NORTHWESTERN UNIVERSITY

The Island (In)Sensitivity of Stripping

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Linguistics

By

David Potter

EVANSTON, ILLINOIS

September 2017

ABSTRACT

The Island (In)Sensitivity of Stripping

David Potter

Ellipsis plays a key role in the understanding of the relationship between what is uttered and what is understood. In elliptical utterances, less is said than is understood. How is it that a mere fragment of a sentence can be understood as something more?

This dissertation explores the syntax of one type of elliptical phenomena, Stripping. A speaker, Alexi, might utter (1), to which Tim might reply with a Stripping continuation, as in (1a). Tim's reply, apparently just a *no* along with a single word correction, is nonetheless understood as if it were a complete sentence, that in (1b). Does this interpretation arise as a result of some process which does not rely on syntactic structure within the ellipsis site, or is there hidden syntactic structure associated with the remnant *Ardbeg*?

(1) Alexi: James likes *Glenlivet*.

 a. Tim: No, *Ardbeg*.

 b. Tim: No, James likes *Ardbeg*.

I argue in this dissertation that Stripping remnants are indeed associated with silent syntactic structure, and that this structure is *isomorphic* to the structure associated with the antecedent. Evidence for this analysis comes from two empirical domains: Islands and Binding Condition C Effects. In Chapter 2, I present the results of four large scale acceptability judgment experiments which reveal that Stripping is insensitive to island constraints. That is, examples like (2b) are grammatical, contrary to various claims in the literature.

(2) a. Tim met [$_{island}$ the researcher who likes *Glenlivet*].

b. No, *Ardbeg*.

The acceptability of these examples leads us to examine just how much structure is contained within the ellipsis site, if any. If there is no structure within the ellipsis site (3a), or if the structure is non-isomorphic to the antecedent (3b), we would expect such examples to be acceptable, as they contain no island. If the ellipsis site contains structure isormorphic to the antecedent, such examples would only be expected to be acceptable if it were possible to repair island constraints through ellipsis.

(3) a. No, *Ardbeg*.

b. No, *Ardbeg* ~~it was e~~.

c. No, *Ardbeg* ~~Tim met [$_{island}$ the researcher who likes e]~~.

These three approaches are evaluated in Chapter 3 on the basis of two large scale plausibility judgment experiments. These experiments concern examples like (4), which were found to be as implausible under the given indexations as typical Binding Condition

C violations. This is to be expected only there is structure within the ellipsis that contains material isomorphic to the antecedent, including the island and material outside it. If the ellipsis site contains no hidden material (4a), or some non-isomorphic material (4b), the pronoun of the antecedent would not be recovered within the ellipsis site. If the full isomorphic structure were recovered, that pronoun would appear within the ellipsis site, yielding a Binding Condition C violation. The implausibility of (4) reflects this violation.

(4) a. He$_i$ met the researcher who likes Mary.

b. *No, James$_i$.

(5) a. No, *James*$_i$.

b. No, *James*$_i$ ~~it was e~~.

c. No, *James*$_i$ ~~he$_i$ met [$_{island}$ the researcher who likes e~~].

In Chapter 4, I argue that the sum of this evidence implicates that, even in instances of island-violating Stripping, the ellipsis site contains isomorphic structure. In turn, I argue that various alternative mechanisms by which the remnant might escape the ellipsis site, without also incurring an island violation, are inadequate, and instead argue that the remnant overtly moves out of the island to escape the ellipsis site. The ill-effects of this island-violating movement, I argue, take the form of an unlinearizable structure, a structure, and the violation associated with it, which is subsequently elided. Thus, the island insensitivity of Stripping is due to island repair.

Acknowledgements

I'd like to first share what I think is the most important thing about doing science that I've learned about during my Ph.D. This is that, to do science, and especially to research and write a dissertation, you are not required to know all the answers. In fact, it seems that you are required to not know all the answers. Completing this dissertation has required me to acknowledge that I don't, won't, and actually can't know all the answers, let alone write about them. My friends outside academia often don't understand this. Experts are fallible, but they are still experts. *I* didn't really understand this until far too long into my graduate studies. The trick, I guess, is to understand that progress is incremental, which implies imperfection at every step. But with every step, progress.

The second thing that I'd like to share is the old saying that you get what you pay for. My graduate studies have been immensely personally rewarding. Enriching, intellectually challenging, satisfying. I have bought a beautiful experience. But I have also paid dearly for it. I've spent years, at serious personal and opportunity cost, practicing science, to the exclusion of much else that life has to offer. However, that I was able to practice science in the way that I have, because I wanted to, that I could pay that price, speaks to the tremendous position of privilege I find myself in. I hope my contributions to academia go some way towards repaying that privilege.

First thanks must go to Masaya Yoshida. From my first visit to Northwestern, Masaya has been a tremendous mentor and friend. He is demanding and compassionate, and when you find these qualities together in a person, I think it means that the person respects themselves and that they respect you. In a supervisor, I'm not sure what higher praise there is. In a friend, too.

My committee members Alexis Wellwood and Kyle Johnson deserve special thanks as well. Alexis came only recently to Northwestern, but she's off soon to California. I'm incredibly lucky our time at NU overlapped, as Alexis's intellectual candor not only improved the dissertation, in absolutely critical respects, but my understanding of the field as well. I have nothing but the highest regards for Kyle Johnson, as he can combine a damn sharp critique with helpful suggestions and a bit of witty banter.

I was lucky to have found my way to Northwestern. Brady Clark served on one of my qualifying paper committees, and the published version wouldn't have happened without his comments and suggestions. He's also a fantastic teacher and teaching mentor. The rest of the faculty at NU, especially Matt Goldrick and Ann Bradlow, and Stefan Kaufmann while he was here, consistently provided the sort of engaged community crucial for the success of their students, me included. Irene Sakk is the foundation of the department, and I wouldn't be finishing without years of her help. The experiments reported in this dissertation could not have been pulled off without the technical wizardry of Chun Liang Chan. I also gratefully acknowledge NSF grant BCS-1323245 awarded to Masaya Yoshida, under which this work was partially supported.

It's hard to overstate the impact that my fellow grads have had on my research and my ability to complete this dissertation. My brilliant collaborator Michael Frazier is an outstanding human being. I'm proud of the work we do together and I'm proud to call him my friend. My lab mates, especially Kat Hall, Peter Baumann, and Lauren Ackerman, proved constantly willing to poke holes in my work, an undervalued quality. My cohort, Michael Blasingame, Angela Fink, Julie Matsubara, and Samantha (Gett) Heidenreich, gave me a ton of needed and appreciated support, especially in the early days. As did the rest of the grads, especially Jeremy Needle, Jordanna Heller, and Tommy Denby. To Angie Cooper, I owe a special debt of gratitude.

Looking beyond Northwestern Linguistics, I realize that my research community is full of supportive and sharp researchers, many of whom have directly contributed in one way or another to my successes. At Simon Frasier University, Chung-hye Han, who supervised my MA thesis, and Nancy Hedberg both helped broaden my understanding of the field, and were fun collaborators and co-presenters. Tim Hunter, a co-author on a paper encompassing part of Chapter 3, was the first person outside my university that I collaborated with. Let's do it again sometime. At conferences and lab meetings, Matt Barros, Brooke Larson, Klaus Abels, Bob Frank, Jason Merchant, and Ming Xiang have all provided insightful critiques and comments. I appreciate you all.

In my capacity as an English as a second language instructor, I am grateful for the opportunity to teach wtih NU's English Language Program, with its incredibly talented administrative and teaching staff. Julia Moore has been not only a tremendous educational

mentor, helping me become a better teacher, but also is probably the kindest hearted person I've ever worked with. Brian Hampson, Erin Leddon, Lisa Del Torto, and Kathleen Geraghty are all pretty incredible educators, and good people. I'm lucky to have worked with them.

I want to also mention my time as an undergrad at Portland State University. I had some great teachers there, including Kim Brown, Lynn Santelmann, Susan Conrad, Janet Cowal, and Thomas Dieterich. I got my first linguistics training, and my first teacher training, at PSU, and given how close both of these are to my heart now, I'd say y'all did well. I was also a McNair scholar at PSU, a program that I think not many people know about, unless you need it. The McNair Scholars program helps low-income first generation college students, and members of under-represented groups, gain the skills and experience necessary to succeed in graduate studies. There is no way that I would have gone to grad school, let alone finished, if it hadn't been for this program.

For my friends and family, there aren't enough words. Mom and Mike, Spencer, and Tina; you're my people. Grandma and Grandpa Blue, I know I told you when I was five that I'd be a paleontologist, but Linguistics is pretty cool too. Foroogh, hurry home! Tim and Dillon, the next one is on me. Grandma and Grandpa Clifton, Grandma and Grandpa Jennings, Tom, and Alexi, rest in peace.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Introduction

Elliptical phenomena, in their seeming simplicity of form, pose a difficult challenge to linguistic analysis. A sentence fragment, often as little as a noun phrase, is uttered as a correction, an elaboration, or as an answer to a question, and it is unremarkable. But reflection on the message conveyed by these fragments would reveal an interpretation approaching that of a complete sentence. How do listeners fill in what hasn't been said?

Consider an example of the type of ellipsis that will be the focus of this dissertation, Stripping. Imagine the good friends Alexi and Tim are discussing which type of whisky their friend James prefers. Alexi might say something like that in (6), to which Tim might reply as in (6a). Tim corrects Alexi with little more than the word *no*, and the correction. Nevertheless, we understand Tim to be saying something along the lines of the complete sentence in (6b). Why is it that we don't understand Tim's correction to mean something like that in (6c)?

(6) Alexi: James likes *Glenlivet*.

a. Tim: No, *Ardbeg*.

b. Tim: No, James likes *Ardbeg*.

c. Tim: No, James loves *Ardbeg*.

Researchers break this problem, a mismatch between what is said and what is understood, down into three parts, as lucidly discussed in Merchant (2016). One part of the puzzle is identifying the constraints on where ellipsis is found. Ellipsis is highly productive, but not just anything can go missing. The question here is why. This is the licensing question on ellipsis. A second part is establishing in what way does the interpretation of a fragment correspond to the preceding linguistic context. Why is (6a) understood as (6b) but never (6c)? This is the identity question on ellipsis. A third part of the problem of ellipsis, and the focus of this dissertation, is whether these fragments are structurally as simple as they seem. Is this structural simplicity linguistic reality, with some extra-linguistic reasoning processes at play to permit us to understand the intended, but linguistically absent, message? Or is this simplicity a mirage, with a hidden but otherwise ordinary linguistic structural present, just out of earshot? This is the structural question of ellipsis.

Approaches to the structural question of ellipsis fall into two broad categories. The non-structural approaches posit that the linguistic structure associated with the fragment, known as the *remnant* in the case of Stripping, is simply what we would expect of the remnant (Culicover and Jackendoff, 2005; Ginzburg and Sag, 2000). There is no hidden structure and so what you see is what you get. The others are structural approaches, which, as their name suggests, posit that there is some hidden structure associated with the remnant (Barros et al., 2013; Lasnik, 1972; Lobeck, 1995; Merchant, 2001; van Craenenbroeck, 2010). This hidden structure might be *isomorphic*, corresponding closely in form, to the antecedent, as in (7a). Here the hidden structure, what I will call the elided material, is indicated by the ~~strikethrough~~ diacritic. Or the elided material might be *non-isomorphic*, as in (7b), where the elided material diverges in form from that of the antecedent.

(7) a. No, *Ardbeg* ~~James loves~~

b. No, *Ardbeg* ~~it is~~

These approaches to the structural question allocate the complexity of ellipsis differently to the various aspects of the cognitive and linguistic systems. The non-structural approaches maintain a simpler syntax, without the need to posit hidden structure, but complicate the interpretive system, which must therefore generate the sorts of meanings typically associated with linguistic structure. The structural approaches complicate the syntax, with various mechanisms for ensuring the structure associated with the remnant is neither seen nor heard, mechanisms which seem not to form a natural class with the various other mechanisms which derive silent elements, such as the generation of traces or the deletion of copies.

The structural and non-structural approaches can also be distinguished on empirical grounds, as they make different predictions as to how the remnant should, or should not, be connected to the elided material. The non-structural approaches predict, broadly, that the remnant should not behave as if it were connected to any elided material, as there is no elided material with with it could be connected under these approaches. The structural approaches do predict some connectivity effects, although just which connectivity effects the remnant should show will depend on just what sort of material is proposed to be present within the ellipsis site. In this dissertation, I explore these approaches with respect to two empirical domains: the locality constraints known as *islands*, and the constraint on the distribution of referring and anaphoric expressions, Binding Condition C.

Islands are structural domains which block certain long distance dependencies. Example (8a) illustrates a typical licit long distance dependency between a *wh*-phrase, *which whisky*, and the structure associated with the verb *perfers*. I will assume that this long distance dependency is expressed as movement: the *wh*-phrase originated within the embedded clause, and moved to its final position at the edge of the matrix clause. Notice that this dependency spans the edge of the embedded clause without difficulty. Islands, in contrast, do block dependencies which cross their edges. Consider (8b), which includes a definite relative clause, the sort of island domain which will be the focus of this dissertation. Here, the long distance *wh*-dependency is blocked, even through it spans roughly the same distance as that in the preceding example. Island domains constrain such long distance dependencies.

(8) a. Which whisky did Tim claim that James prefers e?

b. *Which whiskey did Tim meet [$_{island}$ the researcher who prefers e]?

Now consider example (9), an instance of *island-violating Stripping*. In this example, the correlate, *Glenlivit*, is contained within a definite relative clause island. The various approaches to the structural question on ellipsis make different predictions about the acceptability of examples such as these. Non-structural and non-isomorphic structural approaches generally predict that examples like (9b) should be acceptable. A non-structural approach to ellipsis posits that there is no structure structure within the ellipsis site at all, and therefore there is no island from within which the remnant would have to move. As in (10a), there is no island, so there is no problem. A non-isomorphic structural approach posits that there is structure within the ellipsis site, but this structure need not contain the island. The non-isomorphic construal illustrated in (10b) merely contain a reduced cleft structure. If the remnant originates within the ellipsis site, to escape ellipsis it need not violate any island constraints, and so the examples are predicted to be acceptable. The isomorphic approaches posit that the structure within the ellipsis site would contain the island, as illustrated in (10c). For these approaches, the predictions made vary with the theory of islands. If island constraints are the result of constraints on derivations, the mechanisms by which syntactic structures are composed, then such examples should be unacceptable. Given a structure like that in (10c), the remnant must escape an island to escape the ellipsis site, which would be derivationally impossible. Alternatively, one could imagine that island-violating movement is derivationally possible, but typically yields a representation that is fatally flawed. For example, island-violating movement might ordinarily cause the words in a sentence to be unlinearizable. If this were the case, then we would expect that ellipsis might be able to repair such a flaw, rendering examples of island-violating Stripping acceptable.

(9) a. Tim met [$_{island}$ the researcher who likes *Glenlivet*].

b. No, *Ardbeg*.

(10) a. No, *Ardbeg*.

b. No, *Ardbeg* ~~it was e~~.

c. No, *Ardbeg* ~~Tim met [$_{island}$ the researcher who likes e~~].

Examining the literature reveals that the matter of evaluating the acceptability of examples like (9b) is a complex matter, involving a number of confounding factors. The result is a number of seemingly contradictory acceptability judgments, which are presented and discussed in Chapter 2.

In an attempt to establish a clear empirical basis for the evaluation of the predictions made by the various approaches to the structural question of ellipsis, I conducted a series of four large scale acceptability judgment experiments, presented in Chapter 2. The results of these experiments indicate that Stripping is indeed largely insensitive to definite relative clause islands, and that examples like (9b) are, from a grammatical perspective, as acceptable as their non-island-violating counterparts. This is true of not only the corrective type of Stripping used as illustrated so far, but also what I call elaborative and non-contrastive Stripping as well. Examples and discussion of these other types follow later in this introduction.

The island insensitivity of Stripping lets us set the isomorphic derivational-island approaches to ellipsis aside, as they predict that Stripping should be sensitive to islands. This leaves us with three approaches, each positing a distinct syntax, or no syntax at all, within the ellipsis site. In Chapter 3, I argue that these alternatives can be distinguished with respect to how Stripping interacts with Binding Condition C (BCC).

A typical contrast captured by BCC is given in (11). A pronoun may precede an name and refer to the same person, as in (11a), so long as the pronoun does not C-command the name, as in (11b).

(11) a. His$_i$ mother loves James$_i$.

b. *He$_i$ loves James$_i$.

Chapter 3 presents the results of two large scale plausibility judgment experiments which indicate that Stripping, in both island and non-island-violating guises, is sensitive to BCC. That is the co-indexation indicated of examples like (12), and similar non-island counterparts, was implausible. This implausibility only follows if the ellipsis site is populated with structure that is isomorphic to the antecedent. In the non-structural, e.g. (13a), and non-isomorphic, e.g. (13b), analyses, the structure within the ellipsis site is impoverished such that the C-command relationship between the pronoun and name is never established. Consequently, no BCC effect would be predicted. Under an isomorphic analysis, however, the pronoun is construed within the ellipsis site, predicting a BCC violation.

(12) a. He$_i$ met the researcher who likes Mary.

b. *No, James$_i$.

(13) a. No, *James*$_i$.

b. No, *James*$_i$ ~~it was e~~.

c. No, *James*$_i$ ~~he$_i$ met [$_{island}$ the researcher who likes e~~].

In Chapter 4, I take the results of Chapters 2 and 3 together and argue in favor of an isomorphic structural approach to ellipsis. I then turn to the issue of how the remnant escapes the ellipsis site without resulting in an island violation. I argue that the remnant does indeed move, overtly, from within the island to escape ellipsis, thereby implicating a type of island repair. This I argue is best thought of in terms of a repair of an unlinerizable structure.

1.1. Background on Ellipsis

Stripping is just one of a large family of elliptical configurations. These include Sluicing (14a), Verb Phrase Ellipsis (VPE, 14b), Noun Phrase Ellipsis (NPE, 14c), Verbal Gapping (14d), Nominal Gapping (14e), Left Peripheral Ellipsis (LPE, 14f), PseudoGapping (PsG, 14g), and Fragment Answers (FAs, 14h), among others.

(14) a. James like a certain whisky, but I don't know which whisky ~~James likes~~.

b. James likes Laphroaig, and Mary does ~~like Laphroaig~~ too.

c. I prefer Scotland's style of whiskey to Canada's ~~style of whiskey~~

d. Alexi likes Jim Bean and Tim ~~likes~~ Jack Daniels.

e. James likes Scotland's most famous style of whiskey and Canada's least famous ~~style~~ of gin. .

f. James brought whiskey from Scotland and ~~brought~~ gin from Portland.

g. Alexi likes Jim Bean more than he does ~~like~~ Jack Daniels.

h. Which kind of whiskey does Mike prefer?

~~He prefers~~ Jameson.

These various types of ellipsis can be distinguished descriptively on the basis of what categories seem to be elided, what categories serve as remnants, and the information structural relationship between the remnants and their correlates. Certain configurations seem to elide an entire constituent, such as VPE and NPE. In the rest, it seems that a constituent has been elided, except for one or more elements, the remnants, which would occupy a position within the ellipsis site at some level of representation, yet survive. Stripping, for example, appears to elide a TP, yet the remnant, which would originate within that TP, survives ellipsis. Gapping, both nominal and verbal, PsG, LPE, PsG, and FAs all fall into this category or remnant leaving ellipsis. Sluicing can be considered a special case of remnant leaving ellipsis, if we assume that the *wh*-remnant in such cases originated within the elided IP.

Among the remnant leaving ellipses, some configurations are characterized by the survival of more than one remnant, among these the two types of Gapping, LPE, Multiple Sluicing, and, marginally, PsG (see Bowers (1998)). In Stripping, the remnant can co-occur with certain focus sensitive operators, such as *only* or *not*, as in (15) and (16), though I set these cases aside in this dissertation.

(15) a. James met someone yesterday.

b. Only Bill.

(16) a. James met Suzanne yesterday.

b. Yeah, but not Jane.

Thus Stripping appears to be most similar, among ellipses, to Sluicing and Fragment Answers. Both apparently target a TP for ellipsis, and both can leave just a single remnant. Of course, unlike Stripping and Fragment Answer remnants, the remnant in Sluicing configurations is a *wh*-phrase, or a constituent pied-pied by a *wh*-phrase. Nevertheless, the similarities between these these constructions, and the other types of remnant leaving ellipsis, have lead several researchers to posit mechanisms which derive these various ellipses in fundamentally similar ways. See, for example, in the generative tradition, Van Craenenbroeck and Lipták (2006, 2013); Frazier et al. (2012); Merchant (2001, 2004); Potter et al. (2017).

(17) a. James went to the brewery

b. No, to the distillery.

(18) a. James went quickly through the tour.

b. No, slowly.

(19) a. James enjoyed the partial tour.

b. No, slowly.

Despite these similarities to other elliptical phenomena, and the conceptual attractiveness in treating them with similar analyses, in this dissertation, I focus on Stripping, with occasional reference to Fragment Answers. This is primarily due to the constraints of conducting experimental research, for which it is difficult to cover as much empirical ground as one could do only reporting the results of small scale informal judgments. The analysis presented in Chapter 4 is very much in line with research suggesting a common derivational core to many of these types of ellipsis, but, as I discuss in Chapter 5, I leave exploring this avenue to further work.

Even within the core cases of Stripping, we find substantial variation. Chapter 2 of this dissertation is devoted in part to an exploration of the effects of one locus of variation between Stripping types: contrastivity. Prior research has distinguished two broad categories of ellipsis, contrastive ellipsis and non-contrastive ellipsis, and it has been claimed at various times that the derivations of contrastive and non-contrastive ellipses differ as a function of contrastivity (e.g. Romero (1998); Merchant (2001); Griffiths and Lipták (2014)). I investigate the empirical basis for some of these claims in Chapter 2.

I will understand non-contrastive ellipsis to be those cases in which the correlate introduces a set of relevant alternatives, of which the remnant is one relevant alternative. Non-contrastive correlates are typically an indefinite, a disjunction, a *wh*-phrase, or an implicit argument, as illustrated in (20)-(23). These non-contrastive ellipses have in common that the correlate introduces a set of alternatives, either as a result of the meaning of the correlate itself, being a *wh*-phrase, an indefinite, or an implicit argument, or as a result of the meaning of the structure of the correlate, being a disjunction. Thus, of the types of scotch, James likes Cao Ila. Between Coal Ila and Jameson, James likes Caol Ila.

(20) a. James likes a particular type of scotch.

b. Yeah, Caol Ila.

(21) a. James likes either Caol Ila or Jameson.

b. Caol Ila.

(22) a. Which scotch does James likes?

b. Caol Ila.

(23) a. Does James smoke?

b. Yeah, cigarettes.

I will understand the correlate and remnant in contrastive ellipsis to each be a member of some contextually salient set. For example Kilchoman and Caol Ila are both types of scotch whisky. In such cases, the correlate and remnant are often contrastive with each other. Caol Ila contrasts with Kilchoman, in that each is a different type of scotch. We can further distinguish types of contrastive ellipsis, with corrective cases, as in (24b) and Elaborative cases, as in (24c). However, the remnant and correlate apparently need not contrast with each other. Consider (24d). Here the Stripping continuation seems to simply offer confirmation that, indeed, James like Kilchoman. In light of such examples, *contrastive* seems less than appropriate as a category label, but, as I will be focused on the elaborative and corrective subtypes, I will nevertheless continue to use the term contrastive.

(24) a. James likes Kilchoman.

b. No, Caol Ila.

c. Yeah, and Caol Ila too.

d. Yeah, Kilchoman.

Thus, in how I will use these terms, the primary distinction between contrastive and non-contrastive ellipsis is how the set of alternatives is introduced. In non-contrastive ellipsis, the correlate introduces the alternatives, while in contrastive ellipsis the alternatives are simply contextually salient.

Notice that these definitions differ from how these terms are used by Griffiths and Lipták (2014). There, the term *contrastive ellipsis* is used only to designate corrective type contrasts, as in (24b), while elaborative contrasts, as in (24d), are a type of non-contrastive relation. It is orthogonal to my primary concerns, identifying the syntactic behavior of the various types of Stripping, whether there is any principled reason for categorizing elaborative types with corrective types, though I do believe there are. However, in the studies described in Chapter 2 I distinguish between, and directly compare, these three subtypes. And, as we will soon see, it turns out that elaborative and corrective Stripping pattern together in exhibiting a partial sensitivity to island constraints, in contrast with what I call non-contrastive Stripping, in which no island sensitive was found at all. This stands as justification for the categorizations described above.

Finally, I'd like to point out that I do not assume that every superficially non-sentential fragment should be treated as an instance of of a uniform class of a single linguistic phenomena, e.g. ellipsis. It seems to me that the possibility of fragments which are not elliptical in nature is not problematic to the notion that some fragments, and Stripping in particular, are elliptical. This is, however, a controversial issue, see in particular Stainton (2006); Merchant (2010) for an overview.

CHAPTER 2

The Island (In)Sensitivity of Stripping

2.1. Introduction

In this chapter I explore the question of Ellipsis Island Sensitivity. I first examine what is meant by the phrases *Island sensitivity* and *Island amelioration*. Then, I review the variety of prior claims about the insensitivity of ellipsis to Island effects, showing there is no clear consensus on the acceptability of these constructions. I then present a series of four experiments designed to determine whether any of a range of Stripping configurations exhibit sensitivity to one particular type of island domain: definite relative clauses.

These experiments largely substantiate those claims that Stripping, of both contrastive and non-contrastive varieties, is insensitive to definite relative clause islands (DRCIs). In doing so, several complications arise. For one, an apparent partial sensitivity of Contrastive Stripping to DCRIs is argued to not result from a grammatical factor, but rather from the resolution and subsequent reanalysis of an ambiguous antecedent. Additionally, an unplanned confound in the experimental design, the presence of factive verbs, a type of weak island, alongside non-factive verbs, yielded unexpected fruit. First, the data suggests that overt movement out of the clausal complements of factive verbs yields a reductions in acceptability, for movement of both DP and PP type elements. This reduced acceptability is compatible with the complements of factive verbs being islands for movement.

Second, the data shows that Stripping show no trace of this reduction in acceptability. If we understand these effects to be factive island effects, then this data shows that Stripping is insensitive to weak island effects.

2.2. Ellipsis Island Sensitivity

In this section, I argue that the acceptability of examples of island-violating Stripping, so called because the correlate is contained with an island, as found in (25), cannot be established against a single baseline. Neither a non-island violating Stripping baseline, nor a non-elliptical example involving overt, island-violating movement is sufficient, as each introduces confounding factors. Instead, multiple baselines must be used, in a factorial design.

(25) a. Speaker A: James met [$_{Island}$ the student who speaks *German*].

b. Speaker B: Yeah, and *French*, too.

It is widely accepted, following Ross (1967), that certain types of island-violating ellipsis, such as the Sluicing example in (26), are more acceptable than their overt movement counterparts. In contrast, the long distance dependencies in (27a) and (27b), are severely degraded when crossing a definite relative clauses island.

(26) James met the student who speaks a certain European language, but I don't know which.

(27) a. *Which language did James meet [$_{Island}$ the student who speaks t]?

b. *It was *French* that James met [$_{Island}$ the student who speaks t].

The contrast in acceptability between examples like (26) and (27a)-(27b) has been called, variously, amelioration (Diogo and Yoshida, 2007), repair (Fox and Lasnik, 2003; Merchant, 2004), and island insensitivity (Merchant, 2001; Fukaya, 2007; Griffiths and Lipták, 2014). In discussing this phenomenon, I will use the term *ellipsis island insensitivity*, or some transparent shortened version of this.

There are several reasons to avoid terms like amelioration or repair when discussing the phenomenon. These are theory specific terms, implying both that the island is present at some level in the representation or derivation, that there is a dependency between the remnant and the verb within the island, that such a dependency induces an island violation, and that this violation is ameliorated or repaired, at some point in the derivation, as a result of the violation being elided. The terms repair and amelioration serve our presently descriptive purposes poorly in that they conflate the empirical phenomenon with one of many possible explanation for that phenomenon, in what seems to be a longstanding vice for syntacticians (e.g. Across-The-Board movement phenomenon, *wh*-movement). An unwanted effect of such conflation is that it can be cumbersome and confusing to discuss alternative explanations using a term which implicates the validity of an alternative explanation[1].

In contrast, I intend the term ellipsis island insensitivity to be a relatively theory neutral term, and will used it to signify those configurations, such as (25) or (26), in which the correlate is contained within an island, yet where the example is more acceptable than some

[1]This point holds despite that the central claim of this dissertation is that a particular form of repair best explains the island insensitivity of Stripping.

relevant baseline. Such a conception of the phenomenon at issue avoids the conflation of the description of the phenomenon with the explanation for it.

An added benefit of this terminology is that it becomes more natural to discuss two aspects of the phenomenon that warrant further discussion. The first issue this formulation helps us to discuss is the possibility that ellipsis be *partially* sensitive to island violations. The term island repair makes it easy to imagine that island violations are either present or repaired, with no easy middle ground. A partial island sensitivity, on the other hand, would yield a somewhat reduced acceptability compared with no island violation, but a greater acceptability than for a *completely* island sensitivity. If ellipsis were partially sensitive to island violations, island-violating ellipsis would be neither as unacceptable as a comparable example of overt movement escaping an island, nor as as acceptable as comparable overt movement from a non-island domain. Rather, the acceptability of island-violating ellipsis be intermediate between these two baselines.

The second issue that this terminology, ellipsis island insensitivity, facilitates the discussion of is what the relevant baseline is against which island-violating ellipsis ought to be compared. If island-violating ellipsis is acceptable, it is acceptable compared to which other phenomena? The term *repair* could be taken to implicate overt island-violating movement as the relevant baseline. The violation of such illicit overt movements can be repaired. However, I will argue that overt movement is insufficient as the sole baseline against which we should judge these island-violating ellipsis cases.

Best practices in syntactic investigations involve the judgement of an example not in isolation, but in the context of a minimal pair, against a baseline example (Johnson, 2004;

Sportiche et al., 2013). In isolation, it is difficult to determine whether a judgment is the result of the phenomenon in question or some other confounding factor. However, even if we explicitly use a baseline for comparison, we need carefully consider just what confounding factors we can eliminate.

Take for example, a *wh*-dependency which crosses an island boundary, as in (28). For an example like this, a natural baseline for comparison would be (29), in which there is comparable long distance movement, but from within a non-island domain. Against a baseline like (29), examples such as (28) are typically judged be unacceptable.

(28) *Which language did James hear [$_{Island}$ the student who speaks t]?

(29) Which language did James hear [$_{Non-Island}$ that the student speaks t]?

But let us ask: what is the source of that unacceptability? In Sprouse et al. (2012), in conversation with a large literature on the nature of island constraints, including the influential Kluender and Kutas (1993), the authors ask whether the unacceptability of (28) could be attributed to additive effects of two independent factors inherent to examples like these. Island-violating movement implicates both the presence of an island and long distance movement from within the island. It could be that the unacceptability of examples like (28) is simply the result of the cumulative unacceptability of these two characteristics. It has been claimed that the mere presence of an island in an example, as in (30), will reduce the acceptability judgement of that example, even if there is no extraction from within that island, as compared to a similar example which does not include an island, e.g. (31) (see Sprouse et al. (2012) for discussion). Long-distance dependencies, as in (29), have also been claimed to be less acceptable than comparable short-distance dependencies, as in (31). Thus, the unacceptability of these island-violating examples could be reduced to the simple additive effect of these two factors.

(30) Who met [$_{Island}$ the student who speaks German]?

(31) Who heard [$_{Non-Island}$ that the student speaks German]?

Notice that (28) and (29) do not form a perfect minimal pair, which would be the case if they were only different in the dimension of interest: extraction from within an island. Both contain long distance movement, but only (28) contains an island. Thus, from this pair alone, it is not appropriate to conclude that it is extraction from within an island which is the source of the reduced acceptability of (28), as tempting as this may be; it could simply be the presence of an island which is to blame.

Now, Sprouse et al. (2012) found that the reduction in acceptability for such island-violating examples was greater than the added effects of the presence of an island and long distance dependencies. There was an interaction between these factors, which would be unexpected if the effect of islandhood were reducible to the independent effects of long-distance extraction and presence of an island. Thus, the authors reasoned that some independent factor must be at play; the concluded that this was a grammatical factor.

Returning to ellipsis, what baseline is most appropriate when judging island-violating ellipsis? I would like to suggest that each of the three natural alternatives are independently inadequate. In the above presentation of the example of Sluicing in (26), I used a non-elliptical overt island-violating movement as the baseline. For an example of Stripping like (32), we could take overt *wh*-movement (33a) or it-clefts (33b) as our non-elliptical baseline.

(32) a. Speaker A: James met [$_{Island}$ the student who speaks a European language].

b. Speaker B: Yeah, *French*.

(33) a. Speaker B: Which language did James meet [$_{Island}$ the student who speaks t]?

b. Speaker B: Yeah, it was *French* that James met [$_{Island}$ the student who speaks t].

A baseline in which an element has overtly moved from within an ellipsis site might seem particularly natural if we are considering a hypothesis about the derivation of such examples in which they derive from full syntactic structures, followed by the movement of the remnant from the ellipsis site. On the other hand, if we were interested in testing the hypotheses that fragment ellipsis involved the deletion of a non-constituent, in which the remnant remained in-situ, we might take a non-elliptical non-movement example like (34) as the baseline.

(34) Speaker B: Yeah, James met [$_{Island}$ the student who speaks *French*].

However, both of these non-elliptical baselines suffer from several important difference in comparison to their ellipsis counterparts; differences which could confound the intended comparison. First, these non-elliptical baselines are significantly longer than their elliptical counterparts. This is a problem because it has been shown that shorter examples are rated, in general, more acceptable than longer examples (Kluender and Kutas, 1993; Sprouse et al., 2012). Consequently, the short length of the ellipsis examples could serve to increase their acceptability relative to the baseline.

A second difference between these baselines and the elliptical items is that the non-elliptical baselines contain a substantial amount of repeated material. Assuming that a non-elliptical control were presented in the context of an antecedent phrase, just as the elliptical example must be, this repetition of material, where a reduction, e.g. ellipsis, were also possible, might serve to reduce the acceptability of the baseline, relative to the elliptical examples (see Wellwood et al. (2017) for discussion of this possibility).

A third issue with these baselines concerns whether elliptical configurations involve movement of the remnant in their derivation. In some analyses of ellipsis, the remnant overtly moves out of the ellipsis site, as is the case in the *move and elide* of Merchant (2004). In other analyses, the remnant doesn't move overly at all, as is the case in the analysis of e.g. Bruening (2015) and Weir (2014). Whether movement is involved in the derivation of fragment ellipsis matters because, as mentioned above in the context of *wh*-movement, long distance dependencies yield a reduction in acceptability, relative to short distance dependency baselines. With the non-elliptical controls it is clear whether a given example involves movement. With ellipsis this is not immediately obvious, but rather a still lively contested issue. But imagine that ellipsis did involve the movement of the remnant, this could then plausibly introduce a reduction in acceptability of ellipsis examples relative to in-situ non-elliptical controls. Or, if ellipsis did not involve movement, then the ellipsis examples might be expected to be improved relative to overt movement controls.

