment 3 NPE Maze: Bayesian model estimates for critical region, No Ellipsis conditions, including semantic similarity measure.

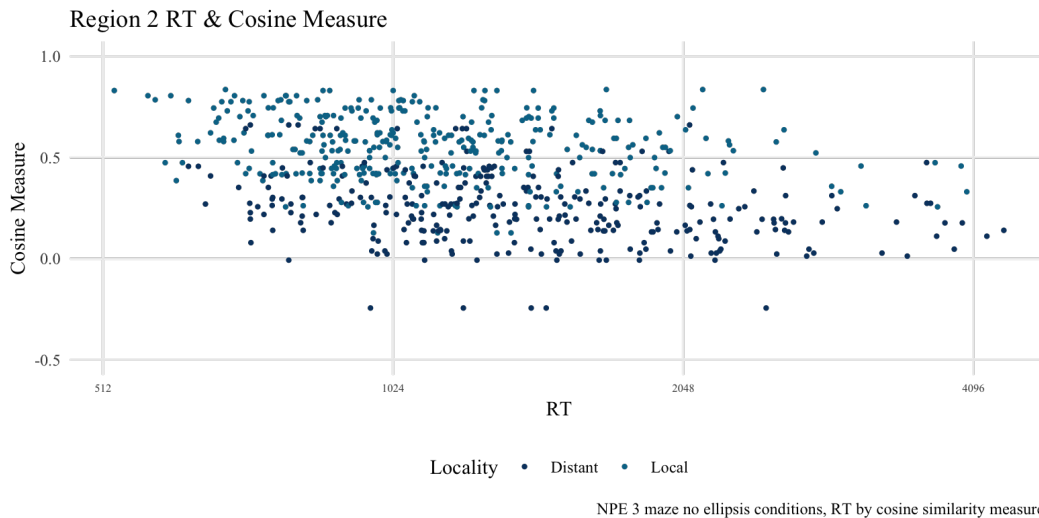


Figure 4.5: Experiment 3 NPE Maze: Critical region reading times and Cosine semantic similarity measures. y-axis is shown in log2 scale.

interval of -1.40 and -0.62. The estimate of -1.00 reflects a greater reading time on nouns that are less semantically related. The model also shows that Locality did not reach significance. The failure of Locality to reach significance in this model tells us that the Cosine measures are accounting for the model variance. Figure 4.5 plots reading times against Cosine measurements, visualizing the relationships given in Tables 4.4 and 4.5.

These findings support the hypothesis that the slow-down in the No Ellipsis condition is driven by the semantic relatedness between the antecedent noun and the critical region

noun. Specifically, I propose that the more distant the semantic relationship between the two nouns, the greater the challenge of integrating the critical noun into the existing discourse context-framework at the point of encountering the noun. This explanation for the RT slow-down is analogous to the argument made for the slow-down in the Ellipsis condition. However, the discourse context in the Ellipsis condition does not have a noun in the first clause to evoke a salient set of entities; therefore, the slow-down in the Ellipsis condition cannot be caused by the difficulty of integrating the noun into the discourse context based on a salient set of nouns. Instead, I have proposed that the integration difficulty is generated by the implausibility of the noun as a postcedent; that is, the difficulty of integrating that noun into the gap in the discourse context created by the cataphoric ellipsis site. Section 4.3 follows up on this hypothesis by altering the properties of the Distant nouns. The experiment shows that the semantic relatedness of the critical nouns in the Distant condition to their corresponding nouns in the first clause modulates the reading time on the critical noun, as hypothesized in this discussion.

4.1.4.2 Indexing processing difficulty in the maze task

Because the maze task is a relatively novel methodology, it is worth digging a bit deeper into the participants' performance on the task. Figure 4.6 shows the total number of trials missed on the critical items broken down by region. Regions 0-11 are plotted because not all conditions in every item extended past region 11. The plot shows that a positive but slowly decreasing slope until region 11. The positive slope of the failure rate is expected given that the failure rates are cumulative. Recall that this is because a failure on one region necessitated failing the remaining regions in that sentence. The general flattening of the slope after region 3 indicates that participants are in general less likely to fail on later regions than they are on the earlier regions. Indeed, the greatest jumps are seen in between region 0-1, region 1-2, and region 2-3. This means that participants are most

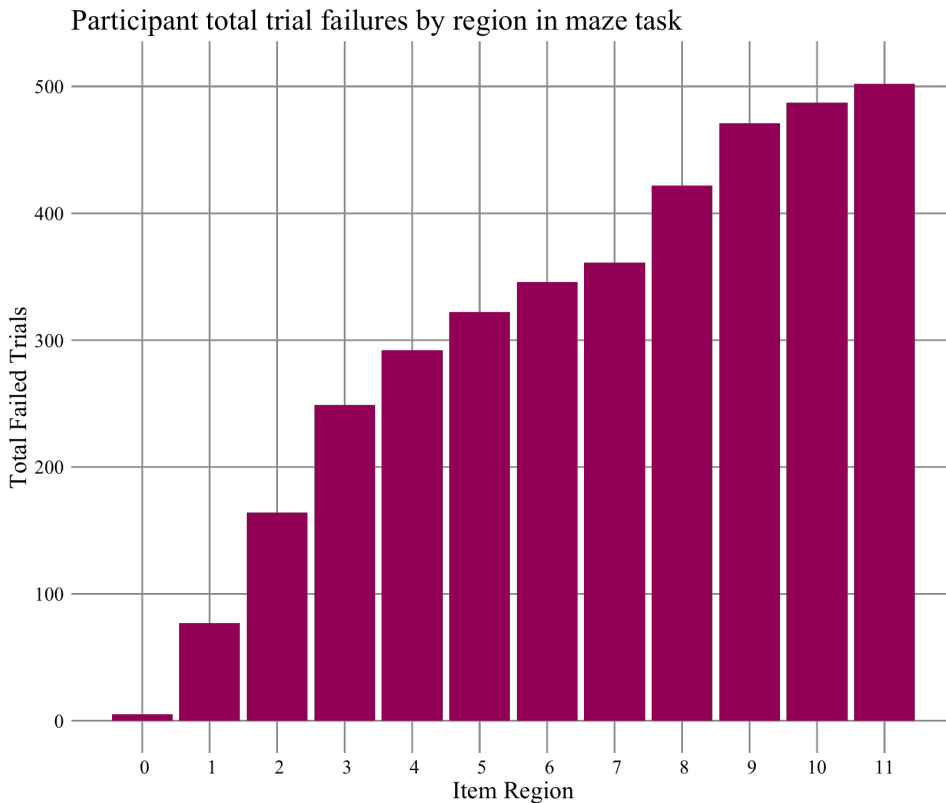


Figure 4.6: Total trial fails by region in NPE maze task. N = 3,698

likely to fail on the second, third, or fourth trial of the sentences overall (see Boyce et al. 2020 and Sloggett et al. 2020 for similar results). Additional increases in failure rate are seen in region 8 and region 9, which correspond to the critical noun in the Ellipsis and No Ellipsis conditions, respectively. This suggests that integrating the critical noun resulted in an increased trial failure rate.

Figure 4.7 shows a cumulative density function plot of region 1-11 of the critical items, by conditions. The plot includes only those regions on which a foil was chosen instead of the correct word. The conditions largely track together through region 2. The plot shows that for each condition, the probability of the failure occurring at or before region 2 is about 30-35%. Recall that this is not the overall chance of failure, it is the




Figure 4.7: Cumulative distribution function, trial fails by region for all conditions in NPE maze task. Counts indicates the region where trial failure occurred.

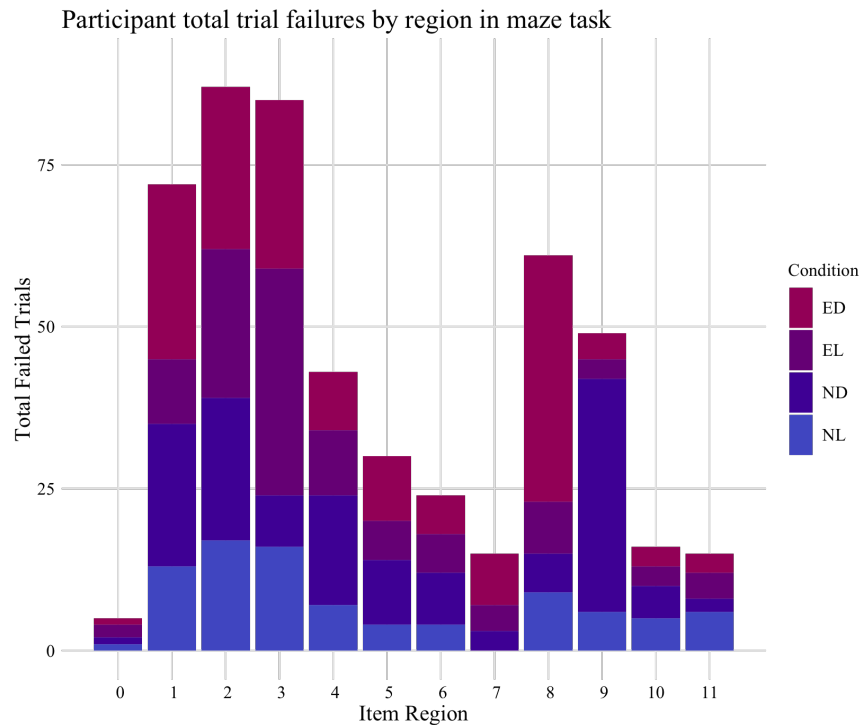


Figure 4.8: Trial fails count by region for all conditions in NPE maze task. Counts indicates the region where trial failure occurred. N = 502.

chance of failing at those particular regions out of overall failures (that is, if you fail, how likely you are to have failed at or before region 2). The rate of failure on the first trial was almost 0%, as expected given the procedural set-up of the experiment. Setting region 0 aside, since it did not involve a true word discrimination task, the rate of failure was substantially higher in the early trials than in the later trials. About 30-35% of the total failures occurred within region 1 and 2. The early failure rate seen here tracks the high cumulative failure rate we saw in Figure 4.6. This pattern is also seen in Figure 4.8, which plots the total count of trials failed for each region.

Region 3 shows a substantial increase in the failure rate in all conditions except for the No Ellipsis, Distant condition. The Ellipsis conditions separate from the No

Ellipsis conditions, showing a generally higher failure rate. This may be attributable to the difficulty in resolving the cataphoric ellipsis at this region; for many items, this region required recognizing that a noun was missing and that the correct continuation was a verb. However, the items were split in whether this resolution was required at region 2 or region 3, suggesting that this cannot be the full explanation. There are only 85 failures in this region total, so I am hesitant to draw any theoretical conclusions from the data. It is possible that idiosyncratic foil choices account for some of the failure rates here, although no obvious patterns are observable in the data.

The failure rate tapers off after region 3 until a large spike in the Ellipsis/Distant condition on region 8 and in the No Ellipsis/Distant condition on region 9. These regions are those in which participants encountered the critical noun in each condition (recall that the regions are offset by one because of the NPE that occurred in the Ellipsis condition and not in the No Ellipsis condition). No comparable increase in failure rates is seen in the Local conditions. The increase in the Distant conditions but lack of similar increase in the Local conditions suggests that integrating the implausible noun lead to a greater failure rate than integrating the plausible noun. This discrepancy in the failure rates between the two conditions further suggests that processing difficulty is indexed not only by an increase in reading time but also by an increase in failure rates on the challenging region.⁸

4.2 Self-paced reading does not show a plausibility effect

The previous section discussed several differences between the Maze methodology and more widely-used experimental measures of incremental processing, such as self-paced

⁸Note that the trials given here were not included in the RT analysis; only correct trials were included in the analysis.

reading. I also discussed in that section some arguments that have been presented in favor of and against using one methodology or the other. These arguments raise questions about which kinds of processing effects can be captured in which methodologies, and whether the choice of methodology can affect whether or not a particular processing property manifests in experimental results. To expand upon this discussion, this section seeks to replicate the results from §4.1 using the self-paced reading methodology.

Because the plausibility effect found in the previous section is one that is not reliably established in the literature, this section also serves as an attempt to replicate those results. I show, however, that the results from the Maze task are not replicated in the self-paced reading task, despite the use of identical items and fillers. I discuss in this section three possible explanations for the failure of replication: (i) experimental procedural differences, (ii) methodological differences, and (iii) Type I error. I argue that the lack of replication is due to general methodological differences between Maze and SPR tasks, and discuss the implications of this finding for studies related to the current design.

4.2.1 Method

4.2.1.1 Methodology

This experiment used the self-paced reading methodology. Words were presented centered on the computer screen, and participants could move through the sentence one word at a time by pressing the space bar at their own pace. Note that this was not a moving window task, so participants did not have visual cues to the length of the sentence. This choice was made to mask the local/distant manipulation from participants.

4.2.1.2 Materials

Materials were identical to those used in Experiment 3, as described in §4.1. An example item is repeated in Table 4.6.

Ellipsis	Plausibility	Sentence
No Ellipsis	Plausible	While many pre-meds are occasionally excited and energetic, most undergraduates are generally depressed.
No Ellipsis	Implausible	While many pre-meds are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.
Ellipsis	Plausible	While many are occasionally excited and energetic, most undergraduates are generally depressed.
Ellipsis	Implausible	While many are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.

Table 4.6: Example item from self-paced reading plausibility experiment.

4.2.1.3 Participants

Fifty-six undergraduate students from the University of California, Santa Cruz participated in the experiment. All students self-identified as native speakers of English. All students received course credit for their participation.

4.2.2 Procedure

The experiment used the IbeXFarm experimental software (Drummond, 2014). As in Experiment 3, the 48 critical items were split into two experimental lists of 24 items each. Each list was then run as a concurrent experiment and the results were combined upon their completion. Within each experimental list, the critical items were distributed into four Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, in an effort to ensure balanced responses across lists. Participants

were trained on six practice items and were given two cartoon breaks half and 2/3 of the way through the experiment. Participants could complete each trial at their own pace.

Fillers were identical to those in Experiment 3. At the completion of the experiment both an online debriefing form and a paper debriefing form was given to participants in which they were asked to record any difficulties they had encountered.

Fifty-three participants participated in the experiment in the psycholinguistics lab in the Linguistics department at UCSC. Participants were seated at a desktop computer in the lab and completed the experiment at their own pace. Due to the unexpected closure of campus due to COVID-19, three additional participants took the experiment online from a location of their choosing. These participants were asked to take the experiment in a quiet place with few distractions.

4.2.3 Results

The comprehension questions of the experimental items were used as an attention check; data from participants who scored less than 70% on these questions were removed. This criterion removed five participants from the analysis. The remaining data were trimmed in the following manner. The first region from each item was removed, then the longest .3% of remaining RTs was removed, as well as anything below 100ms. The data were separated into eight regions, resulting in the same region labelling as in §4.1. Mean RTs computed over items and then within each region are given in Figure 4.9.

The critical region, *NP*, was analyzed using a Bayesian mixed-effects linear regression. The analysis was implemented in the *brms* package for R (Bürkner, 2017) and was fitted with default priors. The model failed to converge without divergences using the *exgaussian* family; therefore, the *shifted_lognormal* family was used, following the recommendation given for response times in Bürkner 2020. The model also included the full fixed effect structure, including all interactions. Table 4.7 provides the posterior




Figure 4.9: Self-paced reading task reading times by region. Error bars indicate standard error of the mean over items. N items = 48, N participants = 56

estimates of the fixed effects as well as 95% credible intervals.

I predicted that the critical region would show a slow-down for the Distant (implausible) condition compared to the Local (plausible condition). This prediction was not borne out, as shown by the lack of a significant effect of **Locality** in Table 4.7. The model also shows a failure to find a significant main effect of **Ellipsis**, or a significant interaction of **Locality:Ellipsis**.

Because self-paced reading often shows delayed ‘spill-over’ or ‘hold-over’ effects, an equivalent analysis was run on the region immediately following the critical region,

	estimate	est.error	2.5%	97.5%
Intercept	5.69	0.06	5.57	5.82
Locality	-0.02	0.04	-0.09	0.05
Ellipsis	0.04	0.03	-0.01	0.10
Locality:Ellipsis	0.00	0.06	-0.12	0.12

Table 4.7: Bayesian model estimates for Region *NP*, self-paced reading.

which for all items contained a tensed version of the verb *to be*. The model estimates are provided in Table 4.8. As in the critical region, no main effects or interactions were found to be significant.

	estimate	est.error	2.5%	97.5%
Intercept	5.82	0.05	5.72	5.92
Locality	-0.03	0.02	-0.07	0.02
Ellipsis	0.02	0.02	-0.03	0.06
Locality:Ellipsis	0.03	0.04	-0.05	0.10

Table 4.8: Bayesian model estimates for Region *be*, self-paced reading.

4.2.4 Discussion

The experiment in this section failed to replicate the effects found in §4.1. There are three possible explanations for the null effect result found here. The first explanation is that the significant results in §4.1 were a Type I error, in which case the null hypothesis in that experiment was incorrectly rejected. This explanation appears unlikely, as the experiment in §4.3 replicates the main findings of §4.1 with slightly altered items and different participants. The second possible explanation for the difference in findings is procedural differences between the two experiments. The experiment in §4.1 was run on an online crowdsourcing platform, while the current experiment was run (mostly)

in-lab with undergraduate participants. It is possible that attentional differences exist across these two sub-populations that account for the differences. A direct comparison could satisfactorily respond to this possibility, for example, re-running the Maze task with the same population in a laboratory setting.⁹ However, undergraduate populations are regularly used in linguistic experiments, and I have no a priori reason for believing these sub-population differences should result in divergent findings.

I believe that the most plausible explanation for the lack of a significant result in the current experiment is methodological in nature. As discussed in §4.1, the Maze is argued to have several advantages over SPR, including forcing incremental interpretation and obviating concerns of "strategic" uniform button-pressing. As the Maze is still a relatively new methodology, there are few comparisons of the findings of each methodology for particular processing procedures (Witzel et al. 2012 being a notable exception). The findings here suggest that the processing mechanisms involved in integrating a discourse entity into the preceding text are sensitive to the differences between the nature of the tasks in SPR and the Maze.

I hypothesize, though cannot conclude with certainty, that the explanation for this stems from the incremental interpretation forced by the Maze task. This could be borne out in several ways that are not mutually exclusive. One possibility is that the Maze encourages participants to make stronger predictions about what they are going to read next; these predictions are then sensitive to the plausibility manipulation of the current experiment. The second is that the Maze simply forces deeper incremental comprehension of the sentence, such that participants are more actively searching for a postcedent in the Maze task, but reading more passively in the SPR task. The third possibility is that the effect in the current experiment is too small to overcome the uniform button-pressing strategy available in SPR, and is therefore obscured in the SPR task. Because the Maze

⁹This type of comparison is outside the scope of the current project.

prevents such a button-pushing strategy, the effect is observed in this task.

The next section presents the results of a follow-up experiment. The experiment slightly alters the items of the experiments in the current section and in §4.1, and uses the Maze task to replicate the main finding of §4.1.

4.3 Semantic relatedness drives reading times: Revisiting the Maze

This chapter is so far in conflict. Section 4.1 found evidence that the plausibility of an available postcedent effects slow-downs in reading times on the noun, suggesting the presence of an active search after encountering an ellipsis gap. Section 4.2, however, failed to replicate this finding using identical items but a different methodology. Furthermore, the strength of the findings of §4.1 for the proposed hypothesis are attenuated because the effect was also found in sentences containing no ellipsis, suggesting that the integration of the critical noun generated processing difficulty regardless of whether it was under consideration as a postcedent of an ellipsis site.

This section seeks to alleviate both of these tensions. First, it provides a replication of the main findings in §4.1 with slightly modified materials. I argue that the failure to find an effect in §4.2 is due to methodological differences between the Maze task and self-paced reading. Second, the experimental materials in this section were modified to follow up on the findings of §4.1. The nouns used in the distant condition in some items were altered to more closely match in semantic relatedness and frequency with the nouns used in that item's local condition. If the reading time slow-down in the No Ellipsis conditions of the experiment in §4.1 was driven by the semantic relationship between the two nouns, as I suggested in that section, then more closely matching the semantic relatedness of both the distant and local nouns to the congruent noun in the first clause should cause this effect to decrease or disappear. I show that the results in this

section do in fact provide evidence in favor of the hypothesis that semantic relatedness was driving the reading-time slowdown in the No Ellipsis conditions in §4.1.

4.3.1 Method

4.3.1.1 Methodology

This experiment used the Maze methodology, as described in §4.1.

4.3.1.2 Participants

Seventy-two participants from Prolific participated in the experiment. Participants self-identified as native speakers of English, and locations were restricted to the United States, Canada, the United Kingdom, Ireland, Australia, and New Zealand. Participants were required to successfully complete three of five practice trials in order to qualify for the experiment. All participants who successfully completed the experiment were paid \$6. Participants were given 87 minutes to complete the entire experiment; it took participants about 30 minutes on average.

4.3.1.3 Materials

Materials were identical to those of the experiment in §4.1 except for the following modifications. First, by-item means were calculated for the experimental items in §4.1. Because the current design matches the distant noun with the local noun, I focus only on the means within the No Ellipsis condition. These means are given in Figure 4.10.

Recall that the hypothesis stated in §4.1 is that the reading time of the noun in the critical region was driven by the semantic relationship between the nouns in the critical region and their analogue noun in the first clause. It follows from this hypothesis that the items that show a greater reading time difference between their local noun and their

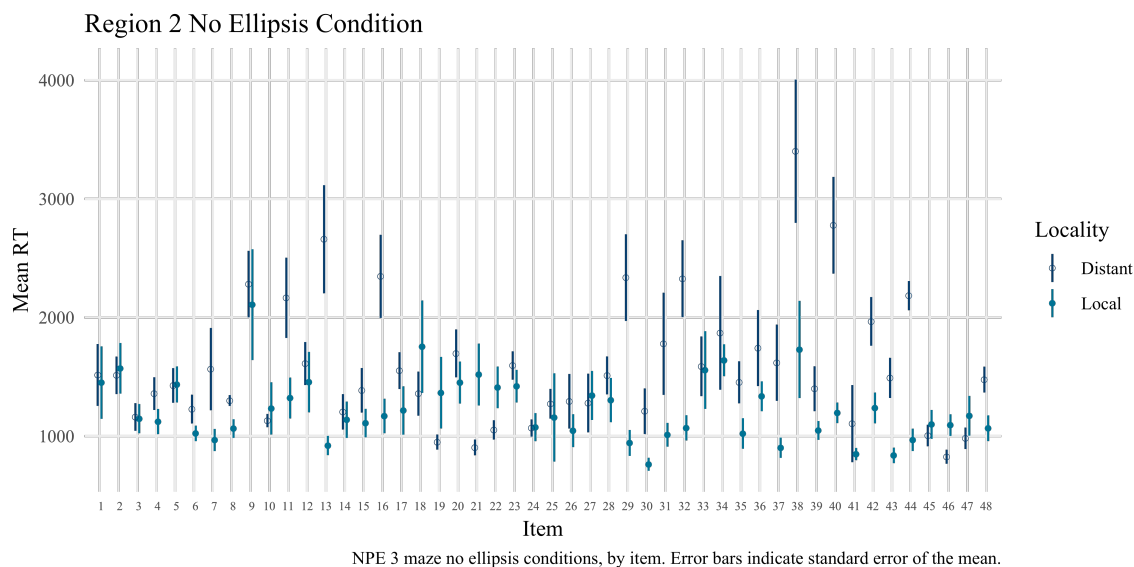


Figure 4.10: Section 4.1 Maze task mean reading times by item in the No Ellipsis condition. Error bars indicate standard error of the mean computed over items.

distant noun also have a greater semantic semantic relatedness difference between each of those nouns and their analogue noun. This is indeed what was found in §4.1. Therefore, the items that were altered for the current experiment were those items that showed a significant difference in reading times between the Local and Distant conditions, as shown in Figure 4.10. For these items, I found a semantic relatedness score for the critical nouns in the Local condition and their analogue noun, using the results given by the University of Colorado, Boulder’s Latent Semantic Analysis.¹⁰ I then selected new nouns

¹⁰<http://lsa.colorado.edu/>

for the critical region in the Distant condition based on the following constraints, best thought of as a partially ordered set of generally decreasing precedence:

1. Maintain the plausibility manipulation of the Local and Distant conditions;
2. Match the semantic relatedness score of the Local and Distant nouns to their analogue noun (so that both nouns were as equally matched as possible to the analogue noun);
3. Match the frequency of the critical nouns in the Local and Distant conditions;
4. Match the length of the critical nouns in the Local and Distant conditions.

The following items were altered: 7, 8, 13, 16, 18, 19, 22, 29, 30, 31, 32, 35, 37,38, 39, 40, 42, 43, 44, 46, 48. Examples of original and altered items are given in (133) and (134). All items are given in Appendix C.1.

(133) Original Item #13

Despite a few vegetables dying from frost and bugs, many books are thankfully saying that tomatoes are surviving to the end of the season.

Altered Item #13

Despite a few vegetables dying from frost and bugs, many gardeners are thankfully saying that tomatoes are surviving to the end of the season.

(134) Original Item #38

Even though few insects have noticeably long antennae, all pre-teens are certainly aware that mayflies are comprised of multiple body sections.

Altered Item #38

Even though few insects have noticeably long antennae, all zoologists are certainly aware that mayflies are comprised of multiple body sections.

I proposed in §4.1.4.2 that the failure rate in particular conditions in the first Maze task experiment may have been due to insufficiently clear foils on some trial. All foils were therefore checked again to ensure that the foil was unambiguously an ungrammatical or highly infelicitous continuation. As in the first maze experiment, all foils within an item were held constant across conditions in critical regions.

4.3.2 Procedure

The procedure was identical to that in §4.1.2.

4.3.3 Results

Maze trials were filtered so that only correct trials were included in the analysis. Following the suggestion in Forster et al. 2009, the first trial of each sentence was removed before any analysis; these trials were uninformative because participants saw the first word of the sentence and the non-word foil 'x-x-x'. After removing the first trials, the longest .3% of remaining RTs was removed, as well as any RT below 200ms. Of the remaining data, each item was separated into eight regions, within which means were calculated for each condition for each item, and then the mean over items within each condition. Means for each region are given in Figure 4.11. The means for the first experiment are repeated for comparison in Figure 4.12.

The critical region, *NP*, was analyzed using a Bayesian mixed-effects linear regression. The analysis was implemented in the *brms* package for R (Bürkner, 2017) and was fitted with default priors. The model failed to converge without divergences using the exgaussian family; therefore, the *shifted_lognormal* family was used, following the recommendation given for response times in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure for Ellipsis and Locality, including their interaction.

A model including the fixed effects of Word Length and Experimental List showed that these effects did not reach statistical significance. The effects were therefore not included in the final model. Word frequency, again expressed as a Zipf score provided by the SUBTLEXus corpus, did reach significance and was therefore included in the final



Figure 4.11: Maze task reading times by region. Error bars indicate standard error of the mean over items. N items = 48, N participants = 72

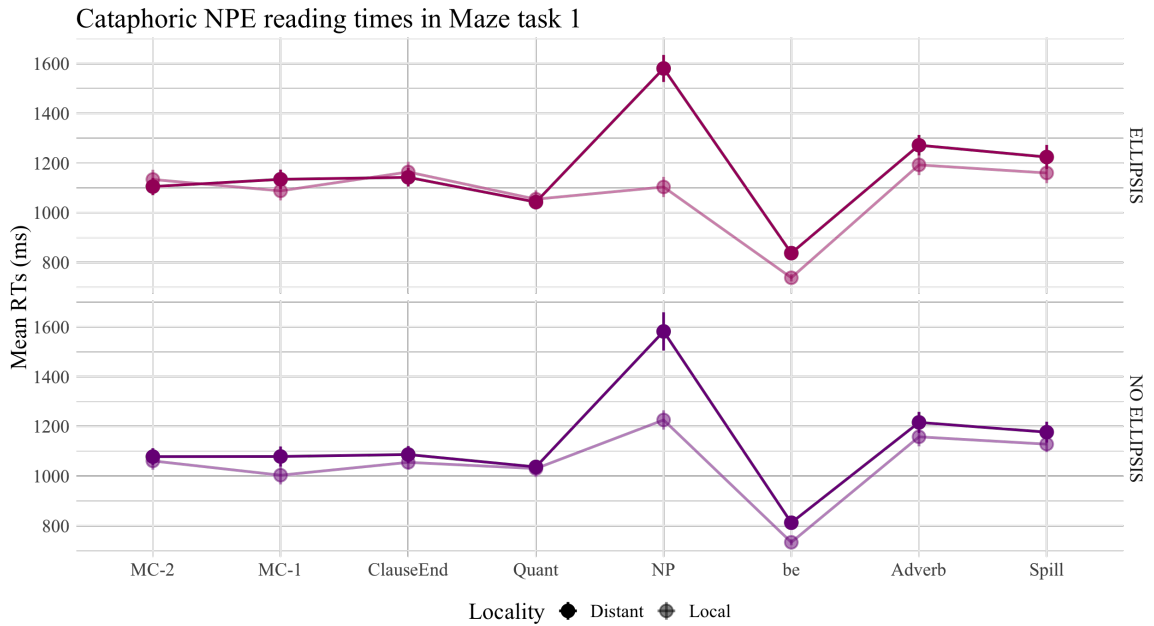


Figure 4.12: Maze task reading times by region in first Maze experiment. Error bars indicate standard error of the mean over items. N items = 48, N participants = 72

model as a fixed effect. Table 4.9 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

	estimate	est.error	2.5%	97.5%
Intercept	7.07	0.10	6.88	7.26
Locality	-0.28	0.04	-0.37	-0.20
Ellipsis	-0.01	0.04	-0.10	0.07
Locality:Ellipsis	-0.18	0.07	-0.31	-0.05
Frequency	-0.14	0.03	-0.20	-0.08
sigma	0.47	0.02	0.44	0.50
ndt	401.59	17.35	363.78	431.04

Table 4.9: Bayesian model estimates for Region *NP*, Maze task revisited.

To check model fit, I ran a posterior predictive check. The posterior predictive check plots the observed data of the experiment against simulated data generated from the posterior predictive distribution (Gabry et al., 2019). If the model is a good fit to the observed data, then the simulated data from the posterior distribution, y_{rep} , should closely match the data we observed, y . The graphical posterior predictive check revealed good model fit, as shown in Figure 4.13.

I predicted that the critical region would show a slow-down for the Distant (implausible) condition compared to the Local (plausible condition). This prediction was borne out, as reflected in the main effect of **Locality** shown in Table 4.9. This finding replicates the finding in §4.1. However, unlike in §4.1, the analysis also found a significant interaction between **Locality:Ellipsis**, in which the reading time difference between the Local and Distant conditions was greater in the Ellipsis condition than in the No Ellipsis condition. Because there is uncertainty in arguing about main effects in the presence of an interaction, I focus on the interaction for the remainder of the discussion.

Because a spill-over effect was found in the *be* region in §4.1, I ran an analysis for that region as well. The analysis procedure was identical to that described above for the

NP region. Table 4.10 provides the posterior estimates of the fixed effects as well as 95% credible intervals. Note that word frequency was not included as an effect because all words in the region consisted of tensed versions of the verb *to be*. To check model fit, I also ran a posterior predictive check for the region. Once again, the posterior predictive check revealed good model fit, as shown in Figure 4.14.

	estimate	est.error	2.5%	97.5%
Intercept	6.00	0.05	5.91	6.09
Locality	-0.06	0.03	-0.12	0.00
Ellipsis	0.04	0.03	-0.02	0.09
Locality:Ellipsis	-0.13	0.05	-0.24	-0.03
sigma	0.42	0.01	0.39	0.45
ndt	332.72	10.07	311.47	350.87

Table 4.10: NPE Maze: Bayesian model estimates for Region *be*.




Figure 4.13: NPE Maze: posterior predictive check generated by Bayesian model for Region *NP*.




Figure 4.14: NPE Maze: posterior predictive check generated by Bayesian model for Region *be*.

The model for the *be* region shows a moderately significant main effect of **Locality**, and an interaction of **Locality:Ellipsis**. Although the mean reading times of each condition in this region are significantly faster than those of the critical region, the model effects replicate those of the critical region, revealing a slight spill-over effect from the critical region onto the following region.

As a final point, note that the presence of an apparently significant interaction in the Spill region should not be taken to be of theoretical import; this region encompasses vastly different words and

word classes across items and conditions, and is the sentence-final word for some items/conditions, but not for others.

4.3.4 Discussion

4.3.4.1 Noun plausibility modulates reading times to a greater degree under postcedent search

The findings in this section replicate and expand upon the findings of §4.1. As in §4.1, the current experiment sought evidence of whether ellipsis resolution is more similar to syntactic dependencies that are actively resolved by the parser, such as those created in *wh*-questions, or whether ellipsis resolution is more similar to pronoun resolution, which is often resolved using general reasoning strategies after sentence parsing is complete.