Let's then consider what we would be able to conclude on the basis of various hypothetical judgments. Suppose that the island-violating ellipsis in (32b) were judged to be more acceptable than the overt island-violating movement example in, for example, (33b). We might be tempted to conclude that island-violating ellipsis is island insensitive. Notice, however, that each of the three differences between our control and elliptical examples could in principle serve to improve the ellipsis example relative to the baseline. The ellipsis fragment is much shorter than the non-elliptical control, it does not include any repeated material, and it is possible that it does not involve any long distance dependencies. Thus, even if there were an effect of the island in the elliptical cases, such an effect could have been obviated by the shorter length of the ellipsis cases, or the reduced acceptability of the baseline due to repetition or long distance movement. So a conclusion that ellipsis is island insensitive in these cases would seem to be premature. On the other hand, if the island-violating ellipsis example were judged to be just as (un)acceptable as the non-elliptical island-violating controls, then it seems safe to conclude that examples of island-violating ellipsis are sensitive to islands, given that none of the differences between elliptical and baseline examples should *decrease* the acceptability of just the ellipsis example.

Now consider the state of affairs when using a non-elliptical in-situ baseline, like (34), for the judgement of the island-violating example of Stripping in (32b). If the Stripping example were judged to be worse than the baseline, we shouldn't immediately conclude that Stripping is island sensitive. After all, if the remnant participated in long-distance movement, this could reduce the acceptability of the Stripping example, relative to the baseline, in which no movement has taken place. Is is also not clear what we could conclude if the Stripping examples were judged to be just as acceptable as the baseline, given the differences in length and repetition between the Stripping and baseline examples. Thus both types of non-elliptical controls are, on their own, inadequate baselines for the evaluation of island-violating Stripping.

A third potential baseline would be a non-island *elliptical* control, directly comparing non-island and island-violating elliptical configurations. This would pit an island-violating ellipsis example, e.g. (35), against an example in which the correlate were contained within a domain from which extraction is possible, such as a complement clause, e.g. (36). In such a scenario, both the test item and the baseline would be of comparable lengths, thereby avoiding length as a possible confound. However, the use of non-island elliptical configurations as controls opens up the possibility of another confound: the presence of an island. As noted above, it has been shown that the mere presence of an island within an example, even in cases where there is no extraction from within the island, results in reduced acceptability.

(35) a. Speaker A: James heard [$_{Island}$ the student who speaks a European language].

b. Speaker B: Yeah, *French*.

(36) a. Speaker B: James heard that the student speaks a European language.

b. Speaker B: Yeah, *French*.

Consequently, we are in a similar situation with the non-island elliptical controls as we found ourselves in with the non-elliptical island-violating controls. Here, if we found that island-violating elliptical cases were judged to be as acceptable as non-island-violating cases, it seems reasonable to conclude that the relevant elliptical configurations were insensitive to island effects. On the other hand, if we were to find that cases of island-violating ellipsis were less acceptable than instances where there was no island in the antecedent at all, then we might be tempted to conclude that the ellipsis is sensitive to island constraints after all. However, such an effect could simply be due to the presence of an island in the antecedent, instead of a genuine sensitivity to island effect.

This discussion has lead us to the conclusion that each of these baselines has limitations when used to draw conclusions about the relative acceptability of island-violating ellipsis, and that which limitations are relevant depend on how the judgements turn out. Accordingly, we should, when assessing claims about the acceptability of island-violating ellipsis, take into consideration the type of baseline used when an author asserts such claims. In particular, we should be wary of conclusions drawn of data where an island-violating ellipsis example has been reported to be more acceptable than an overt island-violating movement non-elliptical control, conclusions drawn of data comparing island-violating ellipsis to non-elliptical in-situ baselines, and conclusions drawn of data in which island-violating ellipsis examples were reported to be less acceptable than than a non-island elliptical controls.

Notice that, even if we were to compare the island-violating ellipsis examples to each of these baselines in turn, we could still find ourselves with results that would be difficult to interpret. For example, data in which the ratings of island-violating elliptical cases are both better than a non-elliptical baseline as well as worse than a non-island elliptical baseline are compatible, as described above, with both a state of affairs in which ellipsis is sensitive to island constraints and a state of affairs in which it is insensitive to island constraints.

To move past this seeming impasse, we could use 2x2 factorial design in conjunction with numerical judgments. In such a design, we would compare non-elliptical configurations with extraction from within an island and from within a non-island, complement clause, against elliptical configurations with the correlate within an island and within a complement clause. These four conditions are described in table (2.1) and illustrated in (37)-(40).

Table 2.1. Ellipsis Island Insensitivity 2x2 Factorial Design

	Island	Non-Island
Ellipsis	ex. (37)	ex. (38)
Non-Ellipsis	ex. (39)	ex. (40)

(37) Stripping, Island

a. Speaker A: James heard [$_{Island}$ the student who speaks a European language].

b. Speaker B: Yeah, *French*.

(38) Stripping, Non-Island

a. Speaker B: James heard that the student speaks a European language.

b. Speaker B: Yeah, *French*.

(39) It-Cleft Island

a. Speaker A: James heard [$_{Island}$ the student who speaks a European language].

b. Speaker B: Yeah, it was *French* that James heard [$_{Island}$ the student who speaks].

(40) It-Cleft Non-Island

a. Speaker B: James heard that the student speaks a European language.

b. Speaker B: Yeah, it was *French* that James heard that the student speaks.

Such a factorial design can be used to identify whether the elliptical configurations are sensitive to islands, by comparing the magnitude of the difference between the island and non-island ellipsis conditions against the magnitude of the difference between the non-elliptical island and non-island conditions. If the magnitude of the difference between the island and non-island conditions is comparable in both elliptical and non-elliptical cases, then, regardless of the overall acceptability of these four conditions, we could conclude ellipsis to be island sensitive. Likewise, if there were no difference between island and non-island elliptical conditions, we could conclude ellipsis to be island sensitive.

Importantly, comparing the magnitude of the differences allows us to abstract away from any differences between elliptical and non-elliptical configurations that might arise due to length, as we are ultimately not interested in the differences in overall acceptability between ellipsis and non-ellipsis items, but only the magnitude of the differences between the island and non-island conditions. Likewise, it allows us to factor out any possible effect of presence of an island in the examples, in that both island elliptical and non-elliptical conditions contain an island, but true island effects yield reductions in unacceptability greater than found with the mere inclusion of an island.

Also, it should be noted that this sort of analysis would be impossible if we were to maintain the simple binary judgements typically reported in the syntax literature. Such a judgment system, by design, washes out any differences between items judged to be unacceptable, or between items judged to be acceptable, even if important differences between items within these categories could be captured with a more sophisticated judgment system.

Suppose that both Island Ellipsis and Island Non-Ellipsis conditions were judged to be unacceptable, and the Non-Island cases acceptable. Such an effect could be due to the island sensitivity of both configurations, or it could simply be due to the presence of an island in the Island conditions. Or imagine that just the Island Non-Ellipsis condition was judged to be unacceptable, and all others to be acceptable. This could be due to positive effect on acceptability of shorter example length, or to the island insensitivity of ellipsis. With simple binary judgements, it is impossible to distinguish between these alternative explanations.

Numerical judgments, in contrast, provide not only the level of granularity necessary to distinguish between arbitrary acceptability levels, but also the means to compare relative acceptability levels. This is the key to distinguishing between the various alternative explanations illustrated above.

In turn, the reporting of judgments along a numerical scale opens the possibility of refining the notion of island sensitivity. Implicit in the literature on ellipsis island sensitivity are two options, at the opposite ends of a spectrum. On the one hand, we would imagine that elliptical configurations are completely insensitive to islands. In such a scenario, an elliptical configuration in which the correlate is contained within an island would be rated no differently than a similar elliptical configuration in which the correlate is not contained within an island. Regardless of the overall acceptability differences between the ellipsis and non-ellipsis conditions, in such a scenario, whatever island effect might be found in non-elliptical controls is completely absent in the elliptical configurations. On the other hand, we could imagine that the island effect is present in full force. This is not to say that in such a scenario the island-violating elliptical configurations need be rated as unacceptably as island-violating non-elliptical controls. As discussed above, it is possible that elliptical configurations are rated, more generally, more acceptably than their non-elliptical controls. Rather, such a complete island sensitivity scenario simply entails that the magnitude of the difference between the island and non-island elliptical conditions be the same as the magnitude of the difference between the island and non-island non-elliptical controls. Thus, although the island-violating elliptical conditions may be rated better than the non-elliptical controls, in this scenario, the island effect would be the same across elliptical and non-elliptical conditions.

A third option is possible, however: that elliptical configurations are only partially sensitive to island violations. In such a partial sensitivity scenario, the island elliptical configurations would be rated worse than the non-island elliptical configurations, but the magnitude of that difference would be smaller than that between the island and non-island non-elliptical configurations. Thus, there would be an island effect in both elliptical and non-elliptical configurations, but this effect would be muted in the ellipsis cases. Such a scenario fits well with the judgments for island-violating Sluicing reported in Ross (1969), which were described as reflecting “violations of lesser severity” (pg. 39). I will pick this discussion up again in the experimental overview. Of course, such a partial island sensitivity would be impossible to describe using only binary acceptability judgments.

2.2.1. Prior Empirical Claims

In this section, I will review the literature on the island sensitivity of Stripping and Fragment Answers and conclude that there is no established consensus on the acceptability of island violating Stripping and Fragment Answers. The judgments in the literature vary along a number of factors, including the type of information structural relationship between the remnant and correlate, the grammatical function the remnant is to serve, the ambiguity of what element would serve as correlate, and the finality of the correlate in the antecedent.

Before we get started I wanted to take note of the distinction between *acceptability* and *grammaticality*. I will understand the former term to describe the judgments which constitute the core type of data employed in this dissertation, and by syntacticians more generally. It is intended as a pre-theoretical term, describing whether a given utterance is judged by a speaker to be acceptable, or natural. Presumably, many factors can influence acceptability judgments, such as plausibility, utterance length or complexity, familiarity or frequency of lexical items or grammatical structures. A mature theory of the language faculty and its interaction with the broader cognitive architecture would identify some of these factors as non-linguistic, and some as non-grammatical. Other factors would be identified as grammatical factors, and utterances which violate one of these factors could be said to be ungrammatical. Unacceptability is a broader term, as an unacceptable utterance could be so judged due to a grammatical factor, a non-grammatical factor, or even a non-linguistic factor.

However, which factors count as a grammatical factors is a theory internal issue and subject to lively debate. As an example, consider island constraints. Do such constraints on locality reflect the workings of the Grammar, long taken to be the working assumption by syntacticians? Or, as proposed in Kluender and Kutas (1993) and subsequent work, do such constraints reflect ordinary grammatical processes yoked to a limited pool of processing resources?

In the upcoming discussion on the sensitivity of ellipsis to island constraints, it is clear that, in nearly every case, the authors intend the given judgments to reflect grammaticality. This is the case when, for example, the authors attribute differences in judgments between minimal pairs of examples to differences in their grammatical properties. However, at least one author, Winkler (2013), has argued that the apparent island insensitivity of Sluicing is a product of extra-grammatical factors, a so-called *grammatical illusion*. At this point in the discussion, I will use the term acceptability in place of grammaticality when discussing the island insensitivity of ellipsis so that I can discuss the judgments of various types of examples, by various authors from an assortment of theoretical backgrounds, abstracting over the degree to which those judgments are the result of just grammatical factors, or an interaction of grammatical and extra-grammatical factors.

To begin the discussion, I would like to point out two well-known sets of data, that of Reinhart (1991) and Drubig (1994), which turn out to be orthogonal to the issue at hand.

Rienhart's influential paper present an analysis of a broad set of elliptical constructions, which she collectively calls Exceptive Constructions (ECs), including Stripping, also known as Bare-Argument Conjunctions, and Comparative Ellipsis. Reinhart proposes a non-structural approach to these ECs; there is no hidden material associated with the remnant. Rather, the correlate quantifier raises (QR), at LF, to form a constituent with the remnant. As the remnant and correlate are linked through a movement operation, Reinhart is interested in whether such an operation is subject to the sorts of locality constraints that QR is, in general, subject to. To this effect, Reinhart reports the judgments in (41), (43), and (45), concerning complex NP islands, *wh*-islands, and subject islands.

(41) Complex NP island

a. The fact that [$_{IP_1}$ some politician has resigned] got much publicity, but (I forgot) which one [$_{IP_2}$ e]. (ex. 36, pg. 372)

b. *The fact that [$_{IP_1}$ some politician has resigned] got much publicity but not the defense minister [$_{IP_2}$ e]. (ex. 37, pg. 372)

c. *The fact that [$_{IP_1}$ all politicians have resigned] got much publicity except the defense minister [$_{IP_2}$ e]. (ex. 38, pg. 372)

d. *More rumors that [$_{IP_1}$ the education minister resigned] were spread than the defense minister [$_{IP_2}$ e]. (ex. 39, pg. 372)

(42) a. The fact that some politician, but not the defense minister, has resigned got much publicity.

b. The fact that some politician has resigned got much publicity, but the fact that the defense minister has resigned has not got much publicity.

c. The fact that some politician has resigned got much publicity, but the fact that the defense minister not has resigned got much publicity (also).

Note however, that the interpretation Reinhart targets in these judgments is not actually an island-violating one; rather the judgments concern an interpretation where the correlate and remnant are arguments of the predicate introduced by the island itself. Consider the judgments given in (41). Reinhart annotates with indices the empty categories with which the remnants would be associated under the intended interpretation. The remnants in the Stripping example, (41b), the exceptive construction, (41c), and the comparative construction, (41d), are all co-indexed with the IP from within the island. That is, Reinhart intends the judgments concerning example (41b) to reflect the interpretation given in (42a), rather than that in (42b), or (42c). In (42a), only one fact is considered: that which concerns both a politician and the defense minister. In the island-violating interpretations in (42b) and (42c), which Reinhart does not consider, there are two facts at issue: that introduced by the antecedent phrase, concerning a politician, and that introduced in associate with the remnant, concerning the defense minister.

Similar facts hold of Reinhart's judgments for the *wh*-island and subject island examples. The intended interpretation for (43b) is given in (44a), and can be paraphrased, roughly, that the thing that Bush said, which Dukakis didn't say, people remembered it in the last poll. Here, there is one thing said, and Bush said it while Dukakis didn't. The island-violating interpretations in (44b) and (44c) both allow multiple things to have been said. Likewise for the subject island cases. Reinhart's target interpretation for her judgments on an example like (45c) entails a single relevant letter, addressed to Max and not to

Felix. An island-violating interpretation would implicate two letters; one to Max and one to Felix.

(43) *Wh*-island

a. [M]ore people remember what was said by Bush than by Dukakis. (ex. 49a, pg. 376)

b. [P]eople remembered what was said by Bush, in the last poll, but not by Dukakis. (ex. 49b, pg. 376)

c. I will tell you what I think about everyone, if you insist, except my boss. (ex. 49c, pg. 376)

(44) a. People remembered what was said by Bush but not by Dukakis, in the last poll.

b. People remembered what was said by Bush, in the last poll, but people (also) remember what was not said by Dukakis, in the last poll.

c. People remembered what was said by Bush, in the last poll, but people don't remember what was said by Dukakis, in the last poll.

(45) Subject Island

a. *[W]ho do [jokes about e] amuse you? (ex. 51a, pg. 376)

b. ?[[J]okes about everyone] amuse me except Felix. (ex. 51b, pg. 376)

c. ?[[T]he letter to Max] got lost, but not (to) Felix. (ex. 51c, pg. 376)

d. ?[M]ore jokes about Max were told in the party than (about) Felix. (ex. 51d, pg. 376)

e. *[Lucie's jokes about *every woman*] amuse me, except Lili. (ex. 52a, pg. 376)

f. *[T]he man [$_{PP}$ near *every woman*] amuses me, except Lili. (ex. 52b, pg. 376)

The question that Reinhart raises with these data is an interesting one. She observes that it would be odd, if these were uniformly elliptical configurations, for the ellipsis recovery mechanism to be constrained by Subjacency. "…it is well-known for all types of constituent ellipsis (Sag 1976; Williams 1977), that the procedure of copying a (well-formed) LF predicate into an empty projection is essentially unrestricted by the grammar - the two can lie as far apart as considerations of discourse coherence permit, and relating them is certainly not sensitive to subjacency or adjacency requirements" (Reinhart, 1991, 372). If a Sluicing antecedent can be contained within an island, why isn't this likewise possible for Stripping? However, this issue of where the antecedent may be found is orthogonal to the question of the sensitivity of ellipsis to islands, and so I set it aside for now.

The data presented by Drubig (1994) similarly initially seems relevant to the discussion at hand, yet, upon closer reflection is actually orthogonal. Drubig is interested in whether association with focus is sensitive to island boundaries intervening between a focus sensitive operator and a focus. The key data are as follows.

In (46) we have an antecedent phrase with the negation *not* acting as a focus sensitive operator. The phrase *black* within a complex NP is stressed. In (46a) through (46e), we have various fragment continuations, of diminishing sized constituents. Drubig reasoned that if focus were island insensitive, then the fragment continuation ought to be able to consist in merely the stressed word itself. The data Drubig reports suggest this is not the case; the smallest possible continuation is the island which contains the stressed word. If we understood the continuations in these examples as being derived from a sentential source, we could take them to be evidence that ellipsis is sensitive to, in this case, complex noun phrase islands.

(46) Mary didn't invite [the man in a black$_F$ suit]$_{FP}$ to the party,

a. but she invited the man in a purple$_F$ suit.

b. but the man in a purple$_F$ suit.

c. *but in a purple$_F$ suit.

d. *but a purple$_F$ suit.

e. *but purple$_F$.

However, there are reasons to suspect that cases like these, and other Exceptive Ellipsis cases, are fundamentally different from the canonical ellipsis cases. One reason is that they appear to conform to different conditions on identity than do canonical cases of ellipsis. The acceptable elliptical continuation of (46), (46b) is interpreted more or less as in (46a). Note that (46a) is of positive polarity, while the antecedent is negative.

In canonical cases of Stripping, as in (47a), the missing material is interpreted more or less identically to the antecedent. Thus the interpretation of (47a) is given in (47b), with missing material ~~struckout~~. Crucially, the polarity of the antecedent and missing material must match; the Stripping phrase in (47a) cannot be interpreted as in (47c). One might notice that in the examples in (46), the continuations are introduced by *but*, a conjunction which can implicate a polarity contrast between the conjuncts. One could then reason that this conjunction allows the listener to infer that the polarity of the antecedent and stripping phrases must differ. However, if this were responsible for the polarity mismatch between antecedent and stripping phrases, we would expect the same to be possible if the antecedent phrase were positive, as in (48a). That is, we would expect the stripping phrase of (48a) to have an interpretation like that in (48b), which it does not appear to have.

(47) a. James went to the store yesterday, and to the mall.

b. and ~~James went~~ to the mall.

c. and ~~James didn't go~~ to the mall.

(48) a. James went to the store yesterday, but to the mall.

b. but ~~James didn't go~~ to the mall.

A second reason to be suspicious that examples like these are instances of the same sort of phenomena as canonical instances of Stripping is that they are infelicitous across interlocutors, unlike other types of Stripping. Consider the contrast between (49) and (50), and set aside for the moment the contrast in how the missing material is interpreted. Assume that the two sentences were uttered by two speakers, one in reply to the other. (49), an instance of corrective Stripping, is perfectly acceptable across interlocutors, while (50) is plainly unacceptable.

(49) a. Mary didn't invite Max_F to the party.

b. No, Alexi_F.

(50) a. Mary didn't invite Max_F to the party.

b. *But Alexi_F.

A third reason is that the resolution of the missing material in such examples is clause bound in a way that Stripping is not. Consider (51), which can have an interpretation roughly like that in (51c). Notice that the matrix clause is recovered: Alexi is the person such that James denied that Mary invited him to the Party. Now compare with (52). The continuation in (52b) cannot have such a long distance interpretation; and though the example is already difficult to judge, much more so than (51b), it seemingly can only have the short interpretation illustrated in (52d).

(51) a. James denied that Mary invited Max$_F$ to the party.

b. No, Alexi$_F$.

c. No, James denied that Mary invited Alexi$_F$ to the party.

(52) a. James denied that Mary didn't invite Max to the party ...

b. ...but Alexi

c. $!=$...but James denied that Mary invited Alexi$_F$ to the party.

d. $=$...but Mary invited Alexi$_F$ to the party.

Given these differences between the sorts of examples Drubig presents and other Stripping configurations, I would like to set these examples aside for the time being[2].

[2]Krifka (2006) discusses these and similar examples which are subject to the same suspicions. See Von Fintel (1994) for an early discussion of the semantics of these constructions. I suspect that something like Reinhart (1991)'s analysis, in which overt movement, but no ellipsis occurs, holds of these constructions, but not of Stripping constructions more broadly.

2.2.2. Claims of Island Sensitivity

The earliest discussion I can find concerning the island sensitivity of an elliptical phenomena other that Sluicing is due to Morgan (1989). Morgan examined instances of corrective fragment answers, as illustrated in (53)-(55). Morgan reported that, if the response to such a polar question were to correct an element within an island, then merely responding with that element would be less acceptable as compared to a full, non-fragment answer. Thus, the question in (53a) can be responded to negatively, with a correction to an element within the coordinate structure, as in (53c). Yet, Morgan reports, if just the correction serves as an answer, as in (53b), the example is degraded. Morgan gives similar judgments for definite relative clause islands, (54), and examples with complex noun phrases and subject islands together, as in (55).

(53) a. Did John and Bill leave this morning?

b. *No, Harry.

c. No, John and Harry left this morning.

(54) a. Was the man who shot Lincoln a Marine?

b. *No, Kennedy.

c. No, the man who shot Kennedy was a Marine

(55) a. Did John's seeing Martha upset the President?

b. *No, Thelma.

c. No, John's seeing Thelma upset the President.

Merchant (2004) reports similar judgments for these types of corrective fragment answers. When the correlate to the fragment answer is contained within an island, Merchant reports that a fragment answer, in which the fragment corresponds just to the correlate, is unacceptable, as compared to a full, non-fragment, answer.

(56) a. Does Abby speak the same Balkan language that Ben$_F$ speaks? (ex. 87, pg. 688)

b. * No, Charlie$_F$.

c. No, she speaks the same Balkan language that Charlie$_F$ speaks.

(57) a. Did Ben leave the party because Abby wouldn't dance with him? (ex. 88, pg. 688)

b. * No, Beth$_F$.

c. No, he left the party because Beth$_F$ wouldn't dance with him.

(58) a. Did Abby vote for a Green Party candidate? (ex. 89, pg. 688)

b. * No, Reform Party$_F$.

c. No, she voted for a Reform Party$_F$ candidate.

(59) a. Did Abby get 'The Cat in the Hat'$_F$ and 'Goodnight Gorilla' for her nephew for his birthday? (ex. 90, pg. 689)

b. * No, 'The Lorax'$_F$.

c. No, she got 'The Lorax'$_F$ and 'Goodnight Gorilla' for her nephew for his birthday.

Merchant also discusses cases of pair-list responses to multiple *wh*-questions, in which an in-situ *wh*-phrase is located within an island. He finds multiple fragment answers to these cases to be less acceptable than full non-fragment answers, as in (60), while comparable non-island, long distance multiple fragment answer versions are fully acceptable, as in (61) and (62).

(60) Which committee member wants to hire someone who speaks which language? (ex. 92, pg. 689)

a. Abby wants to hire someone who speaks Greek and Ben wants to hire someone who speaks Albanian.

b. * Abby Greek, and Ben Albanian.

(61) a. Who's more likely to be influencing who?

b. The CIA John Foreman, or John Foreman the CIA?

(62) a. Which lawyer said he was representing which war criminal?

b. Cochran Milosevic, and Dershowitz Sharon.

Krifka (2006) reports judgments similar to Merchant (2004)'s about multiple fragments answers to multiple island-violating *wh*-questions, as in (63). If the second fragment answer corresponds to the entire island, here a complex noun phrase island, then the response is reported to be acceptable. However, if that second fragment is merely a portion of the island, the the example is reported to be unacceptable.

(63) Who introduced [the author of which novel] to Sue?

a. John, the author of Ulysses.

b. *John, Ulysses.

Krifka also expands the empirical domain by considering two additional type of question-fragment answer pairs. In one, Krifka employs a *wh*-in-situ constituent question, as in (64a). Here, a *wh*-phrase, under either an echo or quiz-question interpretation, is contained within a complex noun phrase island. Krifka reports that a non-contrastive fragment answer response cannot merely consist in an element corresponding to the *wh*-phrase. An answer such as (65b) is unacceptable. Instead, the data reported indicate that the entire complex noun phrase, which itself contains the element corresponding to the *wh*-phrase in the question, is the minimal acceptable fragment answer. Krifka also discusses another type of non-contrastive fragment answer, one in which the correlate is a disjunction, as in (64b). Here, if we take the book titles *Ulysses* and *Moby-Dick* to be focused, an island insensitive fragment answer could be expected to contain just an element corresponding to the relevant focus. Krifka reports that the island insensitive fragment answer such as (65b) isn't possible in this case either.

(64) a. John introduced [the author of which novel] to Sue?

b. Did John introduce [the author of Ulysses or Moby-Dick] to Sue?

(65) a. The author of Ulysses.

b. *Ulysses.

2.2.3. Claims of Island Insensitivity

The judgments reported in the literature are far from unanimous that ellipsis is island sensitive. Lappin notes: "In fact, it is not at all clear that bare argument ellipsis is sensitive to subjacency" (Lappin, 1996, pg. 161). Lappin reports several unacceptable examples of island-violating Stripping, also known as bare argument ellipsis, in (66), but also provides several examples of elaborative focus sensitive operator (FSO) stripping, in (67), which he reports to be acceptable. "In each of these cases it is possible to interpret the negated fragments as outside the scope of the syntactic island containing its corresponding NP in the antecedent. The most natural reading of [(67a)], for example, is that John enjoyed reading the articles which appeared in the *New York Times* last week, but he did not enjoy reading the articles which appeared in the Daily Telegraph last week." (Lappin, 1996, pg. 161).

(66) a. *We have interrogated the burglar who stole the car already, but not the diamonds.

b. *A musician who loved Bach arrived, and Mozart too.

(67) a. John enjoyed reading the articles which appeared in the *New York Times* last week, but not the *Daily Telegraph*.

b. Dancing with Mary in the garden is a pleasure, but not the park.

c. John agreed to the request that he submit articles to the journal, but not book reviews.

Culicover and Jackendoff (2005) provide a host of island-violating elliptical examples which they judge to be acceptable. In (68), they give examples of indefinite relative clause island (IRCI) insensitive non-contrastive Sluicing and Stripping. In (69), an IDRCI insensitive contrastive Stripping example. In (70), a complex noun phrase insensitive non-contrastive Stripping example. In (71), a coordinate structure constraint insensitive non-contrastive Stripping example. And, in (72), a sentential subject insensitive Sprout-Stripping example.

(68) a. John met a guy who speaks a very unusual language.

b. i. Which language? [*Which language did John meet a guy who speaks t?]

ii. Yes, Albanian. [*Albanian, John met a guy who speaks t.]

(69) a. John met a woman who speaks French.

b. And Bengali? [*And Bengali, did John meet a woman who speaks French t?]

(70) a. The Administration has issued a statement that it is willing to meet with one of the student groups.

b. Yeah, right — the Gay Rifle Club. [*The Gay Rifle Club, the administration has issued a statement that it is willing to meet with t.]

(71) a. They persuaded Kennedy and some other senator to jointly sponsor the legislation.

b. Yeah, Hatch. [*Hatch, they persuaded Kennedy and t to jointly sponsor the legislation.]

(72) a. For John to flirt at the party would be scandalous.

b. Even with his wife? [*Even with his wife, would for John to flirt t at the party be scandalous?]

Culicover and Jackendoff also note that they arrive at different judgments than those given by Merchant (Culicover and Jackendoff, 2005, 244, fn. 10). For all of Merchant's examples from (56) to (59) listed above and repeated here, Culicover and Jackendoff note that they find them acceptable, if the correlate in the antecedent is properly focused.

(73) a. Does Abby speak the same Balkan language that Ben_F speaks?

b. No, Charlie_F.

c. No, she speaks the same Balkan language that Charlie_F speaks.

(74) a. Did Ben leave the party because Abby wouldn't dance with him?

b. No, Beth_F.

c. No, he left the party because Beth_F wouldn't dance with him.

(75) a. Did Abby vote for a Green Party candidate?

b. No, Reform Party_F.

c. No, she voted for a Reform Party_F candidate.

(76) a. Did Abby get 'The Cat in the Hat'$_F$ and 'Goodnight Gorilla' for her nephew for his birthday?

b. No, 'The Lorax'$_F$.

c. No, she got 'The Lorax'$_F$ and 'Goodnight Gorilla' for her nephew for his birthday.

Culicover and Jackendoff also provide non-contrastive Fragment Answer and Sprout-Stripping examples which are insensitive to left branch islands.

(77) a. What kind of scotch does Harriet drink?

b. Expensive

(78) a. Let's get some pizza.

b. Pepperoni?

Casielles refers to the examples reported by Morgan (1989), instances of corrective fragment answers, and, contrary to Morgan, judges them to be acceptable: "[t]hese improve if the contrasted element is focused in the question..." (Casielles, 2006, pg. 121).

(79) a. Did John and Bill leave this morning?

b. No, Harry.

c. No, John and Harry left this morning.

(80) a. Was the man who shot Lincoln a Marine?

b. No, Kennedy.

c. No, the man who shot Kennedy was a Marine

(81) a. Did John's seeing Martha upset the President?

b. No, Thelma.

c. No, John's seeing Thelma upset the President.

Likewise Fukaya reports both (82), originally from Merchant (2004), who judged it to be unacceptable, and (83), to be acceptable. Fukaya on the other hand finds examples in which the fragment corresponds to a subject, e.g. (84), to be unacceptable.

(82) a. Did Abby like the candidate who referred to *Chomsky*?

b. No, *to Bresnan*.

(83) a. Did John meet a man who had been *to Paris*?

b. No, *to London*.

(84) a. Does Abby speak the same Balkan language that Ben$_F$ speaks?

b. * No, Charlie$_F$.

c. No, she speaks the same Balkan language that Charlie$_F$ speaks.

Fukaya also reports that the following examples of elaborative and corrective stripping are fully acceptable [3].

(85) a. John talked to the professor who had recommended Mary (at the faculty meeting).

b. Susan as well. / Susan, too.

(86) a. John met a man who had been to Paris.

b. To London as well.

c. No. To London.

(87) a. Microsoft hired a linguist who is on good terms with Chomsky.

b. With Bresnan, too.

Valmala (2007) likewise finds non-contrastive and corrective fragment answers to be acceptable. In the case of corrective fragments, the correlate must be contrastively focused. Valmala suggests that the degraded acceptability of corrective fragments is due to "parsing difficulties, not to island effects" (Valmala, 2007, pg. 12). As evidence Valmala notes that PP fragments, and repetition of the correlate, e.g. with *not*, are acceptable. That this

[3]Fukaya claims, however, that these examples only have with what he calls the local interpretation, a result of a non-isomorphic parse. I discuss such non-isomorphic approaches to island insensitivity in Chapters 3 and 4.

explanation, that a possible ambiguity yields unacceptability is in direct contrast with the explanation given by Winkler (2013), where it was claimed that a potential ambiguity actually gives rise to island insensitivity in English.

(88) a. A: I heard that Ben left the party because some girl wouldn't dance with him.

b. B: Yeah, Susan.

(89) a. A: I heard that Abby speaks the same Balkan language that some linguist speaks.

b. B: Yeah, Zeljko Boskovic.

(90) a. A: I heard they want to hire a linguist who speaks a/some Balkan language.

b. B: Yeah, Albanian.

(91) a. Q: Did Bill leave the party because he wanted to avoid having to talk to Susan?

b. A: No, to Mary.

c. A': No, not to Susan, to Mary. (ex. 58, pg. 12)

Barros et al. (2013) report corrective fragment answers to be acceptable just in case the fragment was an object, on the basis of a small scale acceptability judgment experiment. Examples in which the fragment was a subject were found to be unacceptable.

(92) a. Did they hire someone who works on French (last year)?

b. No, German.

(93) a. Did they leave because you offended Mary?

b. ?No, Sarah.

(94) They hired a multilingual person who speaks Greek, but I don't know which other languages.

(95) a. Did Ben leave the party because Abby wouldn't dance with him?

b. *No, Sally.

(96) a. Does Abby speak the same Balkan language that Ben speaks?

b. *No, Charlie.

Weir (2014) reports a few examples of island-violating corrective fragment answers to be acceptable, especially if the context, such as that in (98a), explicitly supports the relevant contrastive focus information structure.

(97) a. Q: Do they grant scholarships to students that study Spanish$_F$?

b. A: No, French. (relative clause) (ex. 403a, pg. 203)

(98) a. Q: Do you take milk and honey$_F$ in your tea?

b. A: No, sugar. (coordinate structure) (ex. 403b, pg. 203)

a. Context: We have before us lots of people. We know that these people are made up of lots of pairs of people who speak the same language as each other and who do not speak the same language as anyone else. (I.e. John and Mary both speak English and nothing else, Jan and Peter both speak Dutch and nothing else, etc.) A and B are playing a game where A is trying to guess which people belong to which pair. A's just trying to guess the right pairings, though; the actual languages they speak is irrelevant to him, all that's relevant is that the people in the pair speak the same language. B knows the pairings and will answer A's questions. A had already worked out that Abby and Charlie were a pair a while ago, but had forgotten this.

b. Q: Does Abby speak the same language that Ben$_F$ speaks?

c. A: No, Charlie. (You'd already worked that pairing out, remember?) (ex. 420, pg. 208)

Stainton (2006) reports acceptable non-contrative fragment answers to in-situ *wh*-questions, in which the *wh*-phrase is contained within a disjunction.

(99) *Tomato juice the Pope likes beer and.

(100) a. The Pope likes beer and what?

b. Tomato Juice.

Finally, Valmala (2007) and Barros (2012b) independently report that fragment answers in response to an island-violating matrix sluicing questions are acceptable, as in (101) or (102), due to Barros, or as in (103), due to Valmala[4]. Such examples constitute non-contrastive fragment answers, as the correlate is an indefinite, thereby introducing a set of alternatives, and not an element contrastive with the fragment.