The hypothesis I proposed was that if the cataphoric ellipsis site catalyzed an active search process for a postcedent, then the plausibility of a candidate postcedent would modulate the reading time on the postcedent such that an implausible postcedent would be read more slowly than a plausible postcedent. The results in this section support this hypothesis, showing that participants read the critical noun at a significantly slower rate when it was an implausible postcedent (the Distant condition) than when it was a plausible postcedent (the Local condition). The findings therefore overall support a view under which comprehenders were engaging in an active search process to resolve the cataphoric ellipsis site.

Unlike the results of the first maze experiment, in §4.1, the current experiment also found an interaction of the Locality condition and the Ellipsis condition, reflecting the finding that the reading-time difference between the Local and Distant conditions in the No Ellipsis condition was smaller than the difference between the Local and Distant conditions in the Ellipsis condition. This result supports the hypothesis proposed in §4.1.4.1, which stated that the slow-down in the No Ellipsis conditions was driven by the processing difficulty of integrating the critical nouns into a discourse representation with the noun in the first clause. I proposed that the processing difficulty of the critical nouns was driven at least in part by the semantic relatedness between the critical noun and its antecedent noun in the first clause. Table 4.3.4.1 shows summary statistics for the Cosine similarity measures of the items in the Local and Distant conditions in the first and second Maze experiment. The measures confirm that the nouns in the Local condition were more closely semantically related to each other than the nouns in the Distant condition in both experiments. The measures also show that the difference in semantic relatedness is less in the second Maze experiment than in the first, as discussed in §4.3.1.3.

Experiment	Condition	Cosine mean	Cosine s.d.	Cosine median
First	Distant	0.236	0.172	0.213
First	Local	0.553	0.170	0.570
Second	Distant	0.264	0.192	0.270
Second	Local	0.541	0.169	0.550

Table 4.11: NPE Maze Experiments: Cosine semantic similarity measures, No Ellipsis conditions.

To again probe the influence of the semantic similarity of the nouns on reading times I ran a Bayesian linear regression predicting the RTs of the No Ellipsis conditions from Locality and semantic similarity (cosine). Table 4.12 provides the posterior estimates of the fixed effects of Locality and Cosine as well as 95% credible intervals. The model used the shifted_lognormal family, as recommended for response times in Bürkner 2020. The model included random intercepts for Item and Participant, and random slopes for Item and Participant for each main effect; interactions between fixed effects were not included.

	estimate	est.error	2.5%	97.5%
Intercept	6.84	0.09	6.65	7.03
Locality	0.03	0.07	-0.11	0.18
Cosine	-0.63	0.19	-1.01	-0.26
sigma	0.49	0.03	0.44	0.54
ndt	448.39	25.65	392.60	492.26

Table 4.12: NPE second Maze Experiment: Bayesian model estimates for critical region, No Ellipsis conditions, including semantic similarity measure.

I predicted that, as in §4.1.4, semantic similarity between the two nouns would impact the reading time of the second noun such that the greater the semantic distance between the two nouns, the greater the reading time on the second noun. This prediction was borne out. Table 4.12 shows a significant main effect of Cosine, reflected in the 95%

credible interval of -1.01 and -0.26. The model also shows that, as in the first Maze experiment, Locality did not reach significance. We can conclude that this failure to reach significance is again due to the Cosine measure accounting for the model variance. Because the current experiment is identical to the experiment in §4.1 other than the changes in the nouns described in this section, we can conclude that these changes are what drove the decreased RT difference in the No Ellipsis condition.

The current experiment therefore serves two main functions. It confirms the main finding of the Ellipsis condition that we found in the first Maze task experiment: the plausibility of a candidate postcedent noun modulates the reading-time, and thus the processing difficulty, of the noun. Second, the semantic relatedness of the critical noun and the noun in the first clause also modulates the reading time of the critical noun, such that the less the semantic relatedness of the two nouns, the greater the reading time. I believe that this effect on its own is worth further investigation. The constraints of the current experiment prohibited exploring the nuances of this relationship, as it is orthogonal to the main research questions. However, it would be interesting to investigate which semantic properties modulate the reading time of the critical noun. For example, one way to establish a baseline measurement might be repeating the noun in the first clause and in the critical region, but varying the quantifiers to make a super/sub-set relation. One could also test hypernym and hyponym relationships, as well as the givenness of the noun. I leave these questions to be answered by future investigations.

4.3.4.2 Indexing processing difficulty in the maze task

Figure 4.15 shows the total number of trials missed on the critical items broken down by region. As in Figure 4.6, only regions 0-11 are plotted. The plot shows a positive and fairly steady increase in total failure rates until peaking at region 11. The positive slope of the failure rate is once again expected because that the failure rates are cumulative.

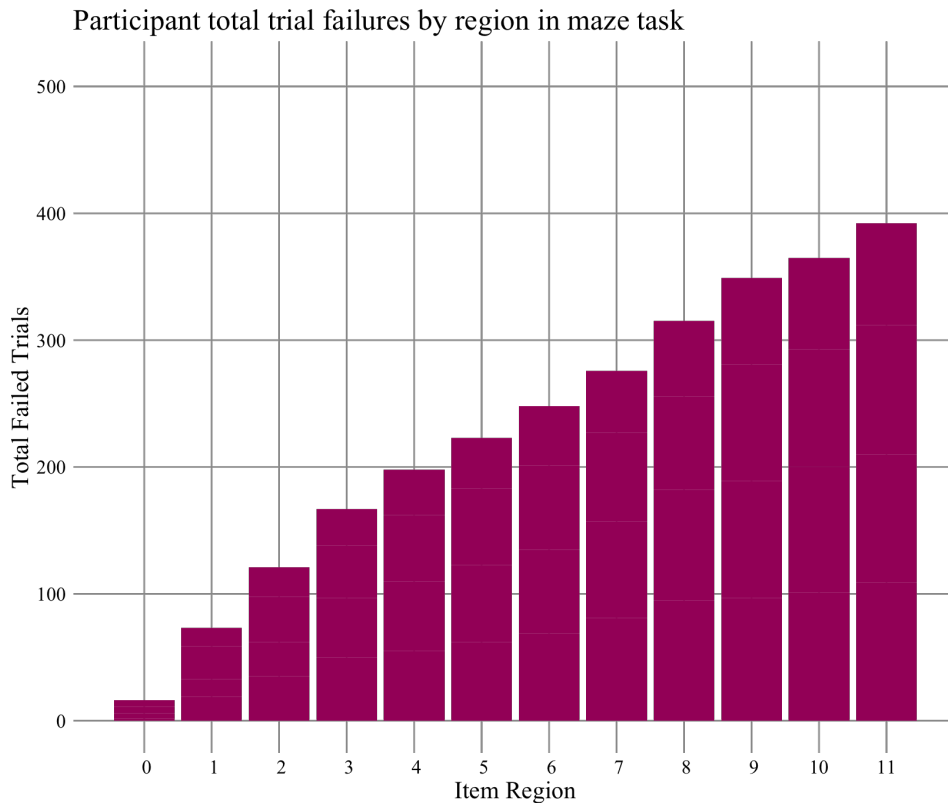


Figure 4.15: Total trial fails by region in NPE maze task. N = 2, 743

The biggest increases in failure are again seen in regions 1, 2, 3, and in region 7 and 8. The explanation for these increases is identical to the analogous increases in Figure 4.15: foils are in general chosen at a greater rate in early trials (Boyce et al., 2020; Sloggett et al., 2020), and the failures in regions 8 (for the Ellipsis conditions) and region 9 (for the No Ellipsis conditions) reflect difficulty in integrating the critical nouns.

The main difference between the failure rate in the first Maze task experiment and the current experiment is the overall failure rate. While the first maze experiment had 3, 698 failures total, the current experiment had only 2, 743. Comparing Figures 4.8 and 4.16 suggest that much of the overall difference is driven by a higher failure rate in earlier trials in the first experiment, which drove up the cumulative failure rate. The

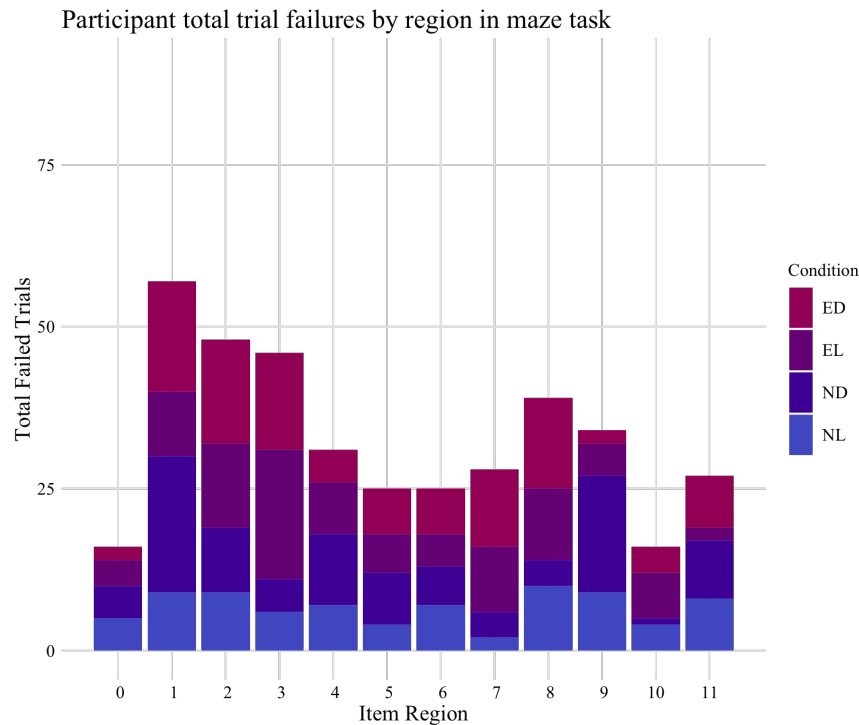


Figure 4.16: Trial fails count by region for all conditions in NPE maze task. Counts indicates the region where trial failure occurred. N = 392

failure rates in regions 1-3 are substantially greater – 72, 87, 85, respectively – for the first experiment than for the current experiment – 57, 48, 46, respectively. The other notable regions of greater failure in the first experiment are regions 8 and 9, which were 61 and 49 in the first experiment and 39 and 34 in the current experiment, respectively.

Figures 4.16 and 4.17 both visualize the failure rate on individual regions. Figure 4.17 is a cumulative density function plot of region 1-11 of the critical items, by conditions. Both plots show that the conditions largely track together through each region, with the No Ellipsis local condition having the overall lowest failure rate.

Overall, I believe that we can draw only tentative conclusions from the failure rate in the maze tasks. The results provide additional evidence that failure rates are greatest in




Figure 4.17: Cumulative distribution function, trial fails by region for all conditions in NPE maze task. Counts indicates the region where trial failure occurred.

early trials, corroborating the findings of previous research. The failure rates do not map clearly onto reading-time changes, and only provide suggestive evidence that processing difficulty can be indexed by trial failure as well as by reading-times. The results of the two experiments also suggest that the failure rate is sensitive to the properties of the foils. Recall that the foils for the current experiment were checked an additional time before the experiment was run, resulting in clearer discrimination between the foils and the grammatical continuations on certain trials. I believe that this clearer discrimination accounts for the difference in failure rates between the two experiments. A more tightly controlled experiment is needed to fully address the effect of foils on failure rate, and in turn the relationship between failure rate in the Maze task and processing difficulty.

4.4 General Discussion

This chapter presented three experiments testing whether language comprehenders employ an active search strategy when resolving cataphoric ellipsis dependencies. Two experiments were conducted using the Maze task, and one was conducted using self-paced reading methodology. The theoretical and methodological implications of these experiments are discussed in this section.

4.4.1 Theoretical implications

Two experiments conducted in the Maze task provided evidence that comprehenders engage in an active search strategy when resolving cataphoric ellipsis dependencies. These findings support the Active Gap Strategy hypothesis: upon encountering an ellipsis gap, comprehenders launched an active search for an antecedent to fill the interpretive gap. Evidence for the search manifested in a significant reading-time slow-down when comprehenders encountered an implausible candidate postcedent for the ellipsis site,

compared to when they encountered a plausible candidate postcedent. This reading-time slow-down indexes processing difficulty, suggesting that the implausible antecedent was either more difficult to integrate into the comprehender's linguistic and non-linguistic representation they had built up to the point of encountering the bottom-up information of the critical noun, and/or that the implausible antecedent was considered for postcedent status and ultimately rejected. The plausibility effect was also found to be greater when the critical noun was under consideration as a postcedent for an ellipsis site than when it was not, suggesting that at least part of the processing difficulty came from attempting to integrate the noun into the gap in the representation left by the cataphoric ellipsis site.

The results support the hypothesis stated in Chapter 3 and at the beginning of the current Chapter: cataphoric ellipsis and anaphoric ellipsis sites both trigger a search for an appropriate antecedent/postcedent to fill the missing representation of the ellipsis site. The difference between anaphoric and cataphoric ellipsis is the nature of the search that is launched. In anaphoric ellipsis, the search takes place over the available representations in focal attention and in the discourse space. In cataphoric ellipsis, comprehenders launch an active, forward-looking search for a postcedent representation. This search process is distinct from the search process of anaphoric ellipsis because it operates over incoming representations as parsing of a sentence happens in real-time. Therefore, the search in cataphoric ellipsis is restricted to those representations that are in focal attention: representations that have been shuttled out of focal attention into declarative memory are not searched.

Besides the difference in the directionality of the search, I have argued that there is no reason to posit any fundamental differences between the two types of search other than those that sensibly fall out from the different nature of the search process. I have also proposed what I have called the Active Gap Strategy, which hews closely to the established Active Filler Strategy. I would like to discuss here differences between what

we know about active filler processes, like those proposed in the Active Filler Strategy, and what we know (and do not know) about the Active Gap Strategy.

First, I have said that my experimental design is inspired by previous experiments examining long-distance dependencies using plausibility manipulations, discussed in the introduction to this chapter. However, one major design difference between the experiments here and these previous experiments is that the plausibility manipulation in the current experiments comes from the predicates of the first clause, and not from the link between the gap and the noun itself. For example, recall the items of Yoshida et al, an example of which is repeated in (135). Under this design, the plausibility effect arises from a mismatch between the thematic role assigned by the verb *notify* to the noun in the *wh*-phrase *which book* (i.e. you cannot notify a book). This is also similar to the plausibility manipulation we saw in Traxler and Pickering (1996), in which there was a thematic role mismatch between a topicalized argument and its verb.

- (135) a. I don't know which book [], but the editor notified the publisher about a new book.
- b. I don't know which book [~~#the editor notified~~], but the editor notified the publisher about a new book.

However, the experiments in this chapter rely on adjectives in the first clause to manipulate the plausibility of the candidate antecedent noun, as shown in (136). I believe that there is no inherent reason to believe that this difference undermines the evidential status of the findings here as supporting an active syntactic dependency formation between the ellipsis gap and its noun. As in previous studies, the semantic plausibility manipulation here serves as an indirect measure of the syntactic search; I can think of no confound introduced by changing the location of the information that creates the semantic plausibility manipulation; however, that does not mean one does not exist. To

test whether this manipulation mattered we would need to test another form of ellipsis that allows us to manipulate semantic plausibility directly, such as sluices. However, such an experiment will run into the problems discussed earlier in this chapter.

- (136) a. While many [] are occasionally **excited** and **energetic**, most **students** are generally depressed.
- b. While many [] are occasionally **excited** and **energetic**, most **reports** are generally saying that students are depressed.

Another possibility is to change what properties are used to evaluate incoming representations for goodness-of-fit. We might think, for example, that the cues used to match a candidate postcedent are the same cues used to search candidate antecedent representations in memory in anaphoric ellipsis. The current experiments manipulated plausibility using mostly animacy and thematic role assignment violations, but further studies could manipulate phi features such as gender/number for candidate postcedents. For example, we saw in Chapter 3 that offline interpretations are sensitive to the specificity of candidate antecedents; therefore, to the extent that specific, definite nouns form bad postcedents for NPE, we predict that a specific candidate antecedent will trigger processing disruption in the same manner as the semantically implausible candidate postcedents in the current experiments.

I believe that the best practice in this situation, as in most situations, is to run multiple experiments and look for convergent evidence across different ellipsis constructions and different items. Several additional follow-up studies suggest themselves based on the current findings. First, due to the experimental design I do not have reliable cross-condition data on what happened more than three words downstream from the critical noun. I can therefore conclude only that it is likely that participants failed to integrate the implausible critical noun as the postcedent of the ellipsis site, and continued their

search by searching every structurally available position until they encountered the plausible noun downstream. This is probable because all sentences in the experiment were grammatical; there was therefore never a case where participants saw a globally ungrammatical or globally pragmatically infelicitous sentence. Participants who persisted correctly to the end of the critical sentences would therefore always end up with a fully felicitous representation of the sentence.

Another follow-up study that immediately presents itself is a version of the current experiment either without a quantifier in either clause or without the quantifier immediately preceding the candidate postcedent. This experiment would address questions about the role that the quantifier plays in structural predictions about the location of the postcedent in the experiments here. Specifically, the quantifier provides a clear structural cue that allows the parser to predict that a noun will immediately follow. The parser may therefore attempt to assign this noun to the missing interpretation of the ellipsis site before confirming its plausibility as an antecedent. If the parser is searching every grammatically available structural position for a postcedent, as predicted by the Active Gap Strategy, then we predict that any DP licensing an NP will, based on knowledge of phrase structure rules, be sufficient to trigger the parser to posit a subsequent NP postcedent. We therefore expect the plausibility manipulation to survive the removal of the quantifier and its replacement with any other head that takes an NP complement. However, if the quantifier supplies key information for the parser, then its removal will modulate the effect. Concerns regarding the role of parallelism with the first clause can also be addressed by varying whether the quantifier appears in both, one, or neither of the two clauses.

A closely related question is at what point in the sentence the reader forms an expectation of the upcoming postcedent. I proposed in the preceding paragraph that early structural evidence for a DP taking an NP complement might trigger structural

expectations for the postcedent. It is also possible that the parser uses more general cues, such as subjecthood, to generate structural expectations about an upcoming postcedent. Recent relevant findings for this question are those of Omaki et al. (2015). Omaki et al. found that, in filler-gap constructions, comprehenders do not appear to wait for information about the transitivity of a verb before positing a gap in the location of the verb's direct object. This effect was manifested in early reading disruptions at intransitive verbs in a potential gap-filling environment, such as (137b), relative to transitive verbs, such as (137a).

- (137) a. The book that the author wrote regularly about was named for an explorer.
b. The book that the author chatted regularly about was named for an explorer.

Omaki et al. argue that this effect is evidence that the parser creates expectations about the location of a gap *before* it reaches the main verb in sentences like (137). The observed slow-down on the transitive verb indicates that comprehenders were positing the gap as the object of the upcoming verb before receiving evidence of the properties of the verb; positing the gap as the object of the upcoming verb generated expectations about the kind of verb that the reader would encounter, namely a certain class of transitive verbs. The slow-down on the verb then reflects a violation of lexical expectations driven by these structural and lexical predictions. Interestingly, Omaki et al. also hypothesize that predictions about the upcoming verb may constrain the verb set not only by the property of transitivity, but also to those that are 'semantically compatible' with the filler NP. Under this account, filler-gap dependency resolution expectations are intertwined with lexical and semantic expectations.

These results are relevant to the current discussion because they suggest that dependency resolution predictions can happen very early in a sentence parse, before even preliminary evidence about the structure that remains to be built is seen. The findings

also suggest that semantic expectations are generated during the dependency resolution prediction process. Together, these results suggest that comprehenders may have been forming early predictions about when they expected to find a postcedent for the ellipsis site and what the form of the postcedent would look like. Further studies can address at what point in the sentence comprehenders form these expectations by varying the structural location of the postcedent.

Omaki et al. acknowledge that their results are also fundamentally compatible with a view of sentence processing in which the parser predicts upcoming structural information by using a wide variety of linguistic and contextual information (Kimball, 1975; Gibson, 1998; Hale, 2003; Staub and Clifton, 2006; Levy, 2008). Similarly, the results here do not allow me to conclusively adjudicate between a theory of active search and an account of the current results that relies on notions of prediction in language processing. I chose to frame the project here in terms of active search because the dominant view in the existing literature on ellipsis processing relies on such a theory. However, the proposal here is not intended to rule out the possibility of a prediction account, and I believe that the results here are completely compatible with the predictions of such an account.

For example, in a review of the role of prediction in language processing, Kuperberg and Jaeger (2016) argue that existing research supports the idea that a comprehender's internal representation of linguistic and non-linguistic context encodes multiple different types of information at different levels of representation. Because of the interactive nature of language processing, any or all of these levels of information can be used to facilitate the processing of incoming information at any other level of representation (Kuperberg and Jaeger 2016, *inter alia.*). They suggest, however, that lower-level information, such as phonological and syntactic information, may be shorter-lived than higher-level information, such as semantic and pragmatic information (see also Sachs 1967 for evidence that semantic 'gist' information persists in memory longer than syntactic information). If

such a model of predictive comprehension is correct, it suggests that the search for a postcedent may be integrating information from multiple levels of representation: syntactic knowledge of structural constraints on the placement of NPs, semantic knowledge of theta roles, and pragmatic information about topicality and discourse relations.

More controversially, Kuperberg and Jaeger propose that information at these levels of comprehension can flow downwards, where high-level top-down information is used to pre-activate information at lower levels of representation (see citations within for extensive discussion). While the evidence for such a theory remains mixed, the pre-activation of the initial noun's semantic neighbors would explain why the reading times on the critical noun in the No Ellipsis conditions were modulated by the nouns' semantic relationship with the corresponding noun in the first clause. Under such an account, comprehenders would be using contextual information given by the discourse relations discussed in §4.1.4 to predict semantic information about the critical noun before they received bottom-up information specifying its lexical context. The predictions regarding the semantic information of the upcoming noun are likely all or partially violated in the condition with the implausible noun, leading to processing difficulty manifesting in reading-time increases and an increased failure rate (Forster, 1981; Schwanenflugel and Schoben, 1985; Schwanenflugel and LaCount, 1988; Kuperberg and Jaeger, 2016).

Alternatively, the reading-time modulation could be due only to the initial activation of the first noun spreading to its neighbors; the activation of the critical noun would therefore not be pre-emptive, but a lingering effect of the activation of the noun in the first clause, also known as *priming effects*. One way to distinguish between the possibilities of pre-activation and spreading-activation is to increase the distance between the two nouns, without adding any intervening candidate nouns. Because priming is short-lived, if semantic relatedness is modulating the reading-time on the second noun due to priming from the first noun, then we expect this difference to decrease as the time between the two

nouns increases. The increased time should, however, have no effect on the reading-time of the second noun if the processing facilitation is due to pre-emptive activation.

One possible way to adjudicate between an active search strategy and prediction is to look for anti-locality effects. Anti-locality effects describe the observation that increasing the length of a dependency does not always lead to decrements in the quality of processing. In fact, sometimes we see the exact opposite effect, in which increased length actually facilitates processing (Konieczny, 2000; Konieczny and Doring, 2003; Lewis and Vasishth, 2005; Levy, 2008; Levy and Keller, 2013). For example, Konieczny and Vasishth & Lewis found that reading times for longer dependencies were faster at the point of dependency integration than reading times for shorter dependencies. This is counter to the result we would expect if longer dependencies always created a drop in processing quality. Generally speaking, the facilitatory effect of longer dependencies are explained by a positive relationship between the probability of encountering the end of the dependency and the length of the dependency: as the length of a dependency increases, the probability of reaching the end of the dependency also increases. As the probability of reaching the integration point of the dependency increases, there is a boost in the activation of the constituent at which the dependency was formed in anticipation of resolving the dependency. This anticipatory boost in activation facilitates processing at the integration site, which manifests in, for example, faster reading times.

For cataphoric ellipsis, we might expect that the ease of integration of the postcedent could be affected by both the length between the ellipsis site and the postcedent, and the informativity of the cues preceding the postcedent. If prediction is indeed playing a role in identifying the postcedent, we might expect that increasing the distance between the ellipsis site and the postcedent will facilitate processing of the postcedent due an increased perceived probability of encountering the postcedent NP. That is, as comprehenders reach the end of a sentence, the probability that the next encountered NP will be the postcedent

increases. These prediction effects should also be modulated by the informativity of the cues preceding the postcedent. I leave the investigation of these interesting questions to future work.

4.4.2 Methodological considerations: SPR vs. Maze

The two Maze task experiments conducted in this chapter found evidence of a forward-looking search process in cataphoric NPE. Evidence of this search process was found in the form of reading-time slow-downs on the experimental critical nouns, indexing processing difficulty. Analogous effects were also found on the verb immediately following the critical noun, which acted as a spill-over region that mirrored the significant findings of the critical region. The self-paced reading experiment, however, failed to reveal these significant effects, or indeed any significant effects at all. I proposed in §4.2.4 that the failure to find an effect in the SPR experiment was due to methodological differences between the Maze and SPR tasks.

This methodological divergence strongly suggests that researchers investigating semantic and pragmatic processing should avail themselves of various experimental methodologies when testing a particular hypothesis. Because psycholinguistic research on semantic and pragmatic processing is a relatively young field,¹¹ the field does not yet have a definitive grasp of which methodologies are sensitive to which semantic and pragmatic processing effects. I hope these findings serve as both a cautionary note about the importance of methodological cognizance, but also an encouraging example of what can be achieved with relatively simple tools.

¹¹cf. syntactic investigations of this nature. See also the methodological debate over sensitivities in SPR and eye-tracking when investigating pronoun resolution in Kazanina et al. (2007) and Patterson and Felser (2019).

4.4.3 Further questions

What remains an interesting and open question is whether the locality bias found here and in previous investigations of ellipsis processing interacts with a preference to interpret an ellipsis site as anaphoric or cataphoric, in cases in which both interpretations are licensed (Liversedge and van Gompel, 2003). For example, in (138), if there is a preference to interpret the ellipsis site as anaphoric, then we expect the interpretation of the first clause of the second sentence to be [if you see any cookies] at the point of encountering the comma (evidence of the end of the clause). However, there is a licit but structurally most distant candidate antecedent, ‘snacks.’ If a pure locality effect is in effect, then we expect that comprehenders will maintain their initial parse even after encountering an alternative antecedent candidate. However, if strict locality is overridden by other factors, then we expect that participants might change their parse, or consider changing their parse, when they encounter the second candidate antecedent.

(138) I want some cookies. If you see any, can you also buy me some snacks from the store?

This question is closely related to an additional question, which is how much of the effect here can be extended to cataphoric processes more generally. This question can be investigated by extending the experimental paradigm used here to additional types of ellipsis as well as to pronominal anaphora, as in (139).

- (139) (a) While they are occasionally excited and energetic, most reports are generally saying that under-graduates are depressed.
- (b) While they are occasionally excited and energetic, most undergraduates are generally depressed.

Chapter 5

Implicit Causality

The last two chapters have been an in-depth case study on the role of proximity and linear positioning in ellipsis interpretation. This dissertation, however, is not entitled ‘The role of proximity in ellipsis interpretation.’ One could write such a dissertation, but it would suffer from a failure to see the forest for the trees. There are many factors that affect our understanding of an ellipsis site, and while I hope I have so far convinced my reader that proximity is an important player, it is certainly not the whole story.

No theory is built from scratch, and all empirical investigations need a starting point. An uncontroversial property of ellipsis (perhaps the only uncontroversial property) is that it is anaphoric. A natural starting point for a theory of ellipsis comprehension, then, is to compare it with better-studied phenomena that are also uncontroversially anaphoric. I have argued that anaphor resolution is such a phenomenon. For my reader who is tempted to dismiss the utility of studying ellipsis comprehension when we have so thoroughly studied anaphoric resolution: it is not a given that ellipsis comprehension and anaphor resolution should be identical. For starters, anaphors are syntactically and phonologically realized morphemes that point (in an a-theoretical sense) to a discourse representation that is introduced via a linguistic antecedent or exophorically. Ellipsis is the absence of

any phonological form; it is not a pointer, it is a gap that must be filled. The processing of ellipsis therefore requires identifying this gap in syntax, and then finding and filling in that representation.

Pronominal anaphors provide featural cues to their antecedent (person, number, animacy), and also to the form of their discourse representation (pronominal, event, propositional). Ellipsis provides only a structural cue via the syntactic content surrounding the gap (is a TP missing? an NP?). Anaphors are also notoriously flexible in their ability to be exophoric and to refer to representations of non-contiguous linguistic structure. The licensing conditions of ellipsis, however, have been shown to be much stricter than those of anaphors. For example, Hankamer and Sag (1976) famously discuss differences between ellipsis, as a surface anaphor, and anaphors, as deep anaphors. Propositional anaphors have also long been known to be liberal in their interpretations (Webber, 1978, 1988; Asher, 1993), but the restrictions on sluicing are significantly more constrained (see overview in Chapter 2). All of which is to say, the pursuit of understanding ellipsis comprehension is not a vapid one; it remains unknown how closely ellipsis hews to better-known and a priori related processes, such as anaphor and filler-gap resolution.

Nevertheless, anaphor resolution has been well-studied and its similarities with ellipsis make it a reasonable starting point of comparison for ellipsis comprehension. Because the bulk of this dissertation has focused on NP ellipsis as a starting point for a theory of ellipsis comprehension, this chapter takes advantage of a well-established property of pronominal anaphors as a way to expand the domain of properties we can investigate for ellipsis. This property is called *implicit causality*. Implicit causality is an off-the-shelf factor that manipulates the salience of pronominal interpretations. The phenomenon was first observed by Garvey and Caramazza (1974). Garvey and Caramazza noticed that certain verbs biased the reading of subsequent pronouns toward either a subject or object interpretation. For example, consider the minimal pair in (140).

(140) Rose frightened Donna because she...

Rose loved Donna because she...

Comprehenders have been shown to overwhelmingly prefer a subject interpretation for the pronoun in the first sentence, but an object interpretation for the pronoun in the second sentence. Since Garvey and Caramazza's initial observation, two dominant explanations for the bias have emerged. Early observations suggested that the bias was not a rule of syntax, as the preferences were not strict, appeared to depend on real-world knowledge, and could be overridden, albeit at some cost (Caramazza et al., 1977; Stewart et al., 2000; Koornneef and van Berkum, 2006). The dominant view was then that the preference was a pragmatic preference, attributed to high-level, non-linguistic information (Brown and Fish 1983; Semin and Fiedler 1991; Corrigan 2001, 2003; Pickering and Majid 2007, a.o.). Kehler 2002; Kehler et al. 2008; Kehler and Rohde 2013 influentially argue that implicit causality is attributable to discourse structure and coherence relations. Recent work has argued, however, that implicit causality should be tied to syntactic verb class, and is therefore not due to pragmatic inference (Hartshorne, 2013; Hartshorne and Snedeker, 2013). The current experiments are not intended to weight in on this debate, and I will not make claims regarding which level of representation implicit causality should be located at.

Instead, this chapter consists of two experiments providing a direct comparison of implicit causality for nominal anaphors and NPE. While implicit causality has been extensively investigated with pronominal anaphors, I know of no studies investigating its effect in ellipsis. As NPE is the closest analogue to pronominal anaphors, that is what is investigated here. It is difficult to make specific predictions about how ellipsis might pattern differently in regards to implicit causality than anaphors, because the locus of implicit causality is not a settled question. However, it is worth exploring if and where the processing of ellipsis and anaphors diverge.

In keeping with the methodological comparisons conducted so far, the first experiment was conducted using the Maze methodology, and the second using self-paced reading. I show that, unlike in the previous chapter, no interesting methodological differences between the two experiments were found. I argue from the experimental findings that the implicit causality bias does not influence pronominal anaphors and NPE differently. I also show that the experiments reveal an object-interpretation bias similar to that seen in the offline comprehension judgments of NPE in §3.4.

5.1 NPE shows an implicit causality and object bias in the Maze task

5.1.1 Method

5.1.1.1 Methodology

This experiment used the *maze* task methodology, as discussed in §4.1.

5.1.1.2 Participants

Seventy-nine workers from Amazon’s Mechanical Turk participated in the experiment. Seven workers were excluded because they were unable to successfully complete any sentence. The data from the remaining seventy-two workers was included in the experimental analysis. Participants were paid \$5, and were awarded a \$2 bonus if they successfully completed 80% of the experimental sentences.

5.1.1.3 Materials

Materials consisted of 24 critical items comprising four conditions. The design was a 2 x 2 with the factors *bias* (Object bias, Subject bias) and *dependency* (Anaphora, Ellipsis). All items consisted of three clauses. The first clause contained a subject and object NP

with a verb that was either a subject-biased verb or an object-biased verb. All verbs were taken from Hartshorne and Snedeker (2013). Object bias verbs were chosen among those reported by Hartshorne and Snedeker as having an object bias of 76% or greater, subject bias verbs were chosen among those reported as having an object bias of 24% or lower. The second clause was embedded under the verb *because*, followed by the expletive subject *it* and a verb embedding the final clause, which contained either an elided noun or an anaphoric noun.