(101) a. They hired someone who speaks a Balkan language.

b. Which one?

c. Albanian.

(102) a. Ben left the party early because someone wouldn't dance with him.

[4]Barros concurs with the judgements of Merchant (2004), however, that corrective fragment answers are island sensitive.

b. Who?

c. Christine.

(103) a. A: I bought a CD which contains several songs by Ottmar Liebert.

b. B: Who?

c. A: Ottmar Liebert.

2.2.4. Claims of Mixed Island Sensitivity

Griffiths and Lipták (2014) report a broad split across types of ellipsis. They claim that non-contrastive ellipsis is island insensitive, while contrastive ellipsis is island sensitive. The contrastive cases, which are limited to the corrective subtype, are given in (104)-(108), and the non-contrastive cases in 109-(114). Among the many instances of non-contrastive cases Griffiths and Liptak give in which the correlate is an indefinte, they also include one example where the correlate is a "quiz question", involving a genuine question with an in-situ *wh*-question (Griffiths and Lipták, 2014, fn. 15). They also claim that fragment answers to echo questions are island insensitive, though without giving an example.

(104) a. A: I heard that a biography of the youngest$_F$ Marx brother is going to be published this year.

b. A': Is the biography of the youngest$_F$ Marx brothers going to be published this year?

c. B: *No, of the oldest$_F$. (ex. 45, pg. 14)

(105) a. A: I imagine John wants a short$_F$ list.

b. A': Does John want a short$_F$ list?

c. B: *No, long$_F$. (ex. 46, pg. 14)

(106) a. A: I heard that Irv and John$_F$ were dancing together last night.

b. A': Were Irv and John$_F$ dancing together last night?

c. B: *No, Bill$_F$. (ex. 47, pg. 14)

(107) a. A: I heard they hired someone who speaks Bulgarian$_F$ fluently.

b. A': Did they hired someone who speaks Bulgarian$_F$ fluently?

c. B: *No, Serbo-Croatian$_F$. (ex. 48, pg. 14)

(108) a. A: I hear that Abby is likely to get mad if Ben$_F$ speaks to Mary.

b. A': Is Abby likely to get mad if Ben$_F$ speaks to Mary?

c. B: *No, Susan$_F$. (ex. 49, pg. 14)

(109) a. A: I heard that a biography of one of the Marx brothers is going to be published this year.

b. B: Yeah, of Groucho

(110) a. A: I imagine John wants a detailed list.

b. B: I'm afraid he does. Very detailed.

(111) a. A: I heard that Irv and a certain someone from your syntax class were dancing together last night.

b. B: Yeah, Bill.

(112) a. A: I heard they hired someone who speaks a Balkan language fluently.

b. B: Yeah, Serbo-Croatian.

(113) a. A: I hear that Abby is likely to get mad if Ben speaks to one of the guys from your syntax class.

b. B: Yeah, John.

(114) a. A: John F. Kennedy was killed in the city that which baseball team calls home?

b. B: The Texas Rangers.

Finally, (Merchant, 2004, pg.709), after discussing his claims of fragment island sensitivity, remarks briefly "that the above island sensitivity does not hold for a range of otherwise similar seeming construction types, such as correctives and multi-speaker cooperative sentence construction and certain confirmatory, clarificational, and elaborative fragments (as noted for elaborative fragments in Hoji and Fukaya (2001, p.12))." Unfortunately, no examples or further description of these types of constructions was given.

2.2.4.1. Summary of Prior Empirical Claims. To conclude this section, let me broadly summarize the range of claims from the literature, tabulated in (2.2)[5]. There is far from a general consensus on the island sensitivity of ellipsis. Neither Stripping nor Fragment Answers have been judged to be uniformly (in)sensitive. Nor is it the case that there is a clear consensus that non-contrastive elliptical configurations are insensitive and contrastive configuraitons sensitive. Nor has there been a shift over time towards considering these island-violating examples as acceptable or unacceptable. The takeaway from this

[5] Barros = Barros (2012b) BET = Barros et al. (2013) Casielles = Casielles (2006) CJ = Culicover and Jackendoff (2005) Fukaya = Fukaya (2007) GL = Griffiths and Lipták (2014) Krifka = Krifka (2006) Lappin = Lappin (1996) Merchant = Merchant (2004) Morgan = Morgan (1989) Stainton = Stainton (2006) Valmala = Valmala (2007) Weir = Weir (2014)

body of work is that a more formal, accountable, approach to data collection and analysis is called for.

Table 2.2. Summary of prior claims concerning Fragment Island Insensitivity

Ellipsis Type	Remnant Grammatical Function	Acceptable Judgment	Unacceptable Judgment
Elaborative Stripping	Subject	ø	ø
	Non-Subject	Lappin, CJ, Fukaya	ø
Corrective Stripping	Subject	ø	GL
	Non-Subject	Fukaya	GL
Non-Contrastive Stripping	Subject	GL, Valmala	ø
	Non-Subject	CJ, GL, Valmala	ø
Elaborative Fragment Answers	Subject	ø	ø
	Non-Subject	ø	ø
Corrective Fragment Answers	Subject	CJ, Casielles, Valmala, Weir	Morgan, Merchant, Barros, Fukaya, BET, GL
	Non-Subject	CJ, Casielles, Fukaya, Valmala, BET, Weir	Morgan, Merchant, GL
Non-Contrastive Fragment Answers	Subject	ø	ø
	Non-Subject	CJ	Krifka
Echo/Quiz Fragment Answers	Subject	GL	ø
	Non-Subject	Stainton	ø
Fragment answers to Sluice	Subject	Barros	ø
	Non-Subject	Barros, Valmala	ø
Multiple Fragment Answers	Subject	ø	ø
	Non-Subject	ø	Merchant, Krifka
Sprout-Stripping	Subject	CJ	ø
	Non-Subject	CJ	ø

2.2.5. Experimental Overview

As a first step towards taking a more formal approach to the status of island-violating ellipsis, I have conducted a series of four experiments, designed to establish a clear empirical baseline for the island (in)sensitivity of Stripping. I will first briefly describe the overall experimental design for these four experiments, and then justify a number of the choices made in the design of these experiments.

Experiments 1-4 involved acceptability judgement tasks, wherein participants were presented with a series of pairs of sentences, the *antecedent item* and the *target item*, and asked to rate the target item on a scale of one to seven points. The pseudo-randomly presented stimuli that participants saw included a number of filler items, to mask the intent behind the experiment, as well as systematically manipulated experimental stimuli which varied along three dimensions. The first dimension was whether the target items were Stripping remnants, preceded by a sentential polarity operator, *Yeah,* or *No,*, or non-reduced it-clefts, e.g. "It was James who bought the last round." In two of the experiments, the pivots and remnants consisted of proper names, and in the other two experiments, they were prepositional phrases which contained proper names, e.g. "to Mary". The second dimension along which the stimuli varied was whether the antecedent sentences contained a complement clause or a definite relative clause island, and it was within these embedded clauses that the correlate to the Stripping remnant or the cleft pivot could be found. The third dimension of differences between the stimuli was the information structural relationship between the correlate and remnant or pivot. In each experiment, half of the target items were an Elaborative continuation of the antecedent. The other half of the target items were either Corrective or Non-Contrastive continuations, depending on the experiment. The details of the design, including an example stimuli set, for each experiment are included below.

The premise of these experiments requires establishing a baseline for island violations, against which the Stripping data can be compared. It-clefts were chosen as that baseline, for a few reasons, instead of plausible alternatives like *wh*-movement or topicalization. First, and most importantly, using it-clefts as the baseline would allow the use of the same words or phrases as both Stripping remnants and cleft pivots. The use of *wh*-movement as the baseline, for example, would have required us to compare *wh*-phrases against nominal or prepositional phrases. Secondly, if we used *wh*-movement to establish a baseline, then if the question and Stripping phrases were not embedded, we would end up comparing constituent questions against Stripping examples, with the presumed differences in frequency being a cause for concern. Or if we did embed the question and Stripping phrases, we would have to contend with the difference in naturalness between embedded questions and embedded Stripping. Although I am convinced the latter is possible, as discussed by Weir (2014), it seems to me that the differences in naturalness between the two are substantial. Third, I avoided using topicalization to establish a baseline because pilot data suggested that topicalizations were generally rated relatively poorly, even when an appropriate linguistic context is given[6].

[6]Unfortunately, the intuition that non-island-violating It-Clefts would be rated better than topicalizations was not born out, as will soon be evident. I discuss possible reasons for these low ratings in the general discussion of these four experiments.

Using it-clefts to establish the baseline may not have been entirely unproblematic, however. It has been said that it-clefts introduce a particular information structural requirement of their own: exhaustivity (Kiss, 1998). This requirement might clash with an aspect of the experimental design for the elaborative experimental conditions, which require the use of the focus operator like *also*. Example (115) illustrates the relevant configuration. As we will see, however, it turns out that such clefts which seemingly violate the exhaustivity implication were, in general, rated only slightly worse than those which did not, such as those in the corrective or non-contrastive conditions. Furthermore, despite these differences in acceptability between elaborative and, e.g., corrective It-Clefts, we will see that such elaborative It-Clefts still demonstrate a sensitivity to islands. So it seems this worry is of no great concern.

(115) a. James met Mary.

b. Yeah, and it was also Bill who James met.

Another choice point in the design of these studies concerned which types of elliptical configurations to examine. Given the attention that Fragment Answers had received in the literature, addressing the relatively less studied Stripping constructions seemed warranted. As for different types of Stripping, it was clear that both non-contrastive and contrastive stripping were important to test, since, for example, Griffiths and Lipták (2014) explicitly argued that the former was island insensitive while the latter was not. Then, given that we can break down contrastive Stripping types into at least two subtypes, elaborative and corrective, as discussed in Chapter 1, and that others have distinguished between these two types in terms of island sensitivity, e.g. Merchant (2004), it was important to test these two types as well.

The choice to use definite relative clause islands, instead of another type of islands, also warrants brief discussion. Definite relative clauses are a type of "strong" or absolute island, from which neither PPs nor DPs can be extracted (Cinque, 1990; Den Dikken and Szabolcsi, 2002a). By using an absolute island, in combination with not just DP remnants, but PP remnants as well, we can be sure that any island insensitivity detected is not merely the result the permeability of the island. A potential drawback of using definite relative clause islands is that they are, in principle, ambiguous between restrictive and non-restrictive interpretations, even if in careful written text the two are distinguished in text by the use of a comma. For example, in a restrictive relative clause reading of (116) there are multiple students under discussion, of which the relevant one speaks Thai. In a non-restrictive reading, there is just one student under discussion, who speaks Thai.

(116) James met the student who speaks Thai.

While both types of relative clauses are said to be islands for extraction Chomsky (1977), Potts observes that non-restrictive relatives never permit extraction, whereas restrictive relatives sometimes do (Potts, 2005, pg. 135). Thus, the island conditions were, in principle at least, ambiguous between two parses, from one of which extraction may be more difficult. If participants freely varied in their interpretations of these relative clauses, then this free variation between parses could have in principle added to the noise present in the ratings.

Finally, a comment on the grammatical categories of the remnants. In two experiments, the remnants and pivots were prepositional phrases, and in the other two, proper names. It is said that prepositional phrases are more sensitive to extraction than are DPs, and that definite relative clause islands block extraction of both DPs and PPs (Den Dikken and Szabolcsi, 2002a). If we are to place the most difficult challenge to the notion of ellipsis island insensitivity, and using PPs would allow us to do so. DP remnants were also used under the suspicion that they might have proven somewhat more insensitive to islands than their PP counterparts. Finally, it is sometimes claimed that apparent island insensitivity is due to an ambiguity in which element serves as correlate (Winkler, 2013). By using PP remnants, and carefully controlling the argument structure of the verbs in the experimental items, we can be sure that these items are not ambiguous. To the degree that the DP and PP items reveal the same island sensitivity or insensitivity, we can be confident that the results for the DP remnant experiments were not due to such an ambiguity resolution.

2.2.6. Predictions

Finally, before we begin discussing the experiments and results in detail, we should consider what kinds of results we would expect to find if any of the Stripping configurations were sensitive or insensitive to islands. Consider figure (2.1). In it are two graphs, which both illustrate scenarios in which Stripping is island sensitive. In the graph on the left, we see a what appears to be a main effect of islandhood. That is, for both the it-cleft and Stripping conditions, the data show a lower mean rating for the island conditions than for the non-island conditions. There appears to be no other effects, with the ratings for the cleft and Stripping conditions otherwise comparable. We would expect this data if Stripping were sensitive to islands, as the difference between the island and non-island conditions doesn't differ between it-cleft conditions, which we assumed would show an island sensitivity, and Stripping conditions.

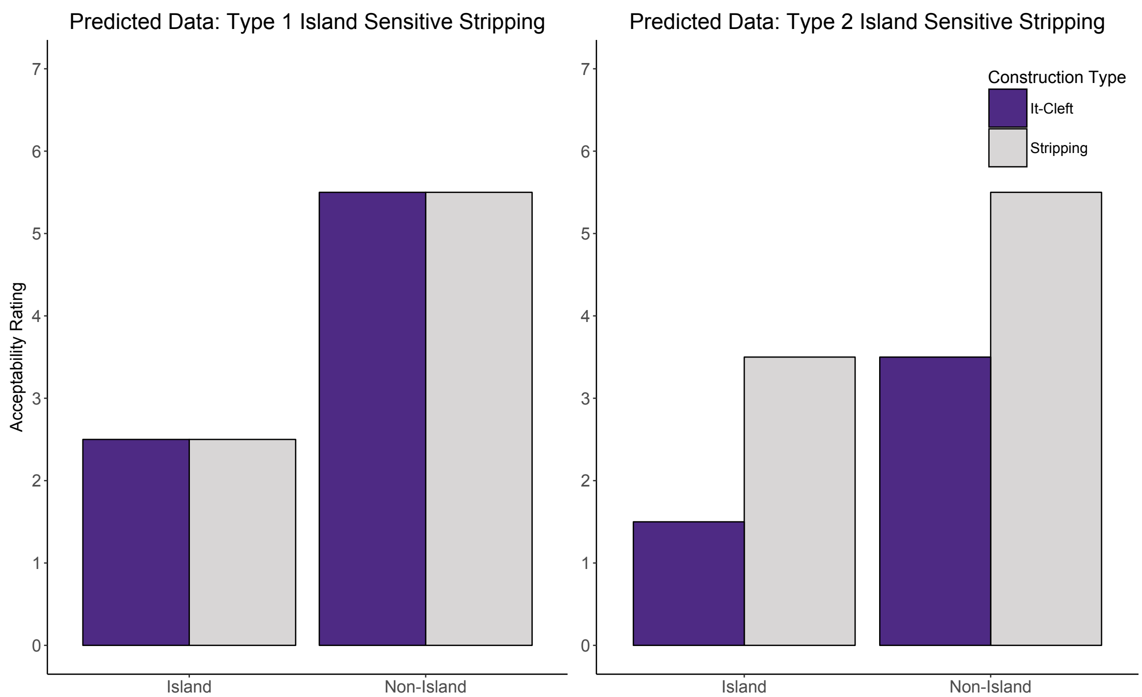

Figure 2.1. Two Types of Stripping Island Sensitivity

The graph on the right also shows what appears to be a main effect of islandhood, with both island conditions rated lower than non-island conditions. There also appears to be a main effect of construction type: the cleft conditions are rated lower than the Stripping conditions. Notice however that there doesn't appear to be an interaction between these two effects. The magnitude of the difference between the island and non-island conditions is the same for both construction types, which would be the result if there were no effect of length, reducing the acceptability of the longer, It-Cleft conditions.

Thus, barring alternative explanations and confounding factors, we would also conclude that the data in the graph on the right is consistent with Stripping being island sensitive. Of course, if we expected the non-island cleft conditions to be as acceptable as the non-island Stripping conditions, this chart might cause us pause. But what shouldn't dissuade us that this chart reflects an island sensitivity in the Stripping data is that the island Stripping cases are rated as acceptable as the non-island cleft conditions. We are not concerned here with absolute ratings, which I assume to be influenced by a variety of non-linguistic factors. Nor is it of primary concern here why there is a main effect of construction type. For the immediate purposes, we are concerned with whether or not the two factors we manipulated, construction type and islandhood, produce an interaction. Only if there is an interaction is it possible for the island effect in the Stripping cases to be different from that in the cleft conditions. In both of the graphs in figure (2.1), the data are consistent with Stripping being island sensitive in a way that is comparable to the island sensitivity of our it-cleft baseline.

If there is an interaction between construction type and islandhood, we could imagine two different ways the data could turn out. In figure (2.2), we again have two graphs, which here each illustrate a possible interaction between our two factors. On the left we see data that likely show main effects of both construction type and islandhood. But notice that the magnitude of the effect differs between the cleft and Stripping conditions; the island effect is smaller, but still present, in the Stripping conditions. This data would therefore

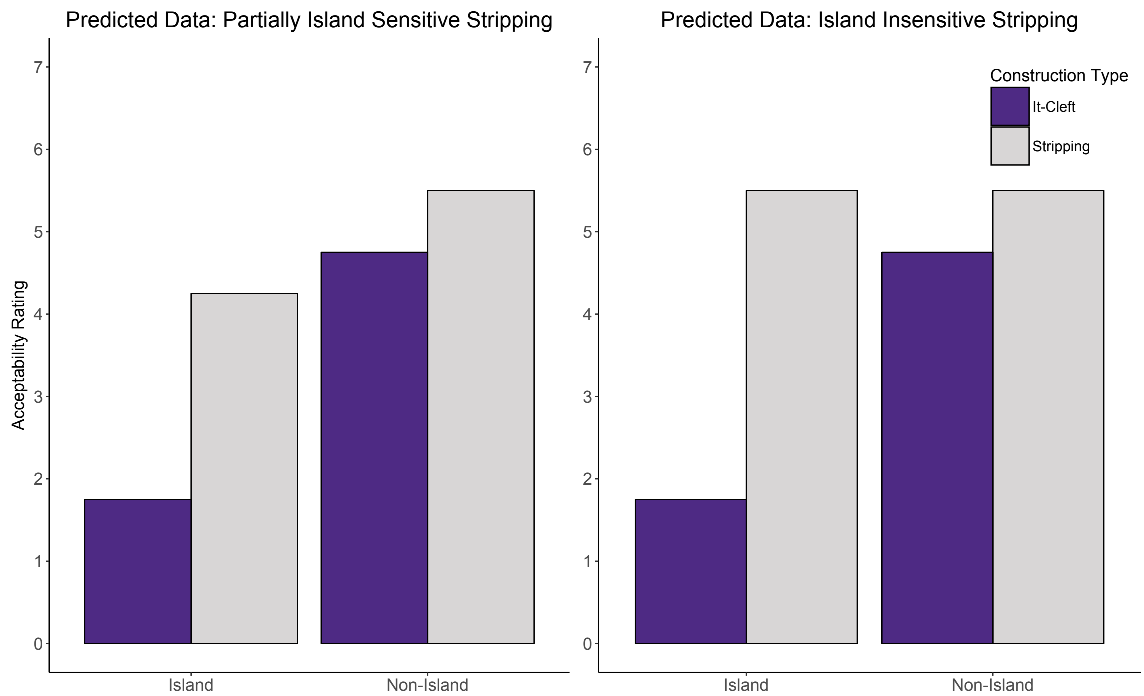

Figure 2.2. Partially and Completely Insensitive Stripping

be consistent with a partially island sensitive Stripping. In the graph on the right, we again see two main effects and an interaction. But here, there is no difference between the island and the non-island Stripping conditions. Here, we would say the data is consistent with an island insensitive Stripping.

2.3. Experiment 1

Experiment 1 investigates the island sensitivity of Elaborative and Non-Contrastive PP remnant Stripping. As discussed in more detail below, the analysis was made somewhat more complicated, due to a design flaw which introduced a confounding factor: some of the matrix verbs were factive, and so potentially introduced island effects into the non-island conditions. Ultimately, however, these results reveal a sensitivity to definite relative clause islands for elaborative Stripping, and island insensitivity for non-contrastive Stripping. Stripping, of all stripes, appears to be insensitive to Factive islands, unlike it-clefts.

2.3.1. Methods

2.3.1.1. Participants. 47 participants were recruited through Amazon's Mechanical Turk[7]. Participants were limited to IP addresses within the US, were only permitted to participate in the experiment once, and were compensated $2 USD. Six participants failed to complete the survey, and so their results were excluded from analysis. A further two participants were excluded on the basis of their rating of filler material. T-tests revealed these participants did not reliably rate the high acceptability fillers differently from the low acceptability fillers. Taking these exclusions into account, the data of 39 participants were analyzed.

2.3.1.2. Stimuli. The stimuli, of which a sample set is illustrated in (117) and (118) consisted of two-turn dialogues between 'Joe' and 'Bill'. The second phrase, Bill's, was the target item, which the participants were asked to rate. The stimuli conformed to a 2x2x2 factorial design: i. Islandhood; ii. Contrast Type; iii. Construction Type. The Islandhood factor manipulated whether the antecedent contained a complement clause (=Non-Island) or a definite relative clause (=Def. Island). The Contrast Type factor manipulated the information structural relationship between the correlate and the remnant or pivot, and the

[7] An online crowdsourcing marketplace platform (www.mturk.com)

antecedent and test phrases. In all conditions the remnant or pivot was a ‘to’ prepositional phrase containing a proper name. In Non-Contrastive conditions , the correlate was a PP containing an indefinite common noun phrase, and the test phrase was preceded by ‘Yeah’. In Elaborative conditions, the correlate was a PP containing another proper name, and the test phrase was preceded by ‘No’. The Construction Type factor manipulated whether the target phrase was Stripping construction (=Stripping) or a non-reduced it-cleft construction (=It-Cleft). The correlate, the Stripping remnant, and Cleft pivot were italicized, to implicitly indicate focus to the participants, and the target item was underlined, so as to indicate to the participants which phrase they were to rate.

Table 2.3. Experiment 1: Factors, Example Stimuli

Construction Type		Pronominal Status: Elaborative	Pronominal Status: Non-Contrastive
Stripping	Def. Island	Example (117a-i)	Example (117a-ii)
	Non-Island	Example (117b-i)	Example (117b-ii)
It-Cleft	Def. Island	Example (118a-i)	Example (118a-ii)
	Non-Island	Example (118b-i)	Example (118b-ii)

(117) Stripping

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: <u>Yeah, and also *to Ashley*.</u>

ii. Non-Contrastive:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: <u>Yeah, *to Ashley*.</u>

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and also *to Ashley*.

ii. Non-Contrastive:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: Yeah, *to Ashley*.

(118) It-Cleft

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and it was also *to Ashley* that Jordan mocked the salesman who sold a car.

ii. Non-Contrastive:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: Yeah, it was *to Ashley* that Jordan mocked the salesman who sold a car.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and it was also *to Ashley* that Jordan implied that the salesman sold a car.

ii. Non-Contrastive:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: Yeah, it was *to Ashley* that Jordan implied that the salesman sold a car.

Forty lexicalizations of the eight conditions were constructed. The embedded verbs, whether in the complement clause, or within the definite relative clause, were ditransitive verbs, with animate DP subject, inanimate DP direct object, and 'to-proper name' indirect objects. Seventy unique matrix verbs were used; ten of which were repeated across island and non-island conditions. Thirty of the remaining sixty verbs were used only in island conditions, and thirty only in non-island conditions.

During the analysis of the experimental results, it was discovered that twenty of the matrix verbs in the non-island conditions were *factive* verbs[8]. Factive verbs are distinguished from non-factive verbs in that the former, unlike the latter, presuppose the truth of their complements (Karttunen, 1971; De Cuba, 2007). Thus (119a) presupposes the truth of the complement. We can test that the complement is presupposed by embedding the complement under negation, as in (119b), with the result that complement must still be true in order for the example to have a truth value. Non-factive verbs, as in (119a), do not pattern in this way. Under negation, the embedded clause need not be true for the whole sentence to have a truth value. Such factive verbs are said to induce island effects for the

[8]Karttunen (1971) distinguishes between several types of Factive and Factive-like verbs. See also Hooper and Thompson (1973). I ignore these distinctions here.

extraction from the complement of the factive verb (Ross 1967, pg. 449; Erteschik-Shir and Lappin 1979 Melvold 1991; Krifka 2006). As a result of this unintentional potential confound, two additional analyses, as described below, were conducted to complement the original.

(119) a. James knows that Mary is dancing tonight.

b. James doesn't know that Mary is dancing tonight.

a. James said that Mary is dancing tonight.

b. James didn't say that Mary is dancing tonight.

Seventy filler items dialogues were also constructed, balanced equally between high and low acceptability items. The high acceptability items consisted of ten dialogues in which the item rated was a complement clause, five with *wh*-movement and pied-piped prepositional phrase items, ten with a *wh*-movement non-pied-piping items, eight with passive items, and two with a pseudo passive items. The low acceptability items consisted of fifteen items with some sort of island violation: 3 tokens of five island types (Adjunct, CNPC, Definite Relative Clause, Subject, and *wh*-Islands). A further ten low acceptability items involved other characteristics designed to lower acceptability, including items with sentential subjects which lacked 'that' (n=3), errors in number agreement (n=4), and errors in case (N=3). The final ten low acceptability filler items were composed of dialogues for which item rated was a Stripping construction, configured in such a way so as to implicate either a voice or an argument structure mismatch between the Stripping and antecedent phrases. In addition to these filler items, six practice items were constructed, two of which were high acceptability, two of which were medium, and two of which were low acceptability items. These items were not marked as practice items in the instructions, nor were they marked for their intended acceptability. This yielded a 1:1.9 target item to non-target item ratio.

Eight lists were intended to be constructed using a Latin square design, each list containing five distinct lexicalizations of each of the eight conditions. An error resulted in each list containing the intended distribution of conditions, five test items from each of the eight conditions, but incorrectly containing eight items from only five lexicalizations. Thus, participants saw five basic sentence frames in eight variations. Each list began with the six aforementioned practice items. In summary, each list contained 40 test items, 70 filler items, and 6 practice items, for a total of 116 items. Each of the eight lists was pseudo-randomized, which ensured that sequential items were not from the same experimental condition.

2.3.1.3. Procedure. The experiment was accessed by and presented to participants through the Amazon Mechanical Turk web portal, via the web-browser of their choice. They were instructed to read the presented dialogues and to rate the underlined phrases for naturalness on a scale of 1 (unnatural) to 7 (natural) in relation to the whole dialogue. The entire dialogue and the rating scale appeared on the same screen. Complete instructions are included in the appendices.

2.3.2. Experiment 1 Results

2.3.2.1. Experiment 1 Preliminary Analysis. Condition means and standard errors are reported in Table (2.4). The data[9] was then analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. (2008)) with rating as the dependent variable, using the lme4 package in R. Contrast-coded fixed effects included Islandhood (Island, Non-Island), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. The maximal random effects structure that would converge was employed, which included random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item. Model comparisons were performed to determine whether the inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

Table 2.4. Experiment 1: Definite Island v. Non-Island mean judgments, with standard errors in parentheses

		Def. Island	Non-Island
Stripping	Elaborative	4.25 (0.13)	4.96 (0.12)
	Non-Contrastive	4.49 (0.13)	4.92 (0.12)
It-Cleft	Elaborative	2.48 (0.12)	3.08 (0.13)
	Non-Contrastive	2.59 (0.12)	3.34 (0.13)

[9]I used the raw rating data in these analyses, instead of using a z-score transformation. Z-score transformations are often employed to control for variance in how e.g. rating scales are used across participants. However, there LMER models I have employed here are designed to account for such item and participant variance, making such data manipulations unnecessary.

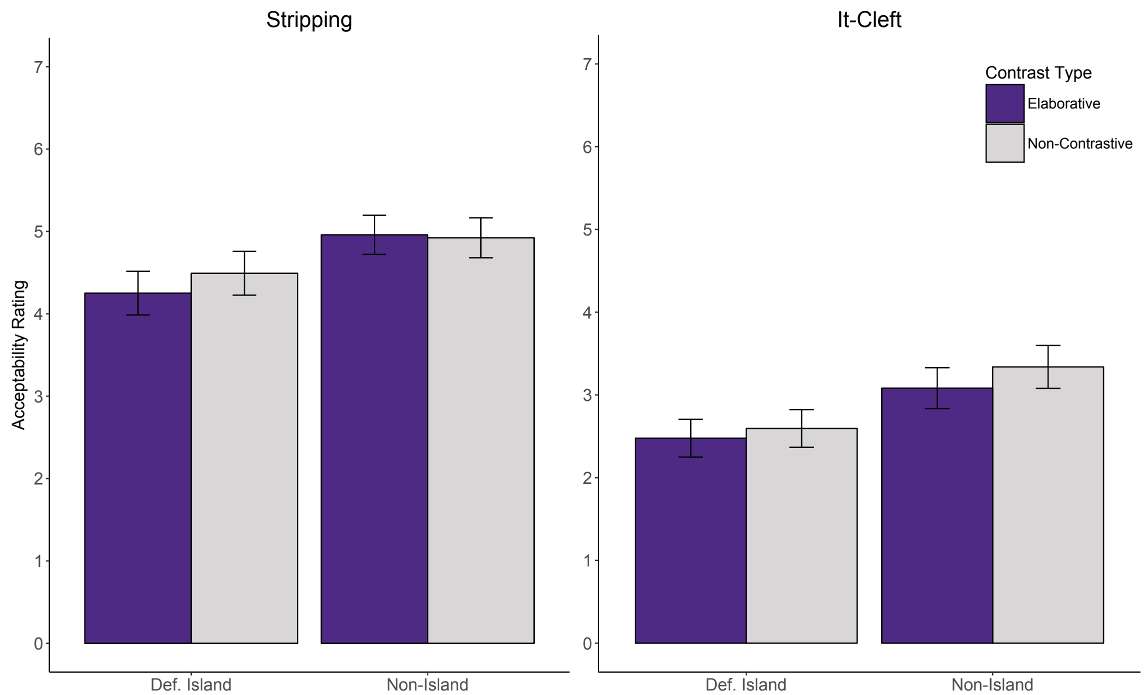

Figure 2.3. Experiment 1: Definite Island vs. Non-Island

The results of these analyses revealed significant main effects of Islandhood (β=0.63, SE β=0.11, $\chi^2(1)$=23.29, p<.001), and Construction Type (β=1.79, SE β=0.08, $\chi^2(1)$ =309.42, p<0.001). No other effects or interactions were observed (χ^2<2.54, p>0.11).

These results indicate that Stripping conditions were more acceptable than It-Cleft conditions, Non-Island conditions were more acceptable than Island conditions, and that there were no differences between Elaborative and Non-Contrastive conditions . Crucially, there were no interactions between any conditions. The lack of interaction between Islandhood and Construction Type indicates that the magnitude of the difference between Island and Non-Island conditions did not differ by construction type or by contrastivity type. Thus, these results would indicate that Stripping is island sensitive.

2.3.2.2. Experiment 1 Definite Island vs. Non-Factive Results. However, as mentioned above, the (non-)factivity of the matrix verbs in the Non-Island conditions could potentially have confounded the results. It could be that the acceptability of the Non-Island conditions was reduced as a result of the presence, and extraction from, the complement of, the factive verbs in these conditions. Consequently, two additional analyses were conducted, in an attempt to see what effect (non-)factivity played in this experiment.

In one analysis, the results from those items containing Factive verbs were eliminated from the analysis, allowing a comparison between those conditions containing just Definite Islands and those containing only Non-Factive verbs. In another, described in (2.3.2.3), the results from those items containing Definite Islands were omitted from the analysis. Instead, in this analysis, the Islandhood condition consisted of a comparison between those conditions containing Factive verbs and those containing Non-Factive verbs. We apply these two additional analyses here in experiment 1, as well as below for experiments 2-4.

In this analysis we exclude those items from the Non-Island conditions which contain Factive Verbs from the analysis. In every other respect, the methods of analysis remain identical to that of the main analysis: the ANNOVA comparison of multiple logistics linear mixed effects regress models, with raw rating as dependent variable. The contrast-coded fixed effects included Islandhood (Island, Non-Factive), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. As with the main analysis, the maximal random effects structure that would

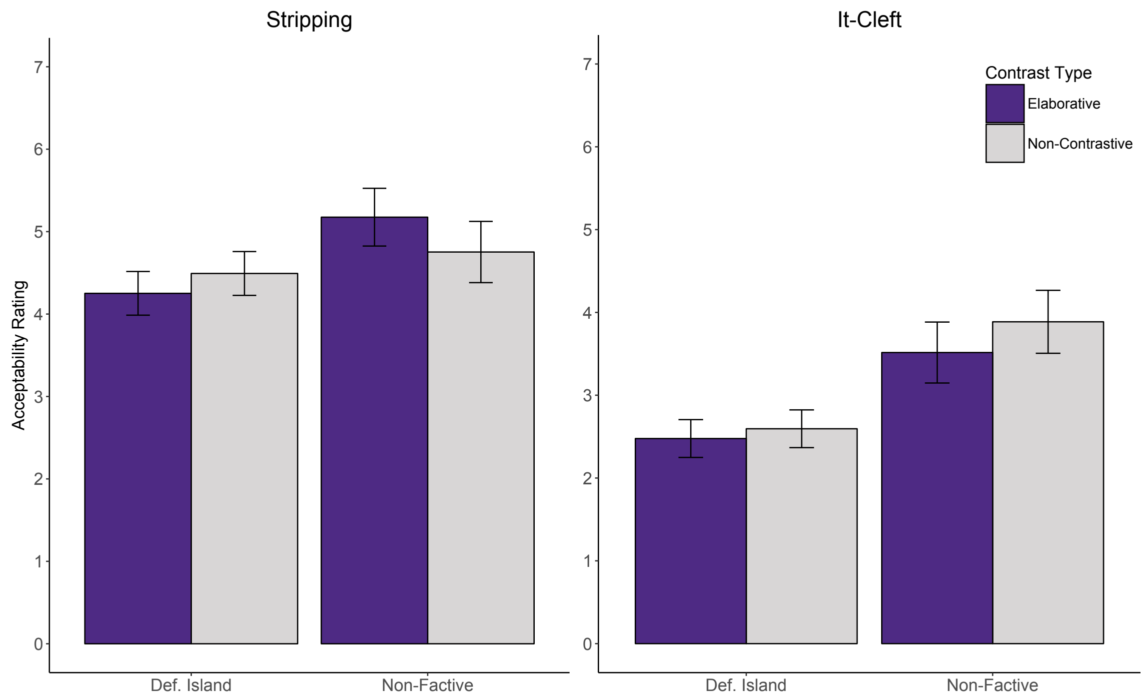

Figure 2.4. Experiment 1: Definite Island vs. Non-Factive

converge was employed, which included at most random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item.

Table 2.5. Experiment 1: Definite Island v. Non-Factive mean judgments, with standard errors in parentheses

		Def. Island	Non-Factive
Stripping	Elaborative	4.25 (0.13)	5.18 (0.18)
	Non-Contrastive	4.49 (0.13)	4.75 (0.19)
It-Cleft	Elaborative	2.48 (0.12)	3.52 (0.19)
	Non-Contrastive	2.59 (0.12)	3.89 (0.19)

Means and standard errors are reported in Table (2.5) and Figure (2.4). The results of the statistical analyses revealed significant main effects of Islandhood (β=0.74, SE β=0.16, $\chi^2(1)$=17.3, p$<$0.001) and Construction Type (β=-1.56, SE β=0.23, $\chi^2(1)$=31.08, p$<$.001). Also significant were the two way interactions between Islandhood and Construction type (β=0.56, SE β=0.18, $\chi^2(1)$=8.95 p=0.003), and Contrast Type and Construction type (β=0.34, SE β=0.17, $\chi^2(1)$=3.84 p=0.05), and the three way interaction between Islandhood, Contrast Type and Construction Type (β=0.92, SE β=0.35, $\chi^2(1)$=6.88 p=0.009). No other effects or interactions were observed ($\chi^2(1)$=1.21 p=0.272). To investigate the nature of these interactions, I analyzed subsets of the full data set.

Analysis of the Stripping subset revealed a main effect of Islandhood (β=0.48, SE β=0.16, $\chi^2(1)$=8.11 p=0.004), and an interaction between Islandhood and Contrast Type (β=-0.62, SE β=0.27, $\chi^2(1)$=4.96 p=0.026), but no other effects ($\chi^2(1)$=0.11 p=0.745). The Elaborative Stripping subset revealed a main effect of Islandhood (β=0.78, SE β=0.23, $\chi^2(1)$=9.96 p=0.002), while Islandhood in the Non-Contrastive Stripping subset did not reach significance (β=0.19, SE β=0.2, $\chi^2(1)$=0.86 p=0.355).

Analysis of the It-Cleft subset of the data revealed a main effect of Islandhood (β=1.03, SE β=0.23, $\chi^2(1)$=15.25, p$<$.001). The main effect of Contrast Type approached significance (β=0.24, SE β=0.12, $\chi^2(1)$=3.51 p=0.061), and the interaction between the two did not reach significance (β=0.24, SE β=0.18, $\chi^2(1)$=1.74 p=0.187).

Analysis of the Elaborative subset of the whole set revealed main effects of Islandhood (β=0.82, SE β=0.18, $\chi^2(1)$=15.65, p$<$.001) and Construction Type (β=-1.7, SE β=0.24, $\chi^2(1)$=33.69, p$<$.001), but no interaction between the two ($\chi^2(1)$=0.36 p=0.546).

Analysis of the Non-Contrastive subset revealed main effects of Islandhood (β=0.67, SE β=0.18, χ^2(1)=12.25, p<.001) and Construction Type (β=-1.41, SE β=0.27, χ^2(1)=21.77, p<.001), with an interaction between these factors (β=0.97, SE β=0.27, χ^2(1)=11.41, p<.001). The Stripping Non-Contrastive subset showed no effect of Islandhood (β=0.19, SE β=0.2, χ^2(1)=0.86 p=0.355), while the It-Cleft subset did show an effect of Islandhood (β=1.19, SE β=0.27, χ^2(1)=15.64, p<.001).

These results reveal an island effect for Corrective and Elaborative It-Cleft conditions and for Elaborative Stripping conditions. Note that the Stripping subset revealed an interaction between Islandhood and Contrast type, which indicates that Islandhood had a different effect in the Elaborative conditions than it had in the Non-Contrastive conditions. Additionally, the Elaborative subset revealed no interaction between Islandhood and Construction Type, indicating that the island effect did not differ between Elaborative Stripping and Elaborative It-Cleft conditions. In contrast, the Non-Contrastive subset did reveal an interaction between Islandhood and Construction type, indicating that the island effect did differ between Non-Contrastive Stripping and Non-Contrastive It-Cleft conditions. Finally, the examination of the Non-Contrastive Stripping subset revealed no Islandhood effect, while the corresponding Non-Contrastive It-Cleft and Elaborative Stripping subsets did reveal an effect of Islandhood. We conclude from this data that both Corrective and Elaborative It-Clefts demonstrate an island effect, and that the Elaborative Stripping conditions show a comparable island effect. The Non-Contrastive Stripping conditions, on the other hand, show no detectable island effect.

2.3.2.3. Experiment 1 Factive vs. Non-Factive Results. In the Factive vs. Non-Factive analysis, we set aside the items from the Island conditions from the analysis, and compare those items in the Non-Island conditions which contain Factive matrix verbs with those containing Non-Factive verbs. In every other respect, the methods of analysis remain identical to that of the main analysis: the ANNOVA comparison of multiple logistics linear mixed effects regress models, with raw rating as dependent variable. The contrast-coded fixed effects included Islandhood (Factive, Non-Factive), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. As with the main analysis, the maximal random effects structure that would converge was employed, which included at most random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item.

Table 2.6. Experiment 1: Factive v. Non-Factive mean judgments, with standard errors in parentheses

		Factive	Non-Factive
Stripping	Elaborative	4.74 (0.16)	5.18 (0.18)
	Non-Contrastive	5.09 (0.16)	4.75 (0.19)
It-Cleft	Elaborative	2.65 (0.16)	3.52 (0.19)
	Non-Contrastive	2.8 (0.16)	3.89 (0.19)

Table (2.6) and Figure (2.5) depict the means and standard errors. The results of the stastical analyses revealed significant main effects of Construction Type (β=-1.73, SE β=0.24, χ^2(1)=33.7, p<.001), the two-way interactions between Islandhood and Contrast Type (β=0.53, SE β=0.21, χ^2(1)=6.03 p=0.014), Islandhood and Construction Type

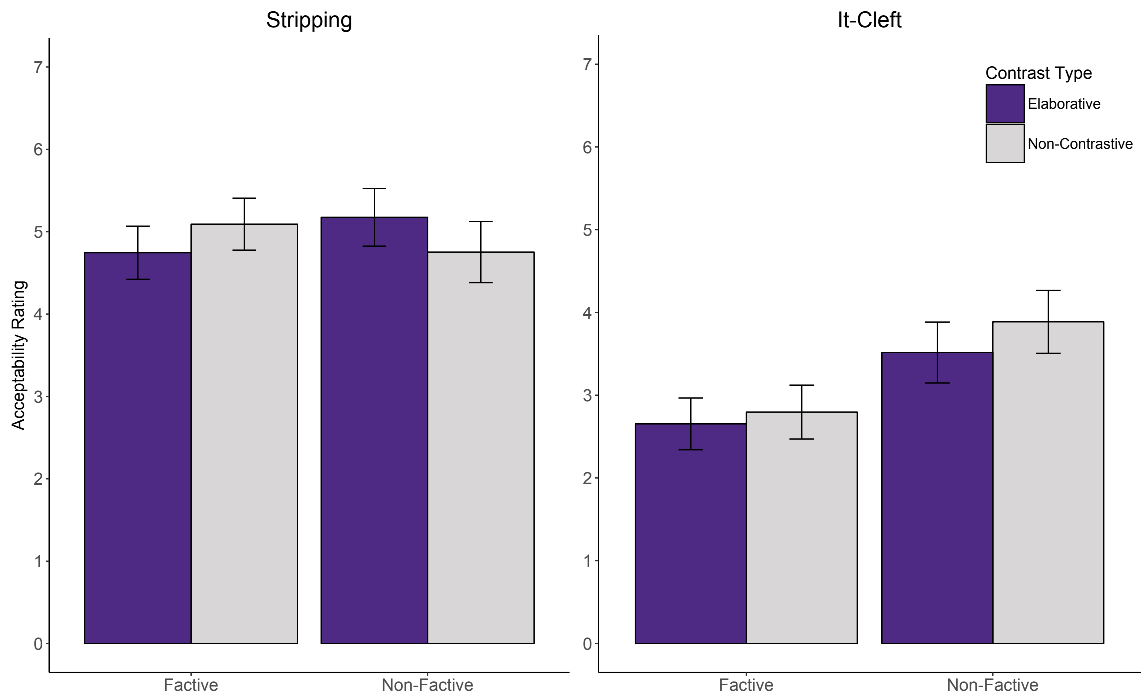

Figure 2.5. Experiment 1: Factive vs. Non-Factive

(β=0.87, SE β=0.26, χ^2(1)=10.28 p=0.001), and the three-way interaction between Islandhood, Contrast Type, and Construction Type (β=1, SE β=0.41, χ^2(1)=5.89 p=0.015). These interactions were investigated with subset analyses of the Stripping and It-Cleft conditions.

The Stripping subset revealed no effects of, or interactions between, Islandhood and Contrast Type, although the interaction between these factors did approach significance (β=-0.69, SE β=0.38, χ^2(1)=3.13 p=0.077). The It-Cleft subset revealed main effects of Islandhood (β=0.56, SE β=0.23, χ^2(1)=5.22 p=0.022) and Contrast Type (β=0.26, SE β=0.12, χ^2(1)=4.43 p=0.035), but no interaction between these factors (χ^2(1)=1.07 p=0.3).

Analysing the Elaborative conditions revealed main effects of Construction Type (β=-1.88, SE β=0.23, χ^2(1)=38.09, p<.001), and while the interaction between the two approached significance, (β=0.57, SE β=0.29, χ^2(1)=3.62 p=0.057), there was no main effect of Islandhood (β=0.1, SE β=0.16, χ^2(1)=0.41 p=0.524).

Pairwise comparisons between Island and Non-Island conditions reveal no main effect of Islandhood in the Elaborative Stripping subset (β=-0.04, SE β=0.28, χ^2(1)=0.02 p=0.881), while the Elaborative It-Cleft subset does exhibit a main effect of Islandhood (β=0.43, SE β=0.21, χ^2(1)=4.11 p= 0.043).

Analysis of the Non-Contrastive conditions revealed a main effect of Construction Type (β=-1.58, SE β=0.28, χ^2(1)=24.09, p<.001), and an interaction between Islandhood and Construction Type (β=1.24, SE β=0.37, χ^2(1)=10.29 p=0.001), but no main effect of Islandhood (χ^2(1)=0.29 p=0.593). The Stripping Non-Contrastive subset showed no effect of Islandhood (β=-0.5, SE β=0.33, χ^2(1)=2.2 p=0.138 5), while the It-Cleft subset did show an effect of Islandhood (β=0.75, SE β=0.33, χ^2(1)=4.92 p=0.026).

These results reveal that there is no effect of factivity in the Stripping conditions: The Factive Stripping conditions are no worse than the Non-Factive Stripping conditions. In contrast, the It-Cleft conditions did show an effect of Factivity: Factive It-Cleft conditions are worse than the Non-Factive conditions.

It is important to note, however, that, in contrast to extraction from relative clause islands, it has not been previously established that extraction from within the complement of factive verbs results in an island effect, in the sense of Sprouse et al. (2012). Extraction from relative clause islands results in a reduction in acceptability greater than the sum of reduction that would be expected given long distance extraction and presence of an island. It has been claimed that the complements of factive verbs are islands for extraction, but it hasn't been shown that the reduced acceptability for extraction from factive complements is greater than the sum of penalties for presence of a factive verb, if such exists, and long distance dependency. Whether the effect of factivity observed for the It-Cleft conditions is a genuine island effect or not, Stripping is completely insensitive to it.

2.3.3. Experiment 1 Discussion

Experiment 1 was intended to investigate the island sensitivity of Non-Contrastive and Elaborative Stripping. The preliminary analysis, in Section (2.3.2.1), revealed island sensitivity for both types of Stripping. However, that analysis also uncovered an unintended confound: the presence of a number of factive verbs in the non-island conditions. Given that factive verbs have been claimed to be islands for movement, the results of Experiment 1 were reanalyzed in two further steps. This reanalysis revealed that the factivity confound did indeed have an impact on the main analysis.

In the Island vs Non-Factive analysis presented in Section (2.3.2.2), we compared Island conditions against those non-island items which did not include a factive verb. Excluding the factive verbs revealed that Non-Contrastive Stripping was completely insensitive to the island effects, while Elaborative Stripping showed full island effects.

In Section (2.3.2.3), the Factive vs Non-Factive analysis compared those items from the Non-Island condition which contained a factive verb against those items which lacked a factive verb. We found that overt extraction from the complement of factive verbs, e.g.

the It-Cleft conditions, resulted in reduced acceptability, compared with extraction from non-factive verb complements. In the Stripping conditions, however, we found no trace of this reduction in acceptability in the Factive conditions. While it's not clear that the reduced acceptability in the It-Cleft Factive conditions is an *island* effect, in the sense that it is a reduction beyond the additive effects of long-distance movement and presence of the island, we do see that the effect is absent in both Elaborative and Non-Contrastive Stripping conditions.

Comparison of the results of the preliminary analysis against the results of the Island vs. Non-Factive and Factive vs. Non-Factive analyses reveals that the complete sensitivity of Elaborative Stripping found in the preliminary analysis can be attributed to the presence of factive verbs in the Non-Island conditions. When these factive verbs were excluded from the analysis, as in the Island vs. Non-Factive analysis, Elaborative Stripping showed only partial sensitivity. And when the Factive and Non-Factive items of the Non-Island conditions were compared, we found a Factivity effect only in the It-Cleft conditions. Thus the presence of factive verbs reduced the acceptability of the It-Cleft Non-Island conditions such that the magnitude of the 'island effect' was comparable to that in the Stripping conditions. However, once we eliminate the Factivity confound, the results of this experiment show that Non-Contrastive PP Stripping is completely insensitive to DRC islands, and Elaborative PP Stripping island sensitive.

In addition to the issue of Factivity, which affected Experiments 1-4, it is important to reiterate that Experiment 1, but only Experiment 1, was subject to an error in the distribution of sentence frames in the materials. Each participant saw the correct number of items from each condition, but the type of sentence frames used to present these conditions were limited. Each participant saw only five sentence frames, each in each of the eight conditions. Thus it is possible that the participants became attuned to the experimental manipulation, thereby affecting their responses. However, as will be discussed in the general discussion, the results of Experiment 1 were largely in line with those of the other experiments. In particular, Experiment 3 also found non-contrastive Stripping, with DP remnants instead of PP remnants, to be island insensitive. Likewise, experiments 1-4 all found Stripping to be insensitive to Factive islands. On the other hand, Experiment 1 was the only experiment in which the Island vs Non-Factive analysis found Elaborative Stripping to be completely sensitive to the definite relative clause island effect. Elaborative Stripping, a factor which was present in each of the other experiments, was found to be *partially* sensitive in each of experiments 2-4, as will be discussed shortly in the General Discussion.

2.4. Experiment 2

Experiment 2 builds on Experiment 1; again the island sensitivity of PP remnant Stripping was investigated. In Experiment 2, Elaborative Stripping was tested, but, here, in comparison to Corrective Stripping. The same factivity issues affected Experiment 2 as affected Experiment 1, the analysis of which reveal the same amelioration of factivity in the Stripping conditions. As concerns the effects of definite relative clause islands, in contrast to the results of Experiment 1, Experiment 2 revealed partial amelioration for not only Corrective Stripping, but Elaborative Stripping as well.

2.4.1. Methods

2.4.1.1. Participants. 43 participants were recruited through Amazon's Mechanical Turk. Participants were limited to IP addresses within the US, were only permitted to participate in the experiment once, and were compensated $2 USD. Five participants failed to complete the survey, and so their results were excluded from analysis. A further three participants were excluded on the basis of their rating of filler material. T-tests revealed these participants did not reliably rate the high acceptability fillers differently from the low acceptability fillers. Taking these exclusions into account, the data of 35 participants were analyzed.

2.4.1.2. Design, Stimuli, and Procedures. The design, stimuli and procedures of Experiment 2 were identical to those of Experiment 1 with the exception of the details of the Contrast Type factor. In Experiment 2, the Contrast Type factor compares Elaborative conditions against Corrective condition. The Corrective and Elaborative conditions are distinguished by the sentential polarity marker which precedes the target phrase. In Elaborative conditions, the target phrase is preceded by "Yeah, . . . ", while Corrective conditions are preceded by "No, . . . ". These contrast types are discussed in more detail in Chapter 1. The factorial design is depicted in Table (2.7) and a sample stimuli paradigm is given in (120).

The same forty lexicalizations as used in Experiment 1 were again employed in Experiment 2, modulo the Corrective Contrast Type replacing the Non-Contrastive type. One consequence of the parallel design of Experiments 1 and 2 is that the same factive verb confound, as described in detail in the section (2.3.1.2), of also affected this experiment. Consequently the same two additional analyses were conducted.

Finally, the same seventy filler items and six practice items described in Experiment 1 were utilized in this experiment in eight pseudorandomized Latin square lists.

Table 2.7. Experiment 2: Factors, Example Stimuli

Construction Type		Pronominal Status Elaborative	 Corrective
Stripping	Def. Island	Example (120a-i)	Example (120a-ii)
	Non-Island	Example (120b-i)	Example (120b-ii)
It-Cleft	Def. Island	Example (121a-i)	Example (121a-ii)
	Non-Island	Example (121b-i)	Example (121b-ii)

(120) Stripping

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and also *to Ashley*.

ii. Corrective:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: No, *to Ashley*.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and also *to Ashley*.

ii. Corrective:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: No, *to Ashley*.

(121) It-Cleft

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and it was also *to Ashley* that Jordan mocked the salesman who sold a car.

ii. Corrective:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: No, it was *to Ashley* that Jordan mocked the salesman who sold a car.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and it was also *to Ashley* that Jordan implied that the salesman sold a car.

ii. Corrective:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: No, it was *to Ashley* that Jordan implied that the salesman sold a car.

2.4.2. Experiment 2 Results

2.4.2.1. Experiment 2 Preliminary Analysis. This preliminary analysis compares Stripping against It-Cleft conditions, Elaborative against Corrective conditions, and DCRI against Non-Island conditions. However, as noted above in Experiment 1, the Non-Island conditions in these preliminary analyses included both Factive and Non-Factive verbs, which we saw above yielded a reduction in the acceptability for just the It-Cleft conditions. This confound in turn influenced the pattern of results; such that it appeared that the Elaborative conditions were completely island sensitive. However, further analysis, in which the factive verbs of the Non-Island conditions were excluded revealed only partial sensitivity for the Elaborative conditions. Given that the same confound affects Experiment 2, the results of the preliminary analysis should be likewise understood with this confound in mind.

Condition means and standard errors can be found in Table (2.8) and Figure (2.6). The data[10] was then analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. (2008)) with rating as the dependent variable, using the lme4 package in R. Contrast-coded fixed effects included Islandhood (Island, Non-Island), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. The maximal random effects structure that would converge was employed, which included random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item.

[10]As in experiment 1 and elsewhere, I used the raw rating data in these analyses, instead of using a z-score transformation, as the LMER models employed were constructed to account for by participant variance.

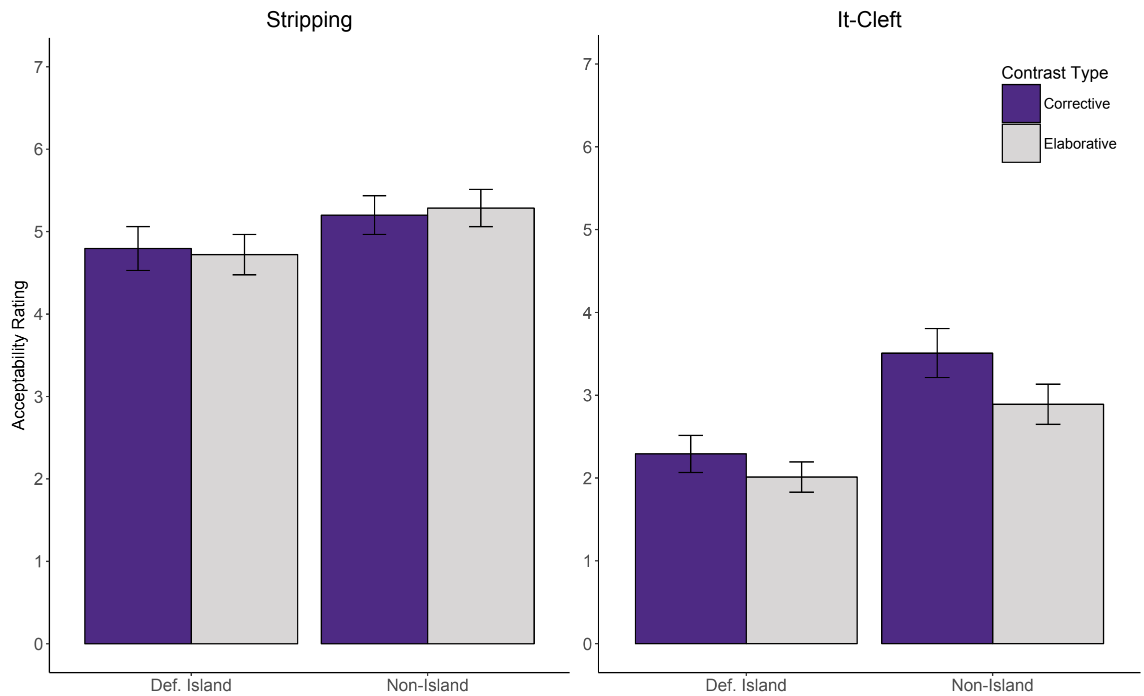

Figure 2.6. Experiment 2: Definite Island vs. Non-Island

Model comparisons were performed to determine whether the inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

Table 2.8. Experiment 2: Definite Island v. Non-Island mean judgments, with standard errors in parentheses

		Def. Island	Non-Island
Stripping	Elaborative	4.72 (0.12)	5.29 (0.11)
	Corrective	4.79 (0.13)	5.2 (0.12)
It-Cleft	Elaborative	2.01 (0.09)	2.89 (0.12))
	Corrective	2.29 (0.11)	3.51 (0.15)

The results of these analyses revealed significant main effects of Islandhood (β=0.76, SE β=0.1, χ^2(1)=35.97, p<.001), Contrast Type (β=0.22, SE β=0.09, χ^2(1)=6.23 p=0.013), and Construction Type (β=-2.32, SE β=0.2, χ^2(1)=57.2, p<.001). Also significant were the two way interactions between Islandhood and Construction Type (β=0.56, SE β=0.14, χ^2(1)=15.3, p<.001) and between Contrast Type and Construction Type (β=0.44, SE β=0.14, χ^2(1)=9.47 p=0.002). The three way interaction approached but did not reach significance (β=0.5, SE β=0.28, χ^2(1)=3.09 p=0.079) and no other effects or interactions were observed (χ^2(1)=0.39 p=0.53). To investigate the nature of these interactions, I analyzed subsets of the full data set.

The Stripping subset revealed a main effect of Islandhood (β=0.48, SE β=0.11, χ^2(1) =14.78, p<.001), with no other effects or interactions (χ^2(1)=0.84 p=0.361). The It-Cleft subset revealed main effects of Islandhood (β=1.05, SE β=0.16, χ^2(1)=28.46, p<.001) and Contrast Type (β=0.44, SE β=0.12, χ^2(1)=13.18, p<.001). No interaction was found (χ^2(1)=2.47 p=0.116).

Analysis of the Elaborative subset of the whole dataset revealed main effects of Islandhood (β=0.72, SE β=0.13, χ^2(1)=24.56, p<.001) and Construction Type (β=-2.55, SE β=0.19, χ^2(1)=64.03, p<.001), but no interaction between the two (χ^2(1)=2.61 p>0.1). Analysis of the Corrective subset revealed both main effects of Islandhood (β=0.81, SE β=0.13, χ^2(1)=28.98, p<.001) and Construction Type (β=-2.1, SE β=0.23, χ^2(1)=44.4, p<.001), as well as an interaction between the two (β=0.8, SE β=0.21, χ^2(1)=13.38, p<.001). Both Corrective It-Cleft (β=1.21, SE β=0.22, χ^2(1)=24.24, p<.001) and Corrective Stripping (β=0.4, SE β=0.13, χ^2(1)=8.06 p=0.005) subsets revealed main effects of Islandhood.

As in the preliminary analysis of Experiment 1, this preliminary analysis reveals Non-Contrastive Stripping to be island insensitive, while Elaborative Stripping, along with showing both Non-Contrastive and Elaborative It-Clefts to be island sensitive. However, as in Experiment 1, this preliminary analysis contained a confound, in that half of the items in the Non-Island conditions contained Factive verbs. The next two analyses will reveal a similar impact of Factivity was found in Experiment 2, such that, as we will see shortly, both Elaborative and Corrective Stripping are indeed only partially sensitive to definite relative clause island effects.

2.4.2.2. Experiment 2 Definite Island vs. Non-Factive Results. To account for the possibility that Factivity introduced a confound into the experimental design, we next compared the Definite Island conditions against those items in the Non-Island conditions which did not contain a factive matrix verb. Condition means and standard errors can be found in Table (2.9) and Figure (2.7). The results of these analyses revealed significant main effects of Islandhood (β=0.94, SE β=0.11, $\chi^2(1)$=40.29, p<.001), Contrast Type (β=0.33, SE β=0.1, $\chi^2(1)$=9.48 p=0.002), and Construction Type (β=-2.22, SE β=0.21, $\chi^2(1)$=51.64, p<.001). Also significant were the two-way interactions between Islandhood and Construction Type (β=0.78, SE β=0.17, $\chi^2(1)$=20.89, p<.001) and between Contrast Type and Construction Type (β=0.53, SE β=0.17, $\chi^2(1)$=9.84 p=0.002), and the three-way interaction between Islandhood, Contrast Type, and Construction Type (β=0.67, SE β=0.33, $\chi^2(1)$=4.04 p=0.044). The two way interaction between Islandhood and Contrast Type approached, but did not reach, significance (β=0.31, SE β=0.17, $\chi^2(1)$=3.38 p=0.066). To investigate the nature of these interactions, I analyzed subsets of the full data set.

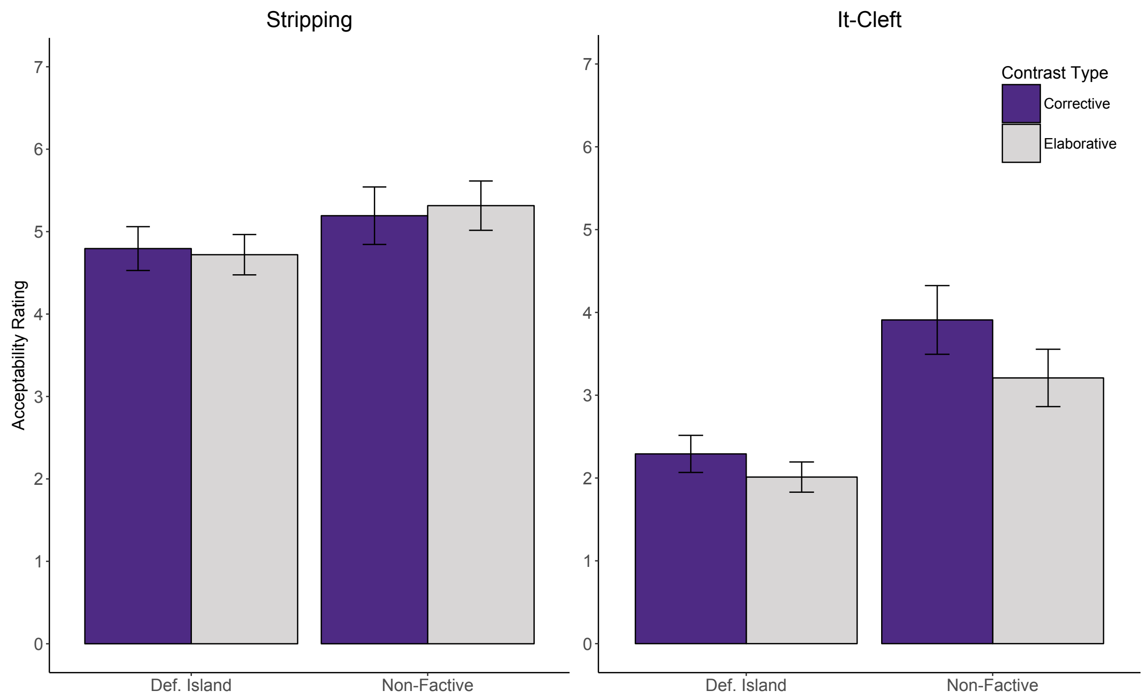

Figure 2.7. Experiment 2: Definite Island vs. Non-Factive

Table 2.9. Experiment 2: Definite Island v. Non-Factive mean judgments, with standard errors in parentheses

		Def. Island	Non-Factive
Stripping	Elaborative	4.72 (0.12)	5.32 (0.15)
	Corrective	4.79 (0.13)	5.19 (0.18)
It-Cleft	Elaborative	2.01 (0.09)	3.21 (0.17)
	Corrective	2.29 (0.11)	3.91 (0.21)

Analysis of the Stripping subset revealed a main effect of Islandhood (β=0.52, SE β=0.12, $\chi^2(1)$=14.29, p<.001), but no other effects ($\chi^2(1)$=0.17 p>0.67). Thus, there is an island effect in the Stripping conditions which doesn't differ between Elaborative and Corrective conditions. Analysis of the It-Cleft subset of the data revealed main effects of Islandhood (β=1.34, SE β=0.19, $\chi^2(1)$=32.41, p<.001) and Contrast Type (β=0.63, SE

β=0.13, $\chi^2(1)$=19.76, p<.001), and an interaction between the two (β=0.71, SE β=0.24, $\chi^2(1)$=8.21 p=0.004). This tells us that there is an island effect in the It-Cleft conditions, but that the magnitude of that effect differs between Elaborative and Corrective It-Cleft conditions, an effect which we can trace to a lower acceptability rating in the Non-Factive Elaborative It-Cleft conditions (3.21), as compared to the ratings of the Non-Factive Corrective It-Cleft conditions (3.91).

Analysis of the Elaborative subset of the whole set revealed main effects of Islandhood (β=0.79, SE β=0.15, $\chi^2(1)$=22.72, p<.001) and Construction Type (β=-2.47, SE β=0.21, $\chi^2(1)$=58.85, p<.001), with an interaction between the two (β=0.48, SE β=0.24, $\chi^2(1)$=3.92 p=0.048). Both Elaborative Stripping conditions (β=0.57, SE β=0.16, $\chi^2(1)$ =10.11 p=0.001) and Elaborative It-Cleft conditions (β=0.99, SE β=0.23, $\chi^2(1)$=15.56, p<.001 and β=1.68) revealed main effects of Islandhood. The interaction between Construction Type and Islandhood, along with inspection of their means, indicates that the magnitude of the effect of Islandhood is reliably larger in the Elaborative It-Cleft conditions than in the Elaborative Stripping Conditions.

Analysis of the Corrective subset revealed main effects of Islandhood (β=1.09, SE β=0.15, $\chi^2(1)$=33.44, p<.001) and Construction Type (β=-1.94, SE β=0.24, $\chi^2(1)$=37.21, p<.001), with an interaction between these factors (β=1.13, SE β=0.25, $\chi^2(1)$=18.89, p<.001). Both Corrective Stripping conditions (β=0.47, SE β=0.17, $\chi^2(1)$=7.23 p=0.007) and Corrective It-Cleft conditions (β=1.68, SE β=0.25, $\chi^2(1)$=30.02, p<.001) revealed Islandhood main effects. Again, the interaction of Islandhood and Construction type within the Corrective subset, along with inspection of condition means, indicates that the island effect is smaller in the Corrective Stripping conditions than in the Corrective It-Cleft conditions.

These results reveal that the Island conditions are rated worse than the comparable Non-Factive conditions. Both It-Cleft and Stripping conditions show an island effect, but the island effect is not uniform across the two contrast types and the two construction types. Between the Corrective It-Cleft and Elaborative It-Cleft conditions, smaller in the latter. The island effect is smaller in the Stripping conditions than in the It-Cleft conditions, results which indicate that Stripping, of both Elaborative and Corrective types, is partially sensitive to Island effects.

2.4.2.3. Experiment 2 Factive vs. Non-Factive Results. Next, we compare those Non-Island items containing Factive verbs with those containing Non-Factive verbs. Condition means and standard errors are reported in Figure (2.8) and Table (2.10). Statistical analysis revealed a main effect of Construction Type (β=-1.73, SE β=0.24, $\chi^2(1)$=33.7, p<.001). Also significant were the two-way interactions between Islandhood and Construction Type (β=0.87, SE β=0.26, $\chi^2(1)$=10.28 p=0.001), and the three-way interaction between Islandhood, Contrast Type, and Construction Type (β=1, SE β=0.41, $\chi^2(1)$=5.89 p=0.015). No other effects or interactions were observed ($\chi^2(1)$=2.29 p=0.13). These interactions were investigated with subset analyses of the Stripping and It-Cleft conditions.

The Stripping subset revealed no effects of or interactions between Islandhood and Contrast Type ($\chi^2(1)$=0.51 p>0.47). The It-Cleft subset revealed main effects of Islandhood (β=0.54, SE β=0.18, $\chi^2(1)$=8.06 p=0.005) and Contrast Type (β=0.62, SE β=0.18, $\chi^2(1)$=10.68 p=0.001), and an interaction between these factors (β=0.8, SE β=0.33, $\chi^2(1)$

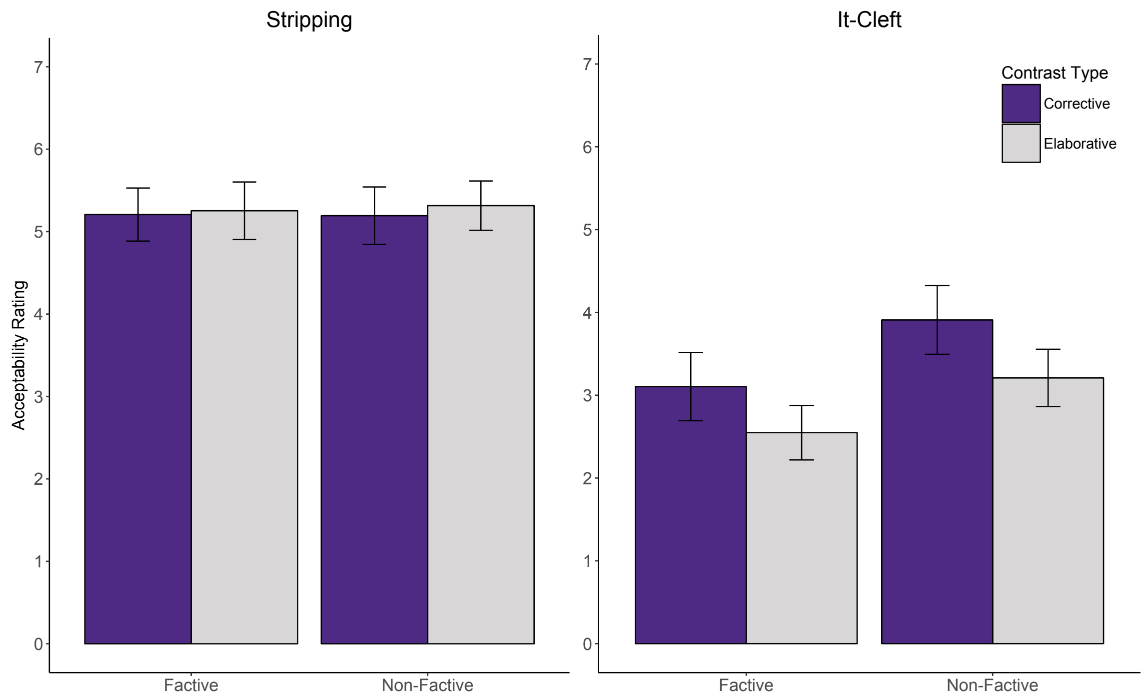

Figure 2.8. Experiment 2: Factive vs. Non-Factive

Table 2.10. Experiment 2: Factive v. Non-Factive mean judgments, with standard errors in parentheses

		Non-Factive	Factive
Stripping	Elaborative	5.32 (0.15)	5.25 (0.18)
	Corrective	5.19 (0.18)	5.21 (0.16)
It-Cleft	Elaborative	3.21 (0.17)	2.55 (0.17)
	Corrective	3.91 (0.21)	3.1 (0.21)

=5.49 p=0.019). The Corrective It-Cleft subset showed an effect of Islandhood (β=0.93, SE β=0.28, $\chi^2(1)$=8.7 p=0.003), while the Elaborative It-Cleft subset revealed no such effect ($\chi^2(1)$=0.52 p>0.46).