Two anaphoric nouns were used in the anaphoric conditions: *ones* and *they*. The use of *they* as an anaphoric noun needs, I assume, no further discussion. The use of *ones* as an anaphor, while potentially eliciting more controversy, has been well-established as an anaphor (Perlmutter, 1970; Jackendoff, 1977). Payne et al. (2013) discuss three uses of *one* in English, briefly described as follows. The pronominal lexical item *one* is a 3rd person singular indefinite pronoun that has no plural form and means along the lines of ‘an arbitrary person’. It can be inflected as *one*, *one’s*, and *oneself*. The determinative lexical item *one* is an indefinite cardinal numeral. It cannot be inflected, has the meaning of single, sole, ‘1’, or ‘some’, and is obligatory when used as a determiner. Finally, the anaphoric lexical item *one* is a regular common count noun. It can be inflected as *one*, *ones*, *one’s*, and *ones’*. This lexical item is anaphoric to some type of class that is referred to in the discourse or is salient in the context of use. Payne et al. give the examples in (141) and (142) of anaphoric *one* (Payne et al., 2013, ex. 6).

(141) The art museum in Bilbao is the most impressive one I’ve seen.

(142) The long, gently curved Victorian railway station building in York is the finest one in the whole of England.

Goldberg and Michaelis (2017) clearly articulate the potential confusion between cardinal *one* and anaphoric *one*. Goldberg and Michaelis point out that when cardinal

one appears without a noun it is nearly indistinguishable from anaphoric *one*. However, they propose two ways to distinguish the two uses. The first way is that the two uses have an interpretation difference, such that cardinal *one* foregrounds the cardinality reading. Second, the two uses are prosodically distinguishable, with the cardinal lexical item carrying a primary accent while the anaphoric lexical item does not.

Given this background, there are several reasons to be reassured that the lexical item used in this experiment is the anaphoric lexical item. First, the anaphoric lexical item is the only version of *one* that can be inflected as *ones*. Second, the lexical item *one* in this experiment does not convey a cardinality reading, and it does not carry a primary accent. We are therefore justified in interpreting the results of the current experiment as reflecting those of anaphoric processing.

Another key feature of the critical materials is the predicate used to describe the candidate antecedents. The predicate is important because its properties could potentially bias the interpretation of the NP elision site or the pronoun, creating an unintended plausibility manipulation. Previous researchers have responded to this design difficulty by using nonce predicates (Hartshorne and Snedeker, 2013) and by using free response items (Kehler et al., 2008). The current experiment used a different strategy, using general properties that could plausibly describe either NP in the first clause. An example item is given in Table 5.1.

Ellipsis	Plausibility	Sentence
Object	Ellipsis	Baritones resent altos because it's hard to find any that aren't competitive.
Subject	Ellipsis	Baritones intimidate altos because it's hard to find any that aren't competitive.
Object	Anaphora	Baritones resent altos because it's hard to find ones that aren't competitive.
Subject	Anaphora	Baritones intimidate altos because it's hard to find ones that aren't competitive.

Table 5.1: Example item from implicit causality experiment.

However, even though predicates were chosen to plausibly apply to either noun, it is still possible that the choice of noun might bias participants toward one interpretation of

the ellipsis site or pronoun rather than the other. For example, participants have a bias for baritones being more competitive than altos. To account for this possibility, NPs were counterbalanced between two experimental lists (see also discussion in 4.1.2), such that half the participants saw a given NP as a subject, and half of the participants saw that same NP as an object. An example item from List 1 and List 2 is given in Table 5.2 and Table 5.3, respectively.

Ellipsis	Bias	Sentence
Object	Ellipsis	Campers trust hikers because it's difficult to find any that are dishonest.
Subject	Ellipsis	Campers delight hikers because it's difficult to find any that are dishonest.
Object	Anaphora	Campers trust hikers because it's difficult to find ones that are dishonest.
Subject	Anaphora	Campers delight hikers because it's difficult to find ones that are dishonest.

Table 5.2: Example item from implicit causality experiment List 1.

Ellipsis	Bias	Sentence
Object	Ellipsis	Hikers trust campers because it's difficult to find any that are dishonest.
Subject	Ellipsis	Hikers delight campers because it's difficult to find any that are dishonest.
Object	Anaphora	Hikers trust campers because it's difficult to find ones that are dishonest.
Subject	Anaphora	Hikers delight campers because it's difficult to find ones that are dishonest.

Table 5.3: Example item from implicit causality experiment List 2.

After each experimental item, participants were given a forced-choice question. The question always targeted the subject of the last clause – either the elided subject or the pronominal subject. The forced-choice question provided both of the NPs from the first clause as options. The order in which the NP options were given was randomized across items and across experimental trials. Participants had eight seconds to respond, otherwise the question timed out and the next item was displayed. Partial items and their corresponding comprehension questions are given in examples (143) and (144).

- (143) Item sentence:
Baritones resent altos because it's hard to find any that aren't competitive.

Item question:
Who are usually competitive?

Forced-choice options:
baritones, altos

- (144) Item sentence:
Campers trust hikers because it's difficult to find ones that are dishonest.

Item question:
Who are rarely dishonest?

Forced-choice options:
campers, hikers

5.1.2 Procedure

The experiment was administered on Amazon's TurkPrime and used the IbexFarm experimental software (Drummond, 2014). Items were run as fillers for the maze task in Experiment 3; materials were therefore presented in two lists. Thirty-six participants saw each list of all 24 items, which were counterbalanced as described above in 5.1.1.3. The results are combined here for the analysis.

Within each experimental list, the critical items were distributed into four Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, ensuring balanced responses across lists. Participants were trained on six practice items and were given two cartoon breaks half and 2/3 of the way through the experiment. Participants could complete each Maze trial at their own pace. If they successfully completed a sentence, they received a screen informing them they were correct and prompting them to press any key to continue to the forced-choice question. Any keypress would then bring up the question screen, which would time out after no activity for eight seconds. If a participant got a Maze trial wrong – if they chose a foil instead of the grammatical/felicitous continuation – the trial was aborted. After an incorrect choice, participants were given a screen telling them that their choice was

incorrect and prompting them to press any key to continue to the question. Note that participants were required to answer the questions for each trial, regardless of whether they successfully completed the Maze task for that trial.

In addition to the 24 critical items, 48 items from Experiment 3 were also included, as well as 24 additional filler items. Filler items included sentences with various types of elliptical and pronominal anaphoric/cataphoric dependencies, including verb phrase ellipsis, sluicing, and pronominal anaphora. At the completion of the experiment a debriefing form was given in which participants were asked to record any difficulties they had encountered.

5.1.3 Results

Trials were filtered so that only completed trials were included in the analysis. This means that only sentences which were correctly completed in full were used; any error (choice of foil word) resulted in the entire trial being discarded. This was done because participants saw the question regardless of whether they successfully completed the entire sentence for that particular item or filler; therefore, failed trials were eliminated in order to exclude occasions when participants were forced to guess with incomplete information. The proportion chosen for each NP, subject or object, is given in Figure 5.1.

The results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the *brms* package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 5.4 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

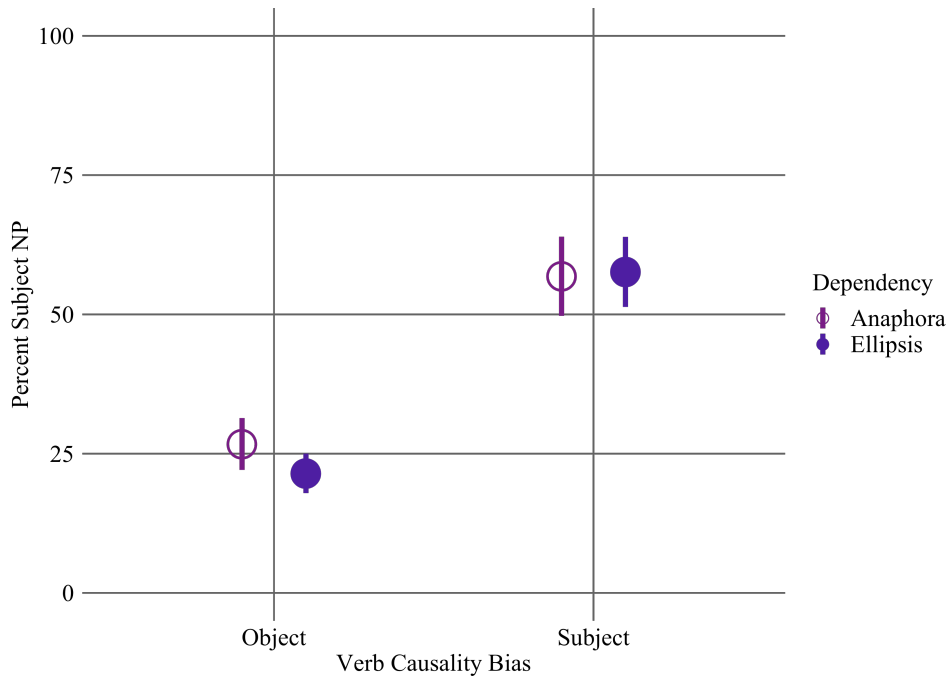


Figure 5.1: Implicit causality proportion subject NP, Maze task. Error bars indicate 95% confidence intervals.

	estimate	est.error	2.5%	97.5%
Intercept	0.47	0.09	0.29	0.66
Bias	1.60	0.21	1.18	2.02
Dependency	0.13	0.15	-0.15	0.42
Bias:Dependency	0.33	0.29	-0.25	0.91

Table 5.4: Bayesian model estimates for implicit causality experiment, Maze task.

To check model fit, I ran a posterior predictive check. The posterior predictive check plots the observed data of the experiment against simulated data generated from the *posterior predictive distribution*, that is, the distribution of the outcome variable provided by the model after using our observed, or experimental, data to update our beliefs about the unknown model parameters θ (Gabry et al., 2019). If the model is a good fit to the

observed data, then the simulated data from the posterior distribution, y_{rep} , should closely match the data we observed, y . The graphical posterior predictive check revealed good model fit, as shown in Figure 5.2.




Figure 5.2: NPE implicit causality posterior predictive check generated by Bayesian model, Maze task.

contained a subject-biased verb, and about 25% of the time when the sentence contained an object-biased verb. This effect is more clearly observed in the bar plots in Figure 5.3, which shows the proportion of object and subject NPs chosen for both the Object-biased verbs and the Subject-biased verbs in the Anaphoric conditions and the Ellipsis conditions, respectively.

An open question was whether the implicit causality bias would affect the Anaphoric and Ellipsis conditions differently. The experiment failed to find any difference between the two conditions, as reflected in the failure to find a significant main effect of **Dependency**, or an interaction of **Dependency: Bias**, in Table 5.4. This effect is also shown

I predicted that the Subject bias condition would see a greater number of subject interpretations – selection of the subject noun – than the Object bias condition, and that the Object bias condition would see a greater number of object interpretations – selection of the object noun – than the Subject bias condition. These predictions were borne out. These findings are reflected in a main effect of **Bias** in Table 5.4. This result is also visible in Figure 5.1, which shows that the subject noun was chosen about 55% of the time when the sentence

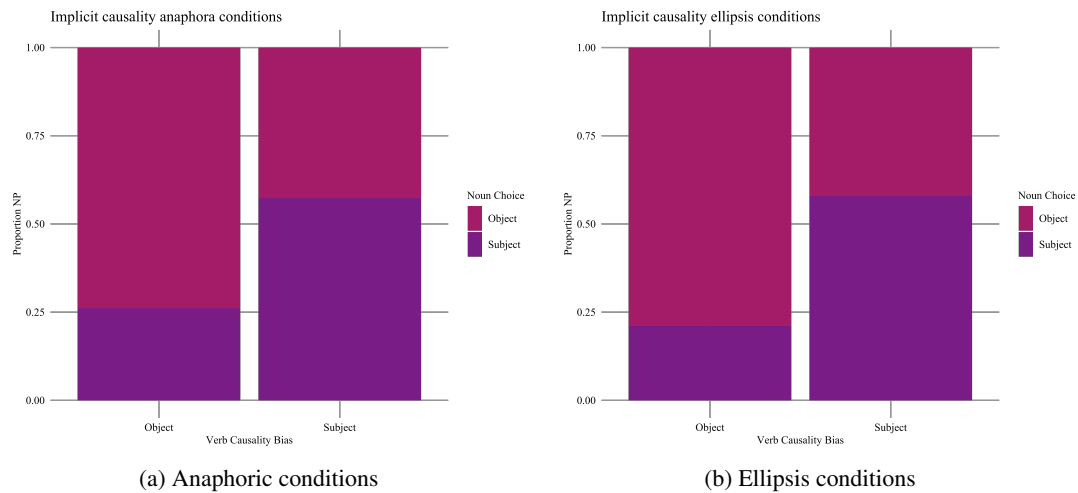


Figure 5.3: Implicit causality proportion subject and object NPs, Maze task.

in Figure 5.1, which shows no significant difference between the proportion of subject interpretations between the Anaphora and Ellipsis conditions.

5.2 NPE shows an implicit causality and object bias in self-paced reading

This section presents a follow-up experiment to that in §5.1. The experiment includes the same items as in §5.1, but uses the self-paced reading methodology. This experiment was run both as a methodological comparison between the self-paced reading and Maze methods, and as a replication of the findings in §5.1.

5.2.1 Methods

5.2.1.1 Methodology

The experiment used the self-paced reading methodology, as described in §4.2.1.1.

5.2.1.2 Materials

Materials were identical to those used in the previous experiment, as described in §5.1.1.3, with the addition of four experimental items. These items were added because the overall pace of self-paced reading experiments is faster than Maze experiments; therefore, additional items could be run while keeping the overall length of the experiment comparable.

5.2.1.3 Participants

Fifty-six undergraduate students from the University of California, Santa Cruz participated in the experiment. All students self-identified as native speakers of English. All students received course credit for their participation.

5.2.2 Procedure

The experiment used the IbexFarm experimental software (Drummond, 2014). Items were run as fillers for the self-paced reading task in §4.2; materials were therefore presented in two lists, counterbalanced as described above in §5.1.1.3. Within each experimental list, the critical items were distributed into four Latin Square lists and were presented in pseudo-random order. Each Latin Square list was run separately, in an effort to ensure balanced responses across lists. Participants were trained on six practice items and were given two cartoon breaks half and 2/3 of the way through the experiment. Participants could complete each trial at their own pace.

Fillers were the items from §4.2 as well as 24 additional fillers. At the completion of the experiment both an online debriefing form and a paper debriefing form was given to participants in which they were asked to record any difficulties they had encountered.

The procedure was otherwise identical to that described in §4.2.2.

5.2.3 Results

The results were analyzed using a Bayesian mixed-effects logistic regression. The analysis was implemented in the `brms` package for R (Bürkner, 2017) and was fitted with default priors. A bernoulli family was used, following the recommendation given for binomial data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure, including all interactions. Table 5.5 provides the posterior estimates of the fixed effects as well as 95% credible intervals.

	estimate	est.error	2.5%	97.5%
Intercept	0.47	0.15	0.18	0.77
Bias	1.93	0.31	1.33	2.56
Dependency	0.02	0.15	-0.28	0.31
Bias:Dependency	-0.06	0.37	-0.77	0.67

Table 5.5: Bayesian model estimates for implicit causality experiment, self-paced reading.

To check model fit, I again ran a posterior predictive check. The posterior predictive check plots the observed data of the experiment against simulated data generated from the posterior predictive distribution (Gabry et al., 2019). If the model is a good fit to the observed data, then the simulated data from the posterior distribution, y_{rep} , should closely match the data we observed, y . The graphical posterior predictive check revealed good model fit, as shown in Figure 5.4.

The proportion of responses for each NP, subject or object, is given in Figure 5.5. I again predicted that the Subject bias condition would see a greater number of subject interpretations – selection of the subject noun – than the Object bias condition, and that the Object bias condition would see a greater number of object interpretations – selection




Figure 5.5: Implicit causality proportion subject NP in self-paced reading. Error bars indicate 95% confidence intervals.

of the object noun – than the Subject bias condition. These predictions were again borne out. The findings are reflected in a main effect of **Bias** in Table 5.5. This result is also visible in Figure 5.5, which shows that the subject noun was chosen about 60% of the time when the sentence contained a subject-biased verb, and about 25% of the time when the sentence contained an object-biased verb.

The experiment also again failed to find any difference between the Anaphora and Ellipsis conditions, as reflected in the lack of a significant main effect of **Dependency**, or an interaction of **Dependency:Bias**, in Table 5.5. This effect can also again be seen in Figure 5.5, which shows no significant difference between the proportion of subject interpretations between the Anaphora and Ellipsis conditions. This null result reveals that the conditions containing NPE and the conditions containing pronominal anaphora

were not treated significantly differently by participants.

5.3 General Discussion

5.3.1 Implicit causality applies across anaphora and ellipsis




Figure 5.4: NPE implicit causality posterior predictive check generated by Bayesian model, self-paced reading.

time.

While this general biasing power of implicit causality verbs has been known for quite some time, the results here are to my knowledge the first experiment directly comparing pronominal anaphora and NPE. The experimental results showed no significant difference between the ellipsis and anaphor conditions, and suggest that the implicit causality bias has an equivalent effect on anaphora and ellipsis processes. This finding gives us two

The results of the two experiments in this section support the the biasing power of implicit causality verbs. In both the Maze task and in the self-paced reading task, participants' interpretations of NPE sites and pronominal anaphors were influenced by the implicit causality bias of the verb.

When the item's main verb had an object-bias norm, participants resolved the NPE sites and pronominal anaphors to an object NP antecedent about 75% of the time. However, when the item's main verb had a subject-bias norm, participants resolved the NPE sites and pronominal anaphors to an object NP antecedent about 45% of the time.

insights. The first regards the breadth of the program of implicit causality. The current results show that implicit causality is not an anaphor-specific phenomenon. While the literature on implicit causality is dominated by anaphora-specific research questions (see Ong (2013) for a notable exception), the current results suggest that ellipsis is another potentially rich area of study for a bias that is well-studied within the domain of anaphora.

The second insight given to us by the current findings regards the nature of ellipsis itself. These experiments show that ellipsis is sensitive to the same type of structural or salience factors that affect anaphoric interpretations. Though the locus of the implicit causality effects is not firmly established (see discussion earlier in this chapter), the impact of these findings now extend to ellipsis as well as to anaphora. For example, let's suppose we live in a world in which implicit causality is a discourse-level phenomenon. The findings here then suggest that this world is also one in which ellipsis interpretation is highly sensitive to discourse-level biases. These findings are then directly predicted by theories such as those given in Kehler 2002.

We can summarize the import of these results as follows. While implicit causality is well established for anaphors, there is less evidence establishing its effect on ellipsis. The current experiments establish that, whatever processing or theoretical factors one ultimately argues is responsible for the implicit causality biasing effect, the explanation must account for the bias applying equally across both anaphora and ellipsis. Additionally, the results show that whatever theory of interpretation one argues is responsible for our comprehension of ellipsis, it must account for the affects of implicit causality.

Finally, there is also a methodological point of interest in the current findings. The implicit causality bias was found in both the self-paced reading experiment and in the Maze experiment. The finding of the effect in both methodologies provides evidence against §4.2's hypothesis that the failure to find any plausibility reading time effects in that experiment was due to participants simply not being attentive. In fact, the

implicit causality results across the two experiments are very consistent. This consistency suggests that the failure to find the plausibility effect in §4.2 was, as discussed, due to methodological factors, and not due to attention factors.

5.3.2 Implicit causality interacts with proximity preference

Unlike in Experiment 1, the current experiments suggest a slight object interpretation bias. In both of the experiments in this section, a subject interpretation was chosen a little over 50% of the time when the verb was a subject-biased verb. However, when the verb was an object-biased verb, participants chose an object interpretation about 75% of the time. This suggests that the subject-biased verb had a harder time ‘pushing’ the interpretation toward a subject bias than the object-biased verbs did for an object interpretation.¹ This subsection discusses this bias by first looking at participants’ response times, and then by discussing several hypotheses about the origin of the response bias.

5.3.2.1 Question response time analysis

The bias toward an object interpretation is also observable in the response times to the experimental questions. A summary of the question response times for both the Maze

¹Note that one possible way of probing this finding is to look at a failure analysis, as done in the previous chapter. However, because no online measures are captured in the current experiments, it’s unclear where we would expect to find relevant failures. The critical responses are conducted offline, and experimental design does not allow us to know when participants are computing the relevant interpretation of the anaphor or ellipsis site. For example, participants may decide on an interpretation once they encounter the ellipsis site/anaphor, or they may wait until the question prompt. Alternatively, participants may have an initial parse of the ellipsis site/anaphor, but then alter it when they choose a response. The experimental design was not set up to answer this questions. Because there is no clear linking hypothesis between failure rates and a theory of implicit causality interpretation for this particular experiment, response time analyses were conducted in lieu of a failure rate analysis.

task and the self-paced reading task are given in Table 5.6, broken down by Bias. The response times were also analyzed using a Bayesian mixed-effects linear regression. The analysis was implemented in the brms package for R (Bürkner, 2017) and was fitted with default priors. A shifted_lognormal family was used, following the recommendation given for response time data in Bürkner 2020. The model included the full random effects structure (random intercepts and slopes for participants and items), following the recommendation in Barr et al. 2013. The model also included the full fixed effect structure. Table 5.7 provides the posterior estimates of the response times plus 95% credible intervals for the Maze task, and Table 5.8 provides the posterior estimates of the response times plus 95% credible intervals for the self-paced reading task.

Task	Bias	mean	med	se	n
Maze	Object	2742	2480	70	56
Maze	Subject	2897	2654	59	56
SPR	Object	3129	2869	71	48
SPR	Subject	3355	3161	64	48

Table 5.6: Question response times computed over items for implicit causality in the Maze task and self-paced reading.

	estimate	est.error	2.5%	97.5%
Intercept	7.99	0.04	7.92	8.06
Dependency	-0.02	0.02	-0.07	0.03
Bias	-0.07	0.02	-0.12	-0.02
Bias:Dependency	-0.03	0.05	-0.13	0.08
sigma	0.35	0.01	0.34	0.37
ndt	26.00	23.54	0.72	87.58

Table 5.7: Bayesian model estimates for implicit causality question response times, Maze task.

The Bayesian models show no significant effect of **Dependency**, revealing that there was no significant difference in the time participants took to respond to questions across

	estimate	est.error	2.5%	97.5%
Intercept	7.80	0.04	7.71	7.88
Dependency	-0.02	0.03	-0.08	0.03
Bias	-0.06	0.03	-0.11	-0.01
Bias:Dependency	-0.01	0.06	-0.13	0.10
sigma	0.44	0.01	0.42	0.47
ndt	70.21	50.35	2.81	185.86

Table 5.8: Bayesian model estimates for implicit causality question response times, self-paced reading task.

the Anaphora and Ellipsis conditions. However, both models show a significant main effect for **Bias**, such that the Subject bias questions were responded to more slowly than the Object bias questions.² This main effect of Bias suggests that not only were subject interpretations less frequent, but that participants were also slightly slower when responding to questions requiring an interpretation of a sentence containing a subject-biased verb.

This finding invites the question of whether the slower response times in the subject bias items was dependent upon the answer a participant chose relative to the bias of the item’s verb. We saw that participants only chose a subject interpretation about 55% of the time when the item contained a subject-biased verb, so it is possible that participants’ response time was dependent upon whether they ultimately chose a subject interpretation or an object interpretation. To probe this question, I ran additional Bayesian models that included participants’ responses. The variable was coded as Congruent, for which 0 indicates a choice incongruent with the bias of the verb in the item (i.e. a subject interpretation for an item containing an object-biased verb, or an object interpretation for an item containing a subject-biased verb), and 1 indicates a choice congruent with

²Bias levels were contrast coded .5 for Object, -.5 for Subject.

the bias of the verb in the item (i.e. a subjective interpretation for an item containing a subject-biased verb, or an object interpretation for an item containing an object-biased verb). The procedure for running the model was identical to that described for in the previous analysis.

The results for the Maze experiment are given in Table 5.9. The effect of **Congruent** failed to reach significance, indicating that participants' response times in the Maze task were not significantly impacted by whether their choice of interpretation was consistent or inconsistent with the verb bias of the sentence.³

	estimate	est.error	2.5%	97.5%
Intercept	8.01	0.04	7.93	8.08
Dependency	-0.02	0.04	-0.10	0.06
Bias	-0.07	0.04	-0.15	0.01
Congruent	-0.03	0.03	-0.08	0.03
Bias:Dependency	-0.04	0.08	-0.19	0.12
Dependency:Congruent	0.00	0.05	-0.09	0.10
Bias:Congruent	0.01	0.05	-0.08	0.10
Dependency:Bias:Congruent	0.01	0.09	-0.17	0.20
sigma	0.35	0.01	0.33	0.37
ndt	27.99	24.83	0.87	92.88

Table 5.9: Bayesian model estimates for implicit causality question response times including participant response, implicit causality in the Maze task.

Analysis of the SPR data revealed a slightly different result. Table 5.10 shows the results of the Bayesian model. Unlike in the Maze task, the results of the SPR task reveal a significant main effect of **Congruent**. This effect tells us that participants were slower to respond to the SPR task questions when their response was incongruent with the bias of the item's verb. No interactions reached significance, including the interaction of

³Choice of NP was also tested (i.e. Subject or Object choice regardless of Bias), but the effect failed to reach significance.

Congruent:Bias, indicating that this slow-down was not greater within the subject or object biased items. That is, participants were slower to respond when their answer was incongruent with the bias of the item’s verb, regardless of the verb bias direction.

Another notable result of the model is that the main effect of Bias no longer reaches significance. This result indicates that the model variance previously captured under Bias is now captured by the main effect of Congruent. This finding is likely due to the fact that the subject-bias items received an incongruent response about 45% of the time, while the object-bias items received an incongruent response only about 25% of the time. Incongruent responses were therefore over-represented in the Subject conditions of the Bias factor. This overrepresentation of incongruent responses in the Subject items relative to the Object items led to the main effect of Bias seen in Table 5.8. However, the current analysis shows that this effect was actually driven by the slow-down in incongruent responses, not by the Subject items themselves.

	estimate	est.error	2.5%	97.5%
Intercept	7.93	0.04	7.84	8.02
Dependency	-0.01	0.05	-0.10	0.08
Bias	0.01	0.04	-0.08	0.10
Congruent	-0.19	0.03	-0.26	-0.13
Bias:Dependency	-0.05	0.09	-0.23	0.13
Dependency:Congruent	-0.02	0.05	-0.13	0.09
Bias:Congruent	-0.07	0.05	-0.17	0.04
Dependency:Bias:Congruent	0.06	0.11	-0.16	0.27
sigma	0.44	0.01	0.41	0.47
ndt	85.22	56.73	4.16	210.46

Table 5.10: Bayesian model estimates for implicit causality question response times including participant response, implicit causality in self-paced reading task.

In conclusion, the analysis shows that participants responded to questions more slowly in the SPR task when their response was incongruent with the bias of the verb.

Because more incongruent responses were given for the Subject-bias items than the Object-bias verbs, the Subject-bias items were overall responded to more slowly than the Object-bias verb items. A similar pattern was seen in the Maze task; however, the effect of Congruent, the effect of providing an incongruent or congruent response, did not reach significance as it did in the SPR task. The Maze analysis therefore only shows a significantly slower response time for the Subject-bias items than the Object-bias items. The failure to find a significant effect of Congruent in the Maze task may be a factor of the number of observations. As discussed earlier, the SPR task contained 4 additional items, leading to a total of 1,514 observations overall compared to the Maze task's 1,329 total observations. It is possible, then, that the Maze task lacks the power to show the additional effect of Congruent. This hypothesis is supported by the fact that the two experiments showed otherwise harmonious results.

Finally, there is the question of what is responsible for the response slow-down. The increase in response time could be reflective of several things. It may reflect greater decision time on the part of participants, meaning that participants wavered on their choice longer in these items. This could potentially indicate that participants did not choose an interpretation of the NPE site/anaphor when it was first encountered, and therefore required greater processing time at the question, when they were required to choose an interpretation of the NPE site/anaphor. It could also indicate that they had chosen an interpretation when they encountered the NPE site/anaphor, but that they reconsidered their choice or altered their interpretation at the question. The current experiment did not capture the online processing data necessary to adjudicate between these possibilities. Another possibility is that the slower response time indicated greater uncertainty in the participants' responses. A study containing confidence ratings could assess this possibility.

5.3.2.2 Origin of object response bias

We have seen so far that participants in both the Maze and the SPR tasks chose object NP interpretations more often than they chose subject NP interpretations. There are several possible explanations for this pattern. One possible explanation is that the results are due to the nature of the specific verbs chosen for subject and object biases. However, this explanation is unlikely as the biases in the verbs were balanced according to the findings of Hartshorne and Snedeker (2013). Another possibility is that the results are due to bias in the nouns chosen. Although the items were balanced so that nouns were swapped between subject and object position between two lists, individual preferences of nouns could be averaging out such that strong preferences in either direction are being masked by the overall proportion reported. For example, perhaps baritones are always interpreted as more competitive than altos, regardless of the noun's position as subject or object. To rule out this possibility, it is worth taking a look at the by-item responses. Figure 5.6 shows the responses of the self-paced reading items separated by list. If the properties of specific nouns for items were driving differences in interpretation, we would expect to see the empty and full circles within each color diverging. However, the patterns across the two lists are overall consistent across items. This suggests that there are no strong interpretation preferences driven by individual nouns.

Another possibility is that the overall proportion is masking stronger preferences within individual items. It is once again worth taking a look at the individual items. Figure 5.7 plots the implicit causality data combined across the Maze and SPR experiments. The plot shows that the object-biased conditions, shown in blue, are overall given a greater interpretation of subject interpretations. The subject-biased conditions, shown in pink, are given overall given a greater interpretation of subject interpretations. While a couple items buck this overall trend by clustering together – notably 107, 109, 116, and 120 –

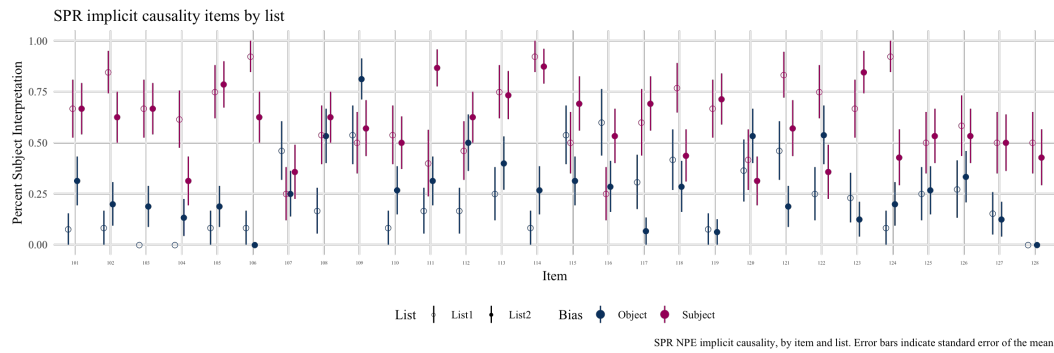


Figure 5.6: SPR implicit causality items, by list. Error bars indicate standard error of the mean.

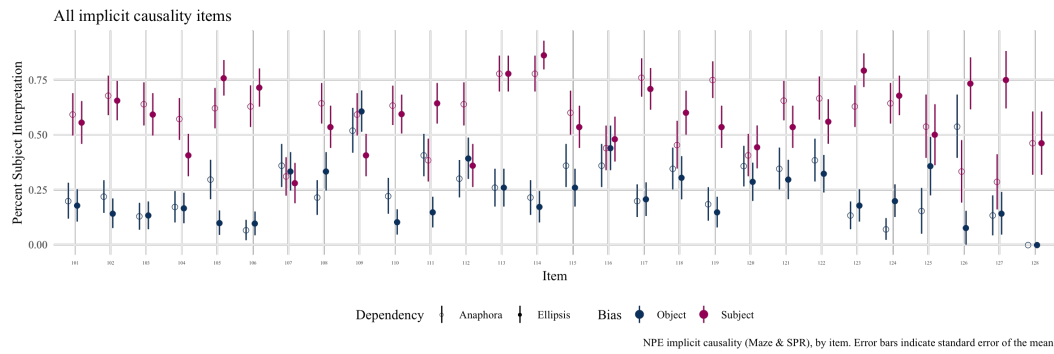


Figure 5.7: Maze and SPR implicit causality items, by list. N=2,852. Error bars indicate standard error of the mean.

there is not overall great item variation.⁴ The generalization that there is a slight object bias therefore remains unexplained by individual item variation.

Yet a further possible explanation is that the location of the critical NPE sites/pronoun in the experimental items affected the bias. The biased verbs taken from Hartshorne and Snedeker (2013) were normed in their experiments using pronouns in subject position, as shown in example (145) (Hartshorne and Snedeker, 2013, ex. 12).