These results reveal an effect of factivity for Corrective It-Cleft conditions, but no comparable effect for either of the Stripping types or for Elaborative It-Cleft conditions.

The Stripping subset showed no effects of either Contrast Type or Islandhood, while the It-Cleft conditions revealed an interaction between these factors. That interaction was due to an effect of Islandhood in the Corrective It-Cleft conditions but not in the Elaborative It-Cleft conditions.

2.4.3. Experiment 2 Discussion

In this experiment, we compared Stripping against It-Cleft conditions, Definite Relative Clause Island against Non-Island conditions, and Elaborative against Corrective conditions. The preliminary analysis revealed full or partial Island sensitivity in the Stripping conditions.

However, as in Experiment 1, this preliminary analysis was confounded by the presence, in half of the items in the Non-Island conditions, of a Factive verb. Factoring out these factive items in the Island vs. Non-Factive analysis, we found partial Island sensitivity for both Stripping types. Then, comparing the Factive and Non-Factive items, we found a factivity effect for the It-Cleft conditions, but no such effect for the Stripping conditions. Thus, in the preliminary analysis, the overall acceptability of the Non-Island It-Cleft conditions was reduced by the presence of Factive islands in this condition. This reduction in acceptability of the Non-Island It-Cleft conditions was sufficient to yield comparable island effects between the Stripping and It-Cleft conditions, in turn yielding what appeared to be partial and full island sensitivity for the Stripping conditions. However, the analysis in which these Factive items were excluded reveals that the magnitude of the island effect in the It-Cleft conditions is greater than the still present effect in the Stripping conditions.

Thus these results support the claim that Elaborative and Corrective Stripping both are partially sensitive to DCRI effects, and completely insensitive to Factivity effects.

2.5. Experiment 3

Experiments 3 and 4 build on Experiments 1 and 2, mirroring the intent and design of those experiments. However in both Experiments 3 and 4, the remnants and cleft-pivots are DPs, in place of the PPs found in Experiments 1 and 2. Here, in Experiment 3, we compared Elaborative Stripping and It-Clefts with Non-Contrastive Stripping and It-Clefts, as we did in Experiment 1. Experiment 4 compares Elaborative and Corrective Contrast types.

2.5.1. Methods

2.5.1.1. Participants. 42 participants were recruited through Amazon's Mechanical Turk. Three participants failed to complete the survey, and so their results were excluded from analysis. Taking these exclusions into account, the data of 39 participants were analyzed.

2.5.1.2. Design, Stimuli, and Procedures. The design, stimuli and procedures of Experiment 3 were identical to those of Experiment 1 with the exception that in Experiment 3 the Stripping remnant and Cleft pivot in the target items were DPs, instead of the PP Stripping remnants and Cleft pivots used in Experiment 1. The experimental factors are given in Table (2.11) and a sample stimuli paradigm is given in examples (122) and (123).

The same forty lexicalizations as used in Experiment 1 were again employed in Experiment 3, modulo the shift from PP to DP Stripping remnants and Cleft pivots. The same factive verb confound, as described in detail in the section (2.3.1.2), of also affected this experiment. Consequently the same two additional analyses were conducted. The same filler and six practice items described in Experiment 1 were utilized in this experiment. Likewise, eight pseudorandomized Latin square lists were constructed.

Table 2.11. Experiment 3: Factors, Example Stimuli

Construction Type		Pronominal Status: Elaborative	Pronominal Status: Non-Contrastive
Stripping	Def. Island	Example (122a-i)	Example (123a-ii)
	Non-Island	Example (122b-i)	Example (123b-ii)
It-Cleft	Def. Island	Example (123a-i)	Example (122a-ii)
	Non-Island	Example (123b-i)	Example (122b-ii)

(122) Stripping

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and also *Ashley*.

ii. Non-Contrastive:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: Yeah, *Ashley*.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and also *Ashley*.

ii. Non-Contrastive:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: Yeah, *Ashley*.

(123) It-Cleft

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and it was also *Ashley* that Jordan mocked the salesman who sold a car to.

ii. Non-Contrastive:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: Yeah, it was *Ashley* that Jordan mocked the salesman who sold a car to.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and it was also *Ashley* that Jordan implied that the salesman sold a car to.

ii. Non-Contrastive:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: Yeah, it was *Ashley* that Jordan implied that the salesman sold a car to.

2.5.2. Experiment 3 Results

2.5.2.1. Experiment 3 Preliminary Analysis. As noted of Experiments 1 and 2, Experiment 3 also suffered a confound, in the presence of factive verbs in what had been intended to be non-island conditions. In those experiments, we found an effect of these factive verbs on just the It-Cleft conditions, a result of which is that, in these preliminary analyses, the island effect in the It-Cleft conditions was diminished to a magnitude comparable to that found in Elaborative and Corrective Stripping. As we will see, these same considerations apply to Experiment 3, and so the results of the preliminary analysis should be understood in this light.

Condition means and standard errors can be found in Table (2.12) and Figure (2.9). The data[11] was then analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. (2008)) with rating as the dependent variable, using the lme4 package in R. Contrast-coded fixed effects included Islandhood (Island, Non-Island), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. The maximal random effects structure that would converge was employed, which included random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item. Model comparisons were performed to determine whether the inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

The results of these analyses revealed significant main effects of Islandhood (β=0.64, SE β=0.11, $\chi^2(1)$=25.64, p<.001), Contrast Type (β=0.25, SE β=0.11, $\chi^2(1)$=5.18 p=0.023),

[11] As in Experiment 1 and elsewhere, I used the raw rating data in these analyses, instead of using a z-score transformation, as the LMER models employed were constructed to account for by participant variance.

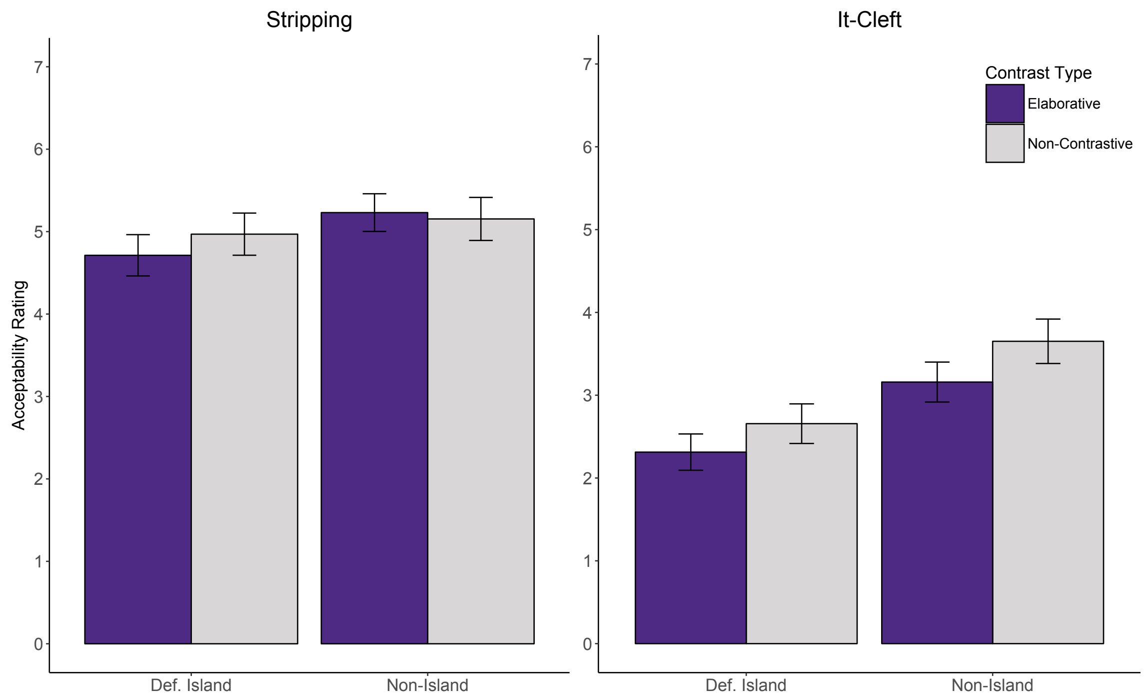

Figure 2.9. Experiment 3: Definite Island vs. Non-Island

Table 2.12. Experiment 3: Definite Island v. Non-Island mean judgments, with standard errors in parentheses

		Def. Island	Non-Island
Stripping	Elaborative	4.71 (0.13)	5.23 (0.12)
	Non-Contrastive	4.97 (0.13)	5.15 (0.13)
It-Cleft	Elaborative	2.31 (0.11)	3.16 (0.12)
	Non-Contrastive	2.66 (0.12)	3.65 (0.14)

and Construction Type (β=-2.07, SE β=0.21, $\chi^2(1)$=48.97, p<.001). Also significant were the two way interactions between Islandhood and Construction Type (β=0.58, SE β=0.16, $\chi^2(1)$=12.86, p<.001) and between Contrast Type and Construction Type (β=0.32, SE β=0.16, $\chi^2(1)$=3.95 p=0.047). No other effects or interactions were observed ($\chi^2(1)$=2.2

p=0.138). To investigate the nature of the two way interactions, I analyzed subsets of the full data set.

The Stripping subset revealed a main effect of Islandhood (β=0.35, SE β=0.1, $\chi^2(1)$ =11.2, p<.001), with no other effects or interactions ($\chi^2(1)$=2.57 p>0.1). The It-Cleft subset revealed main effects of Islandhood (β=0.93, SE β=0.19, $\chi^2(1)$=19.92, p<.001) and Contrast Type (β=0.4, SE β=0.13, $\chi^2(1)$=8.82 p=0.003); no interaction was found ($\chi^2(1)$=0.33 p=0.565).

Analysis of the Elaborative subset of the data revealed main effects of Islandhood (β=0.69, SE β=0.13, $\chi^2(1)$=21.8, p<.001) and Construction Type (β=-2.23, SE β=0.21, $\chi^2(1)$=54.73, p<.001), but no interaction between the two ($\chi^2(1)$=2.87 p=0.09). Further subset analyses reveal a main effect of Islandhood in both Elaborative Stripping conditions (β=0.52, SE β=0.13, $\chi^2(1)$=13.58, p<.001), and in Elaborative It-Cleft Conditions (β=0.85, SE β=0.2, $\chi^2(1)$=15.63, p<.001)

Analysis of the Non-Contrastive subset revealed both main effects of Islandhood (β=0.6, SE β=0.14, $\chi^2(1)$=14.95, p<.001) and Construction Type (β=-1.92, SE β=0.24, $\chi^2(1)$ =38.57, p<.001), as well as an interaction between the two (β=0.82, SE β=0.25, $\chi^2(1)$=10.29 p=0.001). A main effect of Islandhood was found in Non-Contrastive It-Cleft conditions (β=1, SE β=0.25, $\chi^2(1)$=13.94, p<.001), but not in Non-Contrastive Stripping conditions ($\chi^2(1)$=1.41 p=0.235).

These results indicate that both Elaborative and Non-Contrastive It-Cleft conditions, as well as Elaborative Stripping conditions, are subject to an island effect, in which the relevant Island conditions were rated worse than the comparable Non-Island conditions.

Furthermore, the island effect in the Elaborative Stripping conditions didn't differ from that in the Elaborative It-Cleft conditions. Non-Contrastive Stripping conditions, on the other hand, demonstrated no island effect. Thus, as with the preliminary analysis in Experiment 1, these results indicate complete island insensitivity for Non-Contrastive Stripping, and complete sensitivity for Elaborative Stripping both It-Cleft conditions. However, as we will soon see, factoring out the effect of Factivity will reveal that Elaborative Stripping is only partially island sensitive, what we also found in the subsequent analyses in Experiments 1 and 2.

2.5.2.2. Experiment 3 Definite Island vs. Non-Factive Results. To account for the possibility that Factivity introduced a confound into the experimental design, we next compared the Definite Island conditions against those items in the Non-Island conditions which did not contain a factive matrix verb. Condition means and standard errors can be found in Table (2.13) and Figure (2.10). The results of these analyses revealed significant main effects of Islandhood (β=0.78, SE β=0.13, $\chi^2(1)$=26.64, p<.001) and Construction Type (β=-1.84, SE β=0.22, $\chi^2(1)$=41.86, p<.001). Also significant was the two-way interaction between Islandhood and Construction Type (β=1.04, SE β=0.21, $\chi^2(1)$=23.45, p<.001). No other effects or interactions were significant ($\chi^2(1)$=2.8 p=0.094). To investigate the interaction between Islandhood and Construction Type, I analyzed subsets of the full data set.

Analysis of the Stripping subset revealed no significant effects or interactions, although that of Islandhood approached significance (β=0.25, SE β=0.13, $\chi^2(1)$=3.7 p=0.054), as did the interaction between Islandhood and Contrast Type (β=-0.49, SE β=0.26, $\chi^2(1)$=3.48

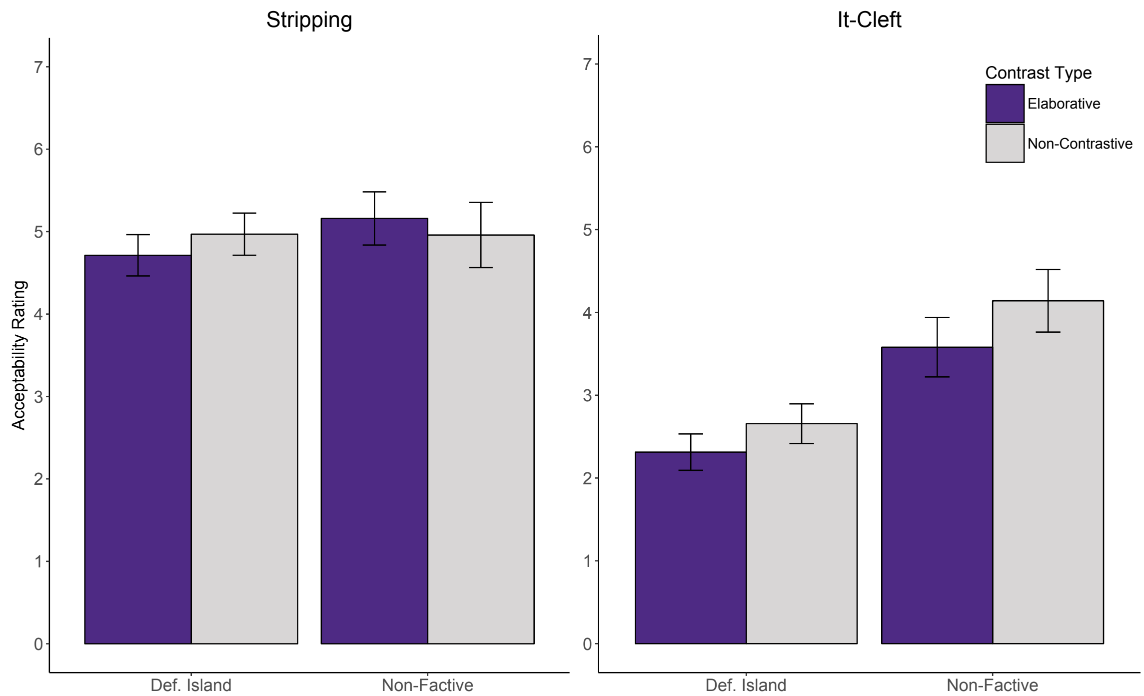

Figure 2.10. Experiment 3: Definite Island vs. Non-Factive

Table 2.13. Experiment 3: Definite Island v. Non-Factive mean judgments, with standard errors in parentheses

		Def. Island	Non-Factive
Stripping	Elaborative	4.71 (0.13)	5.16 (0.16)
	Non-Contrastive	4.97 (0.13)	4.96 (0.2)
It-Cleft	Elaborative	2.31 (0.11)	3.58 (0.18)
	Non-Contrastive	2.66 (0.12)	4.14 (0.19)

p=0.062). Analysis of the Elaborative Stripping data revealed an effect of islandhood (β=0.49, SE β=0.18, $\chi^2(1)$=6.85 p=0.009), while the Non-Contrastive Stripping data revealed no such island effect ($\chi^2(1)$=0.01 $p<0.9$).

Analysis of the It-Cleft subset revealed main effects of Islandhood (β=1.31, SE β=0.21, χ^2(1)=27.05, p<.001) and Contrast Type (β=0.35, SE β=0.14, χ^2(1)=5.69 p=0.017), but no interaction between the two (β=0.01, SE β=0.27, χ^2(1)=0 p=0.96).

Thus, as with the Definite Island vs. Non-Factive analysis in Experiment 1, Experiment 3 reveals both It-Cleft conditions to be sensitive to DCR island effects, Elaborative Stripping to be partially sensitive, and Non-Contrastive Stripping to be insensitive. Likewise, this analysis of Experiment 3 yields results consistent with the results surrounding the Elaborative conditions of Experiment 2, where Elaborative Stripping was also found to be partially island sensitive. Note these similarities hold, despite the differences between Experiments 1 and 2, which examined PP remnants and cleft-pivots, and Experiment 3, which examined DP items.

2.5.2.3. Experiment 3 Factive vs. Non-Factive Results. To determine the effect of Factivity, as with Experiments 1 and 2, an analysis comparing, of the items in the Non-Island condition, those items containing Factive verbs, to those that did not. Condition means and standard errors can be found in Table (2.14) and Figure (2.11). The results of these analyses revealed significant main effects of Islandhood (β=0.31, SE β=0.12, χ^2(1)=6.93 p=0.008), and Construction Type (β=-1.8, SE β=0.24, χ^2(1)=35.69, p<.001). Also significant were the two-way interactions between Islandhood and Construction Type (β=0.94, SE β=0.23, χ^2(1)=15.33, p<.001), and between Contrast Type and Construction Type (β=0.56, SE β=0.22, χ^2(1)=6.25 p=0.012). No other effects or interactions were observed (χ^2(1)=2.33 p>0.127).

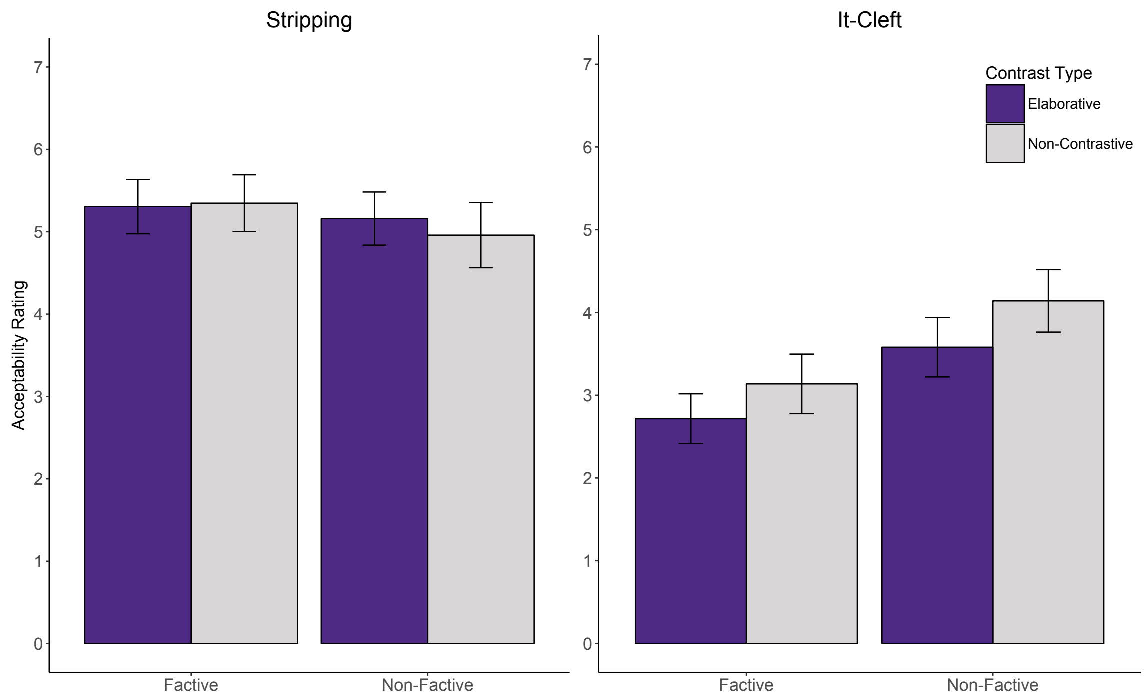

Figure 2.11. Experiment 3: Factive vs. Non-Factive

Table 2.14. Experiment 3: Factive v. Non-Factive mean judgments, with standard errors in parentheses

		Factive	Non-Factive
Stripping	Elaborative	5.31 (0.17)	5.16 (0.16)
	Non-Contrastive	5.35 (0.17)	4.96 (0.2)
It-Cleft	Elaborative	2.72 (0.15)	3.58 (0.18)
	Non-Contrastive	3.14 (0.18)	4.14 (0.19)

These interactions were investigated with subset analyses of the Stripping and It-Cleft conditions. The Stripping subset revealed no effects of or interactions between Islandhood and Contrast Type ($\chi^2(1)$=1.06 p>0.3). The It-Cleft subset revealed main effects of Islandhood (β=0.79, SE β=0.18, $\chi^2(1)$=16.68, p<.001) and Contrast Type (β=0.47, SE β=0.2, $\chi^2(1)$=5.26 p=0.022). The interaction failed to reach significance ($\chi^2(1)$=0.37 p=0.545).

Thus, the It-Cleft conditions reveal a Factivity effect, while the same effect is absent of the Stripping conditions.

2.5.2.4. Experiment 3 Discussion. Experiment 3 yielded results broadly comparable to those of Experiments 1 and 2. As in those earlier experiments, a preliminary analysis found Non-Constrastive Stripping to be completely island insensitive, but Elaborative Stripping to be island sensitive. However, as discussed above, the presence of Factive verbs in the Non-Island conditions introduced a confound into the experiment. Comparing the factive items against the non-factive items revealed a Factivity effect for the It-Cleft conditions, but not for the Stripping conditions. Comparing the Definite Island conditions to the non-factive, non-island items revealed that Non-Contrastive Stripping is island insensitive and Elaborative Stripping is only partially sensitive.

Of particular note is that the broad pattern of results in Experiment 3 did not differ from the comparable results in Experiments 1 and 2, despite the fact that the remnants and cleft pivots in Experiment 3 were DPs while they were PPs in the earlier experiments. We discuss this pattern more in the General Discussion.

2.6. Experiment 4

Experiment 4 builds on Experiments 1-3, mirroring the intent and design of those experiments. An in Experiment 3, the remnants and cleft-pivots in Experiment 4 are DPs, in place of the PPs found in Experiments 1 and 2. Here, in Experiment 4, we compared Elaborative Stripping and It-Clefts with Corrective Stripping and It-Clefts, as we did in Experiment 2.

2.6.1. Methods

2.6.1.1. Participants. 45 participants were recruited through Amazon's Mechanical Turk. Participants were limited to IP addresses within the US, were only permitted to participate in the experiment once, and were compensated $2 USD. Five participants failed to complete the survey, and so their results were excluded from analysis. A further participant was excluded on the basis of their rating of filler material. T-tests revealed this participant did not reliably rate the high acceptability fillers differently from the low acceptability fillers. Taking these exclusions into account, the data of 39 participants were analyzed.

2.6.1.2. Design, Stimuli, and Procedures. The design, stimuli and procedures of Experiment 4 were identical to those of Experiment 2 with the exception that in Experiment 4 the Stripping remnant and Cleft pivot in the target items were DPs, instead of the PP Stripping remnants and Cleft pivots used in Experiment 2. A Sample stimuli paradigm is given in (124). The same forty lexicalizations as used in Experiments 1-3 were again employed in Experiment 4, modulo the DP Stripping remnants and Cleft pivots, and the Corrective contrast type. The same factive verb confound, as described in detail in the section (2.3.1.2), of also affected this experiment. Consequently the same two additional analyses were conducted.

The same seventy filler items and six practice items used in Experiments 1-3 were utilized in this experiment, arranged in eight pseudorandomized Latin square lists.

(124) Stripping

a. Def. Island

Table 2.15. Experiment 4: Factors, Example Stimuli

Construction Type		Pronominal Status: Elaborative	Pronominal Status: Corrective
Stripping	Def. Island	Example (124a-i)	Example (124a-ii)
	Non-Island	Example (124b-i)	Example (124b-ii)
It-Cleft	Def. Island	Example (125a-i)	Example (125a-ii)
	Non-Island	Example (125b-i)	Example (125b-ii)

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and also *Ashley*.

ii. Corrective:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: No, *Ashley*.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and also *Ashley*.

ii. Corrective:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: No, *Ashley*.

(125) It-Cleft

a. Def. Island

i. Elaborative:

Joe: Jordan mocked the salesman who sold a car *to Katherine*.

Bill: Yeah, and it was also *Ashley* that Jordan mocked the salesman who sold a car to.

ii. Corrective:

Joe: Jordan mocked the salesman who sold a car *to an elderly customer*.

Bill: No, it was *Ashley* that Jordan mocked the salesman who sold a car to.

b. Non-Island

i. Elaborative:

Joe: Jordan implied that the salesman sold a car *to Katherine*.

Bill: Yeah, and it was also *Ashley* that Jordan implied that the salesman sold a car to.

ii. Corrective:

Joe: Jordan implied that the salesman sold a car *to an elderly customer*.

Bill: No, it was *Ashley* that Jordan implied that the salesman sold a car to.

2.6.2. Experiment 4 Results

2.6.2.1. Preliminary Analysis. The design of Experiment 4, being based on those of Experiments 1-3, also suffered the factivity confound of those prior experiments. Consequently, the interpretation of the results of this preliminary analysis should be mindful of this confound.

Condition means and standard errors can be found in Table (2.16) and Figure (2.12). The data[12] was then analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. (2008)) with rating as the dependent variable, using the lme4 package in R. Contrast-coded fixed effects included Islandhood (Island, Non-Island), Contrast type (Elaborative, Non-Contrastive), and Construction Type (Stripping, It-Cleft), as well as their 2- and 3-way interactions. The maximal random effects structure that would converge was employed, which included random slopes and intercepts for Islandhood, Contrast type, Construction Type, and the 2- and 3- way interactions by Participant and Item. Model comparisons were performed to determine whether the inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

Table 2.16. Experiment 4: Definite Island v. Non-Island mean judgments, with standard errors in parentheses

		Def. Island	Non-Island
Stripping	Elaborative	4.26 (0.12)	4.7 (0.12)
	Corrective	4.35 (0.12)	4.63 (0.12)
It-Cleft	Elaborative	2.03 (0.09)	3.11 (0.12)
	Corrective	2.22 (0.1)	3.56 (0.13)

The results of these analyses revealed significant main effects of Islandhood (β=0.79, SE β=0.1, $\chi^2(1)$=40.44, p<.001), and Construction Type (β=-1.76, SE β=0.25, $\chi^2(1)$=32.88, p<.001). The main effect of Contrast Type approached, but did not reach, significance (β=0.17, SE β=0.09, $\chi^2(1)$=3.33 p=0.068). Also significant were the two way interactions between Islandhood and Construction Type (β=0.86, SE β=0.14, $\chi^2(1)$=34.11, p<.001) and between Contrast Type and Construction Type (β=0.32, SE β=0.14, $\chi^2(1)$=5.07 p=0.024).

[12]As in Experiment 1 and elsewhere, I used the raw rating data in these analyses, instead of using a z-score transformation, as the LMER models employed were constructed to account for by participant variance.

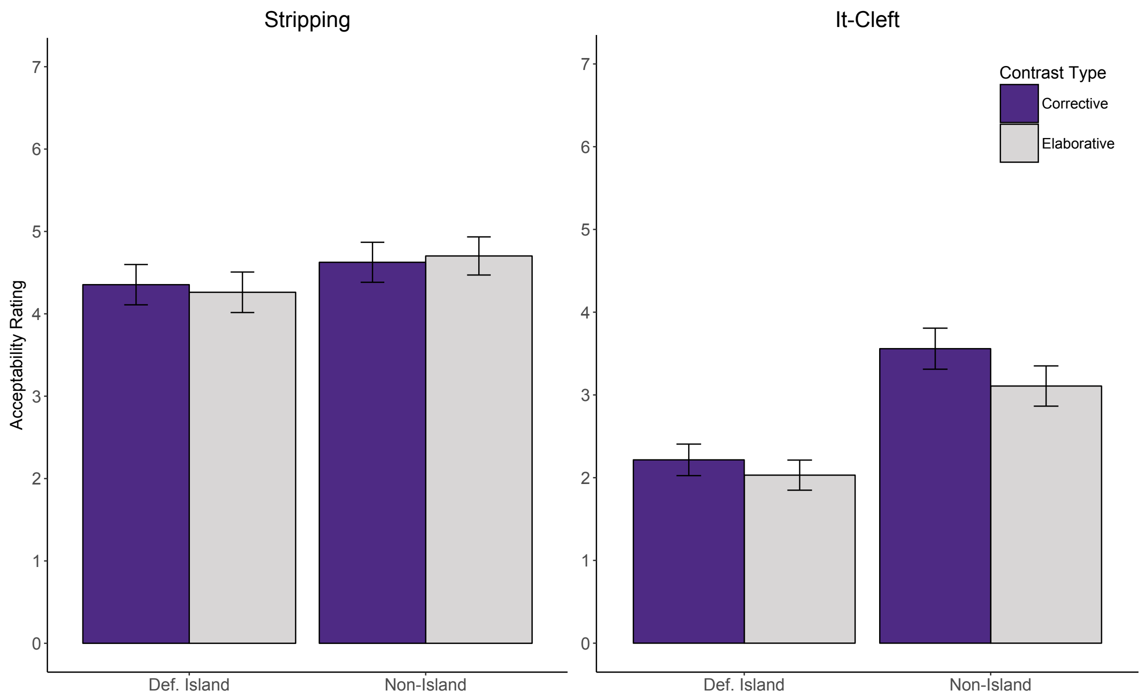

Figure 2.12. Experiment 4: Definite Island vs. Non-Island

No other effects or interactions were observed ($\chi^2(1)$=2.76 p=0.096). These two way interactions were investigated through analysis of subsets of the full data set.

An analysis of the Stripping subset revealed a main effect of Islandhood (β=0.36, SE β=0.09, $\chi^2(1)$=16.24, p<.001), with no other effects or interactions ($\chi^2(1)$=1.16 p=0.282). The It-Cleft subset revealed main effects of Islandhood (β=1.21, SE β=0.17, $\chi^2(1)$=33.7, p<.001) and Contrast Type (β=0.32, SE β=0.12, $\chi^2(1)$=6.78 p=0.009); no interaction was found ($\chi^2(1)$=1.79 p=0.181).

Analysis of the Elaborative subset of the data revealed main effects of Islandhood (β=0.76, SE β=0.14, $\chi^2(1)$=24.26, p<.001) and Construction Type (β=-1.91, SE β=0.24, $\chi^2(1)$=39.21, p<.001), as well as an interaction between the two (β=0.62, SE β=0.2,

$\chi^2(1)$=9.13 p=0.003). Further subset analyses reveal a main effect of Islandhood in Elaborative Stripping conditions (β=0.45, SE β=0.15, $\chi^2(1)$=8.39 p=0.004), and in Elaborative It-Cleft Conditions (β=1.08, SE β=0.22, $\chi^2(1)$=19.9, p$<$.001)

Analysis of the Corrective subset revealed both main effects of Islandhood (β=0.81, SE β=0.11, $\chi^2(1)$=38.58, p$<$.001) and Construction Type (β=-1.6, SE β=0.28, $\chi^2(1)$=23.39, p$<$.001), as well as an interaction between the two (β=1.08, SE β=0.19, $\chi^2(1)$=30.65, p$<$.001). A main effect of Islandhood was found in Corrective It-Cleft conditions (β=1.36, SE β=0.21, $\chi^2(1)$=30.91, p$<$.001), as well as in Corrective Stripping conditions (β=0.27, SE β=0.1, $\chi^2(1)$=7.5 p=0.006).

These results reveal an island effect in both It-Cleft and Stripping conditions, but that both Elaborative and Corrective Stripping conditions are only partially sensitive to this island effect. Note that here, unlike in Experiments 1-3, the preliminary analysis found Elaborative Stripping to be only partially sensitive, and, unlike in Experiment 2, found Corrective Stripping to be only partially sensitive to islands. Nevertheless, as the next two analyses will show, the factivity confound was still present in this experiment.

2.6.2.2. Experiment 4 Definite Island vs. Non-Factive Results. We next compared the Definite Island conditions against those items in the Non-Island conditions which did not contain a factive matrix verb. Condition means and standard errors can be found in Table (2.17) and Figure (2.13). The results of these analyses revealed significant main effects of Islandhood (β=0.93, SE β=0.12, $\chi^2(1)$=38.83, p$<$.001) and Construction Type (β=-1.6, SE β=0.26, $\chi^2(1)$=27.24, p$<$.001). Also significant was the two-way interaction between Islandhood and Construction Type (β=1.17, SE β=0.18, $\chi^2(1)$=34.67, p$<$.001). No other

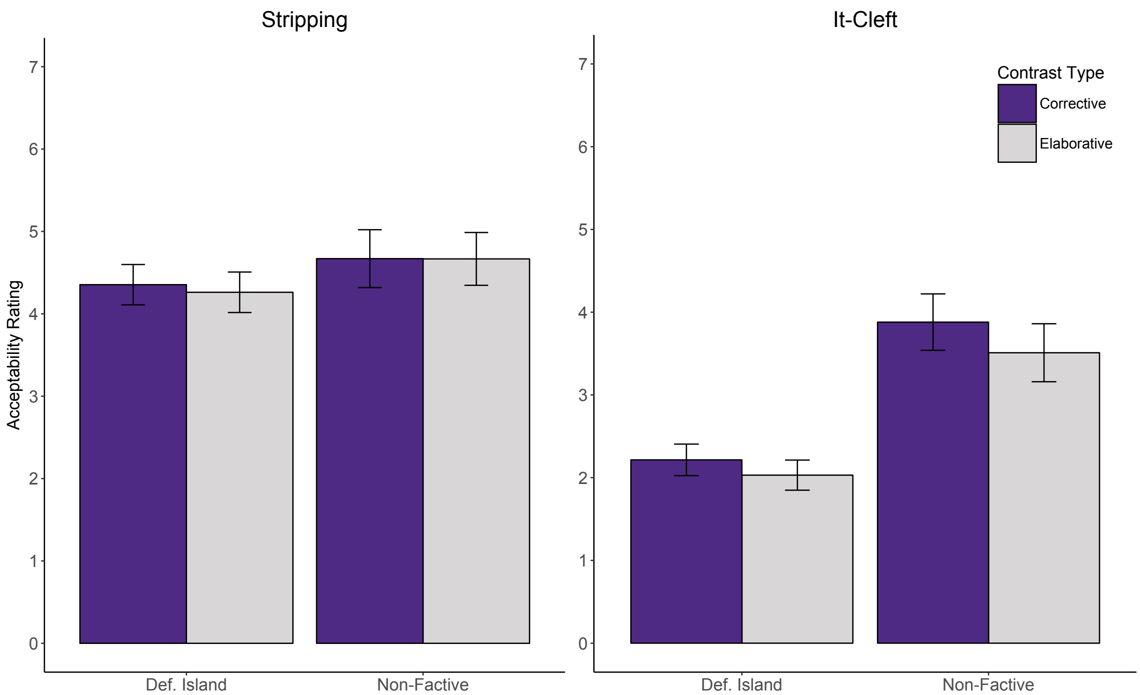

Figure 2.13. Experiment 4: Definite Island vs. Non-Factive

effects or interactions were significant ($\chi^2(1)$=2.6 p=0.107). This interaction between Islandhood and Construction Type was investigated in a comparison of subsets of the full data set.

Table 2.17. Experiment 4: Definite Island v. Non-Factive mean judgments, with standard errors in parentheses

		Def. Island	Non-Factive
Stripping	Elaborative	4.26 (0.12)	4.67 (0.16)
	Corrective	4.35 (0.12)	4.67 (0.18)
It-Cleft	Elaborative	2.03 (0.09)	3.51 (0.18)
	Corrective	2.22 (0.1)	3.88 (0.17)

Analysis of the Stripping subset revealed a significant effect of Islandhood (β=0.34, SE β=0.12, χ^2(1)=7.72 p=0.005); both the effects of Contrast Type and the interaction of the two did not reach significance (χ^2(1)=0.41 p=0.521). Analysis of the It-Cleft subset of the data revealed main effects of Islandhood (β=1.57, SE β=0.22, χ^2(1)=35.94, p<.001) and Contrast Type (β=0.33, SE β=0.14, χ^2(1)=5.21 p=0.022), but no interaction between the two (χ^2(1)=1.18 p=0.277).

Post-hoc pairwise comparisons reveal interactions between Islandhood and Construction type in both Elaborative (β=1.03, SE β=0.24, χ^2(1)=16.66, p<.001) and Corrective (β=1.3, SE β=0.21, χ^2(1)=33.49, p<.001) data subsets. Additional comparisons reaffirm that the Islandhood effect found in the Stripping subset is present in both Elaborative Stripping (β=0.41, SE β=0.2, χ^2(1)=4.12 p=0.042) and Corrective Stripping condition subsets (β=0.41, SE β=0.2, χ^2(1)=4.12 p=0.042).

Thus both Corrective and Elaborative It-Clefts demonstrate an island effect, a reduction in the acceptability of the island conditions relative to the non-island conditions. In contrast, both Elaborative and Corrective Stripping display a partial island sensitivity.

2.6.2.3. Experiment 4 Factive vs. Non-Factive Results. Finally, to determine the effect of Factivity, as with Experiments 1-3, an analysis comparing those Non-Island items containing Factive verbs, to those which did not. Condition means and standard errors can be found in Table (2.18) and Figure (2.14). The results of these analyses revealed significant main effects of Islandhood (β=0.33, SE β=0.11, χ^2(1)=7.74 p=0.005) and Construction Type (β=-1.34, SE β=0.25, χ^2(1)=21.33, p<.001). Also significant were the two-way interactions between Islandhood and Construction Type (β=0.73, SE β=0.22,

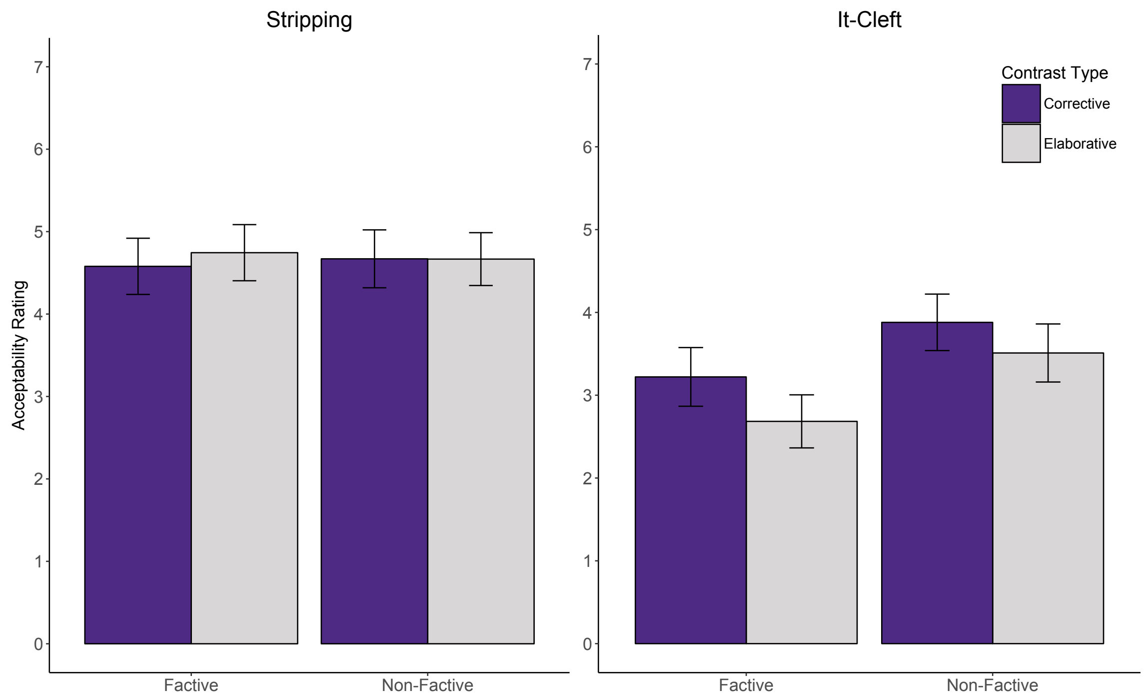

Figure 2.14. Experiment 4: Factive vs. Non-Factive

$\chi^2(1)$=10.67 p=0.001) and between Contrast Type and Construction Type (β=0.56, SE β=0.21, $\chi^2(1)$=6.92 p=0.009). No other effects or interactions were observed ($\chi^2(1)$=2.56 p=0.109). These interactions were investigated with subset analyses.

Table 2.18. Experiment 4: Factive v. Non-Factive mean judgments, with standard errors in parentheses

		Non-Factive	Factive
Stripping	Elaborative	4.67 (0.16)	4.74 (0.17)
	Corrective	4.67 (0.18)	4.58 (0.17)
It-Cleft	Elaborative	3.51 (0.18)	2.68 (0.16)
	Corrective	3.88 (0.17)	3.22 (0.18)

The Stripping subset revealed no main effects or interaction ($\chi^2(1)$=0.38 p=0.538). The It-Cleft subset revealed main effects of Islandhood (β=0.69, SE β=0.18, $\chi^2(1)$=12.5, p<.001) and Contrast Type (β=0.47, SE β=0.2, $\chi^2(1)$=5.08 p=0.024). The interaction failed to reach significance ($\chi^2(1)$=0 p=0.979). These results indicate that an effect of Factivity is present in the It-Cleft conditions, and that this effect is completely absent of the Stripping conditions.

2.7. Experiments 1-4 General Discussion

This set of Experiments 1-4 have attempted to establish a clear empirical base upon which to found studies of the island sensitivity of ellipsis. These studies have compared the effect on acceptability of overt movement from within an island and within a complement clause to the effect on acceptability of island-violating Stripping, operationalized as instances of Stripping in which the correlate for a Stripping phrase is contained within a definite relative clause island.

2.7.1. Preliminaries

Before addressing the island sensitivity of Stripping, some preliminary summary comments on the results are appropriate. First, these experiments have demonstrated that it-clefts were rated significantly worse than the Stripping phrases. This remains true even of a comparison of island-violating Stripping configurations to non-island It-Cleft configurations.

This could be the case simply because the Stripping conditions were much shorter than the It-Cleft conditions, and additional length has been shown to reduce acceptability

(Sprouse et al., 2012). However, length is unlikely to be the sole cause of the lower rating of the It-Clefts, as many of the fillers were multi-clausal yet were rated quite highly, as illustrated in (126) and (127)[13].

(126) Complex Clause filler mean rating: 6.44

a. Joe: I am really nervous about applying for jobs.

b. Bill: Well, I can't imagine that your supervisor wouldn't give you a positive recommendation letter.

(127) Filler Wh-Movement filler mean rating: 6.60

a. Joe: Last I heard, the CIA and FBI both were investigating the cyber attacks on the voting machines.

b. Bill: Well, I doubt we'll ever find out who was responsible.

One might alternatively wonder whether it could be the case that, in contexts where reduction is possible, the failure to reduce yields a reduction in acceptability (Bresnan, 1973; Wellwood et al., 2017). In many of the items considered in these experiments, where a non-reduced cleft was judged for acceptability, a reduced cleft was also possible. This is true of non-contrastive It-Clefts and the reduced cleft continuation in (128) certainly seems to be more acceptable than the non-reduced cleft continuation.

(128) a. James heard that the student gave a present to a teacher.

b. Yeah, it was Mrs. Wilcox that James heard that the student gave a present to.

c. Yeah, it was Mrs. Wilcox.

[13] Here and below, in the discussion of filler means, the means represent the mean of all ratings of fillers of these types across all four experiments, which utilized the same fillers.

However, it has been claimed that contrastive *reduced* It-Clefts are unacceptable when the gap corresponding to the cleft pivot is within an embedded clause or within an island (Barros, 2012a). Thus, Barros (2012a) claims, the reduced cleft continuation in (129a) should be less acceptable than the non-reduced continuation.

(129) James heard that the student gave a present to *Mr. Rock*.

a. i. No, it was Mrs. Wilcox that James heard that the student gave a present to.

   ii. No, it was Mrs. Wilcox.

If these judgments hold, and I will note that it isn't immediately obvious to me that they do, then we would expect for contrastive it-clefts, comprising elaborative and corrective it-clefts, to be more acceptable than the non-contrastive It-Clefts. This is because it is only the non-contrastive It-Clefts for which a reduced form is acceptable. Thus it is only the non-contrastive It-Clefts for which the availability of a reduced form is predicted to reduce the acceptability of the non-reduced form.

However, we found in Experiments 1 and 3 that non-contrastive It-Clefts were rated *more* acceptable than Elaborative It-Clefts. Thus, the prediction does not seem to be born out. I will temper this conclusion in noting that there is reason to believe that Elaborative It-Clefts are reduced in acceptability compared to other types of It-Clefts, due to an exhaustivity constraint on It-Clefts, as discussed below. Thus the reason why elaborative It-Clefts were rated lower than non-contrastive clefts could be due to this exhaustivity effect. However, this leaves the relatively low ratings of the corrective It-Clefts a mystery.

If the sole reason why the non-contrastive It-Cleft conditions were rated to be less acceptable than the Stripping conditions were the availability of a reduced cleft alternative, we would then expect the corrective It-Clefts, for which it was claimed that the reduced form is unacceptable, to be as acceptable as the Stripping conditions, contrary to our findings.

Of course, if it turned out that the cleft continuations in (129a) did not differ in acceptability, or that the reduced form is more acceptable than the non-reduced form, then such a reduction-alternative explanation does make the correct prediction. Both non-contrastive and corrective It-Clefts would have an alternative reduced form, and so, when the reduced form is possible, as in our experimental items, the non-reduced form is of marked acceptability.

A third alternative is that the It-Cleft conditions were rated worse than the Stripping conditions because it was just in these conditions that any repetition at all between the antecedent and test phrases was found in the experiment. This is true not only of the test items, but also of the filler items. I leave it to future research to determine whether the lower acceptability of It-Clefts as compared to Stripping cases is due to item length or to repetition of material given in the linguistic context.

It's important to note the problems for the current experimental design that a reduced acceptability of just the It-Cleft conditions would pose. The ratings of the island It-Cleft conditions was near the bottom of the scale used by participants. Table (2.19) illustrates the ratings for both (definite relative clause) island and non-island non-contrastive It-Clefts from the Definite Island vs. Non-Factive analysis of Experiment 3, found in Section (2.5.2.2), Table (2.13). These represent the highest ratings given to It-Clefts across the four experiments. The table also illustrates the mean ratings of the various fillers, aggregated across Experiments 1-4.

Table 2.19. Various island-violating Filler Means vs. Non-Contrastive It-Cleft Means

	Mean Rating
Filler: Passive	6.70
Filler: *Wh*-Movement	6.60
Filler: Complex Clause	6.44
Filler: Pied-Piping	5.46
Test item: Non-Island Non-Contrastive It-Cleft	4.14
Filler: Sentential Subj Missing *that*	3.54
Filler: PseudoPassive	3.38
Filler: CNPC Island Violation	2.66
Test item: Island Non-Contrastive It-Cleft	2.66
Filler: *Wh*-island Violation	2.62
Filler: Case Violation	2.60
Filler: Definite Relative Clause Island Violation	2.45
Filler: Adjunct Island Violation	2.39
Filler: Subject Island Violation	2.25

Observe that many of the filler items were rated near the very high end of the seven point scale, including many items involving long distance *wh*-movement. However, the non-island non-contrastive It-Clefts were rated much lower than most of these fillers that had intended to be acceptable, and just somewhat more acceptable than instances of sentential subjects missing the complementizer *that*. Then note that the island-violating It-Clefts were rated similarly to other types of island-violating fillers, which all clustered between ratings of 2.66 and 2.25.

This results in the magnitude of difference in means for the It-Cleft items in this case being merely 1.52 points. In contrast, the difference between the mean ratings of examples of *wh*-movement and, e.g., fillers involving complex noun phrase island violations is much larger, at 3.94 points. Given that the ratings of the island-violating It-Cleft examples is more or less in line with the island-violating fillers, we can point to the badness of the non-island-violating It-Clefts for the relatively small difference in mean ratings between the island and non-island It-Clefts.

The magnitude of difference between the baseline island and non-island conditions is, of course, critical for the evaluation of the current experiments. I concluded that island-violating Stripping could be considered to be sensitive to islands if the magnitude of difference between the ratings of the island and the ratings of non-island Stripping conditions were as large as that same difference in the It-Cleft conditions. That is, island-violating Stripping could be considered to be island sensitive if the penalty for an island violation was the same for Stripping as it was for It-Clefts. Similarly, island-violating Stripping could be considered to be insensitive to islands if there were no penalty for an island violation. Or, if there were a penalty for island violations in Stripping but it was less than the penalty observed for It-Clefts, we could consider Stripping to be partially island sensitive.

Now we can see the potential problem that the low ratings for the non-island It-Clefts pose. The unexpectedly low non-island It-Cleft condition ratings resulted in a low delta between the ratings of island and non-island It-Clefts. It was conceivable that the difference between island and non-island Stripping cases could have matched the difference displayed in the It-Cleft cases. We would have then concluded that Stripping was island sensitive. This, despite the relatively small delta for the It-Cleft cases compared with the much larger delta between the island and non-island filler controls.

As it turns out, as we saw above and will discuss in more detail below, each of Experiments 1-4 revealed the various types of Stripping to be either island insensitive or only partially sensitive. Thus, despite the relatively small differences between the island and non-island It-Cleft conditions, the differences between the island and non-island Stripping conditions were either smaller or simply non-existent. The message to take home, however, is that the baseline matters.

A second overall trend that warrants discussion was the reduction in acceptability of elaborative It-Clefts compared with corrective and non-contrastive It-Clefts. In all of the analyses, preliminary, DRCI vs. Non-Factive, and Factive vs. Non-Factive, for all of the experiments except Experiment 1, I found that the elaborative It-Clefts were reliably rated worse than the corrective and non-contrastive It-Clefts. Even in the various analyses of Experiment 1, the elaborative It-Cleft conditions were rated numerically worse than the non-corrective It-Cleft conditions.

It would be tempting to attribute the rating for elaborative It-Clefts as less acceptable than non-contrastive or corrective It-Clefts to an exhaustivity effect. It has been claimed that English It-Clefts introduce an exhaustivity requirement of the cleft pivot, such that the cleft pivot is the only relevant element in the discourse to satisfy the presupposition of the cleft clause Kiss (1998). So, for example, in (130b), the claim is that the cleft structure imposes a requirement that there be one person who James thinks that the lawyer called and that this person be Mary.

(130) a. James thinks that the lawyer called someone.

b. Yeah, it was *Sue* who James thinks that the lawyer called.

If this were true, then we could attribute the lower acceptability of the elaborative It-Cleft conditions, as compared to corrective and non-contrastive It-Cleft conditions, to a violation of this exhaustivity condition. This would explain why examples like (131b) were rated worse than examples like (130b) and (131c). In (131b) it is entailed that there are at least two people who James thinks that the the lawyer called: Mary and Bill. This violates the exhaustivity requirement that is satisfied of the cleft continuations in (130b) and (131c), where it is just Sue, or just Mary and Bill who James thinks the lawyer called.

(131) a. James thinks that the lawyer called *Mary*.

b. Yeah, and it was also *Bill* who James thinks that the lawyer called.

c. No, it was *Bill* who James thinks that the lawyer called.

However, it should be noted that the differences between the elaborative It-Cleft cases and the corrective or non-contrastive It-Cleft conditions are, across Experiments 1-4, numerically and statistically reliably, smaller than the differences between island and non-island conditions. This could suggest that whatever mechanism which induces the reduced acceptability for the elaborative It-Cleft cases is distinct from what induces the reduced acceptability for the island-violating It-Cleft cases. If we take the magnitude of reduction in acceptability in cases of overt island-violating movement to be characteristic of a *grammaticality* violation, then this suggests that the effect of an exhaustivity violation in the It-Cleft cases is distinct from such a grammaticality violation.

For example, it has been claimed that, for example, the focus sensitive operator *only*, exhibits exhaustivity effects due to its semantics (e.g. Beaver and Clark (2003)); exhaustivity is built into the meaning of *only*. An example, such as (132), in which this exhaustivity is violated yields quite a strong unacceptability, perhaps as a result of its contradictoriness. On the other hand, while examples like (131b) are degraded relative to It-Cleft examples which don't violate exhaustivity, they also don't seem to produce a judgment of contradictoriness either. Similar facts have also been observed by Horn (1981). Thus, although this aspect of these results is orthogonal to the main line of inquiry here, they do contribute to the discussion on the varying nature of exhaustivity, suggesting that, unlike the exhaustivity of *only*, that of It-Clefts is defeasible and so likely to originate from pragmatic principles (Horn, 1981).

(132) #James called only Mary, and he called Sue too.

2.7.2. On island (in)sensitivity

Now, on to the island sensitivity of It-Clefts and Stripping. I will first recapitulate the findings of interest from Experiments 1-4 and then discuss possible explanations for the differences in island sensitivity between non-contrastive and contrastive Stripping.

First, we found that It-Clefts of all types, elaborative, corrective, and non-contrastive, show sensitivity to definite relative clause islands. That is to say, that instances of the non-reduced It-Clefts tested here are less acceptable when the gap is found within a definite relative clause than when it is found within a complement clause. This is true both of the preliminary analyses, for which the Non-Island conditions included items that contained so-called factive islands, as well as the subsequent analyses comparing DCR island conditions to non-factive non-island conditions. These results are in line with those reported in Hunter et al., who found that corrective It-Clefts demonstrate patterns of DCR island sensitivity similar to those found by Sprouse et al. (2012) for *wh*-movement for a variety of island types, as noted in the experimental overview.

We also found It-Clefts of all types tested to be sensitive to factivity[14]. That is, It-Cleft items in which the gap was contained within the clausal complement of a factive verb were rated worse than similar It-Cleft items in which the gap was contained within the clausal complement of a non-factive verb. To my knowledge, this is the first formal examination of the effects of factivity on acceptability and the sensitivity of ellipsis to factivity. It should be noted that, this formal examination was, in fact, unintentional on the part of the researcher. The bad fortune of including factive verbs in the Non-Island conditions, of all four experiments no less, was balanced out, in part, by the blind luck of having included exactly 20 factive verbs and 20 non-factive verbs in the Non-Island conditions, thereby permitting easy analysis comparing these two classes of verbs.

Further research into the effects of factivity on acceptability and the sensitivity of ellipsis to factivity should address two aspects of the present unintentional design. First, it has been known since Karttunen (1971) that the category of factive verbs should actually be divided into several types, each with distinct properties. I have ignored these within-category differences here, but investigation into factivity, islandhood, and ellipsis island insensitivity would do well to explore these differences.

Second, as noted in the discussion of Experiment 1, the design of the present experiment limits our ability to determine the source of the reduction in acceptability between

[14]The one exception comes from Experiment 2, where we didn't find an effect of factivity for elaborative It-Clefts. Given that we found such an effect for elaborative It-Clefts in each other experiment, I will attribute this result to statistical power, although I will not attempt here a formal characterization of this conclusion.

Factive and Non-Factive It-Cleft conditions. For comparison, the key characteristic of the acceptability study of Sprouse et al. (2012) was the factorial design, crossing dependency length with islandhood. In such a way, the authors were able to determine that the lower acceptability of a long distance extraction from within an island domain, as compared with long distance extraction from within a non-island domain, was greater than the sum of the detrimental effects of the presence of a long distance dependency and the presence of an island. An island violation is characterized as an interaction between these factors.

The present design simply compares It-Clefts with gaps in the complement of a factive with It-Clefts with gaps in the complement of a non-factive verb. Thus it is not clear whether the effect of factivity observed here is actually any greater than the mere sum of dependency length and presence of a factive verb. I leave it to future research to address these two concerns.

Let us now turn to Stripping and the contexts where it is acceptable. In Experiments 1-4, having factored out the effects of factivity, we found that Stripping, of elaborative, corrective, and non-contrastive types, shows either partial or complete insensitivity to definite relative clause islands. In Experiments 1 and 3, we found that non-contrastive Stripping showed no differences between island and non-island conditions. In Experiment 3, we found that elaborative Stripping showed a reduction in island conditions as compared with non-island conditions, but that these differences were smaller than those found for the comparable elaborative It-Cleft conditions. Experiment 1, on the other hand, revealed full sensitivity for the elaborative Stripping conditions. In Experiments 2 and 4, we found that elaborative Stripping was partially island sensitive, replicating the results of Experiment

3, and that corrective Stripping was as well. Thus, we have what we take to be strong evidence that non-contrastive Stripping is completely insensitive to DRC islands, and that elaborative and corrective Stripping are partially sensitive. Though there is a reduction for the island conditions, to call the island-violating contrastive Stripping instances *unacceptable* seems a stretch here. Such examples were always rated higher than even the non-island-violating it-clefts. In Experiment 4, for example, Corrective Def. Island Stripping was rated on average 4.35, while the Corrective Non-Factive It-Clefts were rated on average 3.88. As illustrated in (2.19), fillers which contained island violations ranged in average acceptability ratings between 2.25 and 2.66.

However, I would speculate that this split between the partial island insensitivity exhibited by the elaborative and corrective types of Stripping, and the fully insensitivity exhibited by the non-contrastive type may provide an explanation for the varying claims in the literature about the island (in)sensitivity of contrastive and non-contrastive ellipsis. Notice that if we had merely compared island-violating contrastive and non-contrastive Stripping with instances of overt island-violating movement, we would likely observe that the Stripping examples were more acceptable, and so probably conclude that Stripping is insensitive to islands.

On the other hand, were we to compare instances of island-violating Stripping with non-island Stripping, we may be able to detect a difference in behavior between the contrastive cases and the non-contrastive cases. Non-contrastive Stripping shows no island effects, while the two contrastive types of Stripping showed some effect. From this we might be conclude that contrastive Stripping is sensitive to islands. But, as discussed earlier, individually, each of these baselines, the overt island-violating movement baseline and the non-island Stripping baseline, is insufficient. Indeed, the results of these experiments bears out that insufficiency. The partial sensitivity of Contrastive Stripping is empirically distinct from the full sensitivity of overt island-violating movement, which could not be detected without the factorial design employed here.

These experiments also provide strong evidence that Stripping is immune to the factivity effects found for It-Clefts. In each of the Factive vs. Non-Factive analyses for Experiments 1-4, no differences were found between the factive and non-factive conditions. In this, the Stripping conditions differ from the It-Clefts, for which a penalty of factivity was found. Whether this effect of factivity constitutes an *island* effect or not, Stripping appears to be completely insensitive to it.

Finally, these experiments found, by and large, no differences in the pattern of results between Experiments 1 and 2, which tested PP remnants and It-Cleft pivots, and Experiments 3 and 4, which tested DPs. This is to say that the sensitivity of It-Clefts to definite relative clause islands and factivity did not differ as a function of whether the pivot was a PP or a DP. Likewise, the partial or complete sensitivity of Stripping to definite relative clause islands and the complete insensitivity to factivity held for both PP and DP remnants.

In summary, It-Clefts of all types are sensitive to definite relative clause islands and factivity, non-contrastive Stripping is insensitive to definite relative clause islands, corrective and elaborative Stripping are partially sensitive definite relative clause islands, and all three types of Stripping are insensitive to factivity. In Chapter 3, I turn from the acceptability of Stripping to an investigation into what structures surround these ellipsis fragments, arguing that there must be structure within the ellipsis site that is isomorphic to the antecedent.

2.7.3. On *partial* island sensitivity

Before moving on, however, let us take a closer look at the *partial* sensitivity to islands that corrective and elaborative Stripping, instances of contrastive Stripping, exhibited in these experiments. I will first look at, and reject, several possible explanations as to why the pattern found in the contrastive cases differs from that of the non-contrastive cases. I would like to ultimately suggest that the partial sensitivity is the result of an ambiguity present just in the antecedents of island-violating contrastive Stripping continuations. This ambiguity is plausibly resolved in a manner inconsistent with the subsequent Stripping continuation, which results in a reduced acceptability for that Stripping continuation.

To start, observe that the factorial design of these studies, along with their parallel designs, allows us to rule out a number of explanations as to why the contrastive cases exhibited partial island sensitivity, rather that full sensitivity or insensitivity. The following points are, with the advantages of the factorial design in mind, obvious, so I will be brief.

One might argue that, in the island conditions, the shortness of the contrastive Stripping items, as compared to the contrastive It-Cleft items, would improve the relative acceptability of the Stripping items. However, this would predict that the non-contrastive Stripping items, which are just as short as the contrastive items, should similarly manifest partial insensitivity, instead of their complete island insensitivity. Or, one could argue that the reduction in the contrastive island-violating Stripping items, as compared with their non-island counterparts, is due to the presence of an island in the antecedent. This isn't likely to be the case because, we would then, incorrectly, predict that the same should be true of the non-contrastive Stripping cases.

Now I would like to explore the idea that this partial insensitivity is due to a violation of one aspect of *Parallelism*, the relevant portion of which I will take to be defined as in (133), following Fiengo and May 1994 (Indices and identity), as cited by Fox and Lasnik 2003, pg.149[15] I will here consider two of ways to distinguishing contrastive and non-contrastive Stripping using Parallelism to derive the partial island sensitivity of the former, but neither of them are tenable. The basic idea for both methods is to induce a Parallelism violation just in the derivation of island-violating contrastive Stripping, and not in island-violating non-contrastive Stripping, nor in non-island contrastive and non-contrastive Stripping.

(133) Variables in the antecedent and the elided clause are bound from parallel positions.

The first way would be to appeal to Non-Isomorphic analyses of ellipsis. The idea would be to assume that Non-Isomorphic construals induce a Parallelism violation, and that just the island-violating contrastive Stripping configurations involve Non-Isomorphic construals. The partial island sensitivity of contrastive Stripping would then follow.

[15] I set aside here the aspect of Parallelism which enforces a version of syntactic identity, as Non-Isomorphic analyses must, obviously, entertain other notions of the identity conditions on ellipsis, and because it is irrelevant to the second of the two options I entertain here.

I will discuss such Non-Isomorphic analyses in more detail in the subsequent chapter, but for the present purposes, a brief description will suffice. Non-Isomorphic analyses of ellipsis are structural analyses, in that they posit that there is silent syntactic material within the ellipsis site, but, crucially, they also posit that this silent syntactic material can substantially differ from the syntactic material of the antecedent such that the ellipsis site and antecedent are non-isomorphic (Abels, 2011; Barros, 2014; Barros et al., 2013; van Craenenbroeck, 2010; Merchant, 2001; Rodrigues et al., 2009). These accounts often appeal to non-isomorphic construals of the ellipsis site to account for claims that ellipsis is, in certain configurations, insensitive to islands. For example, take the island-violating contrastive Stripping continuation in (134b). A non-isomorphic analysis might posit that the ellipsis site is construed as just the material from within the relative clause, as in (134c), or as a reduced It-Cleft, as in (134d). In either of these cases, the remnant has moved, to evacuate the ellipsis site, but, due to the non-isomorphic construal of the ellipsis site, the movement of the remnant has not had to cross island boundary. So the mantra goes: "No island, no island violation".

(134) a. Bill met the student who speaks Thai$_F$.

b. Joe: No, Mandarin$_F$.

c. . . . Mandarin ~~the student speaks t~~.

d. . . . Mandarin ~~it was t~~.

e. . . . Mandarin ~~Bill met the student who speaks t~~.

Let's assume that the remnant in the two non-isomorphic construals have LF structures roughly corresponding to those in (135). Now, Parallelism will be violated under either of the non-isomorphic construals illustrated in (134) if we assume that the correlate in the antecedent scopes either through long distance, island-violating movement, as in (136b), or the correlate could participate in roll-up movement, following Charlow (2017), as in (136c)[16]. However in no combination of the ellipsis phrase LFs in (135) and antecedent phrase LFs in (136) is Parallelism satisfied. For the ellipsis LFs in (135), there is one, short local relationship, while in (136c), there are two and in (136b), one long distance relationship.

(135) a. Mandarin λx the student speaks x.

b. Mandarin λx it was x.

(136) a. Bill met the student who speaks Thai$_F$.

b. Thai λx [Bill met the student who speaks x]

c. [Thai λx [the student who speaks y]] λx [Bill met x]

So, let us attribute a small decrease in acceptability to a violation of Parallelism. Then, assume that only island-violating contrastive Stripping configurations involve non-isomorphic construals. Other configurations involve ellipsis site construals isomorphic to the antecedent. This implies that the island insensitivity of non-contrastive Stripping involves some other mechanism by which these constructions become insensitive to island

[16]In a third potential antecedent LF, the island containing the correlate scopes at the edge of the matrix clause, and the correlate scopes in-situ, within the island, following Krifka (2006), as in (i). Notice here that it is harder here to determine whether Parallelism is satisfied. There is a short dependency, as in the ellipsis LF, but this dependency doesn't directly concern the correlate itself. In any case, if it were to satisfy Parallelism, we come no closer to attributing partial sensitivity to a Parallelism violation, and if it violated Parallelism, the subsequent argument against this general approach holds here too.

(i) [the student who speaks Thai] λx [Bill met x] violations, such as Repair (Ross 1969, Merchant 2001). Thus we can derive the partial insensitivity effect; non-isomorphic construals violate Parallelism, and non-isomorphic construals only occur in the context of contrastive island Stripping. The problem here is that it feels like we've assumed too much. By assumption, to account for island-violating ellipsis, the grammar would have to provide *both* a non-isomorphic construal mechanism, as well as a repair mechanism. But, if repair were possible for the non-contrastive cases, there doesn't seem to be any reason why repair shouldn't be possible for the contrastive cases as well.

Under an island repair analysis, the derivation of island-violating non-contrastive Stripping would involve the remnant participating in island-violating movement, as illustrated in (137). If such movement is possible in these cases, with the island violation repaired by ellipsis presumably, similar movement should possible for the contrastive cases, such as (134e). Thus, such an explanation for the partial island sensitivity of contrastive Stripping relies on the stipulation that repair is not possible for these cases, leaving only a non-isomorphic, and so Parallelism violating, parse.

(137) a. Bill met the student who speaks some asian language.

b. Joe: Yeah, Mandarin$_F$.

c. . . . Mandarin ~~Bill met the student who speaks t~~.

A second way to distinguish the partial island sensitivity of contrastive Stripping from the complete insensitivity of non-contrastive Stripping using Parallelism builds on the proposal of Griffiths and Lipták (2014). Working from the assumption that island-violating contrastive ellipsis is ungrammatical, they argue that this ungrammaticality results from a disparallelism between what the LF of the ellipsis site, under an isomorphic construal, would have to look like in cases of island-violating contrastive ellipsis, and what the LF of the antecedent would look like if we assumed that island-violating covert focus movement were impossible. Of course, the results of Experiments 1-4 show that contrastive ellipsis not generally unacceptable, and the sensitivity to islands is less that what was observed for overt island-violating movement. Rather, the island-violating Stripping contexts are merely slightly less acceptable than the non-island cases. So, let's imagine what this analysis would predict if we recast the effects of a Parallelism not as inducing ungrammaticality but merely a reduction in acceptability.

In essence, Griffiths and Lipták (2014) assume that LF focus movement of the correlate is island sensitive, and, while the island violation resulting from the movement of the remnant in island-violating contrastive ellipsis configurations is repaired, this results in a violation of Parallelism. The difficulty with such a proposal is that, if ellipsis repair is successful because PF island-violating movement has no phonological correlate, due to ellipsis, then it isn't clear why covert movement of the correlate should be island sensitive in the first place. It follows from a repair account of island repair that the relevant islands are PF phenomena. Consequently, movement that has no effect on PF, e.g. covert movement, should induce no violation of a PF constraint.