(145) Sally frightens Mary because she is a dax.

⁴For those interested, these items can be found in the Appendix.

The critical NPE sites/pronouns in the experiments here were either a subject extracted from a relative clause in object position, as in (146), or the subject of an embedded clause under a raising-to-subject verb, as in (147).

(146) Photographers idolize musicians because it’s difficult to find any that don’t respect all artforms.

(147) Anthropologists like geologists because it seems that most are amiable.

Because grammatical role parallelism is known to affect pronominal resolution preferences, it is possible that the location of the NPE sites/pronouns in object position in the current experiment account for the object bias observed (Smyth, 1994; Lappin and Leass, 1994; Chambers and Smyth, 1998). To probe this possibility, I included the position of the NPE site/pronoun as a two-level factor in the Bayesian regression analysis of the response data in the Maze task. The results are shown in Table 5.11.

	estimate	est.error	2.5%	97.5%
Intercept	0.51	0.11	0.29	0.73
Bias	1.60	0.21	1.20	2.03
Dependency	0.13	0.15	-0.15	0.42
SyntRole	-0.10	0.17	-0.44	0.23
Bias:Dependency	0.33	0.30	-0.26	0.93

Table 5.11: Bayesian model estimates for implicit causality experiment, Maze task, including grammatical position of NPE site/pronoun.

As shown in Table 5.11, **SyntRole** failed to reach significance. This result suggests that it was not the syntactic position of the NPE site/pronoun that was driving the object bias. There are two additional reasons to believe it is unlikely that the syntactic position of the NPE site/pronoun is responsible for the object bias. The positioning of the site/pronoun was not completely controlled for in the experimental items, so object items are underrepresented relative to subject items (8 to 16, respectively). Additionally,

grammatical role parallelism is argued to hold of pronouns which have an antecedent not only in the same abstract syntactic position, but also in an analogous constituent structure and attachment site (Smyth, 1994). These parallelism constraints are not satisfied in the experimental items here.

There are some additional experimental designs that could be considered if one wished to follow up on this finding. One design would be to form a cataphoric condition of the items used in the current experiment. Such an experiment could test whether the relative preferences of subject and object interpretations are affected in the cataphoric version of the items, similar to what we observed in Chapter 3. The cataphoric versions of the current items are mildly marked, but not impossible: *It's hard to find {any/ones} that aren't competitive, therefore baritones resent altos*. Another design would be to have the ellipsis site/gap in subject position. This version of the design could directly test whether the positioning of the ellipsis site/anaphor affected interpretation preferences. Such items are again awkward, but maybe not impossible: *Baritones resent altos because {any/ones} that aren't competitive are hard to find*. I leave these possibilities open for future consideration.

To conclude this section, I have discussed several different possibilities for the origin of the object-bias preferences observed in the implicit causality experiments in this chapter. I also concluded that all seem unlikely to be the explanation for the effect. I believe it is therefore most likely that the object bias instead represents the same proximity bias observed in Experiment 2. Although all pronominal anaphora and ellipsis processes in these items were anaphoric and not cataphoric, thus not providing the possibility of showing a preference for the subject, the data pattern similarly to the findings of the **Anaphoric** condition in Experiment 2. The findings therefore provide additional support for a general proximity bias when the nouns are otherwise fairly matched for salience. The findings also show that the implicit causality bias is strong enough to overcome

this bias. However, the implicit causality bias is not strong enough to keep the subject interpretation in the subject-biased verbs from being overall slightly dragged toward an object interpretation. Most important for the current research question, however, is that, regardless of the origin of the object bias observed here, it holds equally of both the anaphoric pronouns and the NPE site interpretations.

Chapter 6

Conclusion

6.1 Why we should look across the aisle

An ellipsis site is, at the most basic level, simply a gap in sound. Ellipsis is arguably unique as an empirical phenomenon in the breadth of research that is undertaken on it. Having said as much, I will simplify slightly and propose that we can bin the bulk of ellipsis research into two camps, that done by formal theoreticians and that done by psycholinguists. Formal theoreticians – e.g. syntacticians, semanticists – are interested in ellipsis largely for what it can tell us about the grammar of natural language. They are interested in investigating questions about when and where ellipsis can occur, and how the grammar accounts for the restrictions on the available interpretations of ellipsis sites. Psycholinguists are interested in ellipsis largely for what it can tell us about the processing mechanisms used to comprehend natural language. They are interested in investigating how we build a meaning representation within an ellipsis site when the input signal is simply the absence of input, and what the mechanisms are that allow us to access those representations in memory. For psycholinguists, ellipsis, like anaphora, provides insights into how language users organize information in memory by telling us something about how we navigate a massive mismatch between overt linguistic form and meaning.

However, an unfortunate reality of the field as it currently stands is that these two camps of research have largely operated independently of one another. There is, perhaps, a common perception that there is no reason for researchers in one camp to consult the other; while both fields are working on the same empirical phenomenon, the goals of the two sides are distinct and therefore non-overlapping. This dissertation is an attempt to bring the two research fields together. I hope to convince researchers that each field can inform the other, and more importantly *should* be informing the other in some way.

There are two main threads running throughout this dissertation. One involves the nature of the relationship that must hold between an ellipsis site and its surrounding linguistic context. This question was covered mostly in Chapter 2, where I argued that bidirectional entailment accounts are too restrictive to account for the full range of possible interpretations of an ellipsis site. I argued against the prevailing view in the literature that the interpretation of an ellipsis site must be restricted to a strict matching relationship with a piece of previous linguistic structure. The requirements of such a theory simply do not account for the polarity reversal data, which I showed do not have a direct antecedent at any level of representation, syntactic, semantic, or pragmatic.

I furthermore presented findings in Chapter 3 that participants freely resolve an NPE site to a noun not mentioned in the discourse context, as long as the discourse context is constrained enough so that the noun is salient in the context. These findings are interesting not in suggesting that ellipsis sites can be freely interpreted, which I do not believe is true, but in what they tell us about the reasoning process language users undertake to resolve ellipsis sites. When no existing structural anaphor was sufficiently salient in the discourse context, participants resolved the ellipsis site to an entity that was salient in the context, e.g. through a relation with a salient QuD, even if that entity was not syntactically present in the context. The experimental findings strongly suggest that comprehenders attempt to resolve ellipsis sites to a salient representation in the discourse,

even if that representation doesn't have a structural antecedent.

The cases I present here are not the only known cases where identity fails to hold between a representation of an ellipsis site and a structural antecedent in the discourse. The literature is replete with counter-examples, which have been discussed throughout the dissertation. For example, see Hardt (1999), Ginzburg and Sag (2001), Dalrymple et al. (1991), Miller and Hemforth (2014), and Poppels and Kehler (2019) for compelling arguments about the challenges accounts relying on structural matching constraints. There is, furthermore, a small cottage industry in experimental work investigating the type and extent of syntactic mismatches that can occur under ellipsis (Frazier and Clifton, 2005; Arregui et al., 2006; Frazier, 2013; Frazier and Clifton, 2015; Frazier and Duff, 2019).

Nevertheless, it is uncontroversially true that bidirectional isomorphism theories account for the majority of ellipsis data. However, as linguists it is the outliers – or the tail end of the distribution – that interests us the most. These cases provide the most interesting insights into the grammatical mechanisms of building a representation of the ellipsis site and the processing mechanisms that support our ability to do so. The focus of the field on bidirectional isomorphism theories risks trapping research at a local minimum, when a more elegant solution may be over the next hill.

This leads me to the second thread running throughout this dissertation, which is the investigation of how comprehenders build their interpretations of ellipsis sites. This question is closely related to the previous one. Let's assume that ellipsis interpretation is constrained by grammatical licensing conditions on its location, e.g., there are grammatical constraints on which syntactic heads and/or grammatical features license an ellipsis site (to take Merchant's E feature as an example). I take this point to be uncontroversial, as there appears to be a limited number of ellipsis constructions available in natural language. Let's also assume, slightly more controversially, that there is also

a grammatical restriction on what the content of the ellipsis site can be relative to the content of the surrounding linguistic context. This grammatical constraint would be an interpretation constraint, given, for example, in Merchant's e-GIVENness, or in the theory of Local Givenness advocated for here.

I believe it is also uncontroversial to say that the process by which language users build a representation of the ellipsis site during comprehension lays outside of grammatical theory, squarely in the domain of processing. Research regarding the properties of the processor and the memory architecture involved in building real-time linguistic representations is exemplified in theories such as those of cue-based retrieval in Martin and McElree (2008, 2009, 2011) and Harris (2015, 2019). Given this picture, the interesting and I believe as-yet-unresolved question is the way in which our grammatical theory of ellipsis interpretation interacts with our theory of comprehension and structure building within the ellipsis site. If we do indeed have a grammatical interpretation rule, then we would expect our structure-building and comprehension processes to be constrained by this rule. For example, let's consider e-GIVENness as our grammatical rule. We could reasonably expect the processor's search, once encountering an ellipsis site, to be constrained to available discourse representations that satisfy the grammatical constraint.¹ Such a theory would predict, for example, that the search mechanisms would fail to provide an interpretation for polarity reversal sluices, as they do not satisfy the constraint of e-GIVENness. Such a theory can even allow for some mild structural repair, as proposed by Arregui et al. (2006).

However, all existing theories of ellipsis acknowledge that the reality is not so simple. As discussed in Chapter 2, all theories of ellipsis include some notion of salience, since it is well known that ellipsis sites are not always felicitous just because they match some

¹Whether the search itself is constrained or whether only the representations returned are constrained is an empirical question that we have yet to answer.

given preceding structure in the linguistic context. For example, recall Barros (2014)'s observation that (25), repeated in (148), is acceptable when the remnant *who* is replaced with the d-linked expression *which man*.

(148) Joe, [who once killed a man in cold blood]_A, doesn't even remember
{#who/which man} [~~he once killed in cold blood~~]_E.

My proposal here is that our grammatical theory of ellipsis interpretation cannot and should not capture this fact. This is why this conclusion does not contain an updated version of Local Givenness, stating something like 'resolve the ellipsis site to the most recent antecedent in anaphoric ellipsis, but the first encountered postcedent in cataphoric ellipsis.' To do so would miss the point, which is that these observations are subsumed under a notion of representational salience. Salience, while almost certainly influenced by grammatical factors (e.g. subjecthood), is not a grammatical rule; it is a product of a general cognitive architecture of which language is only one part.

Here is the overall utility of experimental work for theoreticians. Once we understand the processing facts, these facts can be safely left out of a theoretical account of ellipsis. To make this point more strongly: not everything belongs in the grammar, and not everything is a nail just because we have a hammer. The question of grammatical interpretation is an important one, but to conflate grammatical rules with processing phenomena undermines the goal of understanding the grammatical properties of ellipsis. I am not the first to warn of the consequences of conflating the grammar with parsing. Frazier (2013), for example, argues that the parser, and not the grammar, is a more efficient way of accounting for mismatches between an ellipsis site and an antecedent. Though I disagree with Frazier's particular conclusions regarding the extent of the parser's ability to alleviate mismatches, the conclusions that Frazier argues for, that linguistic data is not always characterized only by the grammar, is analogous to the point being made here.

Despite Frazier's efforts, research teasing apart the parser from the grammar has mostly been concentrated on syntactic phenomenon, and within that largely on island violations specifically. For example, Phillips (1996) covers the history of this line of research and makes specific arguments in favor of collapsing the grammar and the parser into a single mechanism. See Hofmeister (2007); Hofmeister et al. (2013) for arguments that the locus of island effects is found in processing mechanisms, not due in grammatical constraints, and Sprouse et al. (2012) for arguments that island effects are indeed most likely due to grammatical constraints.

Such a debate delineating the role of grammatical constraints and processing constraints in ellipsis does not yet exist in earnest, although it should. To take a concrete example, consider theories of sluicing built around the proposed ungrammaticality of sluicing out of discourse subordinated clauses. Such theories are at best missing the point (which is that the unacceptability of these constructions are the result of processing factors, not grammatical restrictions), and at worst leading the intellectual pursuit astray. In summary, I'm not arguing that theoreticians must care about the exact processing models of ellipsis; but, there are many reasons why an ellipsis interpretation can be unacceptable, and not all are grammatical in origin. Lumping everything into the grammar may only distort the picture and preclude the possibility of creating a unifying, explanatory theory.

6.2 Truth and Beauty?

Is there a way, then, to unify our theories of grammatical licensing, grammatical interpretation, and processing? This section outlines what I believe to be the most promising approach, given what we know about ellipsis at this moment.

One question that I have mostly taken for granted in the discussions here is the presence of syntactic material in the ellipsis site. Merchant's 2001 dissertation has been

hugely influential in moving the field toward an ellipsis theory in which syntactic material is present in the ellipsis site. However, recent cracks have appeared in the unequivocal picture painted by Merchant. For example, the case-licensing facts, which Merchant heavily relies on to show that a continuous syntactic structure is present in an ellipsis site, are not as clear as previously believed. See Thoms 2019 for an overview of the existing literature, but the basic takeaway is that morphological case assigned to a remnant is *not* always identical with the case that would be assigned were there no ellipsis in an otherwise identical clause. The picture therefore seems more complicated than simply chopping off a syntactic clause at PF.

So where do we start? I have already outlined my theory of the grammatical interpretation constraint on ellipsis, in Chapter 2. So here I will focus on the processing picture. Let's start then with Harris (2015)'s proposal of the processing mechanism of ellipsis given in (149), formulated for sluicing, specifically.

(149) Basic tasks of the processor in clausal ellipsis:

1. Parse the remnant by constructing the appropriate phrase structure for the remnant given the input.
2. Locate the correlate, if any, from the antecedent clause and pair it with the remnant.
3. Construct the elided phrase by regenerating or copying a structure at Logical Form.

While steps (1) and (2) are specific to sluicing, it is reasonable to conclude that the first steps of processing any ellipsis site are parsing the overt structure, and then searching for the appropriate representation for the ellipsis site, either through a backward search for anaphoric ellipsis or a forward search for cataphoric ellipsis. For cataphoric ellipsis, I have argued that incoming representations are evaluated in real-time and matched via

cues to the ellipsis site gap. Once an appropriate representation is found, it is immediately integrated into the representational gap in memory. The interpretations for cataphoric ellipsis therefore come as a complete package, with syntactic, semantic, and pragmatic representations together.

However, for anaphoric ellipsis the appropriate representations must be retrieved from memory. While the exact properties of the representations themselves are not yet known, the existing literature supports a view under which ellipsis sites do not involve retrieval of full syntactic representations from memory (Poirier et al., 2010; Martin and McElree, 2011; Miller, 2016). For example, Poirier et al. found evidence for partial reactivation of semantic material 500ms after a sluicing site. They give evidence from a cross modal lexical priming study in which they report priming effects for the object of a sluicing antecedent, but not the subject. These findings unfortunately do not adjudicate on the properties of the syntactic material present in the ellipsis site, but suggest some amount of partial reactivation of previously encountered material. Xiang et al. (2014) show evidence of syntactic reconstruction at a sluicing site by showing that sluicing sites given rise to syntactic priming of dative/double object constructions, therefore providing experimental evidence of syntactic material within the ellipsis site. Miller (2016) investigates the question of syntactic material in the ellipsis site using evidence from agreement attraction. Miller argues based on her findings that there is limited, syntactic reactivation inside NPE sites. These studies provide initial evidence that some syntactic structure is present in the ellipsis site; however, they do not provide evidence that a full syntactic representation is recovered.

I propose the processing mechanism of ellipsis given in (150), formulated in the spirit of Harris's process, but slightly modified and generalized. One important caveat I will make at this point is that this is the process argued to take place on the comprehension end; whether or not a syntactic structure is also constructed in the production of an

ellipsis construction is unknown.

(150) Basic tasks of the processor in anaphoric ellipsis:

1. Identify the ellipsis gap site.
2. Determine the cues necessary for retrieval of the appropriate antecedent from memory, if the appropriate antecedent is not in focal attention.
3. Retrieve the appropriate discourse representation from memory.
4. Construct the elided phrase by regenerating syntactic structure to align with the selected discourse representation.

The remainder of the discussion will focus on steps 3 and 4. I propose that step 3 involves only the retrieval of a discourse representation, e.g. a discourse entity or a discourse segment (Webber, 1988), which does not include a syntactic representation. The idea that syntactic information degrades quickly in memory has been known for some time. Recent theories of verbatim recall propose that language users shuttle information about the syntactic form of incoming material from immediate memory to declarative memory on the scale of hundreds of milliseconds (Broadbent, 1958; McElree et al., 2003). For a linguistic expression of any length, therefore, there is no single encoding in memory, but many smaller bits. When tasked with the immediate recall of an expression, comprehenders may use a “blueprint” that remains in short term memory in order to put back together the various encodings of the expression that have been moved to long-term memory. However, this short-term memory blueprint is quickly overwritten as new, similar sentences are encountered (Wagers, 2014). Therefore, while comprehenders may encode the exact form of utterances they encounter, they lose the ability to recall these exact forms from memory as time elapses and new structures intervene.

Various researchers have provided psycholinguistic evidence on this fallibility of structural memory, also called verbatim memory. Sachs (1967), Jarvella (1971), Bransford

et al. (1972), and Garrod and Trabasso (1973) all report evidence in support of the claim that memory for syntactic structure is short-lived. For example, Sachs found that when test sentences occurred immediately after an original sentence, participants reliably discriminated the original sentence from semantically and syntactically altered test sentences. However, once intervening material was added – at 80 and 160 syllables – participants could discriminate the semantically alternated test sentences but not the syntactically altered test sentences. Sachs concluded from this observation that language users’ memory for syntactic structure fades more rapidly than memory for the meaning of a linguistic expression.

At present, this concerns us only in that it establishes that semantic and pragmatic representations of linguistic expressions are longer-lasting and more accurately retrieved than syntactic representations, showing that the two are not inextricably linked in memory. If we are accessing only discourse representations in ellipsis, then we have explained why ellipsis sites can be relatively far from their antecedents, and why ellipsis interpretations can appear to have non-continuous syntactic antecedents, as in split-antecedents (see discussion in 3.3.4.3). We have also explained why sluices like the polarity reversal sluices are possible, as discourse representations have long been understood to not need strict syntactic analogues (van Dijk and Kintsch, 1983; Webber, 1988). We also get the repair performance of Arregui et al. (2006); Frazier (2013) for free, without a need to posit special operations that take place at Logical Form.

However, I have discussed convincing grammatical and processing evidence that some syntactic structure is inside an ellipsis site. This is where step 4 comes in. I propose, following similar ideas in Belanger (2014) and Harris (2015), that the syntactic structure is not copied or reactivated from an antecedent structure, but is regenerated after the discourse representation is accessed. Arguments for reactivation of syntactic structure in ellipsis sites (Shapiro et al., 2003; Poirier et al., 2010; Miller, 2016) draw from previous

studies looking at the reactivation of antecedent content at the location of a pronoun, and the reactivation of the content of a *wh*-phrase at its corresponding gap (McElree and Bever 1989; Nicol et al. 1994, a.o.). Reactivation in these constructions entails creating a link from a pronoun to its antecedent representation in memory, or from a gap to its corresponding *wh*-phrase in memory. Because the process relies on generating a link from a recently parsed representation to a previously encountered representation, e.g. a pronoun to its antecedent, reactivation necessarily uses an existing representation in memory. This restriction in turn constrains the structure and content of the reactivated material to that of a previously encountered structure. To argue that ellipsis sites contain reactivated syntactic material, then, is to constrain the syntactic material within the ellipsis site to that of a strict structural antecedent, modulo any minor repair processes.

Regeneration is distinct from reactivation or copying procedures because it does not require that the syntactic structure be identical to some previous structure in memory. Belanger draws from work in Potter and Lombardi 1990, 1992, 1998, which shows that, while the verbatim or structural/syntactic representation of a perceived sentence degrades quickly, a conceptual representation persists along with recently activated lexical items from the sentence. Belanger proposes that structure in the ellipsis site contains re-generated material that is primed by these lexical items. Thus, there is no reactivation or copying of syntactic material, which means that the syntactic material in the ellipsis site need not match any syntactic structure in the preceding linguistic context. However, the ellipsis site can draw on and be primed by recently activated lexical items, which can be quickly accessed and used to construct a structural representation of the ellipsis site.

Belanger conducts several lure experiments to test these hypotheses, but they fail to convincingly show empirical evidence for the claims made in the dissertation. However, I believe that, though the theory is challenging to prove experimentally, it is ultimately the

correct story. The regeneration of syntactic structure within the ellipsis site provides the flexibility demanded by ellipsis researchers who have argued for anaphoric accounts, but also explains persistent syntactic connectivity effects that suggest a syntactic relationship between the material inside an ellipsis site and its surrounding linguistic structure. To show this reliably involves further careful investigation of the interaction between our grammatical interpretation constraint on ellipsis and the processing mechanisms responsible for building our representations of ellipsis sites. I believe that only when we look at the full range of evidence – from grammatical theory and experimental investigation – will we build an explanatory theory of ellipsis.

Appendix A

Experiment 1 items

It.	Cond.	Item
1	AA	Honeycrisp apples would taste great with navel oranges for breakfast, but I don't know if we can find any right now.
1	CA	I don't know if we can find any right now, but honeycrisp apples would taste great with navel oranges for breakfast.
1	AS	Honeycrisp apples would taste great with this navel orange I have for breakfast, but I don't know if we can find any right now.
1	CS	I don't know if we can find any right now, but honeycrisp apples would taste great with this navel orange I have for breakfast.
1	AO	This honeycrisp apple I have would taste great with navel oranges for breakfast, but I don't know if we can find any right now.
1	CO	I don't know if we can find any right now, but this honeycrisp apple I have would taste great with navel oranges for breakfast.
2	AA	Acorns would look great with tulips for the head table, but I'm not sure we can get any by tonight.
2	CA	I'm not sure we can get any by tonight, but acorns would look great with tulips for the head table.
2	AS	Acorns would look great with this tulip of yours for the head table, but I'm not sure we can get any by tonight.
2	CS	I'm not sure we can get any by tonight, but acorns would look great with this tulip of yours for the head table.
2	AO	This acorn of yours would look great with tulips for the head table, but I'm not sure we can get any by tonight.
2	CO	I'm not sure we can get any by tonight, but this acorn of yours would look great with tulips for the head table.
3	AA	Balloons would look good next to streamers for the party, but I don't remember if we have any for tomorrow.
3	CA	I don't remember if we have any for tomorrow, but balloons would look good next to streamers for the party.
3	AS	Balloons would look good next to this streamer of mine for the party, but I don't remember if we have any for tomorrow.
3	CS	I don't remember if we have any for tomorrow, but balloons would look good next to this streamer of mine for the party.
3	AO	This balloon of mine would look good next to streamers for the party, but I don't remember if we have any for tomorrow.
3	CO	I don't remember if we have any for tomorrow, but this balloon of mine would look good next to streamers for the party.
4	AA	Eggplants would look cool beside pineapples for the mantelpiece, but I forget if we need any this time.

- 4 CA I forget if we need any this time, but eggplants would look cool beside pineapples for the mantelpiece.
- 4 AS Eggplants would look cool beside my pineapple for the mantelpiece, but I forget if we need any this time.
- 4 CS I forget if we need any this time, but eggplants would look cool beside my pineapple for the mantelpiece.
- 4 AO My eggplant I have would look cool beside pineapples for the mantelpiece, but I forget if we need any this time.
- 4 CO I forget if we need any this time, but my eggplant would look cool beside pineapples for the mantelpiece.
- 5 AA Pears would taste delicious with bananas for dinner, if we can get any in time.
- 5 CA If we can get any in time, pears would taste delicious with bananas for dinner.
- 5 AS Pears would taste delicious with your banana for dinner, if we can get any in time.
- 5 CS If we can get any in time, pears would taste delicious with your banana for dinner.
- 5 AO Your pear would taste delicious with bananas for dinner, if we can get any in time.
- 5 CO If we can get any in time, your pear would taste delicious with bananas for dinner.
- 6 AA Clarinets would sound good with flutes during the reception, if we can find some by this evening.
- 6 CA If we can find some by this evening, clarinets would sound good with flutes during the reception.
- 6 AS Clarinets would sound good with my flute during the reception, if we can find some by this evening.
- 6 CS If we can find some by this evening, clarinets would sound good with my flute during the reception.
- 6 AO My clarinet would sound good with flutes during the reception, if we can find some by this evening.
- 6 CO If we can find some by this evening, my clarinet would sound good with flutes during the reception.
- 7 AA Sofas would go well with tables for the showing, but I don't know if we can buy any by this weekend.
- 7 CA I don't know if we can buy any by this weekend, but sofas would go well with tables for the showing.
- 7 AS Sofas would go well with my table for the showing, but I don't know if we can buy any by this weekend.
- 7 CS I don't know if we can buy any by this weekend, but sofas would go well with my table for the showing.
- 7 AO My sofa would go well with tables for the showing, but I don't know if we can buy any by this weekend.
- 7 CO I don't know if we can buy any by this weekend, but my sofa would go well with tables for the showing.
- 8 AA Calico kittens would look cute with pug puppies for the party, but I don't know how many we'd need.
- 8 CA I don't know how many we'd need, but calico kittens would look cute with pug puppies for the party.
- 8 AS Calico kittens would look cute with your pug puppy for the party, but I don't know how many we'd need.
- 8 CS I don't know how many we'd need, but calico kittens would look cute with your pug puppy for the party.
- 8 AO Your calico kitten would look cute with pug puppies for the party, but I don't know how many we'd need.
- 8 CO I don't know how many we'd need, but your calico kitten would look cute with pug puppies for the party.
- 9 AA Newts can't live with fighting fish in an aquarium, despite my wanting some really badly.
- 9 CA Despite my wanting some really badly, newts can't live with fighting fish in an aquarium.

- 9 AS Newts can't live with this fighting fish of mine in an aquarium, despite my wanting some really badly.
- 9 CS Despite my wanting some really badly, newts can't live with this fighting fish of mine in an aquarium.
- 9 AO This newt of mine can't live with fighting fish in an aquarium, despite my wanting some really badly.
- 9 CO Despite my wanting some really badly, this newt of mine can't live with fighting fish in an aquarium.
- 10 AA Serving spoons would work well with large platters for lunch, but I don't know if we should get any for next week.
- 10 CA I don't know if we should get any for next week, but serving spoons would work well with large platters for lunch.
- 10 AS Serving spoons would work well with your large platter for lunch, but I don't know if we should get any for next week.
- 10 CS I don't know if we should get any for next week, but serving spoons would work well with your large platter for lunch.
- 10 AO Your serving spoon would work well with large platters for lunch, but I don't know if we should get any for next week.
- 10 CO I don't know if we should get any for next week, but your serving spoon would work well with large platters for lunch.
- 11 AA Staplers would work well with notebooks for the meeting, but I don't remember if we can get any this time of day.
- 11 CA I don't remember if we can get any this time of day, but staplers would work well with notebooks for the meeting.
- 11 AS Staplers would work well with my notebook for the meeting, but I don't remember if we can get any this time of day.
- 11 CS I don't remember if we can get any this time of day, but staplers would work well with my notebook for the meeting.
- 11 AO My stapler would work well with notebooks for the meeting, but I don't remember if we can get any this time of day.
- 11 CO I don't remember if we can get any this time of day, but my stapler would work well with notebooks for the meeting.
- 12 AA Horses would live well with donkeys for long periods, but I'm not sure if you want any for the farm.
- 12 CA I'm not sure if you want any for the farm, but horses would live well with donkeys for long periods.
- 12 AS Horses would live well with your donkey for long periods, but I'm not sure if you want any for the farm.
- 12 CS I'm not sure if you want any for the farm, but horses would live well with your donkey for long periods.
- 12 AO Your horse would live well with donkeys for long periods, but I'm not sure if you want any for the farm.
- 12 CO I'm not sure if you want any for the farm, but your horse would live well with donkeys for long periods.
- 13 AA Roosters would get along with cows pretty well, but I forget if you can find any locally.
- 13 CA I forget if you can find any locally, but roosters would get along with cows pretty well.
- 13 AS Roosters would get along with this cow of ours pretty well, but I forget if you can find any locally.
- 13 CS I forget if you can find any locally, but roosters would get along with this cow of ours pretty well.
- 13 AO This rooster of ours would get along with cows pretty well, but I forget if you can find any locally.
- 13 CO I forget if you can find any locally, but this rooster of ours would get along with cows pretty well.

- 14 AA Guitars would sound great with tubas for the quartet, if you can buy any in time.
- 14 CA If you can buy any in time, guitars would sound great with tubas for the quartet.
- 14 AS Guitars would sound great with your tuba for the quartet, if you can buy any in time.
- 14 CS If you can buy any in time, guitars would sound great with your tuba for the quartet.
- 14 AO Your guitar would sound great with tubas for the quartet, if you can buy any in time.
- 14 CO If you can buy any in time, your guitar would sound great with tubas for the quartet.
- 15 AA Penguins would live well with ducks, in case you want to buy some when we go out later.
- 15 CA In case you want to buy some when we go out later, penguins would live well with ducks.
- 15 AS Penguins would live well with the zoo's duck, in case you want to buy some when we go out later.
- 15 CS In case you want to buy some when we go out later, penguins would live well with the zoo's duck.
- 15 AO The zoo's penguin would live well with ducks, in case you want to buy some when we go out later.
- 15 CO In case you want to buy some when we go out later, the zoo's penguin would live well with ducks.
- 16 AA Stuffed animals would look cute alongside dolls for the shelf, but I'm not sure if we have any in the house.
- 16 CA I'm not sure if we have any in the house, but stuffed animals would look cute alongside dolls for the shelf.
- 16 AS Stuffed animals would look cute alongside your doll for the shelf, but I'm not sure if we have any in the house.
- 16 CS I'm not sure if we have any in the house, but stuffed animals would look cute alongside your doll for the shelf.
- 16 AO Your stuffed animal would look cute alongside dolls for the shelf, but I'm not sure if we have any in the house.
- 16 CO I'm not sure if we have any in the house, but your stuffed animal would look cute alongside dolls for the shelf.
- 17 AA Willow trees would look good with rose bushes alongside your house, but I don't know if you want to plant any right now.
- 17 CA I don't know if you want to plant any right now, but willow trees would look good with rose bushes alongside your house.
- 17 AS Willow trees would look good with your rose bush alongside your house, but I don't know if you want to plant any right now.
- 17 CS I don't know if you want to plant any right now, but willow trees would look good with your rose bush alongside your house.
- 17 AO Your willow tree would look good with rose bushes alongside your house, but I don't know if you want to plant any right now.
- 17 CO I don't know if you want to plant any right now, but your willow tree would look good with rose bushes alongside your house.
- 18 AA Parent wolves are apparently always near their cubs at night, even though I haven't seen any yet.
- 18 CA Even though I haven't seen any yet, parent wolves are apparently always near their cubs at night.
- 18 AS Parent wolves are apparently always near their cub at night, even though I haven't seen any yet.
- 18 CS Even though I haven't seen any yet, parent wolves are apparently always near their cub at night.
- 18 AO A parent wolf is apparently always near its cubs at night, even though I haven't seen any yet.
- 18 CO Even though I haven't seen any yet, a parent wolf is apparently always near its cubs at night.
- 19 AA Marshmallows will be served with Hershey bars tonight, despite you forgetting to buy any earlier.

- 19 CA Despite you forgetting to buy any earlier, marshmallows will be served with Hershey bars tonight.
- 19 AS Marshmallows will be served with this Hershey bar I have tonight, despite you forgetting to buy any earlier.
- 19 CS Despite you forgetting to buy any earlier, marshmallows will be served with this Hershey bar I have tonight.
- 19 AO This marshmallow I have will be served with Hershey bars tonight, despite you forgetting to buy any earlier.
- 19 CO Despite you forgetting to buy any earlier, this marshmallow I have will be served with Hershey bars tonight.
- 20 AA Seastars have lived happily with sea turtles before, but I don't remember if there are any in the aquarium right now.
- 20 CA I don't remember if there are any in the aquarium right now, but seastars have lived happily with sea turtles before.
- 20 AS Seastars live happily with our sea turtle before, but I don't remember if there are any in the aquarium right now.
- 20 CS I don't remember if there are any in the aquarium right now, but seastars have lived happily with our sea turtle before.
- 20 AO Our seastar has lived happily with sea turtles before, but I don't remember if there are any in the aquarium right now.
- 20 CO I don't remember if there are any in the aquarium right now, but our seastar has lived happily with turtles before.
- 21 AA Antique xylophones would sound nice played with bassoons for the orchestra, but I forget if you can buy any around town.
- 21 CA I forget if you can buy any around town, but antique xylophones would sound nice played with bassoons for the orchestra.
- 21 AS Antique xylophones would sound nice played with your bassoon for the orchestra, but I forget if you can buy any around town.
- 21 CS I forget if you can buy any around town, but antique xylophones would sound nice played with your bassoon for the orchestra.
- 21 AO Your antique xylophone would sound nice played with bassoons for the orchestra, but I forget if you can buy any around town.
- 21 CO I forget if you can buy any around town, but your antique xylophone would sound nice played with bassoons for the orchestra.
- 22 AA Grizzly bears are usually close to bobcats in the spring, if you want to see any while you're in Yellowstone.
- 22 CA If you want to see any while you're in Yellowstone, grizzly bears are usually close to bobcats in the spring.
- 22 AS Grizzly bears are usually close to the park's oldest bobcat in the spring, if you want to see any while you're in Yellowstone.
- 22 CS If you want to see any while you're in Yellowstone, grizzly bears are usually close to the park's oldest bobcat in the spring.
- 22 AO The park's oldest grizzly bear is usually close to bobcats in the spring, if you want to see any while you're in Yellowstone.
- 22 CO If you want to see any while you're in Yellowstone, the park's oldest grizzly bear is usually close to bobcats in the spring.
- 23 AA Octopuses usually live with eels at the Aquarium, but I don't believe there are any in this tank.
- 23 CA I don't believe there are any in this tank, but octopuses usually live with eels at the Aquarium.
- 23 AS Octopuses usually live with your eel at the Aquarium, but I don't believe there are any in this tank.
- 23 CS I don't believe there are any in this tank, but octopuses usually live with your eel at the Aquarium.
- 23 AO Your octopus usually lives with eels at the Aquarium, but I don't believe there are any in this tank.