A final, more general problem with any account of partial island sensitivity attributing a principal role to Parallelism is that it seems that Parallelism violations, in general, produce much greater reductions in acceptability than we seem to see in cases of island-violating contrastive Stripping. For example, consider the exchange in (138), adapted from

Lappin (2005). The antecedent, taken in isolation, admits two readings. In one, *a supervisor* scopes over *each student*, which is consistent with a state of affairs in which there is one supervisor that each student will consult. In the other, *each student* scopes over *a supervisor*, which is consistent with a state of affairs where there is a supervisor, and possible a different one for each student, consulted by each student. These same interpretations are possible given (138a) in the context of the elaborative continuation in (138b). Note, however, that whatever the scopal interpretation that the Stripping phrase receives, the antecedent must have the same interpretation. Thus, it is quite impossible for (138b) to be interpreted in which *her* is bound by *each student*, resulting in pairs of supervisors and their students, while *a supervisor* scopes over *each* in the antecedent. Such an impossible interpretation constitutes a Parallelism violation, and so, for the relatively weak penalty seen in partially island sensitive Stripping, it seems unlikely that Parallelism can be to blame.

(138) a. A: Each student will consult a supervisor.

b. B: Yeah, her own.

Given these problems with attributing the merely partial island sensitivity of contrastive Stripping to a Parallelism violation, I would like to suggest an alternative account. The basis for this alternative is the observation that the antecedents for island-violating contrastive Stripping continuations, as in (139a), are ambiguous with respect to what constituent is taken to be the focus.

(139) a. James interviewed the student who speaks *Thai*.

b. No, *Mandarin*.

In one parse, illustrated in (140), the italicized phrase is the focus, and in the other, the whole of the island, which contains the italicized phrase, is the focus. This second parse is due to a phenomenon known as focus projection (e.g. Selkirk (1996)), by which a constituent containing a word that is prosodically prominent can sometimes be understood to the focus, rather than just the prominent word itself. Now, notice that in (140a), if the focus is to take wide scope, it must scope out of the island. In (140b) on the other hand, where the whole island is taken to be the focus, the focus does not need to escape an island to take wide scope, as it is the island itself which is the focus.

(140) a. James interviewed the student who speaks [*Thai*]$_{Foc}$

b. James interviewed [the student who speaks *Thai*]$_{Foc}$

Now, let us assume that, in examples ambiguous between a parse involving island violating wide scope and a parse which does not involve any island violating wide scope, the non-island violating parse is preferred. To explore this intuition, consider example (141), which is ambiguous between at least a narrow focus parse, like that in (140a), and a projected focus parse, like in (140b). In the context given in (141a), the continuation in (141b) is only compatible with the narrow focus interpretation. The continuation in (141c) is only compatible with the projected focus interpretation. Disambiguated in context, both readings are available, but the intuition is that the wide scope interpretation is more accessible. For what follows, this bias need not be absolute, nor even overwhelming. All that is required is that the non-island violating parse be preferred over the island violating parse.

(141) James didn't interview the student who speaks *Thai*.

a. Context: the student James was trying to interview today speaks Thai

b. ...so it turns out he interviewed someone who he wasn't trying to interview.

c. ...so he didn't interview anyone today.

Given this assumption, the antecedents of island violating contrastive Stripping continuations, presented to participants as in (142a), would initially be more likely to be parsed with a projected focus, as in (140b) rather than (140a). However, given an island violating Stripping continuation, as in (142b), the only parse of the antecedent that satisfies Parallelism is the dispreferred narrow focus parse in (140a). Arriving at this mismatch between initial parse and the parallel requirements of the Stripping continuation, participants might reanalyze the antecedent from the preferred parse to the dispreferred parse, which would thereby plausibly induce some measure of unacceptability. Or they might simply fail to perform the reanalysis necessary to satisfy Parallelism, which would yield a correlate which mismatched the remnant, rendering the example even more unacceptable, as compared to the non-island cases.

(142) a. James interviewed the student who speaks *Thai*

b. No, *Mandarin*.

The ambiguity represented in (140) does not apply to island-violating non-contrastive Stripping antecedents, as the characteristic by which the indefinite correlate scopes, its 'indefiniteness' does not spread from the italicized element to the whole island, as contrastive focus can. Thus, there is no non-island violating wide scope parse available to such antecedents, which would otherwise be the preferred parse. Consequently, upon arriving at the Stripping continuation, no reanalysis is needed and no mismatch between correlate and remnant is endangered.

(143) James interviewed the student who speaks *an Asian language*

It follows then that just contrastive Stripping should demonstrate any effect of islandhood. Only in these cases does the preference for a non-island violating parse result in a bias towards a misparse of the antecedent. The resulting reanalysis yields a reduction in acceptability as compared to the non-island contrastive Stripping cases. However, given that the bias which leads to the misparse is merely a bias, and not an absolute constrain, the resulting reduction in acceptability is less than that found in instances of overt island movement, such as seen in Experiments 1-4 for island-violating It-Clefts.

If this explanation is on the right track, then it should be possible to override the bias for one focal structure over another by placing the examples in a context which favors the narrow focus parse, over the projected focus parse. Consider the dialogue, illustrated in (144), between three good friends, Alexi, Tim, and Spencer. The indefinite in (144a) raises the issue of which language is spoken by the student that James. It seems intuitive to take Tim's reply to then be addressing that issue, rather than addressing the issue of who James interviewed. Thus, in taking (144b) to be addressing the issue of *which language*, it is parsed with narrow focus, on *Thai*, instead of projected focus, on *the student who speaks Thai*. Consequently, the island-violating Stripping continuation in (144c) should be fully acceptable, as the antecedent has been parsed with a focal structure consistent with the island-violating continuation[17].

(144) a. Alexi: James interviewed the student who speaks *an Asian language*

[17]The same argument would hold if Tim's reply were a non-contrastive Stripping continuation. The focal structure of such an island-violating continuation would be unambiguous, and Spencer's corrective Stripping continuation would have a perfectly appropriately parsed antecedent.

b. Tim: Yeah, he interviewed the student who speaks *Thai*.

c. Spencer: No, *Mandarin*.

To conclude the discussion on the partial island insensitivity of contrastive Stripping, we have examined two possibly ways to explain the partial sensitivity in terms of a Parallelism violation, neither of which held up to closer scrutiny. I then suggested that the partial island sensitivity of contrastive Stripping was not the result of a failure of Parallelism, but rather the result of the reanalysis of an ambiguous antecedent in order to satisfy Parallelism. This reanalysis, where necessary, is the result of biased ambiguity resolution towards a focus structure which does not implicate island-violating wide scope. The projected focus structure favored by this biased ambiguity resolution is, however, disparallel to that of an island-violating Stripping continuation, thereby requiring reanalysis to satisfy Parallelism, but also thereby resulting in a somewhat reduced acceptability.

I would like to conclude the general discussion with a summary of potential shortcomings of Experiments 1-4, as well as some further lines of inquiry. First, Experiment 1 suffered from an issue with the creation of the stimuli lists in which participants saw multiple items from the same sentence frame across experimental conditions. I imagine we should not be seriously concerned here, as the results of Experiment 1 were largely in line with those of Experiments 2-4, with the exception of the complete island sensitivity found for elaborative Stripping in the Definite Island Stripping conditions. However, Experiments 2-4 found only a partial island sensitivity for these conditions, and so I feel comfortable setting this anomalous result of Experiment 1 aside.

A second issue follows from the inclusion of Factive verbs in the non-island conditions across Experiments 1-4. We suspected that this inclusion would confound the results, a suspicion confirmed by subsequent independent comparisons of the non-factive items to the definite island items and to the factive items. However, such an analysis was not built into the design, and could have subsequently reduced the power of the experiments. That the results of Experiments 1-4 yielded a quite similar pattern of results, I set aside a formal power analysis for the time being.

A third issue the attentive reader may have already noticed: although we excluded a certain number of participants on the basis of their ratings of the filler items, we did not include a formal requirement that the participants be native speakers of English. We take the consistently higher ratings of filler items we a priori expected to be good as better than those we expected to be bad, along with the general consistency in experimental results across the four experiments as indicative that any non-native speakers that may have participated in the experiment to not have sufficiently different grammars from the native speakers so as to bias the results in any meaningful manner. If the participation of non-native speakers did indeed influence the results, then, given the consistency of the results, they must have done so in a manner consistent across all four experiments, which we take to be unlikely.

These sets of experiments point to a number of future studies, many obvious. Most obviously, one ought to wonder if these results generalize to more or less similar elliptical configurations. Whether Fragment answers are island sensitive comes to mind immediately, as should Contrastive Sluicing, which has widely been claimed to be island sensitive (e.g. Romero 1998; Merchant 2001). Given the various extant claims in the literature about Stripping and Fragment Answers, it seems pertinent to examine Contrast Sluicing as well.

Similarly, we examined here just one type of island: definite relative clauses. A wealth of other island configurations are well attested in the literature, and it remains to be seen how these alternatives stack up to those tested here. Such a broader island exploration could also serve to inform the source of the partial island sensitivity found here for contrastive Stripping. Other sorts of islands don't seem to require the introduction of additional discourse referents, to which I have attributed what island effect is attested in contrastive island-violating Stripping. For example, in (145), we have foci in a coordinate structure island and in a sentential subject island. Considering alternatives to Bill in these structures, however, doesn't require any additional discourse referents to be accommodated. Thus, if the above hypothesis about the source of the partial island sensitivity of contrastive Stripping is correct, Stripping continuations of such items should not, even without additional supporting context, induce a reduction in acceptability. The ambiguity between a new and a contrastive focus interpretation would remain for antecedents such as these, but, given that such an ambiguity would also be present for the non-island controls, islandhood shouldn't induce any additional unacceptability here.

(145) a. James met Mary and *Bill*.

b. That *Bill* is happy bothers everybody.

From the analysis of factivity across Experiments 1-4 comes the question of whether these are "true" island effect. That is, does extraction of an argument from within the clausal complement of a factive verb induce an reduction in acceptability beyond the additive effects of the factivity of the matrix verb and of long distance extraction? The design of the present experiments does not allow us to settle this issue.

Finally, across these experiments, we found a consistently lower acceptability for elaborative It-Clefts as compared to corrective and non-contrastive It-Clefts. I earlier suggested this was due to an exhaustivity effect. However, given that the reduction in acceptability for these exhaustivity violating examples was only marginally, though reliably, lower than their non-violating counterparts, I speculated that the reduction may indicate that exhaustivity is not encoded in the semantics of It-Clefts, but rather the result of a pragmatic process. Such deserves further investigation.

This concludes the discussion of Experiments 1-4. In the next chapter, I investigate the interpretation of island-violating Stripping, through Binding Condition C reconstruction effects, in an attempt to determine what sort of syntactic structure is contained within the ellipsis site, if any. Taking the results of this chapter along with those of the next, in Chapter 4, I argue that the ellipsis site must be able to contain syntactic structure that is isomorphic to that of the antecedent. A consequence of this, I argue, is that the grammar must be able to perform Island Repair. I there explore how this might be implemented.

CHAPTER 3

Stripping and Binding Condition C

3.1. Introduction

In Chapter 2, I presented the results of a series of large-scale acceptability judgment experiments on the purported island insensitivity of Stripping configurations. Those experiments showed that non-contrastive Stripping is island insensitive and that contrastive Stripping is only partially sensitive to islands. I then suggested that the partial island sensitivity of contrastive Stripping is likely due to a baised parser, leading to the conclusion that, from a grammatical perspective, both are island insensitive.

In this chapter, I present the results of two further large-scale judgment experiments, conducted in colaboration with Tim Hunter and Masaya Yoshida (Hunter et al.). These experiments are designed to directly test the predictions of the myriad approaches to ellipsis, with respect to a particular kind of connectivity effect: Binding Condition C. The results reveal that Stripping patterns identically to non-elliptical controls, in both non-island and island-violating contexts. This pattern is expected only if ellipsis sites are populated with syntactic structure which is both rich enough to resolve the binding conditions and *isomorphic* to the structure of the antecedent. In turn, I argue that these facts are very difficult to account for under approaches to ellipsis which posit that the ellipsis site contains no structure (Ginzburg and Sag 2000, Valmala 2007), a null element anaphoric to the antecedent (Barker, 2013), or material which is non-isomorphic to the antecedent (Merchant 2001,Barros et al. 2013).

In these experiments, we investigated the behavior of Stripping in the context of potential Binding Condition C contexts. The logic was this: if ellipsis sites are populated with elided syntactic structure, isomorphic to that of the antecedent, then a Stripping continuation should be just as sensitive to Binding Condition C contexts as would a comparable overt, unelided counterpart.

We asked participants to rate the plausibility that underlined words in short dialogues referred to the same person. We used the canonical, non-elliptical conditions, illustrated in (146) and (147), as a baseline. In (146), a pronoun and C-commanding R-expression are co-indexed, which participants rated as plausible.

The key finding of Experiment 5 is a contrast in the plausibility of the indicated co-indexation relations in Stripping examples like that in (146) and (147). In (146), the correlate to a Stripping remnant which contains a pronoun is preceded and C-commanded by a co-indexed R-expression. Participants found such co-indexation perfectly plausible. What they found implausible were co-indexation relations in configurations like that in (147). In these examples, the remnant contains an R-expression, co-indexed with a pronoun that precedes and C-commands the correlate. Experiment 6 disentangles the effects of precedence and C-command, keeping precedence relations between pronoun and name constant, while varying whether the pronoun C-commands the base position of the name. Participants exhibited a sensitivity to C-command, such that even if the pronoun precedes the base position of the name, coreference between the two was less plausible if the pronoun C-commanded then name than if it did not.

(146) a. Ann: Mary$_i$ knows the student who sent a text to Bill.

b. Bill: No, to her$_i$.

(147) a. Ann: She$_i$ knows the student who sent a text to Bill.

b. Bill: *No, to Mary$_i$.

3.1.1. Binding condition connectivity effects under Ellipsis

The argument in this chapter hinges on the distribution of R-expressions with respect to co-indexed elements, a distribution known in the literature by the classic constraint which captures this distribution: binding condition C (henceforth BCC, Chomsky (1981)). R-expressions, such as proper names, cannot be C-commanded by co-indexed elements. For example, in (148), *Alexi* is C-commanded by the pronoun *he*; if this pronoun is interpreted as co-indexed with *Alexi*, the example is unacceptable[1]

(148) He$_{\{j/*i\}}$ said the idea bothered Alexi$_i$.

In (148), the pronoun both precedes and C-commands the R-expression, but we can be sure that the unacceptability of the co-indexed interpretation is not due to the precedence relation between the pronoun and R-expression. The examples in (149) demonstrate that

[1] If the the question under discussion is "Who did Alexi say the idea bothered?" a configuration in which an R-expression C-commands another co-indexed R-expression is said to be only somewhat unacceptable: "??Alexi said the idea bothered Alexi$_F$" (Lasnik, 1989). Likewise such could be the response to an echo question. Note, however, that examples with R-expression C-commanded by co-indexed pronouns, such as (148), remain unacceptable in the same contexts. I set aside the issue of why this contrast might hold.

pronouns may precede co-indexed R-expressions, just in case the pronoun does not C-command the R-expression.

(149) a. His$_i$ friends said the idea bothered Alexi$_i$.

b. He$_i$ laughed, but then Alexi$_i$ cried.

Researchers familiar with the literature on ellipsis may be surprised by the claim that it is possible to use the binding conditions to diagnose structure within the ellipsis site. After all, Fiengo and May (1994) demonstrated that Verb Phrase Ellipsis does not appear to respect binding condition C. In (150a), a non-elliptical control illustrates a canonical BCC violation (150). A similar example, with the same intended interpretation, in which the lower VP has been elided, (150b), however, is fully acceptable, with no hint of a BCC violation. Similar pairs of examples can be constructed for Stripping cases as well, (151).

(150) a. *Lara loves Sol$_i$ and he$_i$ thinks that Sally loves Sol$_i$ too.

b. Lara loves Sol$_i$ and he$_i$ thinks that Sally does too.

(151) a. She$_F$ said that Lara admires Bill$_i$. *Yeah, and he$_{F,i}$ said that Lara admires Bill$_i$ too.

b. She$_F$ said that Lara admires Bill$_i$. Yeah, and him$_{F,i}$ too.

Fiengo and May (1994) argue, as part of their structural approach to ellipsis, that these effects result from a process by which elided material may differ in form from the antecedent in just such a way as to avoid the BCC violation, a process known as *Vehicle Change*. Under a vehicle change analysis, the elided syntactic structure associated with the Stripping clause in (151b) would not contain the R-expression *Bill*, but rather a pronoun co-indexed with Bill, as in (152).

(152) ...and $\text{He}_{F,i}$ said that Lara admires him_i too.

Advocates of non-Structural approaches to ellipsis have pointed out that such effects are expected if there is no structure within the ellipsis site (Ginzburg and Sag, 2000, pg.297). We wouldn't need to say anything special about (151b), such as a construction specific rule of vehicle change, because ellipsis sites never contain the material that would trigger BCC effects in the first place. The structure of the Stripping phrase is simply as in (153).

(153) ...and $\text{him}_{F,i}$ too.

Likewise, under accounts of ellipsis in which the ellipsis site is a null element anaphoric to the antecedent, or an anti-constituent thereof, we might expect to find such insensitivity to BCC effects, although, to my knowledge, none of the proponents of these accounts have discussed these issues. Notice that the verbal anaphor *do so* behaves just like VPE with respect to BCC effects. Compare the VPE example in (150b) with the anaphoric variant in (154), both under the interpretation in which Bill admires himself; neither exhibits the unacceptability one would expect to result from a BCC violation. Anaphors, regardless of their interpretation, are structurally insufficient to trigger a BCC violation.

(154) ...and she said that Bill does so too.

However, there are elliptical configurations for which BCC does hold. Morgan (1973) points out that Fragment Answers in which the fragment contains the relevant R-expression do exhibit BCC effects. Morgan provides the paradigm in (155), in which the constituent question in (155) is fragment answered with the unacceptable (155b). Note that this Fragment Answer is just as unacceptable as its unelided counterpart in (155c)[2].

(155) a. Where is he$_i$ staying?

b. *In John$_i$'s apartment.

c. *He$_i$ is staying in John$_i$'s apartment.

We can demonstrate that this BCC effect is due to the C-command relationship, instead of a precedence relationship, between the pronoun and the base position of the correlate. If the pronoun does not C-command the base position of the correlate, as in (156a), the fragment and long answers in (156b) and (156c) remain acceptable.

(156) a. Where is his$_i$ mother staying?

b. In John$_i$'s apartment.

c. His$_i$ mother is staying in John$_i$'s apartment.

Similarly, Hunter and Yoshida (2016) observe the same holds true of Stripping configurations. In (157), a pronoun in the antecedent C-commands the correlate, *that Mary left*, yet an R-expression within that remnant, *him*, can be co-indexed with the pronoun. Again, acceptable controls like (158), in which the pronoun does not c-command the correlate, show that the unacceptability in (157b) can be attributed to BCC.

(157) a. He$_i$ said that Mary$_i$ left.

[2]Note also that (155b) embeds the R-expression within a DP, such that its unacceptability cannot be attributed to a strong crossover effect, of the sort discussed by (Chomsky, 1981; Wasow, 1972), rather than BCC.

b. *Yeah, but not that John$_i$ left.

(158) a. His$_i$ friends said that Mary left.

b. Yeah, but not that John$_i$ left.

As Morgan (1973) and later Merchant (2004) point out, the sensitivity of Stripping to BCC poses problems for non-structural and anaphoric accounts of ellipsis. The apparent insensitivity of elliptical phenomena to BCC illustrated in (151b) and (150b) was explained under these kinds of accounts by the lack of structure within the ellipsis site that would permit the application of BCC. However, this same quality of these accounts likewise predicts that examples like (155b) and (157) should also be insensitive to BCC. There is simply insufficient structure within the ellipsis site, whether because a lack of structure within the ellipsis site entirely, or because that structure is a simple anaphor, to resolve BCC.

These facts follow naturally, by and large, from the remaining structural accounts, the LF copying and deletion accounts. Under both of these accounts, the ellipsis site contains syntactic structure, exactly of the sort that one would expect to yield BCC effects in the appropriate configurations. Such structural resolutions, here of the deletion variety, are depicted in (159) and (160).

(159) a. Where is he$_i$ staying?

b. *In John$_i$'s apartment ~~he$_i$ is staying t~~.

(160) a. He$_i$ said that Mary$_i$ left.

b. *Yeah, but not that John$_i$ left ~~he$_i$ said t~~.

There is a variant of the deletion accounts, however, which makes different predictions for these BCC cases. We've been considering non-island examples, but the island insensitivity of ellipsis has lead to the proposal that, in certain cases, the syntactic structure contained within the ellipsis site is quite different from that of the antecedent (Merchant 2001, Fukaya 2007, Barros et al. 2013). When the correlate is contained within an island, such *non-isomorphic* or *evasion* approaches to ellipsis propose that the structure in the ellipsis site does not contain an island at all. The island insensitivity of island insensitive ellipsis is explained; while the remnant moves to escape ellipsis, because there is no island within the ellipsis site, there is no island violating movement implicated. Thus, for an example like (161b), the ellipsis site would be resolved as in (161c), a *short source* construal. In this short source ellipsis site resolution, the remnant need not move out of an island in order to escape the ellipsis, as the ellipsis site does not contain an island at all.

(161) a. Bill knows the student who speaks Thai$_F$.

b. Joe: No, Mandarin$_F$.

c. . . . Mandarin$_i$ ~~the student speaks t~~.

Above, we mention a short source the non-isomorphic analysis. This is not the only non-isomorphic resolution that has been proposed, however. It has also been proposed that the ellipsis site could instead be resolved as a reduced it-cleft. However, such a resolution would not possible for the kinds of examples under consideration here, if it is correct that contrastive reduced it-clefts are unacceptable when the correlate is contained within an embedded clause, whether a complement clause (162) or an island clause (163) (Barros, 2012a). See also (Griffiths and Lipták, 2014, pg. 31), who claim that clefts are island sensitive, and (Vicente, 2008, pg. 15), who claims that clefts, presumably reduced clefts, are incompatible with contrastivity. Consequently, for the purposes of the present discussion, which centers on contrastive Stripping, I will use *non-isomorphic* to refer to short source analyses.

(162) a. Does Chris think that Jack$_F$ is dating Sandra?

b. *No, it is Bill.

(163) a. Is Chris jealous because Jack$_F$ is dating Sandra?

b. *No, (it is) Bill.

(164) a. Bill knows the student who speaks Thai$_F$.

b. Joe: No, Mandarin$_F$.

c. . . . Mandarin$_i$ ~~the student speaks t~~.

d. . . . Mandarin$_i$ ~~it was t~~.

These non-isomorphic analyses have been proposed to explain a variety of properties of elliptical phenomena, including the availability to strand prepositions in non-P-stranding languages (Vicente, 2008). They have also been proposed in order to account for the island insensitivity of ellipsis without the need to appeal to island repair. The key to how these non-isomorphic analyses account for ellipsis island insensitivity is that the material within the ellipsis site does not contain an island. This trick is accomplished in a wide range of ways, depending on how the ellipsis site is resolved. The short source analysis, for example, construes the ellipsis site as simply the content of the island itself, as in (161c). A consequence of this is that the non-isomorphic construal of (161b) does not contain the matrix subject of the antecedent *Bill*[3] Non-isomorphic analyses also make a clear prediction for the behavior of island violating ellipsis and BCC. Consider then how island insensitive ellipsis in BCC contexts would be predicted to be resolved. In example (165), the antecedent contains a pronoun which C-commands the correlate, and the remnant an R-expression co-indexed with that pronoun. A non-isomorphic construal of the ellipsis site, whether a short source construal, as in (166a), or a cleft construal, as in (166b), however, would not contain an instance of the co-indexed pronoun. Consequently, we would not expect, under either of these non-isomorphic construals, for a BCC violation to arise. Under an isomorphic construal, on the other hand, as depicted in (167), we *would* expect a BCC effect.

(165) a. She$_i$ likes the manager who assigned the job *to Bill*.

b. No, to *Mary*$_i$.

(166) a. ...to *Mary*$_i$ ~~the manager assigned the job t~~

b. ...to *Mary*$_i$ ~~it was~~

(167) ...to *Mary*$_i$ ~~she$_i$ likes the manager who assigned the job t~~

[3]One could also imagine that the subject in the ellipsis site in the *short source* construal depicted in (161c) could contain a relative clause depicting the material from outside of the relative clause island. In such an analysis, what had been the matrix subject of the antecedent no longer C-commands the base position of the remnant, and so, from the perspective of what follows, would make identical predictions as the simpler analysis in (161c)

(i) Mandarin$_i$ ~~the student who Bill knows speaks t~~.

Non-isomorphic approaches to ellipsis island insensitivity all posit that a non-isomorphic parse is available when the correlate is contained within an island. The question then arises whether such non-isomorphic parses are available for any elliptical example, or only when the correlate is contained within an island[4]. For example, in (168), the correlate, *Thai*, is contained within a complement clause. If non-isomorphic parses are generally available, what I will call a *consistently* non-isomorphic analysis, then even in cases of complement clauses, a non-isomorphic parse, like that in (168c) or (168d), would be available. If non-isomorphic parses are only available as a last resort mechanism, in cases of islands, in what I will call a *selectively* non-isomorphic analysis, then such examples as (168b) should only have an isomorphic parse, as in (168e).

(168) a. Bill said that the student speaks Thai$_F$.

b. Joe: No, Mandarin$_F$.

c. . . . Mandarin ~~the student speaks t~~.

d. . . . Mandarin ~~it was t~~.

e. . . . Mandarin ~~Bill said that the student speaks t~~.

These two variants of the non-isomorphic analysis make different predictions in light of the BCC and non-island configurations. A consistently non-isomorphic analysis would permit the non-isomorphic parses of (169b), given in (169c)[5]. Example (169c) thereby

[4]Analyses which appeal to non-isomorphism to account for other phenomena give different licensing conditions for non-isomorphism. For example, Vicente claims that non-isomorphic, it-cleft, parses are possible if such "strengthens antecedent clause" (Vicente, 2008, pg. 6).

[5]It is sometimes said, in the context of Sluicing, that such short source parses are regularly available in non-island contexts and, in certain cases, the ellipsis site is obligatory construed as the short source; see discussion in Van Craenenbroeck (2010). If such a short source were available for Sluicing cases, illustrated by the parse of (ib) as (ic), we might expect BCC effects to be absent in cases such as (ii). However, my

evades the potential BCC violation. A selectively non-isomorphic analysis would, under the assumption that non-isomorphic parses are only available to evade island violations, only permit the isomorphic parse for such examples, given in (169d).

(169) a. She$_i$ said that the manager assigned the job *to Bill*.

b. Joe: No, to Mary$_i$.

c. . . . to Mary$_i$ ~~the manager assigned the job t~~.

d. . . . to Mary$_i$ ~~she$_i$ said that the manager assigned the job t~~.

Investigating BCC in non-island and island contexts allows us to test the predictions of non-structural analyses, isomorphic structural and the two kinds of non-isomorphic structural analyses. These predictions are summarized in table (3.1).

Table 3.1. Predictions concerning potential BCC violations

Analyses	Island Context	Non-Island Context
Non-Structural	acceptable	acceptable
Isomorphic Structural	unacceptable	unacceptable
Consistently Non-Isomorphic Structural	acceptable	acceptable
Selectively Non-Isomorphic Structural	acceptable	unacceptable

informants judge the co-indexation marked there to be unacceptable. This suggests that the judgment that a short source is available might have something other than a syntactic basis.

(i) a. Bill said that the manager assigned the job to an employee.
b. Yeah, but I don't know which.
c. . . . which employee the manager assigned the job to
d. . . . which employee Bill said that the manager assigned the job to

(ii) a. She$_i$ said that the manager assigned the job to one of her$_i$ employees.
b. Yeah, but I don't know which one of Mary$_i$'s employees.
c. . . . which one of Mary$_i$'s employees the manager assigned the job to
d. . . . which one of Mary$_i$'s employees she$_i$ said that the manager assigned the job to

3.2. Experiment 5

We tested the predictions of non-structural analyses, isomorphic structural and the two kinds of non-isomorphic structural analyses with a large-scale judgment study, using a methodology adapted from Kazanina et al. (2007). Results support the predictions of the isomorphic approaches. As expected, the indexation in non-elliptical island and non-island controls with canonical word order, like (170) was rated by participants to be unlikely, showing they were sensitive to BCC. The Stripping conditions, in both island (171) and non-island (172) contexts, patterned just like the non-elliptical controls, in that the co-idexation illustrated here was rated to be unlikely as compared to non-BCC contexts like that in (173).

(170) a. Ann: She$_i$ knows the student who sent a text to Bill.

b. Bill: No, She$_i$ knows the student who sent a text to Mary$_i$.

(171) a. Ann: She$_i$ knows the student who sent a text to Bill.

b. Bill: No, to Mary$_i$.

(172) a. Ann: She$_i$ heard that the student sent a text to Bill.

b. Bill: No, to Mary$_i$.

(173) a. Ann: Mary$_i$ knows the student who sent a text to Bill.

b. Bill: No, to her$_i$.

3.2.1. Participants

44 participants were recruited through Amazon's Mechanical Turk[6]. Participants were limited to IP addresses within the US, were only permitted to participate in the experiment once and were compensated $2 USD. Five participants failed to complete the survey, and so their results were excluded from analysis. The data from one additional subject was excluded, on the grounds that the mean ratings of the low plausibility fillers was not lower than that of the high plausibility fillers. Taking these exclusions into account, the data of 38 participants were analyzed.

3.2.2. Design, Stimuli, Procedures

The stimuli were presented to participants as two sentence dialogues, in which the second sentence, preceded by 'No, . . . ', was intended to be understood as a corrective response to the first sentence. In all conditions, a prepositional phrase found in an embedded clause in the first dialogue sentence was italicized; we call this the PP correlate. The second dialogue sentence also contained an italicized prepositional phrase, which we call the PP focus. The PP focus was intended to be understood as the correction. Across conditions, these PPs were always underlined, as were the matrix subjects of the first dialogue sentence and, in the canonical conditions, the matrix subject of the second dialogue sentence.

The stimuli conformed to a 2x2x2 factorial design: i. Islandhood; ii. Construction Type; iii. Pronominal Status. The Islandhood factor manipulated whether the first sentence in the dialogue contained a complement clause (=Non-Island) or a definite relative

[6]An online crowdsourcing marketplace platform (www.mturk.com)

clause (=Island)[7]. The Construction Type factor manipulated whether the second dialogue sentence was a Corrective Stripping construction (=Stripping) or a non-elliptical canonical word-ordered construction (=Canonical). In the Stripping conditions, the remnant was the PP focus. The use of prepositional phrases ensured that the only possible correlate to the Stripping remnant was the PP correlate in the first dialogue sentence. In the Canonical conditions, the first and second dialogue sentences were identical, except for the PP remnant and PP focus. The final factor, Pronominal Status, manipulated the distribution of pronouns and proper names in the underlined phrases in the dialogues. In the 'pronoun' conditions, the PP focus contained an underlined pronoun and the matrix subjects of the first dialogue, and also the matrix subject of the second sentence in the canonical conditions, were underlined proper names that were gender matched to the pronoun in the PP focus. In the 'name' conditions, the PP focus contained an underlined proper name and the matrix subjects were underlined gender-matched pronouns.

Table 3.2. Experiment 5: Factors, Example Stimuli

Construction Type		Pronominal Status: Pronoun-Name	Pronominal Status: Name-Pronoun
Stripping	Island	Example (174a-i)	Example (174a-ii)
	Non-Island	Example (174b-i)	Example (174b-ii)
Canonical	Island	Example (175a-i)	Example (175a-ii)
	Non-Island	Example (175b-i)	Example (175b-ii)

(174) Stripping

a. Island

[7]The same factivity issue as seen in Experiments 1-4 was present in Experiment 5. However, as will be discussed shortly, there was no effect of Factivity in the current experiment, so the main analysis collapses the analysis of the Factive and Non-Factive items.

i. Pronoun-Name:

Joe: She likes the manager who assigned the job *to Bill*.

Bill: No, *to Mary*.

ii. Name-Pronoun:

Joe: Mary likes the manager who assigned the job *to Bill*.

Bill: No, *to her*.

b. Non-Island

i. Pronoun-Name:

Joe: She said the manager assigned the job *to Bill*.

Bill: No, *to Mary*.

ii. Name-Pronoun:

Joe: Mary said the manager assigned the job *to Bill*.

Bill: No, *to her*.

(175) Canonical

a. Island

i. Pronoun-Name

Joe: She likes the manager who assigned the job *to Bill*.

Bill: No, She likes the manager who assigned the job *to Mary*.

ii. Name-Pronoun:

Joe: Mary likes the manager who assigned the job *to Bill*.

Bill: No, Mary likes the manager who assigned the job *to her*.

b. Non-Island

i. Pronoun-Name:

Joe: She said the manager assigned the job *to Bill*.

Bill: No, She said the manager assigned the job *to Mary*.

ii. Name-Pronoun:

Joe: Mary said the manager assigned the job *to Bill*.

Bill: No, Mary said the manager assigned the job *to her*.

Forty lexicalizations of the eight conditions were constructed. The gender of the names and pronouns used in the Pronominal Status manipulation were balanced across the lexicalizations. 40 filler sentences were constructed to contain a range of grammatical constructions, yielding a 1:1 filler to target item ratio. 10 filler sentences contained underlined phrases, which, if understood to refer to the same person, would result in either a Binding Condition A violation (n=5) or a Binding Condition C violation. The remaining 30 fillers contained underlined words, which, if understood to refer to the same person, would not result in any Binding Condition violation. However, 10 of these 30 fillers contained some other defect that would result in reduced acceptability.

Four lists were constructed using a Latin square design, each containing five distinct lexicalizations of each of the eight conditions. Each list began with six practice items in a variety of grammatical configurations, each containing underlined words. One practice item would yield a Binding Condition A violation and another a Binding Condition B violation, if the underlined words in these items were understood to refer to the same person. The practice items did not vary in lexicalization or order between lists, and were not marked explicitly as practice items. In sum, each list contained 40 test items, 40 filler items, and 6 practice items, for a total of 86 items. Each of the four lists was pseudo-randomized to ensure that sequential items were not from the same condition.

The stimuli were presented to participants in the form of two-turn dialogues between ‘Joe’ and ‘Bill’. Following the procedure of Kazanina et al. (2007), participants were instructed to read the presented dialogues and to rate how plausible it would be for the underlined phrases in the dialogue to refer to the same person, on a scale of 1 (implausible) to 7 (plausible). The entire dialogue and the rating scale appeared on the same screen. Complete instructions are included in the appendices.