- 23 CO I don't believe there are any in this tank, but your octopus usually lives with eels at the Aquarium.
- 24 AA Dirt bikes would be fun with four-wheelers for the trip, but I'm not sure we have any in the garage.
- 24 CA I'm not sure we have any in the garage, but dirt bikes would be fun with four-wheelers for the trip.
- 24 AS Dirt bikes would be fun with our four-wheeler for the trip, but I'm not sure we have any in the garage.
- 24 CS I'm not sure we have any in the garage, but dirt bikes would be fun with our four-wheeler for the trip.
- 24 AO Our dirt bike would be fun with four-wheelers for the trip, but I'm not sure we have any in the garage.
- 24 CO I'm not sure we have any in the garage, but our dirt bike would be fun with four-wheelers for the trip.
- 25 AA Otters often swim with harbor seals in Fall, if you want to spot any from your canoe.
- 25 CA If you want to spot any from your canoe, otters often swim with harbor seals in Fall.
- 25 AS Otters often swim with my favorite harbor seal in Fall, if you want to spot any from your canoe.
- 25 CS If you want to spot any from your canoe, otters often swim with my favorite harbor seal in Fall.
- 25 AO My favorite otter often swims with harbor seals in Fall, if you want to spot any from your canoe.
- 25 CO If you want to spot any from your canoe, my favorite otter often swims with harbor seals in Fall.
- 26 AA Sunflowers won't grow under apple trees in our climate, despite my wanting to plant some in the garden.
- 26 CA Despite my wanting to plant some in the garden, sunflowers won't grow under apple trees in our climate.
- 26 AS Sunflowers won't grow under our apple tree in our climate, despite my wanting to plant some in the garden.
- 26 CS Despite my wanting to plant some in the garden, sunflowers won't grow under our apple tree in our climate.
- 26 AO Our sunflower won't grow under apple trees in our climate, despite my wanting to plant some in the garden.
- 26 CO Despite my wanting to plant some in the garden, our sunflower won't grow under apple trees in our climate.
- 27 AA Stingrays are swimming alongside reef sharks right now, even though it's hard to see any in the pool.
- 27 CA Even though it's hard to see any in the pool, stingrays are swimming alongside reef sharks right now.
- 27 AS Stingrays are swimming alongside the largest reef shark right now, even though it's hard to see any in the pool.
- 27 CS Even though it's hard to see any in the pool, stingrays are swimming alongside the largest reef shark right now.
- 27 AO The largest stingray is swimming alongside reef sharks right now, even though it's hard to see any in the pool.
- 27 CO Even though it's hard to see any in the pool, the largest stingray is swimming alongside reef sharks right now.
- 28 AA Zebras would get along well with antelopes next year, but I don't know if it's possible to get any for the pen.
- 28 CA I don't know if it's possible to get any for the pen, but zebras would get along well with antelopes next year.
- 28 AS Zebras would get along well with our antelope next year, but I don't know if it's possible to get any for the pen.
- 28 CS I don't know if it's possible to get any for the pen, but zebras would get along well with our antelope next year.

- 28 AO Our zebra would get along well with antelopes next year, but I don't know if it's possible to get any for the pen.
- 28 CO I don't know if it's possible to get any for the pen, but our zebra would get along well with antelopes next year.
- 29 AA Orange ottomans would contrast well with blue rugs for the photo shoot, in case you want to buy some before next month.
- 29 CA In case you want to buy some before next month, orange ottomans would contrast well with blue rugs for the photo shoot.
- 29 AS Orange ottomans would contrast well with our blue rug for the photo shoot, in case you want some for the photo shoot.
- 29 CS In case you want to buy some before next month, orange ottomans would contrast well with our blue rug for the photo shoot.
- 29 AO Our orange ottoman would contrast well with blue rugs for the photo shoot, in case you want to buy some before next month.
- 29 CO In case you want to buy some before next month, our orange ottoman would contrast well with blue rugs for the photo shoot.
- 30 AA Ties would look good with hats for work, even though you don't wear any right now.
- 30 CA Even though you don't wear any right now, ties would look good with hats for work.
- 30 AS Ties would look good with your hat for work, even though you don't wear any right now.
- 30 CS Even though you don't wear any right now, ties would look good with your hat for work.
- 30 AO Your tie would look good with hats for work, even though you don't wear any right now.
- 30 CO Even though you don't wear any right now, your tie would look good with hats for work.
- 31 AA Photographs would look elegant in large frames for the wall, but I'm not sure if you have any already.
- 31 CA I'm not sure if you have any already, but photographs would look elegant in large frames for the wall.
- 31 AS Photographs would look elegant in your large frame for the wall, but I'm not sure if you have any already.
- 31 CS I'm not sure if you have any already, but photographs would look elegant in your large frame for the wall.
- 31 AO Your photograph would look elegant in large frames for the wall, but I'm not sure if you have any already.
- 31 CO I'm not sure if you have any already, but your photograph would look elegant in large frames for the wall.
- 32 AA Jade stones would pair well with gold clasps for winter, but I don't remember if the store has any in stock.
- 32 CA I don't remember if the store has any in stock, but jade stones would pair well with gold clasps for winter.
- 32 AS Jade stones would pair well with your gold clasp for winter, but I don't remember if the store has any in stock.
- 32 CS I don't remember if the store has any in stock, but jade stones would pair well with your gold clasp for winter.
- 32 AO Your jade stone would pair well with gold clasps for winter, but I don't remember if the store has any in stock.
- 32 CO I don't remember if the store has any in stock, but your jade stone would pair well with gold clasps for winter.
- 33 AA Vintage books would look cool beside candlesticks for the shelf, but I forget if you have any around.
- 33 CA I forget if you have any around, but vintage books would look cool beside candlesticks for the shelf.
- 33 AS Vintage books would look cool beside our candlestick for the shelf, but I forget if you have any around.

- 33 CS I forget if you have any around, but vintage books would look cool beside our candlestick for the shelf.
- 33 AO Our vintage book would look cool beside candlesticks for the shelf, but I forget if you have any around.
- 33 CO I forget if you have any around, but our vintage book would look cool beside candlesticks for the shelf.
- 34 AA Scallop shells would look pretty strung with opals for the gala, if you want to see any while you're at the store.
- 34 CA If you want to see any while you're at the store, scallop shells would look pretty strung with opals for the gala.
- 34 AS Scallop shells would look pretty strung with your opal for the gala, if you want to see any while you're at the store.
- 34 CS If you want to see any while you're at the store, scallop shells would look pretty strung with your opal for the gala.
- 34 AO Your scallop shell would look pretty strung with opals for the gala, if you want to see any while you're at the store.
- 34 CO If you want to see any while you're at the store, your scallop shell would look pretty strung with opals for the gala.
- 35 AA Cowrie shells would pair well with stone beads for the bracelet, but I don't believe you need any really.
- 35 CA I don't believe you need any really, but cowrie shells would pair well with stone beads for the bracelet.
- 35 AS Cowrie shells would pair well with your stone bead for the bracelet, but I don't believe you need any really.
- 35 CS I don't believe you need any really, but cowrie shells would pair well with your stone bead for the bracelet.
- 35 AO Your cowrie shell would pair well with stone beads for the bracelet, but I don't believe you need any really.
- 35 CO I don't believe you need any really, but your cowrie shell would pair well with stone beads for the bracelet.
- 36 AA Brass handles would go well with ceramic door pulls for on the dresser, but I don't think you have to get any for today.
- 36 CA I don't think you have to get any for today, but brass handles would go well with ceramic door pulls for on the dresser.
- 36 AS Brass handles would go well with my ceramic door pull for on the dresser, but I don't think you have to get any for today.
- 36 CS I don't think you have to get any for today, but brass handles would go well with my ceramic door pull for on the dresser.
- 36 AO My brass handle would go well with ceramic door pulls for on the dresser, but I don't think you have to get any for today.
- 36 CO I don't think you have to get any for today, but my brass handle would go well with ceramic door pulls for on the dresser.
- 37 AA Whales are often by sardine schools, but I don't know how many we'll see today near the beach.
- 37 CA I don't know how many we'll see today near the beach, but whales are often by sardine schools.
- 37 AS Whales are often by this tagged sardine school, but I don't know how many we'll see today near the beach.
- 37 CS I don't know how many we'll see today near the beach, but whales are often by this tagged sardine school.
- 37 AO This tagged whale is often by sardine schools, but I don't know how many we'll see today near the beach.
- 37 CO I don't know how many we'll see today near the beach, but this tagged whale is often by sardine schools.
- 38 AA Strawberry vines won't grow with watermelon vines in winter, despite my fussing over some in the backyard.

- 38 CA Despite my fussing over some in the backyard, strawberry vines won't grow with watermelon vines in winter.
- 38 AS Strawberry vines won't grow with my watermelon vine in winter, despite my fussing over some in the backyard.
- 38 CS Despite my fussing over some in the backyard, strawberry vines won't grow with my watermelon vine in winter.
- 38 AO My strawberry vine won't grow with watermelon vines in winter, despite my fussing over some in the backyard.
- 38 CO Despite my fussing over some in the backyard, my strawberry vine won't grow with watermelon vines in winter.
- 39 AA Lilac bushes would look nice with cherry trees alongside the pond, but I don't believe we can find any on short notice.
- 39 CA I don't believe we can find any on short notice, but lilac bushes would look nice with cherry trees alongside the pond.
- 39 AS Lilac bushes would look nice with our cherry tree alongside the pond, but I don't believe we can find any on short notice.
- 39 CS I don't believe we can find any on short notice, but lilac bushes would look nice with our cherry tree alongside the pond.
- 39 AO Our lilac bush would look nice with cherry trees alongside the pond, but I don't believe we can find any on short notice.
- 39 CO I don't believe we can find any on short notice, but our lilac bush would look nice with cherry trees alongside the pond.
- 40 AA Bean stalks would grow nicely planted with pea vines along the fence, even though we haven't found any in the stores yet.
- 40 CA Even though we haven't found any in the stores yet, bean stalks would grow nicely planted with pea vines along the fence.
- 40 AS Bean stalks would grow nicely planted with our pea vine along the fence, even though we haven't found any in the stores yet.
- 40 CS Even though we haven't found any in the stores yet, bean stalks would grow nicely planted with our pea vine along the fence.
- 40 AO Our bean stalk would grow nicely planted with pea vines along the fence, even though we haven't found any in the stores yet.
- 40 CO Even though we haven't found any in the stores yet, our bean stalk would grow nicely planted with pea vines along the fence.
- 41 AA Gold statuettes would look nice with crystal chandeliers for the foyer, but I don't know how many we can afford.
- 41 CA I don't know how many we can afford, but gold statuettes would look nice with crystal chandeliers for the foyer.
- 41 AS Gold statuettes would look nice with our crystal chandelier for the foyer, I don't know how many we can afford.
- 41 CS I don't know how many we can afford, but gold statuettes would look nice with our crystal chandelier for the foyer.
- 41 AO Our gold statuette would look nice with crystal chandeliers for the foyer, but I don't know how many we can afford.
- 41 CO I don't know how many we can afford, but our gold statuette would look nice with crystal chandeliers for the foyer.
- 42 AA Abstract paintings would look good with marble statues for the entryway, but I don't remember how many the gallery carries.
- 42 CA I don't remember how many the gallery carries, but abstract paintings would look good with marble statues for the entryway.
- 42 AS Abstract paintings would look good with our marble statue for the entryway, but I don't remember how many the gallery carries.
- 42 CS I don't remember how many the gallery carries, but abstract paintings would look good with our marble statue for the entryway.
- 42 AO Our abstract painting would look good with marble statues for the entryway, but I don't remember how many the gallery carries.

42 CO I don't remember how many the gallery carries, but our abstract painting would look good with marble statues for the entryway.

Appendix B

Experiment 2 items

It.	Cond.	Item
1	AN	I saw Mary's dog and Susan's cat yesterday, but I didn't see Jane's at the time.
1	CN	I didn't see Jane's at the time, but I saw Mary's dog and Susan's cat yesterday.
1	AP	I saw Mary's dog and I saw Susan's cat yesterday, but I didn't see Jane's at the time.
1	CP	I didn't see Jane's at the time, but I saw Mary's dog and I saw Susan's cat yesterday.
2	AN	I talked to Chris's partner and Pete's colleague for a while, but I wasn't able to talk to Owen's at the dinner.
2	CN	I wasn't able to talk to Owen's at the dinner, but I talked to Chris's partner and Pete's colleague for a while.
2	AP	I talked to Chris's partner and I talked to Pete's colleague for a while, but I wasn't able to talk to Owen's at the dinner.
2	CP	I wasn't able to talk to Owen's at the dinner, but I talked to Chris's partner and I talked to Pete's colleague for a while.
3	AN	I tried Ashley's steak and Tony's casserole at the potluck, but I didn't get a chance to try Quinn's at all.
3	CN	I didn't get a chance to try Quinn's at all, but I tried Ashley's steak and Tony's casserole at the potluck.
3	AP	I tried Ashley's steak and I tried Tony's casserole at the potluck, but I didn't get a chance to try Quinn's at all.
3	CP	I didn't get a chance to try Quinn's at all, but I tried Ashley's steak and I tried Tony's casserole at the dinner.
4	AN	I put up Josh's streamers and Bert's balloons for the party, but I couldn't find Evan's anywhere.
4	CN	I couldn't find Evan's anywhere, but I put up Josh's streamers and Bert's balloons for the party.
4	AP	I put up Josh's streamers and I put up Bert's balloons for the party, but I couldn't find Evan's anywhere.
4	CP	I couldn't find Evan's anywhere, but I put up Josh's streamers and I put up Bert's balloons for the party.
5	AN	I played Francis's clarinet and Lucy's flute at the rehearsal, but I wasn't able to play Willow's at all.
5	CN	I wasn't able to play Willow's at all, but I played Francis's clarinet and Lucy's flute at the rehearsal.
5	AP	I played Francis's clarinet and I played Lucy's flute at the rehearsal, but I wasn't able to play Willow's at all.
5	CP	I wasn't able to play Willow's at all, but I played Francis's clarinet and I played Lucy's flute at the rehearsal.

- 6 AN I used Emily's chair and Scarlett's table for the apartment, but I wasn't able to use Nicole's in the end.
- 6 CN I wasn't able to use Nicole's in the end, but I used Emily's chair and Scarlett's table for the apartment.
- 6 AP I used Emily's chair and I used Scarlett's table for the apartment, but I wasn't able to use Nicole's in the end.
- 6 CP I wasn't able to use Nicole's in the end, but I used Emily's chair and I used Scarlett's table for the apartment.
- 7 AN I swam with Janice's pug puppy and Kelsey's niece for a while, but I wasn't able to swim with Ivy's for some reason.
- 7 CN I wasn't able to swim with Ivy's for some reason, but I swam with Janice's pug puppy and Kelsey's niece for a while.
- 7 AP I swam with Janice's pug puppy and I swam with Kelsey's niece for a while, but I wasn't able to swim with Ivy's for some reason.
- 7 CP I wasn't able to swim with Ivy's for some reason, but I swam with Janice's pug puppy and I swam with Kelsey's niece for a while.
- 8 AN I bought Liam's newt and Mateo's fighting fish last week, but I couldn't buy Noah's in time.
- 8 CN I couldn't buy Noah's in time, but I bought Liam's newt and Mateo's fighting fish last week.
- 8 AP I bought Liam's newt and I bought Mateo's fighting fish last week, but I couldn't buy Noah's in time.
- 8 CP I couldn't buy Noah's in time, but I bought Liam's newt and I bought Mateo's fighting fish last week.
- 9 AN I borrowed Michael's serving spoon and William's platter for the luncheon, but I couldn't borrow Grant's in time.
- 9 CN I couldn't borrow Grant's in time, but I borrowed Michael's serving spoon and William's platter for the luncheon.
- 9 AP I borrowed Michael's serving spoon and I borrowed William's platter for the luncheon, but I couldn't borrow Grant's in time.
- 9 CP I couldn't borrow Grant's in time, but I borrowed Michael's serving spoon and I borrowed William's platter for the luncheon.
- 10 AN I brought Jim's stapler and Isabella's notebook for the meeting, but I couldn't bring Logan's for us.
- 10 CN I couldn't bring Logan's for us, but I brought Jim's stapler and Isabella's notebook for the meeting.
- 10 AP I brought Jim's stapler and I brought Isabella's notebook for the meeting, but I couldn't bring Logan's for us.
- 10 CP I couldn't bring Logan's for us, but I brought Jim's stapler and I brought Isabella's notebook for the meeting.
- 11 AN I stole Sophia's horse and Emma's donkey last month, but I couldn't steal Mia's for the stable.
- 11 CN I couldn't steal Mia's for the stable, but I stole Sophia's horse and Emma's donkey last month.
- 11 AP I stole Sophia's horse and and I stole Emma's donkey last month, but I couldn't steal Mia's for the stable.
- 11 CP I couldn't steal Mia's for the stable, but I stole Sophia's horse and I stole Emma's donkey last month.
- 12 AN I snapped up Mason's rooster and Benjamin's cow right away, but I couldn't get Elijah's for the farm.
- 12 CN I couldn't get Elijah's for the farm, but I snapped up Mason's rooster and Benjamin's cow right away.
- 12 AP I snapped up Mason's rooster and I snapped up Benjamin's cow right away, but I couldn't get Elijah's for the farm.
- 12 CP I couldn't get Elijah's for the farm, but I snapped up Mason's rooster and I snapped up Benjamin's cow right away.

- 13 AN I purchased Amelia's guitar and Katie's tuba for the orchestra, but I couldn't purchase Evelyn's in time.
- 13 CN I couldn't purchase Evelyn's in time, but I purchased Amelia's guitar and Katie's tuba for the orchestra.
- 13 AP I purchased Amelia's guitar and I purchased Katie's tuba for the orchestra, but I couldn't purchase Evelyn's in time.
- 13 CP I couldn't purchase Evelyn's in time, but I purchased Amelia's guitar and I purchased Katie's tuba for the orchestra.
- 14 AN I pet Jacob's penguin and Abigail's duck a lot, but I didn't get to pet Harper's at all.
- 14 CN I didn't get to pet Harper's at all, but I pet Jacob's penguin and Abigail's duck a lot.
- 14 AP I pet Jacob's penguin and I pet Abigail's duck a lot, but I didn't get to pet Harper's at all.
- 14 CP I didn't get to pet Harper's at all, but I pet Jacob's penguin and I pet Abigail's duck a lot.
- 15 AN I got John's stuffed animal and Emily's book already, but I wasn't able to get Morgan's yet.
- 15 CN I wasn't able to get Morgan's yet, but I got Michael's stuffed animal and John's book already.
- 15 AP I got John's stuffed animal and I got Emily's book already, but I wasn't able to get Morgan's yet.
- 15 CP I wasn't able to get Morgan's yet, but I got John's stuffed animal and I got Emily's book already.
- 16 AN I planted Andrew's tree sapling and Ryan's dandelion seed along my house, but I wasn't able to get Lincoln's in time.
- 16 CN I wasn't able to get Lincoln's in time, but I planted Andrew's tree sapling and Ryan's dandelion seed along my house.
- 16 AP I planted Andrew's tree sapling and I planted Ryan's dandelion seed along my house, but I wasn't able to get Lincoln's in time.
- 16 CP I wasn't able to get Lincoln's in time, but I planted Andrew's tree sapling and I planted Ryan's dandelion seed along my house.
- 17 AN I brought Zoe's marshmallows and Olivia's chocolate bars for the campfire, but I forgot to bring Stella's with me.
- 17 CN I forgot to bring Stella's with me, but I brought Zoe's marshmallows and Olivia's chocolate bars for the campfire.
- 17 AP I brought Zoe's marshmallows and I brought Olivia's chocolate bars for the campfire, but I forgot to bring Stella's with me.
- 17 CP I forgot to bring Stella's with me, but I brought Zoe's marshmallows and I brought Olivia's chocolate bars for the campfire.
- 18 AN I have Hazel's seastar and Ellie's sea turtle already, but I haven't seen Charlotte's in the aquarium yet.
- 18 CN I haven't seen Charlotte's in the aquarium yet, but I have Hazel's seastar and Ellie's sea turtle already.
- 18 AP I have Hazel's seastar and I have Ellie's sea turtle already, but I haven't seen Charlotte's in the aquarium yet.
- 18 CP I haven't seen Charlotte's in the aquarium yet, but I have Hazel's seastar and I have Ellie's sea turtle already.
- 19 AN I heard Aaron's aria and Oliver's sonata last night, but I didn't hear Isaiah's at the try-out.
- 19 CN I didn't hear Isaiah's at the try-out, but I heard Aaron's aria and Oliver's sonata last night.
- 19 AP I heard Aaron's aria and I heard Oliver's sonata last night, but I didn't hear Isaiah's at the try-out.
- 19 CP I didn't hear Isaiah's at the try-out, but I heard Aaron's aria and I heard Oliver's sonata last night.
- 20 AN I bought Joshua's statue and Edward's painting at the show, but I couldn't afford Eric's at all.

- 20 CN I couldn't afford Eric's at all, but I bought Joshua's statue and Edward's painting at the show.
- 20 AP I bought Joshua's statue and I bought Edward's painting at the show, but I couldn't afford Eric's at all.
- 20 CP I couldn't afford Eric's at all, but I bought Joshua's statue and I bought Edward's painting at the show.
- 21 AN I took Caleb's dirt bike and Hunter's four-wheeler from the garage, but I couldn't find Eli's anywhere.
- 21 CN I couldn't find Eli's anywhere, but I took Caleb's dirt bike and Hunter's four-wheeler from the garage.
- 21 AP I took Caleb's dirt bike and I took Hunter's four-wheeler from the garage, but I couldn't find Eli's anywhere.
- 21 CP I couldn't find Eli's anywhere, but I took Caleb's dirt bike and I took Hunter's four-wheeler from the garage.
- 22 AN I can point out Anna's favorite otter and Samantha's favorite seal in the Bay, but I don't think I can find Caroline's anymore.
- 22 CN I don't think I can find Caroline's anymore, but I can point out Anna's favorite otter and Samantha's favorite seal in the Bay.
- 22 AP I can point out Anna's favorite otter and I can point out Samantha's favorite seal in the Bay, but I don't think I can find Caroline's anymore.
- 22 CP I don't think I can find Caroline's anymore, but I can point out Anna's favorite otter and I can point out Samantha's favorite seal in the Bay.
- 23 AN I planted Jonathan's sunflower and Connor's apple tree last week, but I haven't planted Adrian's in the yard yet.
- 23 CN I haven't planted Adrian's in the yard yet, but I planted Jonathan's sunflower and Connor's apple tree last week.
- 23 AP I planted Jonathan's sunflower and I planted Connor's apple tree last week, but I haven't planted Adrian's in the yard yet.
- 23 CP I haven't planted Adrian's in the yard yet, but I planted Jonathan's sunflower and I planted Connor's apple tree last week.
- 24 AN I boarded Leo's zebra and Nate's giraffe last month, but I haven't needed to board Ted's yet.
- 24 CN I haven't needed to board Ted's yet, but I boarded Leo's zebra and Nate's giraffe last month.
- 24 AP I boarded Leo's zebra and I boarded Nate's giraffe last month, but I haven't needed to board Ted's yet.
- 24 CP I haven't needed to board Ted's yet, but I boarded Leo's zebra and I boarded Nate's giraffe last month.
- 25 AN I want Adeline's table and Alexa's rug for the photo shoot, but I don't want Ariana's in the house.
- 25 CN I don't want Ariana's in the house, but I want Adeline's table and Alexa's rug for the photo shoot.
- 25 AP I want Adeline's table and I want Alexa's rug for the photo shoot, but I don't want Ariana's in the house.
- 25 CP I don't want Ariana's in the house, but I want Adeline's table and I want Alexa's rug for the photo shoot.
- 26 AN I'll wear Alexa's ring and Gabriella's earring for the interview, but I won't be able to wear Naomi's at all.
- 26 CN I won't be able to wear Naomi's at all, but I'll wear Alexa's ring and Gabriella's earring for the interview.
- 26 AP I'll wear Alexa's ring and I'll wear Gabriella's earring for the interview, but I won't be able to wear Naomi's at all.
- 26 CP I won't be able to wear Naomi's at all, but I'll wear Alexa's ring and I'll wear Gabriella's earring for the interview.
- 27 AN I bought Hailey's photograph and Autumn's frame yesterday, but I couldn't afford Eva's for my office.

- 27 CN I couldn't afford Eva's for my office, but I bought Hailey's photograph and Autumn's frame yesterday.
- 27 AP I bought Hailey's photograph and I bought Autumn's frame yesterday, but I couldn't afford Eva's for my office.
- 27 CP I couldn't afford Eva's for my office, but I bought Hailey's photograph and I bought Autumn's frame yesterday.
- 28 AN I wore my mother's necklace and my grandmother's ring all day, but I wasn't able to wear my sister's for the ceremony.
- 28 CN I wasn't able to wear my sister's for the ceremony, but I wore my mother's necklace and my grandmother's ring all day.
- 28 AP I wore my mother's necklace and I wore my grandmother's ring all day, but I wasn't able to wear my sister's for the ceremony.
- 28 CP I wasn't able to wear my sister's for the ceremony, but I wore my mother's necklace and I wore my grandmother's ring all day.
- 29 AN I pasted Angela's finger-painting and Shirley's collage on a wall, but there wasn't room to paste Kathleen's at the same time.
- 29 CN There wasn't room to paste Kathleen's at the same time, but I pasted Angela's finger-painting and Shirley's collage on a wall.
- 29 AP I pasted Angela's finger-painting and I pasted Shirley's collage on a wall, but there wasn't room to paste Kathleen's at the same time..
- 29 CP There wasn't room to paste Kathleen's at the same time, but I pasted Angela's finger-painting and I pasted Shirley's collage on a wall.
- 30 AN I put Evan's vintage book and Bryson's candlestick on the shelf, but I couldn't find Maverick's anywhere in the house.
- 30 CN I couldn't find Maverick's anywhere in the house, but I put Evan's vintage book and Bryson's candlestick on the shelf.
- 30 AP I put Evan's vintage book and I put Bryson's candlestick on the shelf, but I couldn't find Maverick's anywhere in the house.
- 30 CP I couldn't find Maverick's anywhere in the house, but I put Evan's vintage book and I put Bryson's candlestick on the shelf.
- 31 AN I have Jason's article and Gavin's cartoon already, but I haven't received Paul's for printing yet.
- 31 CN I haven't received Paul's for printing yet, but I have Jason's article and Gavin's cartoon already.
- 31 AP I have Jason's article and I have Gavin's cartoon already, but I haven't received Paul's for printing yet.
- 31 CP I haven't received Paul's for printing yet, but I have Jason's article and I have Gavin's cartoon already.
- 32 AN I checked Stephanie's lettuce and Patricia's tomatoes the other day, but I didn't have time to look at Jennifer's very carefully.
- 32 CN I didn't have time to look at Jennifer's very carefully, but I checked Stephanie's lettuce and Patricia's tomatoes the other day.
- 32 AP I checked Stephanie's lettuce and I checked Patricia's tomatoes the other day, but I didn't have time to look at Jennifer's very carefully.
- 32 CP I didn't have time to look at Jennifer's very carefully, but I checked Stephanie's lettuce and I checked Patricia's tomatoes the other day.
- 33 AN I planted Elizabeth's lilac and Jessica's tulip in the front yard, but there wasn't room for Karen's in the front.
- 33 CN There wasn't room for Karen's in the front, but I planted Elizabeth's lilac and Jessica's tulip in the front yard.
- 33 AP I planted Elizabeth's lilac and I planted Jessica's tulip in the front yard, but there wasn't room for Karen's in the front.
- 33 CP There wasn't room for Karen's in the front, but I planted Elizabeth's lilac and I planted Jessica's tulip in the front yard.
- 34 AN I watched Daniel's movie and Anthony's sitcom last night, but I haven't had time to watch Matthew's yet this week.

- 34 CN I haven't had time to watch Matthew's yet this week, but I watched Daniel's movie and Anthony's sitcom last night.
- 34 AP I watched Daniel's movie and I watched Anthony's sitcom last night, but I haven't had time to watch Matthew's yet this week.
- 34 CP I haven't had time to watch Matthew's yet this week, but I watched Daniel's movie and I watched Anthony's sitcom last night.
- 35 AN I'm opening my aunt's gift bag and my uncle's card later tonight, but I haven't received my cousin's in the mail yet.
- 35 CN I haven't received my cousin's in the mail yet, but I'm opening my aunt's gift bag and my uncle's card later tonight.
- 35 AP I'm opening my aunt's gift bag and I'm opening my uncle's card later tonight, but I haven't received my cousin's in the mail yet.
- 35 CP I haven't received my cousin's in the mail yet, but I'm opening my aunt's gift bag and I'm opening my uncle's card later tonight.
- 36 AN I tried Andy's beer and Mark's cider last night, but I wasn't able to find Steven's on the menu.
- 36 CN I wasn't able to find Steven's on the menu, but I tried Andy's beer and Mark's cider last night.
- 36 AP I tried Andy's beer and I tried Mark's cider last night, but I wasn't able to find Steven's on the menu.
- 36 CP I wasn't able to find Steven's on the menu, but I tried Andy's beer and Mark's cider last night.

Appendix C

Maze/SPR experimental items

It.	Cond.	Item	Foils
1	ND	While many pre-meds are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.	x-x-x hear mixture thou consequences belongs go textbooks, ends ago net extension effort ring sentimental sin assertion.
1	EL	While many are occasionally excited and energetic, most undergraduates are generally depressed.	x-x-x hear thou consequences belongs go textbooks, ends ago net extension assertion.
1	NL	While many pre-meds are occasionally excited and energetic, most undergraduates are generally depressed.	x-x-x hear mixture thou consequences belongs go textbooks, ends ago net extension assertion.
1	ED	While many are occasionally excited and energetic, most reports are generally saying that undergraduates are depressed.	x-x-x hear thou consequences belongs go textbooks, ends ago net extension effort ring sentimental sin assertion.
2	ND	While most peacocks are naturally aggressive and combative, some studies are periodically finding that swans are friendly.	x-x-x lose ablative why expansion delegation saw dissented, rose beneath sin manufactures amounts poet silks job monetary.
2	ED	While most are naturally aggressive and combative, some studies are periodically finding that swans are friendly.	x-x-x lose why expansion delegation saw dissented, rose beneath sin manufactures amounts poet silks job monetary.
2	NL	While most peacocks are naturally aggressive and combative, some swans are periodically friendly.	x-x-x lose ablative why expansion delegation saw dissented, rose beneath sin manufactures monetary.
2	EL	While most are naturally aggressive and combative, some swans are periodically friendly.	x-x-x lose why expansion delegation saw dissented, rose beneath sin manufactures monetary.
3	ND	Although many manuscripts are very well-written and engrossing, few editors are ultimately convinced that novels are bestsellers.	x-x-x hear elimination nor ship upon ill monarchies, ago custody job prevention molecules ball loaded joy cheerleader.
3	ED	Although many are very well-written and engrossing, few editors are ultimately convinced that novels are bestsellers.	x-x-x hear nor ship upon ill monarchies, ago custody job prevention molecules ball loaded joy cheerleader.
3	NL	Although many manuscripts are very well-written and engrossing, few novels are ultimately bestsellers.	x-x-x hear elimination nor ship upon ill monarchies, ago custody job prevention cheerleader.
3	EL	Although many are very well-written and engrossing, few novels are ultimately bestsellers.	x-x-x hear nor ship upon ill monarchies, ago custody job prevention cheerleader.

4	ND	Although few boats are sturdy or completely watertight, many captains are surprisingly convinced that kayaks are able to withstand hurricanes.	x-x-x saw serum why tenet fat comparison syndicated, rose fruitful ad demonstrates breakfast thou unseat why went age readiness achievable.
4	ED	Although few are sturdy or completely watertight, many captains are surprisingly convinced that kayaks are able to withstand hurricanes.	x-x-x saw why tenet fat comparison syndicated, rose fruitful ad demonstrates breakfast thou unseat why went age readiness achievable.
4	NL	Although few boats are sturdy or completely watertight, many kayaks are surprisingly able to withstand hurricanes.	x-x-x saw serum why tenet fat comparison syndicated, rose fruitful ad demonstrates went age readiness achievable.
4	EL	Although few are sturdy or completely watertight, many kayaks are surprisingly able to withstand hurricanes.	x-x-x saw why tenet fat comparison syndicated, rose fruitful ad demonstrates went age readiness achievable.
5	ND	Though many buses are certainly damaged or destroyed, several investigators are definitely confident that trains are still operative.	x-x-x drop whence no authority exclude sex happiness, granted approximation lot matrix ancestors fish loving bit truth ascertain.
5	ED	Though many are certainly damaged or destroyed, several investigators are definitely confident that trains are still operative.	x-x-x drop no authority exclude sex happiness, granted approximation lot matrix ancestors fish loving bit truth ascertain.
5	NL	Though many buses are certainly damaged or destroyed, several trains are definitely still operative.	x-x-x drop whence no authority exclude sex happiness, granted approximation lot matrix truth ascertain.
5	EL	Though many are certainly damaged or destroyed, several trains are definitely still operative.	x-x-x drop no authority exclude sex happiness, granted approximation lot matrix truth ascertain.
6	ND	Although many ships are already docked and de-rigged, several onlookers are definitely seeing that sailboats are still racing.	x-x-x door thy lot ability curfew joy lemons, degrees multiple bed navigation duties ages racehorse thy taxes torque.
6	ED	Although many are already docked and de-rigged, several onlookers are definitely seeing that sailboats are still racing.	x-x-x door lot ability curfew joy lemons, degrees multiple bed navigation duties ages racehorse thy taxes torque.
6	NL	Although many ships are already docked and de-rigged, several sailboats are definitely still racing.	x-x-x door thy lot ability curfew joy lemons, degrees multiple bed navigation taxes torque.
6	EL	Although many are already docked and de-rigged, several sailboats are definitely still racing.	x-x-x door lot ability curfew joy lemons, degrees multiple bed navigation taxes torque.
7	ND	While almost all blazes are natural and ultimately inevitable, a few firefighters are unfortunately finding that wildfires are from arson each year.	x-x-x nor so dimmer why suppose ask conception objectives, am sweeter red revolutionary respect went triathlon how tell obeys seat thin.
7	ED	While almost all are natural and ultimately inevitable, a few firefighters are unfortunately finding that wildfires are from arson each year.	x-x-x nor so why suppose ask conception objectives, am sweeter red revolutionary respect went triathlon how tell obeys seat thin.
7	NL	While almost all blazes are natural and ultimately inevitable, a few wildfires are unfortunately from arson each year.	x-x-x nor so dimmer why suppose ask conception objectives, am sweeter red revolutionary tell obeys seat thin.
7	EL	While almost all are natural and ultimately inevitable, a few wildfires are unfortunately from arson each year.	x-x-x nor so why suppose ask conception objectives, am sweeter red revolutionary tell obeys seat thin.
8	ND	While nearly all whales are swimming south for winter, many marine biologists are sadly recognizing that great white sharks are becoming stranded in the northern Pacific.	x-x-x sat nor most ago efficacy sales sit agents, went nonviolence old amend liabilities army curate son conflict comedies arm me purposes Prevent.
8	ED	While nearly all are swimming south for winter, many marine biologists are sadly recognizing that great white sharks are becoming stranded in the northern Pacific.	x-x-x sat nor ago efficacy sales sit agents, went nonviolence old amend liabilities army curate son conflict comedies arm me purposes Prevent.

8	NL	While nearly all whales are swimming south for winter, many great white sharks are sadly becoming stranded in the northern Pacific.	x-x-x sat nor most ago efficacy sales sit agents, went nonviolence old amend conflict comedies arm me purposes Prevent.
8	EL	While nearly all are swimming south for winter, many great white sharks are sadly becoming stranded in the northern Pacific.	x-x-x sat nor ago efficacy sales sit agents, went nonviolence old amend conflict comedies arm me purposes Prevent.
9	ND	Although a few groundhogs are ready for winter hibernation, almost all assessments are still showing that black bears are too skinny.	x-x-x ask turboprop mid grace sat partly compressing, reform several influence hotness death hour threw now key me esters.
9	ED	Although a few are ready for winter hibernation, almost all assessments are still showing that black bears are too skinny.	x-x-x ask mid grace sat partly compressing, reform several influence hotness death hour threw now key me esters.
9	NL	Although a few groundhogs are ready for winter hibernation, almost all black bears are still too skinny.	x-x-x ask turboprop mid grace sat partly compressing, reform several influence hotness me esters.
9	EL	Although a few are ready for winter hibernation, almost all black bears are still too skinny.	x-x-x ask mid grace sat partly compressing, reform several influence hotness me esters.
10	ND	Although too many animals are becoming endangered or extinct, numerous scientists are nevertheless surmising that insects are still to be discovered.	x-x-x walk so summary nor sentence shortening hot stained, thousand obligation red consequences womanizer milk goddess car trust add won additional.
10	ED	Although too many are becoming endangered or extinct, numerous scientists are nevertheless surmising that insects are still to be discovered.	walk so nor sentence shortening hot stained, thousand obligation red consequences womanizer milk goddess car trust add won additional.
10	NL	Although too many animals are becoming endangered or extinct, numerous insects are nevertheless still to be discovered.	walk so summary nor sentence shortening hot stained, thousand obligation red consequences trust add won additional.
10	EL	Although too many are becoming endangered or extinct, numerous insects are nevertheless still to be discovered.	x-x-x walk so nor sentence shortening hot stained, thousand obligation red consequences trust add won additional.
11	ND	Even though few people are unprepared or surprised, several statistics are regrettably showing that climbers are killed on Mt. Everest every season.	x-x-x extent arm unable sky chieftain lot diffusion, century assumption sun rattlesnake welfare boat rallying boy sister map Custard
11	ED	Even though few are unprepared or surprised, several statistics are regrettably showing that climbers are killed on Mt. Everest every season.	x-x-x extent arm sky chieftain lot diffusion, century assumption sun rattlesnake welfare boat rallying boy sister map Custard fruit exists.
11	NL	Even though few people are unprepared or surprised, several climbers are regrettably killed on Mt. Everest every season.	x-x-x extent arm unable sky chieftain lot diffusion, century assumption sun rattlesnake sister map Custard fruit exists.
11	EL	Even though few are unprepared or surprised, several climbers are regrettably killed on Mt. Everest every season.	x-x-x extent arm sky chieftain lot diffusion, century assumption sun rattlesnake sister map Custard fruit exists.
12	ND	Even though countless people are killed or injured, numerous forums are genuinely showing that adventurers are obsessed with cliff jumping in Norway.	x-x-x visor coherence heaven boy heaven art cluster, argument basing ice apostolic herself mode concurrence non aromatic drop appoint buy Deputy.
12	ED	Even though countless are killed or injured, numerous forums are genuinely showing that adventurers are obsessed with cliff jumping in Norway.	x-x-x visor coherence boy heaven art cluster, argument basing ice apostolic herself mode concurrence non aromatic drop appoint buy Deputy.
12	NL	Even though countless people are killed or injured, numerous adventurers are genuinely obsessed with cliff jumping in Norway.	x-x-x visor coherence heaven boy heaven art cluster, argument basing ice apostolic aromatic drop appoint buy Deputy.

12	EL	Even though countless are killed or injured, numerous adventurers are genuinely obsessed with cliff jumping in Norway.	x-x-x visor coherence boy heaven art cluster, argument basing ice apostolic aromatic drop appoint buy Deputy.
13	ND	Despite a few vegetables dying from frost and bugs, many gardeners are thankfully saying that tomatoes are surviving to the end of the season.	x-x-x sky thou cable wall buyer arm feat, step fifth sea exuberance decide cent auditors ten sculpture sat nor we bar sky taught.
13	ED	Despite a few dying from frost and bugs, many gardeners are thankfully saying that tomatoes are surviving to the end of the season.	x-x-x sky many wall buyer arm feat, step fifth sea exuberance decide cent auditors ten sculpture sat nor we bar sky taught.
13	NL	Despite a few vegetables dying from frost and bugs, many tomatoes are thankfully surviving to the end of the season.	x-x-x sky thou cable wall buyer arm feat, step fifth sea exuberance sculpture sat nor we bar sky taught.
13	EL	Despite a few dying from frost and bugs, many tomatoes are thankfully surviving to the end of the season.	x-x-x sky many wall buyer arm feat, step fifth sea exuberance sculpture sat nor we bar sky taught.
14	ND	Despite much pollution produced annually by coal plants, most research is regrettably finding that carbon is generated by human activity.	x-x-x cent commander select fragment buy pair length, feel attitude sit capitulated replied neck manner ten discourse his happiness thousand.
14	ED	Despite much produced annually by coal plants, most research is regrettably finding that carbon is generated by human activity.	x-x-x cent select fragment buy pair length, feel attitude sit capitulated replied neck manner ten discourse his happiness thousand.
14	NL	Despite much pollution produced annually by coal plants, most carbon is regrettably generated by human activity.	x-x-x cent commander select fragment buy pair length, feel attitude sit capitulated discourse his happiness thousand.
14	EL	Despite much produced annually by coal plants, most carbon is regrettably generated by human activity.	x-x-x cent select fragment buy pair length, feel attitude sit capitulated discourse his happiness thousand.
15	ND	While many animals are poached in the wild, numerous surveys are tragically revealing that elephants are killed on game reserves.	x-x-x ages whereas job rivalry sum saw grew, document destiny few knighthood sixteenth drop dialectic him virtue non suite.
15	ED	While many are poached in the wild, numerous surveys are tragically revealing that elephants are killed on game reserves.	x-x-x ages job rivalry sum saw grew, document destiny few knighthood sixteenth drop dialectic him virtue non suite.
15	NL	While many animals are poached in the wild, numerous elephants are tragically killed on game reserves.	x-x-x ages whereas job rivalry sum saw grew, document destiny few knighthood virtue non suite.
15	EL	While many are poached in the wild, numerous elephants are tragically killed on game reserves.	x-x-x ages job rivalry sum saw grew, document destiny few knighthood virtue non suite.
16	ND	While most reptiles appear frightening and potentially venomous, almost all turtles are actually positive that snakes are harmless.	x-x-x tube although fourth transmitter saw manuscripts countries, nearly sat smuggled art yourself struggle seat candle hot massacre.
16	ED	While most appear frightening and potentially venomous, almost all turtles are actually positive that snakes are harmless.	x-x-x tube fourth transmitter saw manuscripts countries, nearly sat smuggled art yourself struggle seat candle hot massacre.
16	NL	While most reptiles appear frightening and potentially venomous, almost all snakes are actually harmless.	x-x-x tube although fourth transmitter saw manuscripts countries, nearly sat smuggled art yourself massacre.
16	EL	While most appear frightening and potentially venomous, almost all snakes are actually harmless.	x-x-x tube fourth transmitter saw manuscripts countries, nearly sat smuggled art yourself massacre.
17	ND	While a few people have no electricity or water, most maps are luckily showing that homes are still with power.	x-x-x sum beauty zero thy maintaining thy ratio, lady tail mid bedside teacher goes clerk ad goals says agree.

17	ED	While a few have no electricity or water, most maps are luckily showing that homes are still with power.	x-x-x sum zero thy maintaining thy ratio, lady tail mid bedside teacher goes clerk ad goals says agree.
17	NL	While a few people have no electricity or water, most homes are luckily still with power.	x-x-x sum beauty zero thy maintaining thy ratio, lady tail mid bedside goals says agree.
17	EL	While a few have no electricity or water, most homes are luckily still with power.	x-x-x sum zero thy maintaining thy ratio, lady tail mid bedside goals says agree.
18	ND	Although few raccoons are particularly feral or timid, some neighborhoods are understandably reporting that squirrels are scared of humans.	x-x-x buy scuttled sum relationship overruns ad lotus, anti unconsciously job intermediaries chemistry soul examiner me oceans sin ethics.
18	ED	Although few are particularly feral or timid, some neighborhoods are understandably reporting that squirrels are scared of humans.	x-x-x buy sum relationship overruns ad lotus, anti unconsciously job intermediaries chemistry soul examiner me oceans sin ethics.
18	NL	Although few raccoons are particularly feral or timid, some squirrels are understandably scared of humans.	x-x-x buy scuttled sum relationship overruns ad lotus, anti unconsciously job intermediaries oceans sin ethics.
18	EL	Although few are particularly feral or timid, some squirrels are understandably scared of humans.	x-x-x buy sum relationship overruns ad lotus, anti unconsciously job intermediaries oceans sin ethics.
19	ND	While few octopuses are seen or understood well, many swimmers are actually finding that lobsters are living around Monterey Bay.	x-x-x ago numerous why walk ran difficulty hear, anti pancake bad addition summary thou grapple lot author helped Intrigue.
19	ED	While few are seen or understood well, many swimmers are actually finding that lobsters are living around Monterey Bay.	x-x-x ago why walk ran difficulty hear, anti pancake bad addition summary thou grapple lot author helped Intrigue.
19	NL	While few octopuses are seen or understood well, many lobsters are actually living around Monterey Bay.	x-x-x ago numerous why walk ran difficulty hear, anti pancake bad addition author helped Intrigue.
19	EL	While few are seen or understood well, many lobsters are actually living around Monterey Bay.	x-x-x ago why walk ran difficulty hear, anti pancake bad addition author helped Intrigue.
20	ND	Despite hardly any wetsuits being sold or marketed, most teenagers are somehow thinking that surfboards are costing more than ever.	x-x-x have bed ago ought hour met ancients, poem predictor oil regards hospital term overpasses mid precept send poet mode.
20	ED	Despite hardly any being sold or marketed, most teenagers are somehow thinking that surfboards are costing more than ever.	x-x-x have bed ought hour met ancients, poem predictor oil regards hospital term overpasses mid precept send poet mode.
20	NL	Despite hardly any wetsuits being sold or marketed, most surfboards are somehow costing more than ever.	x-x-x have bed ago ought hour met ancients, poem predictor oil regards precept send poet mode.
20	EL	Despite hardly any being sold or marketed, most surfboards are somehow costing more than ever.	x-x-x have bed ought hour met ancients, poem predictor oil regards precept send poet mode.
21	ND	Because so few mountains sit on major faults, many hikers are actually proposing that volcanoes are buried under sediment.	x-x-x ad why procedure am sum taste mouth, lady dismissing joy district borrowers lose polyester bit intent tried solitude.
21	ED	Because so few sit on major faults, many hikers are actually proposing that volcanoes are buried under sediment.	x-x-x ad why am sum taste mouth, lady dismissing joy district borrowers lose polyester bit intent tried solitude.
21	NL	Because so few mountains sit on major faults, many volcanoes are actually buried under sediment.	x-x-x ad why procedure am sum taste mouth, lady dismissing joy district intent tried solitude.
21	EL	Because so few sit on major faults, many volcanoes are actually buried under sediment.	x-x-x ad why am sum taste mouth, lady dismissing joy district intent tried solitude.
22	ND	Because most deer can outrun fast predator animals, many hunters are surprisingly sure reindeer are able to live long lives.	x-x-x page vein why carrot tone thalamus brother, soul linkage bed depreciation post barked our acid him debt cast woman.

22	ED	Because most can outrun fast predator animals, many hunters are surprisingly sure reindeer are able to live long lives.	x-x-x page why carrot tone thalamus brother, soul linkage bed depreciation post barked our acid him debt cast woman.
22	NL	Because most deer can outrun fast predator animals, many reindeer are surprisingly able to live long lives.	x-x-x page vein why carrot tone thalamus brother, soul linkage bed depreciation acid him debt cast woman.
22	EL	Because most can outrun fast predator animals, many reindeer are surprisingly able to live long lives.	x-x-x page why carrot tone thalamus brother, soul linkage bed depreciation acid him debt cast woman.
23	ND	Because many dinosaurs had feathers and not scales, most paleontologists were actually wrong that triceratops were not colorful.	x-x-x necked compiling am vacation net am utmost, week decommissioning milk occasion drink lose telekinesis fund job twenties.
23	ED	Because many had feathers and not scales, most paleontologists were actually wrong that triceratops were not colorful.	x-x-x necked am vacation net am utmost, week decommissioning milk occasion drink lose telekinesis fund job twenties.
23	NL	Because many dinosaurs had feathers and not scales, most triceratops were actually not colorful.	x-x-x necked compiling am vacation net am utmost, week decommissioning milk occasion job twenties.
23	EL	Because many had feathers and not scales, most triceratops were actually not colorful.	x-x-x necked am vacation net am utmost, week decommissioning milk occasion job twenties.
24	ND	Because few canines attack cattle unless extremely desperate, most ranchers will generally assume that wolves avoid farmland.	x-x-x add magenta leader forget winter reference exchanges, goal interned poet interests temple goes rushes banks turbines.
24	ED	Because few attack cattle unless extremely desperate, most ranchers will generally assume that wolves avoid farmland.	x-x-x add leader forget winter reference exchanges, goal interned poet interests temple goes rushes banks turbines.
24	NL	Because few canines attack cattle unless extremely desperate, most wolves will generally avoid farmland.	x-x-x add magenta leader forget winter reference exchanges, goal interned poet interests banks turbines.
24	EL	Because few attack cattle unless extremely desperate, most wolves will generally avoid farmland.	x-x-x add leader forget winter reference exchanges, goal interned poet interests banks turbines.
25	ND	Even though hardly any businesses remain open long, countless critics are nevertheless loving that restaurants are scattered around town.	x-x-x extent affect job priority thou plan item, manganese diagram lot contribution shapes duty chromosomes fat ministers agency sale.
25	ED	Even though hardly any remain open long, countless critics are nevertheless loving that restaurants are scattered around town.	x-x-x extent affect job thou plan item, manganese diagram lot contribution shapes duty chromosomes fat ministers agency sale.
25	NL	Even though hardly any businesses remain open long, countless restaurants are nevertheless scattered around town.	x-x-x extent affect job priority thou plan item, manganese diagram lot contribution ministers agency sale.
25	EL	Even though hardly any remain open long, countless restaurants are nevertheless scattered around town.	x-x-x extent affect job thou plan item, manganese diagram lot contribution ministers agency sale.
26	ND	While a few senators are immoral and extremely greedy, most polls are unexpectedly showing that politicians are honest.	x-x-x eat catheter ice seminal mid condition larval, boat hers bit chemotherapy herself fell dimensional me hearts.
26	ED	While a few are immoral and extremely greedy, most polls are unexpectedly showing that politicians are honest.	x-x-x eat thy seminal mid condition larval, boat hers bit chemotherapy herself fell dimensional me hearts.
26	NL	While a few senators are immoral and extremely greedy, most politicians are unexpectedly honest.	x-x-x eat catheter ice seminal mid condition larval, boat hers bit chemotherapy hearts.
26	EL	While a few are immoral and extremely greedy, most politicians are unexpectedly honest.	x-x-x eat thy seminal mid condition larval, boat hers bit chemotherapy hearts.

27	ND	Despite many cages rusting badly and sometimes cracking, most fishermen are thankfully certain that crab traps are fairly sturdy.	x-x-x feet groin rafting likes mid interests reagents, arts extremity ten rectangles teacher poet elbow am agency kisses.
27	ED	Despite many rusting badly and sometimes cracking, most fishermen are thankfully certain that crab traps are fairly sturdy.	x-x-x feet been likes mid interests reagents, arts extremity ten rectangles teacher poet elbow am agency kisses.
27	NL	Despite many cages rusting badly and sometimes cracking, most crab traps are thankfully fairly sturdy.	x-x-x feet groin rafting likes mid interests reagents, arts extremity ten rectangles agency kisses.
27	EL	Despite many rusting badly and sometimes cracking, most crab traps are thankfully fairly sturdy.	x-x-x feet been likes mid interests reagents, arts extremity ten rectangles agency kisses.
28	ND	While nearly all skateboards are overpriced and technologically advanced, a few students are occasionally finding that bicycles are affordable.	x-x-x thy sulked why cannonball ask totalitarianism occasion, sat sentence lot conversation summary goes stamping joy criticizes.
28	ED	While nearly all are overpriced and technologically advanced, a few students are occasionally finding that bicycles are affordable.	x-x-x sulked why cannonball ask totalitarianism occasion, sat sentence lot conversation summary goes stamping joy criticizes.
28	NL	While nearly all skateboards are overpriced and technologically advanced, a few bicycles are occasionally affordable.	x-x-x thy sulked why cannonball ask totalitarianism occasion, sat sentence lot conversation criticizes.
28	EL	While nearly all are overpriced and technologically advanced, a few bicycles are occasionally affordable.	x-x-x sulked why cannonball ask totalitarianism occasion, sat sentence lot conversation criticizes.
29	ND	Although most wildcats are highly reclusive and solitary, all national parks are fully clear that mountain lions are protected under environmental regulations.	x-x-x mark much joy notion checkmate me covenant, won voyage fact him coast image hour covenant most non diagnosis mixed journey interesting.
29	ED	Although most are highly reclusive and solitary, all national parks are fully clear that mountain lions are protected under environmental regulations.	x-x-x mark her notion checkmate me covenant, won voyage fact him coast image hour covenant most non diagnosis mixed journey interesting.
29	NL	Although most wildcats are highly reclusive and solitary, all mountain lions are fully protected under environmental regulations.	x-x-x mark much joy notion checkmate me covenant, won voyage fact him coast image mixed journey interesting.
29	EL	Although most are highly reclusive and solitary, all mountain lions are fully protected under environmental regulations.	x-x-x mark her notion checkmate me covenant, won voyage fact him coast image mixed journey interesting.
30	ND	Even though enough homes are planned and built, most renters are unfortunately complaining that houses are increasing in price.	x-x-x review sector urging his counsel me taste, read voyage why mathematics exterior rule sought me philosophy bar lived.
30	ED	Even though enough are planned and built, most renters are unfortunately complaining that houses are increasing in price.	x-x-x review sector his counsel me taste, read voyage why mathematics exterior rule sought me philosophy bar lived.
30	EL	Even though enough are planned and built, most houses are unfortunately increasing in price.	x-x-x review sector his counsel me taste, read voyage why mathematics philosophy bar lived.
30	NL	Even though enough homes are planned and built, most houses are unfortunately increasing in price.	x-x-x review sector urging his counsel me taste, read voyage why mathematics philosophy bar lived.
31	ND	Even though few birds hatch and ultimately survive, several zoos are nonetheless revealing that eagles are repopulating the forest.	x-x-x target sit ratio farce map securities essence, teacher most arm philosopher numerical turn orally why apologizes me walked.
31	ED	Even though few hatch and ultimately survive, several zoos are nonetheless revealing that eagles are repopulating the forest.	x-x-x target sit farce map securities essence, teacher most arm philosopher numerical turn orally why apologizes me walked.

31	NL	Even though few birds hatch and ultimately survive, several eagles are nonetheless repopulating the forest.	x-x-x target sit ratio farce map securities essence, teacher most arm philosopher apologizes me walked.
31	EL	Even though few hatch and ultimately survive, several eagles are nonetheless repopulating the forest.	x-x-x target sit farce map securities essence, teacher most arm philosopher apologizes me walked.
32	ND	Even though nearly all pets are obedient and affectionate, many veterinarians are unsurprisingly writing that guinea pigs are independent animals.	x-x-x editor ask thee mid exponent map indefinitely, poem numerous me rehabilitating vessels wife rude sea publication address.
32	ED	Even though nearly all are obedient and affectionate, many veterinarians are unsurprisingly writing that guinea pigs are independent animals.	x-x-x editor ask all thee exponent map indefinitely, poem numerous me rehabilitating vessels wife rude sea publication address.
32	NL	Even though nearly all pets are obedient and affectionate, many guinea pigs are unsurprisingly independent animals.	x-x-x editor ask thee mid exponent map indefinitely, poem numerous me rehabilitating publication address.
32	EL	Even though nearly all are obedient and affectionate, many guinea pigs are unsurprisingly independent animals.	x-x-x editor ask all thee exponent map indefinitely, poem numerous me rehabilitating publication address.
33	ND	While almost all rocks are discarded and forgotten about, some tourists are actually insistent that pinecones are taken as treasured souvenirs.	x-x-x nor solve ad diligence act programme angle, hour trouser why doctrine genealogy born inductee mid shift thy menopause apertures.
33	ED	While almost all are discarded and forgotten about, some tourists are actually insistent that pinecones are taken as treasured souvenirs.	x-x-x nor thee diligence act programme angle, hour trouser why doctrine genealogy born inductee mid shift thy menopause apertures.
33	NL	While almost all rocks are discarded and forgotten about, some pinecones are actually taken as treasured souvenirs.	x-x-x nor solve ad diligence act programme angle, hour trouser why doctrine shift thy menopause apertures.
33	EL	While almost all are discarded and forgotten about, some pinecones are actually taken as treasured souvenirs.	x-x-x nor thee diligence act programme angle, hour trouser why doctrine shift thy menopause apertures.
34	ND	Despite most buildings being heavily damaged and abandoned, many newspapers are fortunately reporting that lighthouses are being repaired.	x-x-x file indicates moral genetics brigade bit paragraph, knew consisting lot perceptions integrity ring decapitated sun pages hastened.
34	ED	Despite most being heavily damaged and abandoned, many newspapers are fortunately reporting that lighthouses are being repaired.	x-x-x file indicates genetics brigade bit paragraph, knew consisting lot perceptions integrity ring decapitated sun pages hastened.
34	NL	Despite most buildings being heavily damaged and abandoned, many lighthouses are fortunately being repaired.	x-x-x file indicates moral genetics brigade bit paragraph, knew consisting lot perceptions pages hastened.
34	EL	Despite most being heavily damaged and abandoned, many lighthouses are fortunately being repaired.	x-x-x file indicates genetics brigade bit paragraph, knew consisting lot perceptions pages hastened.
35	ND	Even though much crude oil is being captured underwater, much inspection is regrettably finding that petroleum is spilling into the ocean.	x-x-x circle wall sins who her views portions maturation, poem centering me officiating suppose lady doubtless her firewall band did throw.
35	ED	Even though much is being captured underwater, much inspection is regrettably finding that petroleum is spilling into the ocean.	x-x-x circle wall her views portions maturation, poem centering me officiating suppose lady doubtless her firewall band did throw.
35	NL	Even though much crude oil is being captured underwater, much petroleum is regrettably spilling into the ocean.	x-x-x circle wall sins who her views portions maturation, poem centering me officiating firewall band did throw.
35	EL	Even though much is being captured underwater, much petroleum is regrettably spilling into the ocean.	x-x-x circle wall her views portions maturation, poem centering me officiating firewall band did throw.

36	ND	Although all workers are bored and anxious occasionally, many corporations are mostly claiming that executives are happy at work.	x-x-x arm supposes joy spoon eat cabinet conversation, knew illustration hot pounds simplest lady photograph son angle son seen.
36	ED	Although all are bored and anxious occasionally, many corporations are mostly claiming that executives are happy at work.	x-x-x arm unable spoon eat cabinet conversation, knew illustration hot pounds simplest lady photograph son angle son seen.
36	NL	Although all workers are bored and anxious occasionally, many executives are mostly happy at work.	x-x-x arm supposes joy spoon eat cabinet conversation, knew illustration hot pounds angle son seen.
36	EL	Although all are bored and anxious occasionally, many executives are mostly happy at work.	x-x-x arm unable spoon eat cabinet conversation, knew illustration hot pounds angle son seen.
37	ND	While few artworks are sold at high prices, most sculptors are generally saying that paintings are bought for sentimental reasons.	x-x-x ago gridiron wife register me formerly fight, thou stewarding fat existence remove user dismissed joy resist lie theologians changed.
37	ED	While few are sold at high prices, most sculptors are generally saying that paintings are bought for sentimental reasons.	x-x-x ago hers register me formerly fight, thou stewarding fat existence remove user dismissed joy resist lie theologians changed.
37	NL	While few artworks are sold at high prices, most paintings are generally bought for sentimental reasons.	x-x-x ago gridiron wife register me formerly fight, thou stewarding fat existence resist lie theologians changed.
37	EL	While few are sold at high prices, most paintings are generally bought for sentimental reasons.	x-x-x ago hers register me formerly fight, thou stewarding fat existence resist lie theologians changed.
38	ND	Even though few insects have noticeably long antennae, all zoologists are certainly aware that mayflies are comprised of multiple body sections.	x-x-x stages ask destiny text condensing says imagines, am outtake job operation pages wine facelift thy terrorism try sentence arts medicine.
38	ED	Even though few have noticeably long antennae, all zoologists are certainly aware that mayflies are comprised of multiple body sections.	x-x-x stages ask much condensing says imagines, am outtake job operation pages wine facelift thy terrorism try sentence arts medicine.
38	NL	Even though few insects have noticeably long antennae, all mayflies are certainly comprised of multiple body sections.	x-x-x stages ask destiny text condensing says imagines, am outtake job operation terrorism try sentence arts medicine.
38	EL	Even though few have noticeably long antennae, all mayflies are certainly comprised of multiple body sections.	x-x-x stages ask much condensing says imagines, am outtake job operation terrorism try sentence arts medicine.
39	ND	Although few roads are ever uncrowded or picturesque, some groundskeepers are actually arguing that expressways are beautifully decorated.	x-x-x ask most why tell pressings ill computation, copy telepathically thy attitude offense mail screenplay art legislators spreading.
39	ED	Although few are ever uncrowded or picturesque, some groundskeepers are actually arguing that expressways are beautifully decorated.	x-x-x ask why tell pressings ill computation, copy telepathically thy attitude offense mail screenplay art legislators spreading.
39	NL	Although few roads are ever uncrowded or picturesque, some expressways are actually beautifully decorated.	x-x-x ask most why tell pressings ill computation, copy telepathically thy attitude legislators spreading.
39	EL	Although few are ever uncrowded or picturesque, some expressways are actually beautifully decorated.	x-x-x ask why tell pressings ill computation, copy telepathically thy attitude legislators spreading.
40	ND	While most evaluations are inaccurate as psychological tests, some psychiatrists are surprisingly convinced that inkblots are revealing.	x-x-x wall diffraction why subjection lot consciencious watch, dear pressed old continuation affection gave typecast sky collector.
40	ED	While most are inaccurate as psychological tests, some psychiatrists are surprisingly convinced that inkblots are revealing.	x-x-x wall why subjection lot consciencious watch, dear pressed old continuation affection gave typecast sky collector.

40	NL	While most evaluations are inaccurate as psychological tests, some inkblots are surprisingly revealing.	x-x-x wall diffraction why subjection lot consciencious watch, dear pressed old continuation collector.
40	EL	While most are inaccurate as psychological tests, some inkblots are surprisingly revealing.	x-x-x wall why subjection lot consciencious watch, dear pressed old continuation collector.
41	ND	Despite most arrests being unflattering and extremely embarrassing, many police are ultimately convinced that mugshots are decent pictures.	x-x-x cent elevate lower exacerbate sea technique transporting, root oxygen son contribute precisely fact punt my calmly realized.
41	ED	Despite most being unflattering and extremely embarrassing, many police are ultimately convinced that mugshots are decent pictures.	x-x-x cent some exacerbate sea technique transporting, root oxygen son contribute precisely fact punt my calmly realized.
41	NL	Despite most arrests being unflattering and extremely embarrassing, many mugshots are ultimately decent pictures.	x-x-x cent elevate lower exacerbate sea technique transporting, root oxygen son contribute calmly realized.
41	EL	Despite most being unflattering and extremely embarrassing, many mugshots are ultimately decent pictures.	x-x-x cent some exacerbate sea technique transporting, root oxygen son contribute calmly realized.
42	ND	Even though few farm animals are domesticated or trainable, some ranchers are nonetheless finding that horses are kept in pens.	x-x-x punt sit keeping why articulating war scrimmage, tone informed bed inscription account gave crisis lot flow ask bids.
42	ED	Even though few are domesticated or trainable, some ranchers are nonetheless finding that horses are kept in pens.	x-x-x punt sit why articulating war scrimmage, tone informed bed inscription account gave crisis lot flow ask bids.
42	NL	Even though few farm animals are domesticated or trainable, some horses are nonetheless kept in pens.	x-x-x punt sit keeping why articulating war scrimmage, tone informed bed inscription flow ask bids.
42	EL	Even though few are domesticated or trainable, some horses are nonetheless kept in pens.	x-x-x punt sit why articulating war scrimmage, tone informed bed inscription flow ask bids.
43	ND	Although a few rodents are rabid and not domesticated, many arguments are actually made that rats are excellent pets.	x-x-x eat carbide easy respondents my thy cusps, debt axis my struggle auditor gold obey hot agreement pier.
43	ED	Although a few are rabid and not domesticated, many arguments are actually made that rats are excellent pets.	x-x-x eat while respondents my thy cusps, debt axis my struggle auditor gold obey hot agreement pier.
43	NL	Although a few rodents are rabid and not domesticated, many rats are actually excellent pets.	x-x-x eat carbide easy respondents my thy cusps, debt axis my struggle agreement pier.
43	EL	Although a few are rabid and not domesticated, many rats are actually excellent pets.	x-x-x eat while respondents my thy cusps, debt axis my struggle agreement pier.
44	ND	While hardly any beetles survive over six months, most naturalists are amazingly sure that ladybugs are choosy about their mates.	x-x-x please lie assigns choices goal car likely, arts infomercial job affiliate port coal backpacking gas grunge doubt aside enact.
44	ED	While hardly any survive over six months, most naturalists are amazingly sure that ladybugs are choosy about their mates.	x-x-x please lie much goal car likely, arts infomercial job affiliate port coal backpacking gas grunge doubt aside enact.
44	NL	While hardly any beetles survive over six months, most ladybugs are amazingly choosy about their mates.	x-x-x please lie assigns choices goal car likely, arts infomercial job affiliate grunge doubt aside enact.
44	EL	While hardly any survive over six months, most ladybugs are amazingly choosy about their mates.	x-x-x please lie much goal car likely, arts infomercial job affiliate grunge doubt aside enact.
45	ND	Because so many sanctions undermine American foreign policy goals, few economists are ultimately convinced that tariffs are effective.	x-x-x joy wait several detective Provides merely whom sight, eat meditation dog suggestion synthesis seen footing sea discovery.

45	ED	Because so many undermine American foreign policy goals, few economists are ultimately convinced that tariffs are effective.	x-x-x joy wait few Provides merely whom sight, eat meditation dog suggestion synthesis seen footing sea discovery.
45	NL	Because so many sanctions undermine American foreign policy goals, few tariffs are ultimately effective.	x-x-x joy wait several detective Provides merely whom sight, eat meditation dog suggestion discovery.
45	EL	Because so many undermine American foreign policy goals, few tariffs are ultimately effective.	x-x-x joy wait few Provides merely whom sight, eat meditation dog suggestion discovery.
46	ND	Because few monsters are ever seen in sunlight, many myths are mainly concerned with vampires causing mischief in the shadows.	x-x-x joy hampered six duty post joy mentally, seat quote net degree authority mind radially ethical polymers ill me abandon.
46	ED	Because few are ever seen in sunlight, many myths are mainly concerned with vampires causing mischief in the shadows.	x-x-x joy some six post joy mentally, seat quote net degree authority mind radially ethical polymers ill me abandon.
46	NL	Because few monsters are ever seen in sunlight, many vampires are mainly causing mischief in the shadows.	x-x-x joy hampered six duty post joy mentally, seat quote net degree ethical polymers ill me abandon.
46	EL	Because few are ever seen in sunlight, many vampires are mainly causing mischief in the shadows.	x-x-x joy some six post joy mentally, seat quote net degree ethical polymers ill me abandon.
47	ND	Because many baked goods are stuffed throughout with sausage, few bakers are actually able to make biscuits vegetarian.	x-x-x mean chaise unity mid notices settlement wish stumble, thy raster her appeared mind won goes quantify attentions.
47	ED	Because many are stuffed throughout with sausage, few bakers are actually able to make biscuits vegetarian.	x-x-x mean unity thou settlement wish stumble, thy raster her appeared mind won goes quantify attentions.
47	NL	Because many baked goods are stuffed throughout with sausage, few biscuits are actually vegetarian.	x-x-x mean chaise unity mid notices settlement wish stumble, thy raster her appeared attentions.
47	EL	Because many are stuffed throughout with sausage, few biscuits are actually vegetarian.	x-x-x mean unity thou settlement wish stumble, thy raster her appeared attentions.
48	ND	Because hardly any psalms were written in English, most ministers are reasonably embarrassed that hymns are difficult to memorize.	x-x-x notion bed most salt density sin Charged, step adaptability bad capability counterpart land blown sun increases map blisters.
48	ED	Because hardly any were written in English, most ministers are reasonably embarrassed that hymns are difficult to memorize.	x-x-x notion bed meme density sin Charged, step adaptability bad capability counterpart land blown sun increases map blisters.
48	NL	Because hardly any psalms were written in English, most hymns are reasonably difficult to memorize.	x-x-x notion bed most salt density sin Charged, step adaptability bad capability increases map blisters.
48	EL	Because hardly any were written in English, most hymns are reasonably difficult to memorize.	x-x-x notion bed meme density sin Charged, step adaptability bad capability increases map blisters.

Appendix D

Implicit causality items

Item #	Condition	Item
1	OE	Baritones resent altos because it's hard to find any that aren't competitive.
1	SE	Baritones intimidate altos because it's hard to find any that aren't competitive.
1	OA	Baritones resent altos because it's hard to find ones that aren't competitive.
1	SA	Baritones intimidate altos because it's hard to find ones that aren't competitive.
2	OE	Historians loathe curators because it's challenging to find any that are respectful.
2	SE	Historians enrage curators because it's challenging to find any that are respectful.
2	OA	Historians loathe curators because it's challenging to find ones that are respectful.
2	SA	Historians enrage curators because it's challenging to find ones that are respectful.
3	OE	Geologists like anthropologists because it seems that most are amiable.
3	SE	Geologists please anthropologists because it seems that most are amiable.
3	OA	Geologists like anthropologists because it seems that they are amiable.
3	SA	Geologists please anthropologists because it seems that they are amiable.
4	OE	Campers trust hikers because it's difficult to find any that are dishonest.
4	SE	Campers delight hikers because it's difficult to find any that are dishonest.
4	OA	Campers trust hikers because it's difficult to find ones that are dishonest.
4	SA	Campers delight hikers because it's difficult to find ones that are dishonest.
5	OE	Brewers admire bartenders because it seems that most are loyal to employers.
5	SE	Brewers impress bartenders because it seems that most are loyal to employers.
5	OA	Brewers admire bartenders because it seems that they are loyal to employers.
5	SA	Brewers impress bartenders because it seems that they are loyal to employers.
6	OE	Fishermen hate sailors because it seems that many are unfriendly.
6	SE	Fishermen enrage sailors because it seems that many are unfriendly.
6	OA	Fishermen hate sailors because it seems that they are unfriendly.
6	SA	Fishermen enrage sailors because it seems that many are unfriendly.
7	OE	Editors mesmerize directors because it appears that most are easily impressed.
7	SE	Editors worship directors because it appears that most are easily impressed.
7	OA	Editors mesmerize directors because it appears that they are easily impressed.
7	SA	Editors worship directors because it appears that they are easily impressed.
8	OE	Surfers deplore paddleboarders because it seems that most don't want to share waves.
8	SE	Surfers perturb paddleboarders because it seems that most don't want to share waves.
8	OA	Surfers deplore paddleboarders because it seems that they don't want to share waves.
8	SA	Surfers perturb paddleboarders because it seems that they don't want to share waves.

9	OE	Professors fear lecturers because it appears that many worry about job shortages.
9	SE	Professors scare lecturers because it appears that many worry about job shortages.
9	OA	Professors fear lecturers because it appears that they worry about job shortages.
9	SA	Professors scare lecturers because it appears that they worry about job shortages.
10	OE	Consultants enjoy lobbyists because it's difficult to find any that are disagreeable.
10	SE	Consultants amuse lobbyists because it's difficult to find any that are disagreeable.
10	OA	Consultants enjoy lobbyists because it's difficult to find ones that are disagreeable.
10	SA	Consultants amuse lobbyists because it's difficult to find ones that are disagreeable.
11	OE	Musicians idolize photographers because it's difficult to find any that don't respect all artforms.
11	SE	Musicians inspire photographers because it's difficult to find any that don't respect all artforms.
11	OA	Musicians idolize photographers because it's difficult to find ones that don't respect all artforms.
11	SA	Musicians inspire photographers because it's difficult to find ones that don't respect all artforms.
12	OE	Senators dread congressmembers because it appears that most hate talking to elected officials.
12	SE	Senators weary congressmembers because it appears that most hate talking to elected officials.
12	OA	Senators dread congressmembers because it appears that they hate talking to elected officials.
12	SA	Senators weary congressmembers because it appears that they hate talking to elected officials.
13	OE	Surgeons revere nurses because it appears that many value patient satisfaction.
13	SE	Surgeons dazzle nurses because it appears that many value patient satisfaction.
13	OA	Surgeons revere nurses because it appears that they value patient satisfaction.
13	SA	Surgeons dazzle nurses because it appears that they value patient satisfaction.
14	OE	Farmers adore environmentalists because it seems that most care about water policy.
14	SE	Farmers excite environmentalists because it seems that most care about water policy.
14	OA	Farmers adore environmentalists because it seems that they care about water policy.
14	SA	Farmers excite environmentalists because it seems that they care about water policy.
15	OE	Stockbrokers disdain accountants because it seems that many prefer to work alone.
15	SE	Stockbrokers irritate accountants because it seems that many prefer to work alone.
15	OA	Stockbrokers disdain accountants because it seems that they prefer to work alone.
15	SA	Stockbrokers irritate accountants because it seems that they prefer to work alone.
16	OE	Chiropractors venerate opticians because it's hard to find any that don't glorify healthcare workers.
16	SE	Chiropractors mesmerize opticians because it's hard to find any that don't glorify healthcare workers.
16	OA	Chiropractors venerate opticians because it's hard to find ones that don't glorify healthcare workers.
16	SA	Chiropractors mesmerize opticians because it's hard to find ones that don't glorify healthcare workers.
17	OE	Gamblers deplore loan sharks because it's difficult to find any that have sufficient money.
17	SE	Gamblers frustrate loan sharks because it's difficult to find any that have sufficient money.
17	OA	Gamblers deplore loan sharks because it's difficult to find ones that have sufficient money.
17	SA	Gamblers frustrate loan sharks because it's difficult to find ones that have sufficient money.
18	OE	Illustrators treasure authors because it seems that many love collaborations.
18	SE	Illustrators gladden authors because it seems that many love collaborations.
18	OA	Illustrators treasure authors because it seems that they love collaborations.
18	SA	Illustrators gladden authors because it seems that they love collaborations.

19	OE	Newscasters distrust meteorologists because it appears that many are difficult to work with.
19	SE	Newscasters disappoint meteorologists because it appears that many are difficult to work with.
19	OA	Newscasters distrust meteorologists because it appears that they are difficult to work with.
19	SA	Newscasters disappoint meteorologists because it appears that they are difficult to work with.
20	OE	Magicians alienate clairvoyants because it seems that most are fearful.
20	SE	Magicians unsettle clairvoyants because it seems that most are fearful.
20	OA	Magicians alienate clairvoyants because it seems that they are fearful.
20	SA	Magicians unsettle clairvoyants because it seems that they are fearful.
21	OE	Osteopaths fear paramedics because it seems that many are insecure.
21	SE	Osteopaths frighten paramedics because it seems that many are insecure.
21	OA	Osteopaths fear paramedics because it seems that they are insecure.
21	SA	Osteopaths frighten paramedics because it seems that they are insecure.
22	OE	Composers exalt singers because it's challenging to find any that don't work with choirs.
22	SE	Composers spellbind singers because it's challenging to find any that don't work with choirs.
22	OA	Composers exalt singers because it's challenging to find ones that don't work with choirs.
22	SA	Composers spellbind singers because it's challenging to find ones that don't work with choirs.
23	OE	Bodyguards dislike paparazzi because it seems that many are ill-tempered.
23	SE	Bodyguards displease paparazzi because it seems that many are ill-tempered.
23	OA	Bodyguards dislike paparazzi because it seems that they are ill-tempered.
23	SA	Bodyguards displease paparazzi because it seems that they are ill-tempered.
24	OE	Hipsters detest baristas because it appears that most are intolerant.
24	SE	Hipsters infuriate baristas because it appears that most are intolerant.
24	OA	Hipsters detest baristas because it appears that they are intolerant.
24	SA	Hipsters infuriate baristas because it appears that they are intolerant.
25	OE	Golfers respect jockeys because it seems that most have a high opinion of other sports.
25	SE	Golfers charm jockeys because it seems that most have a high opinion of other sports.
25	OA	Golfers respect jockeys because it seems that they have a high opinion of other sports.
25	SA	Golfers charm jockeys because it seems that they have a high opinion of other sports.
26	OE	Astronomers despise chemists because it appears that few respect other sciences.
26	SE	Astronomers miff chemists because it appears that few respect other sciences.
26	OA	Astronomers despise chemists because it appears that they disrespect other sciences.
26	SA	Astronomers miff chemists because it appears that they disrespect other sciences.
27	OE	Comedians alienate producers because it's difficult to find any that are warmhearted.
27	SE	Comedians exasperate producers because it's difficult to find any that are warmhearted.
27	OA	Comedians alienate producers because it's difficult to find ones that are warmhearted.
27	SA	Comedians exasperate producers because it's difficult to find ones that are warmhearted.
28	OE	Nuns hate missionaries because it's hard to find any that are religiously-tolerant.
28	SE	Nuns disturb missionaries because it's hard to find any that are religiously-tolerant.
28	OA	Nuns hate missionaries because it's hard to find ones that are religiously-tolerant.
28	SA	Nuns disturb missionaries because it's hard to find ones that are religiously-tolerant.

Appendix E

GloVe Cosine Measures, NPE Maze Experiments

First Maze Experiment

Item #	NP1	NP2	Cosine	Condition
1	pre-meds	undergraduates	0.357789	NL
1	pre-meds	reports	-0.243899	ND
2	peacocks	swans	0.637104	NL
2	peacocks	studies	-0.007248	ND
3	manuscripts	novels	0.41969	NL
3	manuscripts	editors	0.352049	ND
4	boats	kayaks	0.695117	NL
4	boats	captains	0.4392	ND
5	buses	trains	0.683653	NL
5	buses	investigators	0.137331	ND
6	ships	sailboats	0.421055	NL
6	ships	onlookers	0.079141	ND
7	blazes	wildfires	0.609649	NL
7	blazes	rangers	0.197099	ND
8	whales	sharks	0.745733	NL
8	whales	biologists	0.407072	ND
9	groundhogs	bears	0.33075	NL
9	groundhogs	assessments	0.027649	ND
10	animals	insects	0.581973	NL
10	animals	scientists	0.374125	ND
11	people	climbers	0.278133	NL
11	people	statistics	0.311652	ND
12	people	adventurers	0.261677	NL
12	people	forums	0.334394	ND
13	vegetables	tomatoes	0.785454	NL
13	vegetables	books	0.176886	ND
14	pollution	carbon	0.539953	NL
14	pollution	research	0.295135	ND
15	animals	elephants	0.67496	NL

15	animals	surveys	0.195938	ND
16	reptiles	snakes	0.74472	NL
16	reptiles	toddlers	0.247005	ND
17	people	homes	0.420326	NL
17	people	maps	0.270251	ND
18	raccoons	squirrels	0.836117	NL
18	raccoons	neighborhoods	0.215244	ND
19	octopuses	turtles	0.522853	NL
19	octopuses	swimmers	0.302659	ND
20	wetsuits	surfboards	0.594936	NL
20	wetsuits	teenagers	0.143929	ND
21	mountains	volcanoes	0.562845	NL
21	mountains	geologists	0.227632	ND
22	deer	reindeer	0.533743	NL
22	deer	hunters	0.642896	ND
23	dinosaurs	triceratops	0.550231	NL
23	dinosaurs	paleontologists	0.530592	ND
24	canines	wolves	0.4462	NL
24	canines	ranchers	0.270044	ND
25	businesses	restaurants	0.474479	NL
25	businesses	critics	0.210338	ND
26	senators	politicians	0.622028	NL
26	senators	polls	0.448699	ND
27	cages	crab traps	0.423324	NL
27	cages	fishermen	0.177662	ND
28	skateboards	bicycles	0.577148	NL
28	skateboards	students	0.03813	ND
29	wildcats	tigers	0.578976	NL
29	wildcats	laws	0.047086	ND
30	homes	houses	0.805709	NL
30	homes	voters	0.263081	ND
31	birds	eagles	0.609863	NL
31	birds	investigations	0.140146	ND
32	dogs	cats	0.831176	NL
32	dogs	blogs	0.180939	ND
33	rocks	pinecones	0.256234	NL
33	rocks	tourists	0.256802	ND
34	buildings	lighthouses	0.350564	NL
34	buildings	newspapers	0.2738	ND
35	oil	petroleum	0.707457	NL
35	oil	fieldwork	0.022859	ND
36	workers	executives	0.499764	NL
36	workers	corporations	0.475025	ND
37	artworks	paintings	0.781508	NL
37	artworks	salespeople	0.098742	ND
38	insects	mayflies	0.457987	NL
38	insects	preteens	0.111028	ND
39	roads	expressways	0.531989	NL
39	roads	groundskeepers	0.087084	ND
40	evaluations	inkblots	0.127587	NL

40	evaluations	proctologists	0.01255	ND
41	arrests	mugshots	0.38592	NL
41	arrests	police	0.661076	ND
42	farm animals	horses	0.445406	NL
42	farm animals	teachers	0.194894	ND
43	rodents	rats	0.775844	NL
43	rodents	kids	0.165078	ND
44	beetles	dragonflies	0.585458	NL
44	beetles	scoutmasters	0.132408	ND
45	sanctions	tariffs	0.415315	NL
45	sanctions	economists	0.218686	ND
46	monsters	vampires	0.612695	NL
46	monsters	legends	0.408979	ND
47	baked goods	biscuits	0.702098	NL
47	baked goods	bakers	0.456492	ND
48	psalms	hymns	0.730552	NL
48	psalms	parishioners	0.305006	ND

Second Maze Experiment

Item #	NP1	NP2	Cosine	Condition
7	blazes	firefighters	0.421253	ND
8	whales	great white sharks	0.745733	NL
8	whales	marine biologists	0.407072	ND
13	vegetables	gardeners	0.416086	ND
16	reptiles	turtles	0.656174	ND
19	octopuses	lobsters	0.601382	NL
21	mountains	hikers	0.433492	ND
29	wildcats	mountain lions	0.459178	NL
29	wildcats	national parks	0.138252	ND
30	homes	renters	0.551621	ND
31	birds	zoos	0.379758	ND
32	dogs	guinea pigs	0.616212	NL
32	dogs	veterinarians	0.468964	ND
35	oil	inspection	0.213827	ND
37	artworks	sculptors	0.501209	ND
38	insects	zoologists	0.226370	ND
40	evaluations	psychiatrists	0.375192	ND
42	farm animals	ranchers	0.194894	ND
43	rodents	arguments	0.027147	ND
44	beetles	ladybugs	0.622864	NL
44	beetles	naturalists	0.231841	ND
46	monsters	myths	0.354267	ND
48	psalms	ministers	0.256101	ND

Bibliography

- Abusch, Dorit. 2009. Presupposition triggering from alternatives. *Journal of Semantics* 27:37–80.
- Amaral, Patricia, Craige Roberts, and E. Allyn Smith. 2007. Review of *The Logic of Conventional Implicatures* by Chris Potts. *Linguistics and Philosophy* 30:707–749.
- Anand, Pranav, and Valentine Hacquard. 2013. Epistemics and attitudes. *Semantics & Pragmatics* 6:1–59.
- Anand, Pranav, and Valentine Hacquard. 2014. Factivity, belief and discourse. In *The Art and Craft of Semantics: A Festschrift for Irene Heim*, ed. Luka Crnić and Uli Sauerland, volume 1, 69–90. Cambridge, MA: MIT Working Papers in Linguistics.
- AnderBois, Scott. 2014. The semantics of sluicing: Beyond truth-conditions. *Language* 90:887–926.
- AnderBois, Scott, Adrian Brasoveanu, and Robert Henderson. 2015. At-issue proposals and appositive impositions in discourse. *Journal of Semantics* 32:93–138.
- Anderson, John R. 2007. *How can the human mind occur in the physical universe?*. Oxford University Press.
- Anderson, John R., and Lael J. Schooler. 1991. Reflections of the environment in memory. *Psychological Science* 2.
- Ariel, Mira. 1990. *Accessing noun-phrase antecedents*. Routledge.
- Arnold J.E., Brown-Schmidt S., Eisenband J.G., and Trueswell J.C. 2000. The rapid use of gender information: Evidence of the time course of pronoun resolution from eyetracking. *Cognition* 76:B13–B26.
- Arnon, Inbal, and Neal Snider. 2010. More than words: Frequency effects for multi-word phrases. *Journal of Memory and language* 62:67–82.
- Arregui, Ana, Charles Clifton, Lyn Frazier, and Keir Moulton. 2006. Processing elided verb phrases with flawed antecedents. *Journal of Memory and Language* 55:232–246.

- Asher, Nicholas. 1987. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy* 10:125–197.
- Asher, Nicholas. 1993. *Reference to abstract objects in discourse*. Kluwer Academic Publishers: Dordrecht.
- Asher, Nicholas. 2008. Troubles on the right frontier. Ms.
- Asher, Nicholas, and Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.
- Asher, Nicholas, and Laure Vieu. 2005. Subordinating and coordinating discourse relations. *Lingua* 115:591–610.
- Badecker, William, and Kathleen Staub. 2002. The processing role of structural constraints on the interpretation of pronouns and anaphors. *Journal of Experimental Psychology: Learning, Memory and Cognition* 28:748–769.
- Barker, Chris. 2007. Parasitic scope. *Linguistics and Philosophy* 30:407–444.
- Barker, Chris. 2013. Scopability and sluicing. *Linguistics and Philosophy* 36:187–223.
- Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68:255–278.
- Barros, Matthew. 2014. Sluicing and identity in ellipsis. Doctoral Dissertation, Rutgers University.
- Barros, Matthew, Patrick D. Elliott, and Gary Thoms. 2013. More variation in island repair: Clausal vs. non-clausal islands. Ms.
- Bartsch, Renate. 1973. Negative transportation. *Linguistische Berichte* 27.
- Beaver, David. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.
- Bechhofer, Robin. 1976. Reduced wh-questions. In *Harvard studies in syntax and semantics*, ed. Jorge Hankamer and Judith L. Aissen, volume 2, 31–67.
- Belanger, Suzanne Michelle. 2014. Regeneration in recall and verb phrase ellipsis. Doctoral Dissertation, University of Toronto.
- Boyce, Veronica, Richard Futrell, and Roger P. Levy. 2020. Maze made easy: Better and easier measures of incremental processing difficulty. *Journal of Memory and Language* 111.
- Bransford, J.D., J.R. Barclay, and J.J. Franks. 1972. Sentence memory: A constructive versus interpretive approach. *Cognitive Psychology* 3:193–209.

- Brasoveanu, Adrian, and Jakub Dotlacil. 2020. *Computational cognitive modeling and linguistic theory*. Springer.
- Broadbent, Donald E. 1958. *Perception and communication*. New York: Pergamon Press.
- Brown, R., and D. Fish. 1983. The psychological causality implicit in language. *Cognition* 13:237–273.
- Brown-Schmidt, Sarah, Donna K. Byron, and Michael K. Tanenhaus. 2005. Beyond salience: Interpretation of personal and demonstrative pronouns. *Journal of Memory and Language* 53:292–313.
- Bürkner, PC. 2017. brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software* 80:1–28.
- Bürkner, PC. 2020. Parameterization of response distributions in brms. [Http://paul-buerkner.github.io/brms/articles/brms_families.html](http://paul-buerkner.github.io/brms/articles/brms_families.html).
- Cantor, Sara. 2013. An exploration of non-grammatical sluicing constructions. Master's Thesis.
- Caramazza, Alfonso, E. Grober, Catherine Garvey, and J. Yates. 1977. Comprehension of anaphoric pronouns. *Journal of verbal learning and verbal behavior* 16:601–609.
- Carlson, Katy. 2002. Parallelism and prosody in the processing of ellipsis sentences. Doctoral Dissertation, University of Massachusetts Amherst.
- Carlson, Katy, Michael Walsh Dickey, Lyn Frazier, and Jr. Charles Clifton. 2009. Information structure expectations in sentence comprehension. *Quarterly Journal of Experimental Psychology* 62:114–139.
- Chambers, Craig G., and Ron Smyth. 1998. Structural parallelism and discourse coherence: A test of centering theory. *Journal of Memory and Language* 39:593–608.
- Chomsky, Noam, and George A. Miller. 1963. Introduction to the formal analysis of natural languages. In *Handbook of mathematical psychology*, ed. R. R. Bush and E. Galanter, volume 2, 444–466. New York: Wiley.
- Chung, Sandra. 1994. Wh-agreement and "referentiality" in Chamorro. *Linguistic Inquiry* 15:1–44.
- Chung, Sandra. 2013. Syntactic identity in sluicing: How much and why. *Linguistic Inquiry* 44:1–44.
- Chung, Sandra, William A. Ladusaw, and Jim McCloskey. 1995. Sluicing in logical form. *Natural Language Semantics* 3:239–282.

- Cinque, Guglielmo. 1990. *Types of A'-dependencies*. Cambridge: MIT Press.
- Clark, H.H., and C.J. Sengul. 1979. In search of referents for nouns and pronouns. *Memory Cognition* 7:35–41.
- Clifton, Charles, and Lyn Frazier. 1989. Comprehending sentences with long distance dependencies. In *Linguistic Structure in Language Processing*, ed. G.N. Carlson and M. Tanenhaus, 273–317.
- Clifton, Charles, and Adrian Staub. 2008. Parallelism and competition in syntactic ambiguity resolution. *Language and Linguistics Compass* 2:234–250.
- Collins, Chris, and Paul M. Postal. 2014. *Classical NEG raising: An essay on the syntax of negation*. MIT Press.
- Collins, James, Daria Popova, Ivan Sag, and Thomas Wasow. 2014. Sluicing and the inquisitive potential of appositives. Ms.
- Condoravdi, Cleo, and Sven Lauer. 2011. Performative verbs and performative acts. In *Sinn and Bedeutung 15: Proceedings of the 2010 annual conference of the Gesellschaft für Semantik*, ed. Ingo Reich, Eva Horch, and Dennis Pauly, 149–164. Saarbrücken: Universaar Saarland University Press.
- Condoravdi, Cleo, and Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. In *Empirical issues in syntax and semantics 9*, ed. Christopher Pinon, 37–58. Oxford: Oxford University Press.
- Copley, Bridget. 2014. Causal chains for futurates. In *Future times, future tenses*, ed. Philippe De Brabanter, Mikhail Kissine, and Saghie Sharifzadeh, 72–86. Oxford: Oxford University Press.
- Corbett, A.T., and F.R. Chang. 1983. Pronoun disambiguation: Accessing potential antecedents. *Memory Cognition* 11:283–294.
- Corrigan, R. 2001. Implicit causality in language: Event participants and their interactions. *Journal of Language and Social Psychology* 20:285–320.
- Corrigan, R. 2003. Preschoolers' and adults' attributions of who causes interpersonal events. *Infant and child development* 12:305–328.
- van Craenenbroeck, Jeroen. 2010. Invisible last resort: A note on clefts as the underlying source for sluicing. *Lingua* 1,714–1726.
- Crain, Stephen, and Janet Dean Fodor. 1985. How can grammars help parsers? In *Natural language parsing*, ed. Arnold M. Zwicky David R. Dowty, Lauri Karttunen, 94–128. Cambridge University Press.

- Crawley, Rosalind A., Rosemary J. Stevenson, and David Kleinman. 1990. The use of heuristic strategies in the interpretation of pronouns. *Journal of Psycholinguistic Research* 4:245–264.
- Dalrymple, Mary, Stuart M. Shieber, and Fernando C.N. Pereira. 1991. Ellipsis and higher-order unification. *Linguistics and Philosophy* 14:399–452.
- Dayal, Veneeta, and Roger Schwarzschild. 2010. Definite inner antecedents and wh-correlates in sluices. In *Rutgers working papers in linguistics*, ed. Peter Staroverov, Daniel Altschuler, Aaron Braver, Carlos Fasola, and Sarah Murray, volume 3, 92–114. New Brunswick, NJ: LGSA.
- Declerck, Renaat. 1995. The problem of *not...until*. *Linguistics* 33:51–98.
- Del Gobbo, Francesca. 2007. On the syntax and semantics of appositive relative clauses. In *Parentheticals*, ed. Nicole Dehé and Yordanka Kavalova, 173–201. Amsterdam: John Benjamins.
- van Dijk, Teun A., and Walter Kintsch. 1983. *Strategies of discourse comprehension*. New York: Academic Press.
- Dowty, David. 1979. *Word meaning and Montague grammar*. Dordrecht: Reidel.
- Drummond, Alex. 2014. IbexFarm. <http://spellout.net/ibexfarm>.
- Dyke, Julia A. Van, and Richard L. Lewis. 2003. Distinguishing effects of structure and decay on attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language* 49:285–316.
- Dyke, Julia A. Van, and Brian McElree. 2006. Retrieval interference in sentence comprehension. *Journal of Memory and Language* 55:157–166.
- Elliott, Patrick D., and Yasutada Sudo. 2016. E-type readings of quantifiers under ellipsis: Consequences for the identity condition. In *Sinn und Bedeutung (SuB) 20*, ed. Nadine Bade, Polina Berezovskaya, and Anthea Schöller, 149–164. University of Tübingen.
- Erteschik-Shir, Nomi. 1977. *On the nature of island constraints*. Bloomington: Indiana University Linguistics Club.
- Evans, Gareth. 1977. Pronouns, quantifiers, and relative clauses. *Canadian Journal of Philosophy* 7:467–536.
- Fiengo, Robert, and Robert May. 1994. *Indices and identity*. MIT Press.
- Fillmore, Charles J. 1963. The position of embedding transformations in a grammar. *Word* 19:208–231.

- Fischler, Ira, and Paul A. Bloom. 1979. Automatic and attentional processes in the effects of sentence contexts on word recognition. *Journal of Verbal Learning and Verbal Behavior* 18:1–20.
- Fitzpatrick, Justin. 2005. The whys and how comes of presupposition and NPI licensing in questions. In *24th West Coast Conference on Formal Linguistics (WCCFL)*, ed. John Alderete, 138–145. Somerville, MA: Cascadilla Proceedings Project.
- Fodor, Janet Dean. 1978. Parsing strategies and constraints on transformation. *Linguistic Inquiry* 9:427–473.
- Foraker, Stephani, and Brian McElree. 2006. The role of prominence in pronoun resolution: Active vs. passive representations. *Journal of Memory and Language* 56:357–383.
- Forster, Kenneth I. 1981. Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* 33:465–495.
- Forster, Kenneth I., C. Guerrera, and L. Elliot. 2009. The maze task: Measuring forced incremental sentence processing time. *Behavior Research Methods* 41:163–171.
- Frazier, Lyn. 1978. On comprehending sentences: Syntactic parsing strategies. Doctoral Dissertation, University of Connecticut, Storrs.
- Frazier, Lyn. 1987. Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory* 5:519–559.
- Frazier, Lyn. 2013. A recycling approach to processing ellipsis. In *Diagnosing syntax*, ed. Lisa Lai-Shen Cheng and Norbert Corver, 485–501. Oxford University Press.
- Frazier, Lyn, and Charles Clifton. 1989. Successive cyclicity in the grammar and the parser. *Language and Cognitive Processes* 4:93–126.
- Frazier, Lyn, and Charles Clifton. 1998. Comprehension of sluiced sentences. *Language and Cognitive Processes* 13:499–520.
- Frazier, Lyn, and Charles Clifton. 2001. Parsing coordinates and ellipsis: Copy α . *Syntax* 4:1–22.
- Frazier, Lyn, and Charles Clifton. 2005. The syntax-discourse divide: Processing ellipsis. *Syntax* 8:121–174.
- Frazier, Lyn, and Charles Clifton. 2010. Imperfect ellipsis: Antecedents beyond syntax? *Syntax* 13:279–297.
- Frazier, Lyn, and Charles Clifton. 2015. Without his shirt off he saved the child from almost drowning: Interpreting uncertain input. *Language, Cognition, and Neuroscience* 30:635–647.

- Frazier, Lyn, and Jack Duff. 2019. Repair or accommodation? Split antecedent ellipsis and the limits of repair. *Glossa* 4:78.
- Frederiksen, John R. 1981. Understanding anaphora: Rules used by readers in assigning pronominal referents. *Discourse Processes* 4:323–347.
- Freedman, Sandra E., and Kenneth I Forster. 1985. The psychological status of overgenerated sentences. *Cognition* 19:101–131.
- Gabry, J., D. Simpson, A. Vehtari, M. Betancourt, and A. Gelman. 2019. Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A* 182:389–402.
- Gajewski, Jon R. 2007. Neg-raising and polarity. *Linguistics and Philosophy* 30:289–328.
- Garnham, Alan. 2001. *Mental models and the interpretation of anaphora*. Psychology Press Ltd.
- Garrod, S., and T. Trabasso. 1973. A dualmemory information processing interpretation of sentence comprehension. *Journal of Verbal Learning and Verbal Behavior* 12:155–167.
- Garvey, Catherine, and Alfonso Caramazza. 1974. Implicit causality in verbs. *Linguistic Inquiry* 5:459–464.
- Gazdar, Gerald. 1979. *Pragmatics: Implicature, presupposition, and logical form*. New York: Academic.
- Geach, Peter T. 1962. *Reference and generality*. Ithaca: Cornell University Press.
- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian data analysis*. CRC Press.
- Gernsbacher, M.A. 1989. Mechanisms that improve referential access. *Cognition* 32:99–156.
- Geurts, Bart. 2009. *Quantity implicatures*. Cambridge University Press.
- Giannakidou, Anastasia. 2002. UNTIL, aspect, and negation: A novel argument for two *untils*. In *Semantics and Linguistic Theory (SALT)*, ed. B. Jackson, volume 12, 84–103. Ithaca, NY: CLC Publications, Cornell University.
- Gibson, Edward. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68:1–76.
- Gibson, Edward, N. Pearlmutter, E. Canseco-Gonzalez, and G. Hickok. 1996. Recency preferences in the human sentence processing mechanism. *Cognition* 59:23–59.
- Ginzburg, Jonathan. 1996. Dynamics and the semantics of dialog. In *Language, Logic, and Computation, Volume 1*, ed. J. Seligman. Stanford: CSLI.

- Ginzburg, Jonathan. 2012. *The interactive stance*. New York: Oxford University Press.
- Ginzburg, Jonathan, and Ivan A. Sag. 2001. *Interrogative investigations*. Chicago: University of Chicago Press.
- Givón, Thomas. 1983. *Topic continuity in discourse: A quantitative cross-language study*. Amsterdam: John Benjamins.
- Goldberg, Adele E., and Laura A. Michaelis. 2017. One among many: Anaphoric one and its relationship with numeral one. *Cognitive Science* 41:233–258.
- van Gompel, Roger P.G., M.J. Pickering MJ, J. Pearson, and Simon P. Liversedge. 2005. Evidence against competition during syntactic ambiguity resolution. *Journal of Memory and Language* 52:284–307.
- Gordon, Peter C., and D. Chan. 1995. Pronouns, passives, and discourse coherence. *Journal of Memory and Language* 34:216–231.
- Gordon, Peter C., Barbara J. Grosz, and Laura A. Gilliom. 1993. Pronouns, names, and the centering of attention in discourse. *Cognitive Science* 17:311–347.
- Greene, Steven B., Gail McKoon, and Roger Ratcliff. 1992. Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18:266–283.
- Gribanova, Vera. 2013. Copular clauses, clefts, and putative sluicing in Uzbek. *Language* 89:830–882.
- Groenendijk, Jeroen, and Floris Roelofsen. 2009. Inquisitive semantics and pragmatics. Workshop on Language, Communication, and Rational Agency. Stanford.
- Groenendijk, Jeroen, and Martin Stokhof. 1990. Dynamic Montague grammar. In *Papers from the second symposium on logic and language*, ed. L. Kalman and L. Polos, 3–48. Budapest: Akademiai Kiadoo.
- Grosz, Barbara J., Aravind K. Joshi, and Scott Weinstein. 1995. Centering: A framework for modelling the local coherence of discourse. *IRCS Technical Reports Series* 116:1–30.
- Grosz, Barbara J., and Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics* 12:175–204.
- Gullifer, Jason W. 2004. Processing reverse sluicing: A contrast with processing filler-gap dependencies. In *North East Linguistics Society (NELS) 34*, 1–20.
- Gundel, Jeanette K., Kaja Borthen, and Thornstein Fretheim. 1999. The role of context in pronominal reference to higher order entities in english and norwegian. In *Modeling and Using Context: Proceedings from the Second International and Interdisciplinary Conference*, ed. P. Bouquet.

- Gundel, Jeanette K., Nancy Hedberg, and Ron Zacharski. 1990. Givenness, implicature, and the form of referring expressions. In *Sixteenth Annual Meeting of the Berkeley Linguistics Society*, 442–453.
- Hale, John. 2003. The information conveyed by words in sentences. *Journal of Psycholinguistic Research* 32:101–123.
- Halliday, Michael A.K. 1967. Notes on transitivity and theme in English: Parts 1, 2, 3. *Journal of Linguistics* 3.
- Hankamer, Jorge. 1977. Multiple analyses. In *Mechanisms of syntactic chains*, ed. Charles N. Li, 583–607. Austin: University of Texas Press.
- Hankamer, Jorge, and Ivan Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry* 7:391–426.
- Hardt, Daniel. 1993. Verb phrase ellipsis: Form, meaning, and processing. Doctoral Dissertation, University of Pennsylvania.
- Hardt, Daniel. 1999. Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosopher* 22:187–221.
- Hardt, Daniel, and Maribel Romero. 2004. Ellipsis and the structure of discourse. *Journal of Semantics* 21:375–414.
- Harris, Jesse. 2015. Structure modulates similarity-based interference in sluicing: An eye tracking study. *Frontiers of Psychology* 6.
- Harris, Jesse. 2019. Alternatives on demand and locality: Resolving discourse-linked wh-phrases in sluiced structures. In *Grammatical approaches to language processing*, ed. Katy Carlson, Charles Clifton Jr., and Janet Dean Fodor.
- Harris, Jesse, and Katy Carlson. 2015. Information structure preferences in focus-sensitive ellipsis: How defaults persist. *Frontiers of Psychology* 6.
- Harris, Jesse, and Katy Carlson. 2016. Keep it local (and final): remnant preferences for let alone ellipsis. *Quarterly Journal of Experimental Psychology* 69:1278–1301.
- Hartman, Jeremy. 2009. When e-GIVENness over-predicts identity. Handout presented at the Fourth Brussels Conference on Generative Linguistics.
- Hartman, Jeremy. 2011. The semantic uniformity of traces: Evidence from ellipsis parallelism. *Linguistic Inquiry* 42:367–388.
- Hartshorne, Joshua K. 2013. What is implicit causality? *Language, Cognition, and Neuroscience* 29.

- Hartshorne, Joshua K., and Jesse Snedeker. 2013. Verb argument structure predicts implicit causality: The advantages of finer-grained semantics. *Language and Cognitive Processes* 28:1474–1508.
- Heim, Irene. 1983a. File change semantics and the familiarity theory of definiteness. In *Meaning, use and interpretation of language*, ed. Rainer Baeuerle, Christoph Schwarze, and Arnim von Stechow, 164–189. Berlin: De Gruyter.
- Heim, Irene. 1983b. On the projection problem for presuppositions. 114–126. Palo Alto, CA: CSLI Publications.
- Heim, Irene. 1990. E-type pronouns and donkey anaphora. *Linguistics and Philosophy* 13:137–138.
- van Heuven, W.J., P. Mandera, E. Keuleers, and M. Brysbaert. 2014. SUBTLEX-UK: A new and improved word frequency database for British English. *Quarterly Journal of Experimental Psychology* 67.
- Higginbotham, James. 1975. On assertive predicates. In *Syntax and semantics*, ed. John P. Kimball, volume 4, 91–124. New York : Academic Press.
- Higginbotham, James. 2003. Remembering, imagining, and the first person. In *Epistemology of language*, ed. Alex Barber, 496–533. Oxford University Press.
- Hobbs, Jerry R. 1978. Resolving pronoun references. *Lingua* 44:311–338.
- Hofmeister, Philip. 2007. Retrieval and gradience in filler-gap dependencies. Ms.
- Hofmeister, Philip. 2011. Representational complexity and memory retrieval in language comprehension. *Language and Cognitive Processes* 26:376–405.
- Hofmeister, Philip, Laura Staum, and Ivan A. Sag. 2013. Islands in the grammar? Standards of evidence. In *Experimental syntax and island effects*, 42–63.
- Horn, Laurence R. 1978. Remarks on neg-raising. *Semantics and Pragmatics* 9:129–220.
- Hwang, Chung Hee. 1992. A logical approach to narrative understanding. Doctoral Dissertation, University of Alberta, Edmonton.
- Iatridou, Sabine, and Ivy Sichel. 2011. Negative DP's, A-movement and scope diminishment. *Linguistic Inquiry* 42:595–629.
- Itti, Laurent, Christof Koch, and Ernst Niebur. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20:1254–1259.
- Jackendoff, Ray. 1977. *X-Bar syntax: A study of phrase structure*. Cambridge: MIT Press.

- Jarvella, Robert J. 1971. Syntactic processing of connected speech. *Journal of Verbal Learning and Verbal Behavior* 10:409–416.
- Just, M.A., and P.A. Carpenter. 1980. A theory of reading: From eye-fixations to comprehension. *Psychological Review* 87:329–354.
- Kadmon, Nirit. 2001. *Formal Pragmatics*. Blackwell Publishers Inc.
- Kaiser, Elsi. 2011. Focusing on pronouns: Consequences of subjecthood, pronominalisation, and contrastive focus. *Language and Cognitive Processes* 26:1625–1666.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In *Formal methods in the study of language*, ed. Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 277–322. Amsterdam: Mathematisch Centrum.
- Kamp, Hans, and Uwe Reyle. 1993. *From discourse to logic*. Dordrecht: Kluwer.
- Karttunen, Lauri. 1971. Implicative verbs. *Language* 47:340–358.
- Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4:169–193.
- Karttunen, Lauri. 1974a. Presuppositions and linguistic context. *Theoretical Linguistics* 1:181–194.
- Karttunen, Lauri. 1974b. *Until*. volume 10, 284–297. Chicago Linguistics Society.
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy* 1:3–44.
- Karttunen, Lauri, and Stanley Peters. 1979. Conventional implicature. In *Syntax and Semantics 11: Presupposition*, ed. C.K. Oh and D.A. Dinneen. Academic Press: New York.
- Kaufmann, Magdalena. 2012. *Interpreting Imperatives* (Studies in Linguistics and Philosophy). Dordrecht, London: Springer.
- Kazanina, Nina, Ellen F. Lau, Moti Lieberman, Masaya Yoshida, and Colin Phillips. 2007. The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language* 56:384–409.
- Kehler, Andrew. 2002. *Coherence, reference, and the theory of grammar*. CSLI Publications.
- Kehler, Andrew, Laura Kertz, Hannah Rohde, and Jeffrey L. Elman. 2008. Coherence and coreference revisited. *Journal of Semantics* 25:1–44.
- Kehler, Andrew, and Hannah Rohde. 2013. A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics* 39:1–37.

- Keppel, G., and B.J. Underwood. 1962. Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior* 1:153–161.
- Kimball, John. 1975. Predictive analysis and over-the-top parsing. *Syntax and Semantics* 4:155–179.
- Kintsch, Walter. 1988. The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review* 95:163–182.
- Kizu, Mika. 1997. A note on sluicing in wh-in situ languages. In *9th Student Conference in Linguistics*, ed. Ljuba Nikolova Veselinova, Susan Robinson, and Lamont Antieau, volume 36, 143–159. Cambridge, MA: MIT, MIT Working Papers in Linguistics.
- Konieczny, L. 2000. Locality and parsing complexity. *Journal of Psycholinguistic Research* 29:627–645.
- Konieczny, L., and P. Doring. 2003. Anticipation of clause-final heads: Evidence from eyetracking and SRNs. In *4th International Conference on Cognitive Science and the 7th Conference of the Australasian Society for Cognitive Science*, 330–335.
- Koornneef, A., and J.J.A. van Berkum. 2006. On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language* 54:445–465.
- Kratzer, Angelika. 1991. Modality. In *Semantik: Ein Internationales Handbuch Zeitgenössischer Forschung (Handbücher zur Sprach- und Kommunikationswissenschaft)*, ed. Arnim von Stechow and Dieter Wunderlich, 639–650. Berlin: de Gruyter.
- Krifka, Manfred. 1989. Nominal reference, temporal constitution, and quantification in event semantics. In *Semantics and contextual expressions*, ed. Renate Bartsch, Johan van Benthem, and Peter van Emde Boas, 75–115. Foris: Dordrecht.
- Kroll, Margaret, and Deniz Rudin. 2017. Syntactic and pragmatic constraints on the acceptability of sluicing. In *47th Annual Meeting of the North East Linguistic Society (NELS)*, ed. Andrew Lamont and Katerina Tetzloff, 177–190.
- Kroll, Margaret, and Amanda Rysling. 2019. The search for truth: Appositives weigh in. In *Semantics and Linguistic Theory (SALT) 29*, 180–200.
- Kroll, Margaret, and Matt Wagers. 2017. Is working memory sensitive to discourse status? Experimental evidence from responsive appositives. Poster presented at XPrag.
- Kroll, Margaret, and Matt Wagers. 2018. Working memory resource allocation is not modulated by clausal discourse status. Ms., University of California, Santa Cruz.
- Kuperberg, Gina R., and T. Florian Jaeger. 2016. What do we mean by prediction in language comprehension? *Language, Cognition, and Neuroscience* 31:32–59.

- Kuperberg, Gina R., Martin Paczynski, and Tali Ditman. 2011. Establishing causal coherence across sentences: An ERP study. *Journal of Cognitive Neuroscience* 23:1230–1246.
- Ladusaw, William A. 1979. Polarity sensitivity as inherent scope relations. Doctoral Dissertation, University of Texas, Austin.
- Ladusaw, William A. 1988. Adverbs, negation, and QR. 481–488. The Linguistics Society of Kansas.
- Ladusaw, William A. 1992. Expressing Negation. In *The Proceedings of Semantics and Linguistic Theory (SALT) 2*, ed. Chris Barker and David Dowty.
- Laka, Itziar. 1990. Negation in Syntax: On the Nature of Functional Categories and Projections. Doctoral Dissertation, Massachusetts Institute of Technology.
- Lappin, Shalom, and Herbert J. Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational Linguistics* 20:535–561.
- Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106:1126–1177.
- Levy, Roger, and Frank Keller. 2013. Expectation and locality effects in German verb-final structures. *Journal of Memory and Language* 168:199–222.
- Lewis, Richard L., and Shravan Vasishth. 2005. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science* 29:375–419.
- Liversedge, Simon P., and Roger P.G. van Gompel. 2003. The influence of morphological information on cataphoric pronoun assignment. *Journal of Experimental Psychology: Learning, Memory, Cognition* 29:128–129.
- Lobeck, Anne. 1995. *Ellipsis: Functional Heads, Licensing, and Identification*. Oxford University Press.
- Lobeck, Anne. 2006. Ellipsis in dp. In *The blackwell companion to syntax*, ed. Martin Everaert and Henk van Riemsdijk, 145–173. Wiley.
- Love, Jessica, and Gail McKoon. 2011. Rule of engagement: Incomplete and complete pronoun resolution. *Journal of Experimental Psychology* 37:874–887.
- MacWhinney, B., and M.C. MacDonald. 1990. Measuring inhibition and facilitation for pronouns. *Journal of Memory and Language* 29:469–492.
- Martin, Andrea, and Brian McElree. 2008. A content-addressable pointer mechanism underlies comprehension of verb-phrase ellipsis. *Journal of Memory and Language* 58:879–906.
- Martin, Andrea, and Brian McElree. 2009. Memory operations that support language comprehension: Evidence from verb-phrase ellipsis. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35:1231–1239.

- Martin, Andrea E. 2016. Language processing as cue integration: Grounding the psychology of language in perception and neurophysiology. *Frontiers in Psychology* 7.
- Martin, Andrea E. 2018. Cue integration during sentence comprehension: Electrophysiological evidence from ellipsis. *PLoS One* 13.
- Martin, Andrea E., and Brian McElree. 2011. Direct-access retrieval during sentence comprehension: Evidence from sluicing. *Journal of Memory and Language* 64:327–343.
- McElree, Brian. 2001. Working memory and focal attention. *Journal of Experimental Psychology* 27:817–835.
- McElree, Brian, and Thomas Bever. 1989. The psychological reality of linguistically defined gaps. *Journal of Psycholinguistic Research* 18:21–35.
- McElree, Brian, and B.A. Doshier. 1993. Serial position and set size in short-term memory: Time course of recognition. *Journal of Experimental Psychology: General* 122:291–315.
- McElree, Brian, Stephani Foraker, and Lisbeth Dyer. 2003. Memory and language memory structures that subserve sentence comprehension. *Journal of Memory and Language* 48:67–91.
- Merchant, Jason. 1998. Pseudosluicing: Elliptical clefts in Japanese and English. In *Zas working papers in linguistics*, ed. A. Alexiadou, N. Fuhrhop, P. Law, and U. Kleinhenz, volume 10, 88–112. Zentrum für Allgemeine Sprachwissenschaft: Berlin.
- Merchant, Jason. 2001. *The syntax of silence: Sluicing, islands, and the theory of ellipsis*. Oxford University Press.
- Merchant, Jason. 2007. The syntactic representation of implicit arguments. Handout from ‘Funny indefinites’ workshop, Zentrum für Allgemeine Sprachwissenschaft, July.
- Merchant, Jason. 2013a. Polarity items under ellipsis. In *Diagnosing syntax*, ed. Lisa L.S. Cheng and Norbert Corver. Oxford University Press.
- Merchant, Jason. 2013b. Voice and ellipsis. *Linguistic Inquiry* 44:77–108.
- Messick, Troy, and Gary Thoms. 2016. Ellipsis, economy, and the (non)uniformity of traces. *Linguistic Inquiry* 27:306–332.
- Miller, Chelsea. 2016. Limited, syntactic reactivation in noun phrase ellipsis. MA Thesis, University of California, Santa Cruz.
- Miller, Philip, and Barbara Hemforth. 2014. Verb phrase ellipsis with nominal antecedents. Ms., Université Paris Diderot.
- Mitchell, D.C. 2004. On-line methods in language processing: Introduction and historical review. In *The on-line study of sentence comprehension: Eye-tracking, ERP, and beyond*, ed. Manuel Carreiras and Chuck Clifton, 15–32. Hillsdale, NJ: Erlbaum.

- Mittwoch, Anita. 1977. Negative sentences with *Until*. volume 13, 410–417. Chicago Linguistics Society.
- Murray, Sarah E. 2014. Varieties of update. *Semantics and Pragmatics* 7:1–53.
- Nairne, James S. 2002. The myth of encoding-retrieval match. *Memory* 10:389–395.
- Neath, I. 1993. Contextual and distinctive processes and the serial position function. *Journal of Memory and Language* 32:820–840.
- Neath, I., and A. Knoedler. 1994. Distinctiveness and serial position effects in recognition and sentence processing. *Journal of Memory and Language* 33:776–795.
- Neely, J.H. 1977. Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity intention. *Journal of Experimental Psychology: General* 106:226–254.
- Ng, Shukhan. 2008. An active gap strategy in the processing of filler-gap dependencies in Chinese. In *Proceedings of the 20th north american conference on chinese linguistics (naccl-20)*, ed. Marjorie K.M. Chan and Hana Kang, 943–957.
- Nicol, Janet L., Janet Dean Fodor, and David Swinney. 1994. Using cross-modal lexical decision tasks to investigate sentence processing. *Journal of Experimental Psychology* 20:1–10.
- Nouwen, Rick. 2007. On dependent pronouns and dynamic semantics. *Journal of Philosophical Logic* 36:123–154.
- Nouwen, Rick. 2014. A note on the projection of appositives. In *Formal approaches to semantics and pragmatics: Japanese and beyond*, ed. Eric McCready, Katsuhiko Yabushita, and Kei Yoshimoto, 205–222. Springer.
- O’Brien, Edward J., Gary E. Raney, Jason E. Albrecht, and Keith Rayner. 1997. Processes involved in the resolution of explicit anaphors. *Discourse Processes* 23:1–24.
- Omaki, Akira, Ellen F. Lau, Imogen Davidson White, Myles L. Dakan, Aaron Apple, and Colin Phillips. 2015. Hyper-active gap filling. *Frontiers in Psychology* 6:384.
- Ong, Matthew. 2013. Strict and sloppy reflexives in VP-Ellipsis. University of California, Santa Cruz. MA Thesis.
- Patterson, Clare, and Claudia Felser. 2019. Delayed application of binding condition c during cataphoric pronoun resolution. *Journal of Psycholinguistic Research* 48:453–475.
- Payne, John, Geoffrey K. Pullum, Barbara C. Scholz, and Eva Berlage. 2013. Anaphoric *one* and its implications. *Language* 89:794–829.

- Pearlmutter, N.J., and E. Gibson. 2001. Recency in verb phrase attachment. *Journal of Experimental Psychology: Learning, memory and cognition* 27:574–590.
- Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. URL <http://www.aclweb.org/anthology/D14-1162>.
- Perlmutter, David. 1970. On the article in English. In *Progress in linguistics*, 233–248. The Hague: Mouton.
- Pesetsky, David. 1987. Wh-in-situ: Movement and unselective binding. In *The representation of (in)definiteness*, ed. Eric Reuland and Alice G.B. ter Meulen, 98–129. Cambridge: MIT Press.
- Phillips, Colin. 1996. Order and structure. Doctoral Dissertation, MIT.
- Phillips, Colin, and Matt Wagers. 2007. Relating structure and time in linguistics and psycholinguistics. In *The oxford handbook of psycholinguistics*, ed. M. Gareth Gaskell, 739–756. Oxford University Press.
- Pickering, Martin J., and Asifa Majid. 2007. What are implicit causality and consequentiality? *Language and Cognitive Processes* 22:780–788.
- Poirier, Josee, Katie Wolfinger, Lisa Spellman, and Lewis P. Shapiro. 2010. The real-time processing of sluiced sentences. *Journal of Psycholinguistic Research* 39:411–427.
- Polanyi, Livia. 1986. The linguistic discourse model: Towards a formal theory of discourse structure. TR-6409. Cambridge: BBN Laboratories.
- Pollmann, Thijs. 1975. Een regel die subject en copula deletert? *Spektator* 5:282–292.
- Poppels, Till, and Andrew Kehler. 2019. Reconsidering asymmetries in voice-mismatched VP-ellipsis. *Glossa* 4:1–22.
- Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In *Proceedings of SALT 14*, ed. Kazuha Watanabe and Robert B. Young.
- Posner, M.I., and C.R. Snyder. 1975. Attention and cognitive control. In *Information processing and cognition*, ed. R.K. Solso, 55–85. Hillsdale, NJ: Erlbaum.
- Potsdam, Eric. 2013. CP-negation and the domain of quantifier raising. *Linguistic Inquiry* 44:674–684.
- Potter, Mary C., and Linda Lombardi. 1990. Regeneration in the short-term recall of sentences. *Journal of Memory and Language* 29:633–654.

- Potter, Mary C., and Linda Lombardi. 1992. The regeneration of syntax in short-term memory. *Journal of Memory and Language* 31:713–733.
- Potter, Mary C., and Linda Lombardi. 1998. Syntactic priming in immediate recall of sentences. *Journal of Memory and Language* 38:265–282.
- Potts, Christopher. 2005. *The Logic of Conventional Implicature*. Oxford University Press.
- Prince, Ellen. 1971. Futurate being, or why yesterday morning, I was leaving tomorrow on the midnight special is OK. Presentation at Summer Meeting of the Linguistic Society of America.
- Ratcliff, Roger, and Gail McKoon. 1981. Automatic and strategic priming in recognition. *Journal of Verbal Learning and Verbal Behavior* 20:204–215.
- Ratcliff, Roger, and Gail McKoon. 1988. A retrieval theory of priming in memory. *Psychological Review* 95:385–408.
- Rayner, Keith, Alexander Pollatsek, Jane Ashby, and Charles Clifton. 2012. *Psychology of reading*. Psychology Press.
- Rizzi, Luigi. 1990. *Relativized minimality*. Cambridge: MIT Press.
- Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12:683–721.
- Roberts, Craige. 1996. Information structure in discourse. *OSU Working Papers in Linguistics* 49:91–136.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5:1–69.
- Roberts, Craige. 2015. Modal subordination: *It would eat you first!* Ms.
- Rodrigues, Cilene, Andrew Nevins, and Luis Vicente. 2009. Cleaving the interactions between sluicing and p-stranding. In *Romance languages and linguistic theory 2006*, ed. Danièle Torck and W. Leo Wetzels, 245–270. Amsterdam: John Benjamins.
- Rohde, Hannah, and William S. Horton. 2014. Anticipatory looks reveal expectations about discourse relations. *Cognition* 133:667–691.
- Romero, Maribel. 1997. Recoverability conditions for sluicing. In *Empirical issues in formal syntax and semantics: Selected papers from the Colloque de Syntaxe et de Semantique de Paris*, ed. Francis Corblin, Danièle Godard, and Jean-Marie Marandin, 193–216. New York: Peter Lang.
- Romero, Maribel. 1998. Focus and reconstruction effect in wh-phrases. Doctoral Dissertation, University of Massachusetts, Amherst.

- van Rooij, Robert. 1997/2006. *Attitudes and changing contexts*. New York: Academic Press.
- van Rooij, Robert, and Katrin Schulz. 2004. Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information* 13:491–519.
- Rooth, Mats. 1985. Association with Focus. Doctoral Dissertation, UMass Amherst.
- Rooth, Mats. 1992. A Theory of Focus Interpretation. *Natural Language Semantics* 1:75–116.
- Rosen, Carol. 1976. Guess what about? In *North East Linguistic Society (NELS) 6*, ed. Montreal Working Papers in Linguistics, 205–211.
- Ross, John. 1973. Slifting. In *The Formal Analysis of Natural Languages: The First International Conference*, ed. Mouton: The Hague, 133–169.
- Ross, John R. 1967. Constraints on Variables in Syntax. Doctoral Dissertation, MIT.
- Ross, John R. 1969. Guess who? In *Proceedings from the 5th Meeting of the Chicago Linguistics Society (CLS)*, 252–286. University of Chicago, Illinois: Chicago Linguistic Society.
- Rudin, Deniz. 2019. Head-based syntactic identity in sluicing. *Linguistic Inquiry* 50:253–283.
- Saab, Andres. 2018. Nominal ellipsis. In *The oxford handbook of ellipsis*, ed. Jeroen van Craenenbroeck and Tanja Temmerman. Oxford University Press.
- Sachs, Jacqueline S. 1967. Recognition memory for syntactic and semantic aspects of connected discourse. *Perception Psychophysics* 2:437–442.
- Schlenker, Philippe. 2009. Local contexts. *Semantics and Pragmatics* 2:1–78.
- Schlenker, Philippe. 2010. Local contexts and local meanings. *Philosophical Studies* 151:115–142.
- Schlenker, Philippe. 2011a. DRT with local contexts. *Natural Language Semantics* 19:373–392.
- Schlenker, Philippe. 2011b. Presupposition projection: Two theories of local contexts Part I. *Language and Linguistics Compass* 5:848–857.
- Schlenker, Philippe. 2013. Supplements within a unidimensional semantics II: Epistemic status and projection. 167–182.
- Schwanenflugel, Paula J., and K.L. LaCount. 1988. Semantic relatedness and the scope of facilitation for upcoming words in sentences. *Journal of Experimental Psychology: Learning Memory and Cognition* 14:344–354.

- Schwanenflugel, Paula J., and Edward J. Schoben. 1985. The influence of sentence constraint on the scope of facilitation for upcoming words. *Journal of Memory and Language* 24:232–252.
- Schwarzschild, Roger. 1999. Givenness, AvoidF and Other Constraints on the Placement of Accent. *Natural Language Semantics* 7:141–177.
- Semin, Gun R., and Klaus Fiedler. 1991. The linguistic category model, its bases, application, and range. *European Review of Social Psychology* 2:1–30.
- Shapiro, Lewis P., Arild Hestvik, Lesli Lesan, and A. Rachel Garcia. 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.
- Simons, Mandy. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117:1034–1056.
- Simons, Mandy. 2013. Local pragmatics and structured contents. *Philosophical Studies* 168:21–33.
- Sloggett, Shayne. 2017. When errors aren't: How comprehenders selectively violate Binding Theory. Doctoral Dissertation, University of Massachusetts, Amherst.
- Sloggett, Shayne, Nicholas Van Handel, and Amanda Rysling. 2020. A-maze by any other name. Poster presented at CUNY 2020.
- Smith, Steven Bradley. 1974. *Meaning and negation*. Mouton: The Hague.
- Smyth, Ron. 1994. Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research* 23:197–229.
- Soames, Scott. 1989. Presuppositions. In *Handbook of philosophical logic*, ed. Dordrecht, volume 3, 553–616.
- Sprouse, Jon, Matt Wagers, and Colin Phillips. 2012. A test of the relation between working memory capacity and syntactic island effects. *Language* 88:83–123.
- Stalnaker, Robert. 1973. Presuppositions. *Journal of Philosophical Logic* 2:447–457.
- Stalnaker, Robert. 1974. Pragmatic presuppositions. In *Semantics and Philosophy*, ed. Milton K. Munitz and Peter K. Unger, 197–213. New York University Press.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25:701–721.
- Staub, Adrian, and Charles Clifton. 2006. Syntactic prediction in language comprehension: Evidence from either...or. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32:425–436.

- Sternberg, Saul. 1966. High-speed scanning in human memory. *Science* 153:652–654.
- Stewart, Andrew J., J. Holler, and E. Kidd. 2007. Shallow processing of ambiguous pronouns: Evidence for delay. *The Quarterly Journal of Experimental Psychology* 60:1680–1696.
- Stewart, Andrew J., Martin J. Pickering, and Anthony J. Sanford. 2000. The time course of the influence of implicit causality information: Focusing versus integration accounts. *Journal of Memory and Language* 42:423–443.
- Stowe, Laurie A. 1986. Parsing WH- constructions: Evidence for on-line gap location. *Language and Cognitive Processes* 1:227–245.
- Stowe, Laurie A., Edith Kaan, Laura Sabourin, and Ryan C. Taylor. 2018. The sentence wrap-up dogma. *Cognition* 176:232–247.
- Sturt, P., C. Scheepers, and M. Pickering. 2002. Syntactic ambiguity resolution after initial misanalysis: The role of recency. *Journal of Memory and Language* 33:407–430.
- de Swart, Henriëtte. 1996. Meaning and use of not...until. *Journal of Semantics* 13:221–263.
- Syrett, Kristen, and Todor Koev. 2014. Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics* 0:1–53.
- Tancredi, Christopher. 1992. Deletion, deaccenting, and presupposition. Doctoral Dissertation, MIT.
- Thoms, Gary. 2019. Reassessing case-matching. Talk given at Sluicing@50, University of Chicago.
- Toosarvandani, Maziar. 2014. Contrast and the structure of discourse. *Semantics and Pragmatics* 7:1–57.
- Traxler, Matthew J., and Martin J. Pickering. 1996. Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language* 35:454–475.
- Uegaki, Wataru. 2015. Interpreting questions under attitudes. Doctoral Dissertation, MIT.
- Vetter, David C. 1973. Someone solves this problem tomorrow. *Linguistic Inquiry* 4:104–108.
- Vicente, Luis. 2014. Sluicing and its subtypes. Ms.
- Vicente, Luis. 2018. Sluicing and its subtypes. In *The oxford handbook of ellipsis*, ed. Jeroen van Craenenbroeck and Tanja Temmerman. Oxford University Press.
- Wagers, Matthew. 2014. Syntax in forward and reverse. In *The routledge handbook of syntax*, ed. Andrew Carnie, Yosuke Sato, and Daniel Siddiqi. Routledge.

- Wagers, Matthew, and Brian McElree. 2012. Working memory and language processing: Theory, data, and directions for future research. In *The cambridge handbook of biolinguistics*, ed. Cedric Boeckx and Kleantes K. Grohmann. Cambridge.
- Warren, Tessa, Sarah J. White, and Erik D. Reichle. 2009. Investigating the causes of wrap-up effects: Evidence from eye movements and E-Z reader. *Cognition* 111:132–137.
- Waugh, N.C., and D.A. Norman. 1965. Primary memory. *Psychological Review* 72:89–104.
- Webber, Bonnie Lynn. 1978. A formal approach to discourse anaphora. Doctoral Dissertation, Harvard University.
- Webber, Bonnie Lynn. 1988. Discourse deixis: Reference to discourse segments. In *The 26th Annual Meeting of the Association for Computational Linguistics*, ed. SUNY Buffalo, 113–122.
- Witzel, Naoko, Jeffrey Witzel, and Kenneth Forster. 2012. Comparisons of online reading paradigms: Eye tracking, moving-window, and maze. *Journal of Psycholinguistic Research* 41:105–128.
- Xiang, Ming, Julian Grove, and Jason Merchant. 2014. Ellipsis sites induce structural priming effects. Ms., University of Chicago.
- Yoshida, Masaya, Lauren Ackerman, Morgan Purrier, and Rebekah Ward. 2012. The processing of backward sluicing. Presentation at CUNY.