3.2.3. Predictions

The Canonical conditions serve as a control, to ensure that the participants exhibit the BCC constraint, as has been observed elsewhere (Kazanina et al., 2007). As there is no ellipsis in these conditions, we expect the same results, no matter which analysis of ellipsis we adopt. In the Name-Pronoun conditions, both Island and Non-Island, examples (175a-ii) and (175b-ii), we expect the a high plausibility rating for co-reference between the underlined name and pronoun, as no binding condition is violated. In the Pronoun-Name conditions, again, both Island and Non-Island, examples (175a-i) and (175b-i), we expect a low plausibility rating for co-reference between the pronoun and the name it C-commands, as these configurations violation BCC.

Non-structural analyses predict no difference between the ratings of the Name-Pronoun and the Pronoun-Name conditions. Recall that, in these analyses, the ellipsis site is

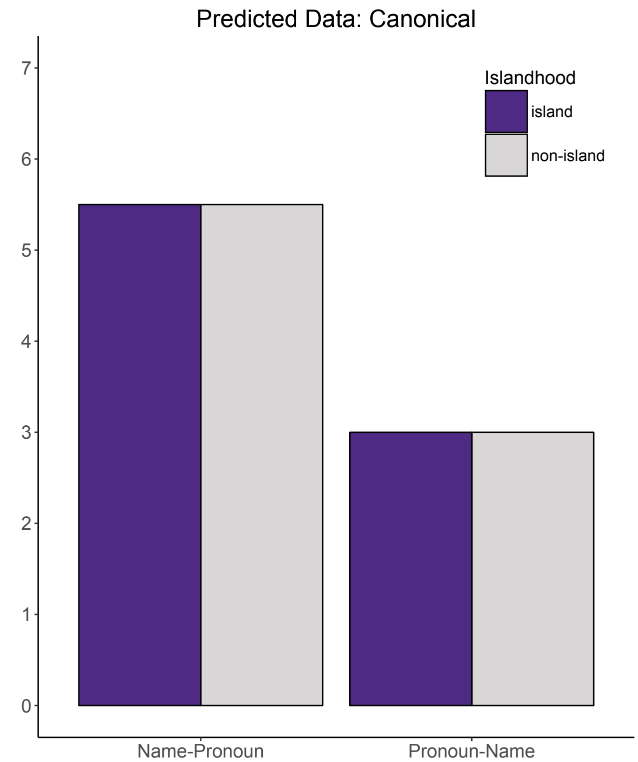

Figure 3.1. Experiment 5: Predicted Canonical Data

not populated with any syntactic structure. Consequently, no BCC is expected to result from configurations like those in the Pronoun-Name conditions; we should find no difference between Pronoun-Name and Name-Pronoun conditions. The indicated co-reference should be plausible in any condition. These predicted results are depicted in the left chart in figure (3.2).

Isomorphic analyses of ellipsis predict that Stripping configurations, just like the Canonical conditions, should yield lower plausibility ratings for the Pronoun-Name conditions than the Name-Pronoun conditions, as illustrated in the right chart in figure (3.2). Under these analyses, in both island and non-island conditions, the ellipsis site is construed

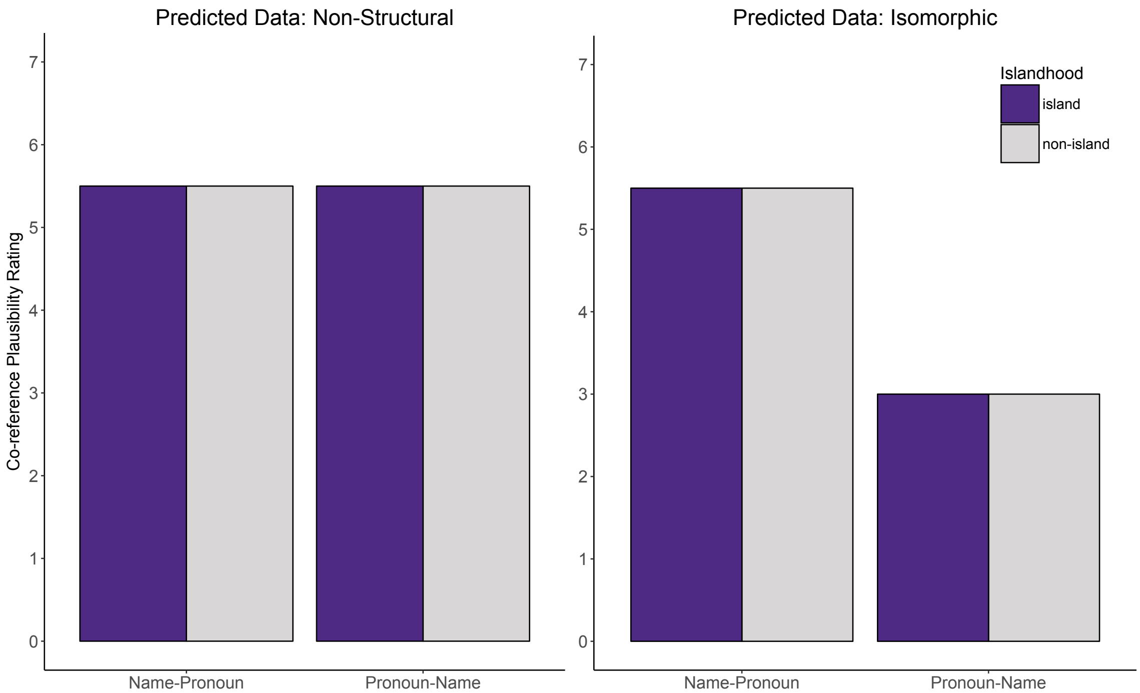

Figure 3.2. Experiment 5: Predicted Data, Non-Structural and Isomorphic Analyses

as containing syntactic material isomorphic to the antecedent. Thus, the base position of the remnant is C-commanded by the matrix subject, and so when the remnant contains a proper name, and the matrix subject is a pronoun, as in the Pronoun-Name conditions, co-indexation of these elements should be rated to be implausible, due to BCC. In the Name-Pronoun conditions, no binding condition is violated, and so co-indexation of the name and pronoun is predicted to be plausible.

In selectively non-isomorphic analyses, island antecedents license non-isomorphic construals of the ellipsis site, such that the matrix subject of the antecedent is not present within the ellipsis site. Consequently, in the Pronoun-Name Island conditions, the ellipsis site will not contain a configuration which would trigger a BCC violation. Thus, Island

Pronoun-Name conditions are expected to be rated as plausible as the Name-Pronoun conditions. In these selectively non-isomorphic analyses however, the non-isomorphic construal is not licensed in non-island conditions. And so in the Pronoun-Name Non-Island conditions, the ellipsis site will contain a BCC violating configuration and these conditions should be rated as less plausible than the Name-Pronoun Conditions. These predictions are illustrated with dummy data in (3.3).

In consistently non-isomorphic analyses, non-isomorphic construals are always possible, and so the potential BCC violations in both Island and Non-Island Pronoun-Name conditions can be avoided by appeal to a non-isomorphic construal of the ellipsis site. These analyses would then predict no differences in plausibility between the Name-Pronoun and Pronoun-Name conditions, as illustrated in (3.3).

3.2.4. Results

The data was analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. 2008), using lme4 in R, with rating as the dependent variable. Contrast-coded fixed effects included Islandhood (Island, Non-Island), Pronominal Status (Pronoun-Name, Name-Pronoun), and Construction Type (Stripping, Canonical), as well as their 2- and 3-way interactions. The maximal random effects structure that would converge was employed, which included random intercepts for Participant and Item, as well as random slopes by participant for Islandhood, Pronominal Status, and Construction Type, and the 2- and 3- way interactions. Model comparisons were performed to determine whether the

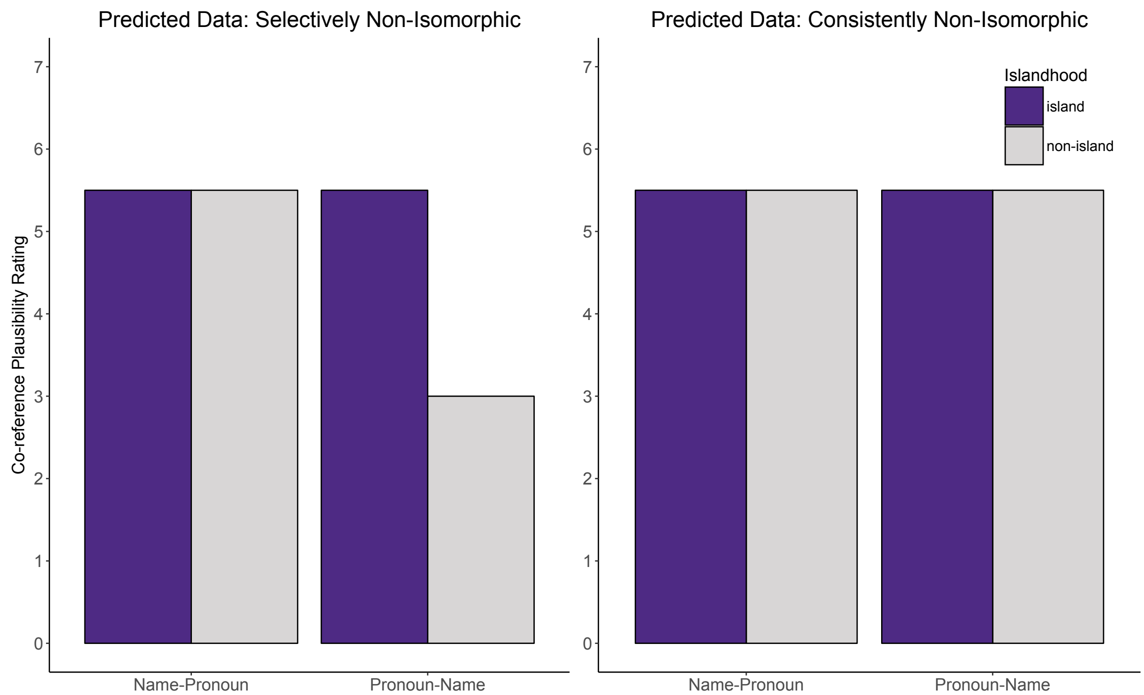

Figure 3.3. Experiment 5: Predicted Stripping Data, Two Non-Isomorphic Analyses

inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

Table 3.3. Experiment 5: Mean Plausibility Results, with standard errors in parentheses

Construction Type		Pronominal Status: Pronoun-Name	Pronominal Status: Name-Pronoun
Stripping	Island	2.64 (0.12)	4.67 (0.13)
	Non-Island	2.68 (0.12)	4.42 (0.14)
Canonical	Island	2.76 (0.14)	5.58 (0.11)
	Non-Island	2.85 (0.14)	5.54 (0.11)

The results of these analyses revealed significant main effects of Pronominal Status (β=2.32, SE β=0.27, $\chi^2(1)$=41.64, p<.001), and Construction type (β=0.58, SE β=0.14,

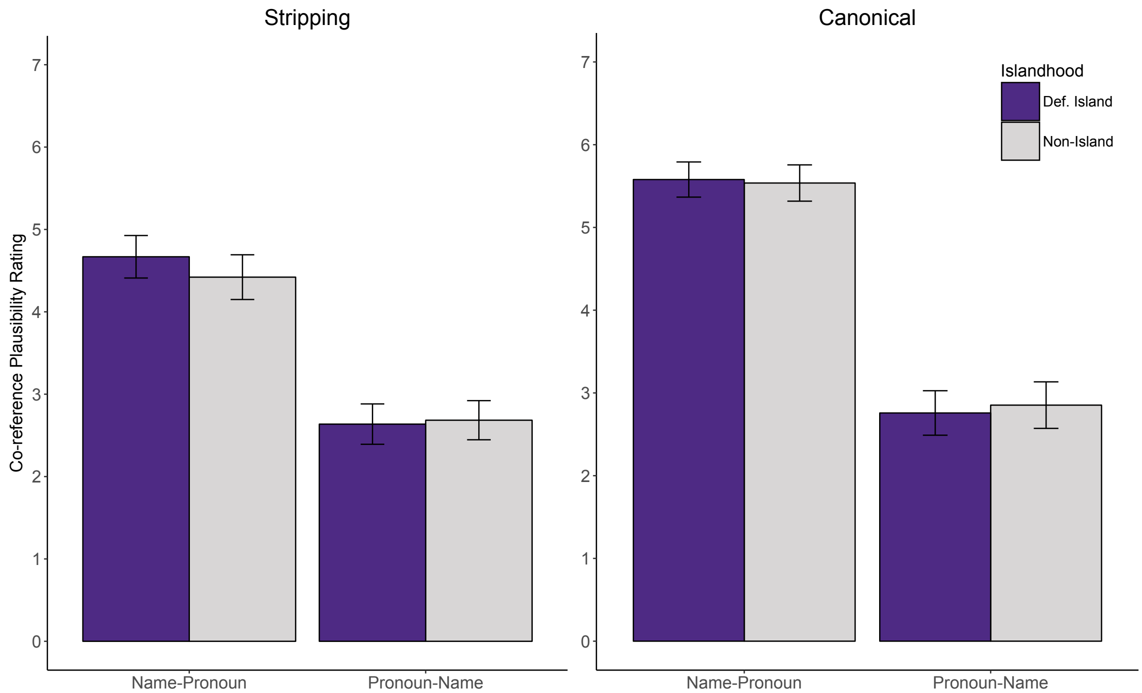

Figure 3.4. Experiment 5: Data

$\chi^2(1)$=15.06, p<.001), where overall participants rated the co-reference relations as less plausible in Stripping Constructions than in Clefts, and less plausible in Name conditions than in Pronoun conditions. Additionally, the 2-way interaction of Pronominal Status x Construction type (β=0.86, SE β=0.15, $\chi^2(1)$=32.7, p<.001), was also significant. No other effects reached significance ($\chi^2(1)$=2.03 p=0.154).

We investigated the significant 2-way interaction with two subset LMER models. We first examined the Stripping subset of the data, which displayed a main effect of Pronominal Status (β=1.89, SE β=0.29, $\chi^2(1)$=27.96, p<.001), such that the co-reference relations in the Name conditions were rated as less plausible than in the Pronoun conditions. No other effects reached significance ($\chi^2(1)$=2.29 p=0.13). Second, the Canonical data subset

Table 3.4. Experiment 5: Estimates of fixed effects

	Coefficient	StandardError	*t*	*p*
Intercept	3.89	0.13	29.96	*n/a*
Islandhood	-0.04	0.07	-0.5	0.616
Pronominal Status (PS)	2.32	0.27	8.67	$<$.001
Construction Type (CT)	0.58	0.14	4.27	$<$.001
Islandhood:PS	-0.21	0.15	-1.43	0.154
Islandhood:CT	0.12	0.15	0.83	0.409
PS:CT	0.86	0.15	5.91	$<$.001
Islandhood:PS:CT	0.16	0.29	0.55	0.582

yielded a main effect of Pronominal Status (β=2.76, SE β=0.28, χ^2(1)=49.34, p$<$.001), such that the co-reference relations in the Name conditions were rated as less plausible than in the Pronoun conditions. No other effects reached significance (χ^2(1)=0.32 p=0.569).

In summary, the Canonical conditions exhibited a classic BCC effect: when an R-expression is C-commanded by a pronoun, participants judged coreference between these elements to be improbable. This was true whether the pronoun and name were separated by a clausal boundary or an island edge. The Stripping conditions showed a very similar pattern. When the remnant contained a name and the correlate was C-commanded by a pronoun, coreference between these items was rated to be improbable. Importantly, the Island and Non-Island conditions, for both Stripping and Canonical conditions, did not differ from each other.

3.2.5. Factive Analysis

The experimental design in Experiment 5 was also subject to the potential factivity confound that afflicted Experiments 1-4. In the subsequent analysis of those experiments, it was found that the factivity of the the matrix verb in the non-island conditions proved to be a true confound. Because the overt movement baseline in those experiments involved a long distance dependency across the factive island boundary, the acceptability of the It-Cleft items involving a factive verb was lower than that of items with a non-factive verb. The Stripping items showed no such sensitivity to factive islands, however.

Experiment 5 differs from the prior experiments in that the baseline in the current experiment involved no long distance movement dependencies. Consequently, no effect of factivity would be expected in the Canonical conditions, an expectation that was born out in the data. However, it has been claimed that ellipsis is indeed sensitive to factive islands, such that acceptable examples of factive island-violating ellipsis involve non-isomorphic structure (Agüero-Bautista, 2007; Sauerland, 1996). If such an analysis were correct, we would expect the factive items to display no sensitivity to BCC, just as the non-isomorphic analyses predict no BCC sensitivity in the definite relative clause island conditions. These predictions were not born out.

Statistical analysis using a logistic linear mixed-effects regression model revealed main effects of Pronominal Status (β=2.22, SE β=0.26, $\chi^2(1)$=40.67, p<.001), and Construction type (β=0.64, SE β=0.15, $\chi^2(1)$=14.55, p<.001), as well as a 2-way interaction of Pronominal Status x Construction type (β=0.93, SE β=0.22, $\chi^2(1)$=16.68, p<.001), was also significant. No other effects reached significance ($\chi^2(1)$=0.71 p=0.401). These results mirror those of the main analysis, in which factive and non-factive items were analyzed together. Table (3.5) details the means and standard errors by condition, figure (3.5) illustrates this data.

Table 3.5. Experiment 5: Factive Analysis Mean Plausibility Results, with standard errors in parentheses

Construction Type		Pronominal Status: Pronoun-Name	Pronominal Status: Name-Pronoun
Stripping	Factive	2.61 (0.17)	4.55 (0.20)
	Non-Factive	2.74 (0.17)	4.31 (0.19)
Canonical	Factive	2.98 (0.22)	5.69 (0.15)
	Non-Factive	2.75 (0.19)	5.42 (0.16)

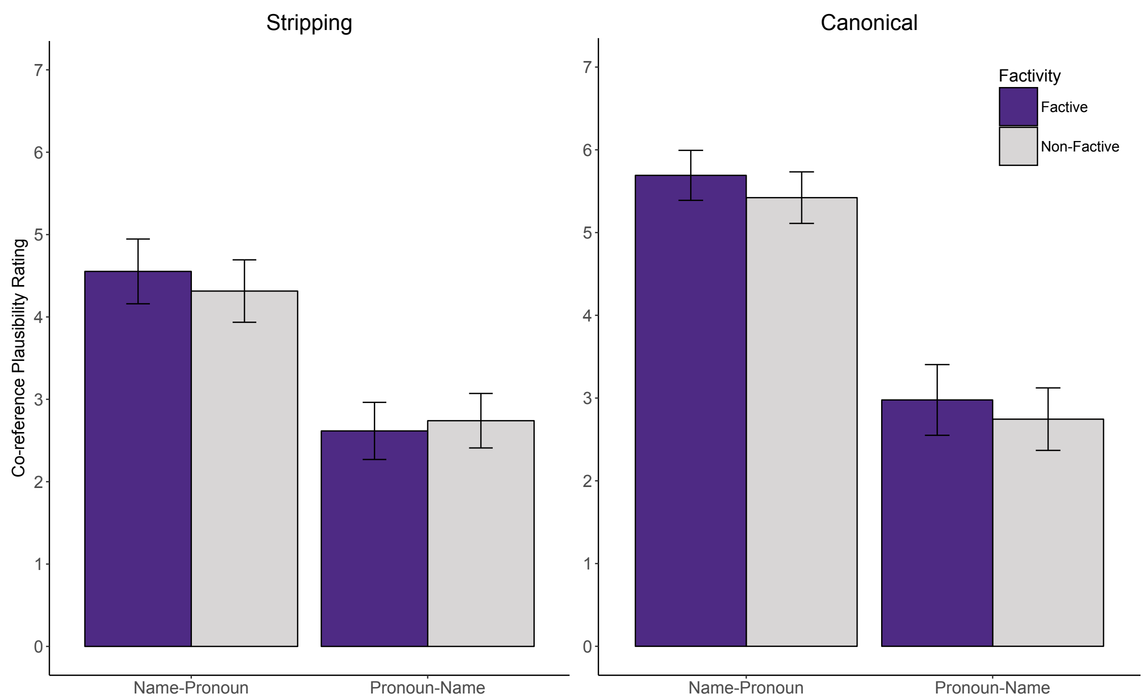

Figure 3.5. Experiment 5: Factive Analysis Data

Although the central focus of this dissertation is on definite relative clause islands, this analysis comparing the sensitivity of non-island and factive island violating Stripping to BCC suggests that the insensitivity of Stripping of not only definite relative clause islands, but factive islands as well, is not due to a non-isomorphic construal of the ellipsis site.

Instead, the sensitivity of Stripping to BCC, when the relevant pronoun would be contained within structure that would not be resolved within the ellipsis site under non-isomorphic analyses, suggests that the full structure isomorphic to the antecedent is resolved within the ellipsis site.

3.2.6. Discussion

In this experiment, we sought to determine how much structure is contained within the ellipsis site of island-violating Contrastive Stripping configurations, using Binding Condition C as a probe. The various approaches to ellipsis, e.g. structural vs. non-structural, make competing predictions about whether island violating Stripping or non-island Stripping or both should exhibit BCC effects on the basis of the syntactic structure and content of the antecedent. If there was no structure within the ellipsis site, there should be no BCC effect. If there was structure in the gap, but if that structure could differ from the antecedent in significant ways, then there should not be a BCC effect when these non-isomorphic structures are available. Finally, if there was structure in the ellipsis site, and if that structure had to be isomorphic to the antecedent at all times, then we should see a BCC effect across the board.

The results bear out the predictions of the isomorphic structural account: the Stripping conditions patterned with the Canonical conditions. The pattern displayed by the non-island Stripping conditions confirm the judgements reported in the literature by Morgan (1973) and others that Stripping configurations are sensitive to BCC effects. This study extends the generalization to remnants originating within definite relative clause islands.

Stripping, in both island and non-island contexts, behaves as if there were syntactic structure, isomorphic to the antecedent, within the ellipsis site.

It is hard to see how these results could be reconciled with the available alternative analyses of ellipsis. Non-structural analyses, by definition, lack the sort of structures which would facilitate the BCC effect observed of Stripping constructions. Non-isomorphic analyses fare no better. Wherever non-isomorphic construals of the ellipsis site are licensed, such construals eliminate the structure necessary for the BCC effect to arise.

Additionally, these results bear on the claim that island-violating corrective Stripping, in which the correlate is final within the antecedent, are derived in a manner distinct from that of other correctives. Griffiths and Lipták remark: "Note that all the cases we construct here for illustration involves correlates that are not sentence-final. This is because sentence-final constituents can be corrected, regardless of their discourse status: (...) We believe this is because sentence-final correction involves a strategy that is distinct from our cases of contrastive fragments" (Griffiths and Lipták, 2014, pg.11, fn.10). Unfortunately, we are given no independent evidence that such sentence final-corrections are distinct from non-sentence final ones, nor what such an alternative strategy ought to look like. The results from Experiment 5, in conjunction from those from Experiments 1-4, however, suggest that the strategy used in these sentence final corrections is one in which the ellipsis site is populated with syntactic structure that is isomorphic to the antecedent, exactly what Griffiths and Lipták propose for non-contrastive ellipsis.

To recap, then, the broad conclusions from the previous chapter and from this one are 1) that Stripping configurations are insensitive to at least definite relative clause islands and

2) that the ellipsis site contains silent syntactic structure, isomorphic to the antecedent. The next step, addressed in Chapter 4, is how we might reconcile the notion that the remnant originates within a domain from which extraction, of the sort which would be required of the remnant if it were to survive ellipsis as a constituent targeting operation, is generally disallowed. However, before we move on, two questions which arose from the results of Experiment 5 need to be addressed.

The first question posed to us about the Experiment 5 concerns whether or not reduced clefts, the structure underlying certain versions of the non-isomorphic analysis, would pattern differently from teh Stripping cases we examined. The question is whether there is a contrast between the two versions of examples like (176), under the indicated co-indexation. Here I use a non-contrastive Stripping example, because, as noted above, it has been claimed that truncated cleft pivots cannot be contrastive with their correlate. Our informants do report a contrast between these two versions, with the version in which the correlate is C-commanded by the pronoun worse. If such facts hold, it seems that the only explanation would be that such reduced clefts are derived form the full clausal structure, which has subsequently been elided. Although Merchant (2001) argues against such a position, it is difficult to imagine how such a BCC effect would otherwise arise.

(176) a. {She$_i$, Her$_i$ friends} said that the manager assigned the job to someone.

b. Yeah, it was *Mary*$_i$.

A second question posed to us about these results concerns the possibility that the observed effect, which we attribute to BCC, is in fact due to not a C-command relationship between the pronoun and name, but rather by the linear order of these two elements. In

Experiment 5, items for which co-reference between the pronoun and name was rated to be implausible placed the pronoun as the matrix subject of the antecedent, and the name within either the Stripping remnant or the embedded indirect object of the canonical conditions, as illustrated in (177) and (178). Thus, no matter whether there was syntactic structure resolved within the ellipsis site or not, in such examples, the pronoun preceded the name.

(177) a. He$_i$ admires the professor who supervises *Mary*.

b. *No, *James*$_i$.

(178) a. He$_i$ admires the professor who supervises *Mary*.

b. …*No, he$_i$ admires the professor who supervises *James*$_i$.

Kyle Johnson (p.c.) suggested that an unacceptability due to this precedence relation might be the source of the effects found in Experiment 5. Across sentences, with no C-command relationship between pronoun and R-expression, a pronoun preceding a co-referential R-expression results in degredation, as illustrated in 179. It could therefore be the case that the results of Experiment 5 were simply be due to what ever causes the effect in (177). If this were the case, then these results would not be evidence of isomorphic structure within ellipsis site, as the effect would be entirely driven by the cross-sentential precedence relation between the pronoun and proper name. However, I do think there are good reasons to reject the relative unacceptability of (179b) as the source for the unacceptability of Stripping examples like (177b).

(179) a. First, he$_i$ walked in.

b. ??And then James$_i$ sneezed.

First, it seems that the unacceptability of examples like (179b) is largely ameliorated through contrastive focus, as in (180b). Now recall that the Stripping examples in Experiment 5 were all instances of corrective Stripping. Thus, whatever the source is of the effect seen in (179b), it should be neutralized in the instances of BCC sensitive Stripping of Experiment 5, just as it is neutralized in (180b). Consequently, it is unlikely that the sensitivity of Stripping to BCC seen here can be attributed to whatever causes the degredation in (179b).

(180) a. First, he$_i$ walked in. And then Mary$_i$ sneezed.

b. No, and then *James*$_i$ sneezed.

Notice also that the the degredation in (179b) seems to be much weaker than that of a typical BCC violation, as in (178b). If the effect in the Stripping conditions in Experiment 5 were due to whatever causes the degredation in (179b), we would expect the Stripping examples like (177b) to be rated better than (178b). This expectation is not borne out however, as the Pronoun-Name Stripping conditions were rated as implausible as the Pronoun-Name canonical conditions.

Finally, if source of the unacceptability of Stripping examples like (177b) was entirely due to what is causing the unacceptability of (179b), then the effect should persist even in continuations which do not contain an instance of the relevant pronoun in a BCC violating configuration. For example, consider the continuation in (179b). Such is the overt counterpart to a short-source non-isomorphic construal of a Stripping ellipsis site. Notice that it also embodies the same basic configuration as that in (179b). Thus, if the source of the unacceptability of the Stripping configuration in (177b) were simply whatever causes the effect in (179b), then it would be possible for the ellipsis site to be resolved with a non-isomorphic construal, the elided equivalent to (181b), while yielding the same effect. We would then expect non-elided instances of such short source resolutions, as in (181b) to also yield the same degredation. This prediction also does not seem to be borne out.

(181) a. He$_i$ admires the professor who supervises *Mary*.

b. No, the professor (that he$_i$ admires) supervises *Bill*$_i$.

Taking these three arguments together, it seems that the results of Experiment 5 can not be attributed to the degredation that results from cross-sentential precedence relations between a pronoun and a name, of the sort seen in (179b). Thus the major conclusion from the results of Experiment 5 still stands: the ellipsis site is populated with structure isomorphic to the antecedent.

However, one might still wonder whether the effect observed is the result of a C-command relationship between the base position of the remnant and the pronoun within the ellipsis site, or whether the effect can be attributed to the precedence relationship between these items within the ellipsis site itself. In the Pronoun-Name conditions, as in (182b), the pronoun both precedes and C-commands the base position of the remnant, while in the Name-Pronoun conditions, as in (183b), the pronoun is the remnant, and so neither precedes nor C-commands the name from its base position. Thus it is in principle possible that the contrast between (182b) and (183b) is due to a precedence relation instead of a C-command relation. Examples like (184) allow us to distinguish between these alternatives. Here, the pronoun precedes, but does not C-command, the base position of the remnant. Consequently, if precedence relations between pronoun and name are the relevant factor for the unacceptability of (182a), examples like (184) should show the same unacceptability. Or, if C-command were the principal factor, (184) should be more acceptable than (182a). Experiment 6 directly compares examples like (182) against examples like (184) to test these predictions.

(182) a. He$_i$ admires the professor who supervises *Mary*.

b. *No, *James*$_i$, ~~he$_i$ admires the professor who supervises e~~.

(183) a. James$_i$ admires the professor who supervises *her*.

b. No, *him*$_i$, ~~James$_i$ admires the professor who supervises e~~.

(184) a. His$_i$ colleague admires the professor who supervises *Mary*.

b. No, *James*$_i$, ~~his$_i$ colleague admires the professor who supervises e~~.

There is reason to think that the contrast between (182) and (183) is due to the C-command relationships between the pronoun and name, rather than their precedence relationships. First, recall that the relevant judgments in the literature explicitly control for the C-command relation between R-expression and its binder. For example, Hunter and Yoshida (2016) present the following minimal pairs, discussed above in (157) and (158), and repeated here as (185) and (186). If the reduced plausibility in the Pronoun-Name conditions above were solely due to the pronoun preceding the co-indexed R-expression, then we would expect examples like (186) to be as unacceptable, under the given co-indexation, as (185), contrary to the judgments in the literature.

(185) a. He$_i$ said that Mary$_i$ left.

b. *Yeah, but not that John$_i$ left.

(186) a. His$_i$ friends said that Mary left.

b. Yeah, but not that John$_i$ left.

Second, in their off-line plausibility rating study, the methodology of which our Experiment 5 utilizes, Kazanina et al. (2007) found co-referentiality between R-expressions and pronouns which preceded but did not C-command them, e.g. (187a), to be more plausible than co-referentiality between R-expressions and pronouns which both precede and C-command the names, e.g. (187b). This data corroborates the judgments in the literature concerning the nature of BCC: linearly preceding but non-C-commanding pronouns do not trigger BCC violations, e.g. (Chomsky, 1981). Taken together, we would expect that a comparison will reveal examples like (182) to be less plausible under the indicated co-reference relations than examples like (184).

(187) a. Because last semester she$_i$ was taking classes full-time while Kathryn was working two jobs to pay the bills, Erica$_i$ felt guilty.

b. Because last semester while she$_i$ was taking classes full-time Kathryn$_i$ was working two jobs to pay the bills, Russell never got to see her.

3.3. Experiment 6

3.3.1. Overview

Experiment 6 tested the whether the implausibility of coreference relations in examples like (188), as observed in Experiment 5, could be attributed to the C-command relationship between the base position of the remnant and the matrix subject within the ellipsis, or the precedence relationship between them. To do so, we compared the plausibility of the coreference relations in such examples against that in examples like 189, for which the pronoun precedes, but does not C-command, the base position of the remnant within the ellipsis site.

(188) a. He$_i$ admired the professor who supervises *Mary*.

b. *No, *James*$_i$, ~~he$_i$ admired the professor who supervises e~~.

(189) a. One of his$_i$ colleagues admired the professor who supervises *Mary*.

b. No, *James*$_i$, ~~one of his$_i$ colleagues admired the professor who supervises e~~.

The results indicate that the coreference relations in examples like (188) were reliably less plausible than those in (189), indicating that the effect observed in Experiment 5 cannot be exclusively attributed to the precedence relations between the pronoun and base position of the remnant.

3.3.2. Participants

57 participants were recruited through Amazon's Mechanical Turk[8]. Participants were limited to IP addresses within the US, were only permitted to participate in the experiment once and were compensated $2 USD. Participants were asked whether they were native speakers of English, and all replied that they were. Twelve participants failed to complete the survey, and so their results were excluded from analysis. The data from eleven additional subjects was excluded, on the grounds that their mean ratings of the low plausibility

[8]An online crowdsourcing marketplace platform (www.mturk.com)

fillers was not lower than that of the high plausibility fillers. Taking these exclusions into account, the data of 36 participants were analyzed.

3.3.3. Design, Stimuli, Procedures

The stimuli were presented to participants as two sentence dialogues, in which the second sentence, preceded by ‘No, . . . ’, was intended to be understood as a corrective response to the first sentence. In all conditions, a prepositional phrase found in an embedded clause in the first dialogue sentence was italicized; we call this the Correlate. The second dialogue sentence also contained an italicized prepositional phrase, which we call the Focus. The Focus was intended to be understood as the correction. Across conditions, the Correlate and Focus were always underlined. Also underlined were the pronouns that were, or were contained within, matrix subjects of the first dialogue sentence and, in the canonical conditions, those of the second dialogue sentence. All of the first dialogue sentences, as well as all of the second sentences of the canonical conditions, contained definite relative clause islands, within which were found the Focus and/or Correlate.

The stimuli conformed to a 2x3 factorial design: i. C-Command Status; ii Construction Type. The C-Command Status factor manipulated the distribution of pronouns and proper names in the underlined phrases in the dialogues. In the ‘C-Command’ conditions, the matrix subjects of the first dialogue, and also the matrix subject of the second sentence in the canonical conditions, was a pronoun, gender matched to the Focus, which C-commanded the Correlate. In the ‘Non-C-Command’ conditions, the matrix subject of the antecedent contained a pronoun which did not C-command the Correlate, as it was a possessive pronoun within a complex DP, such as *one of her friends*. In addition to *one of* . . . , the other quantifiers used were *all, none, few, each, most, half*.

The Construction Type factor manipulated whether the second dialogue sentence was a corrective Stripping construction with a PP remnant (=PP Stripping), a corrective DP remnant Stripping construction (=DP Stripping), or a non-elliptical canonical word-ordered construction (=Canonical). In the Stripping conditions, the remnant was the Focus. In the Canonical conditions, the first and second dialogue sentences were identical, except for the Correlate and Focus.

Table 3.6. Experiment 6: Factors, Example Stimuli

Construction Type		C-Command Status: C-Command	C-Command Status: Non-C-Command
Stripping	DP Remnant	Example (190a-i)	Example (190a-ii)
	PP Remnant	Example (190b-i)	Example (190b-ii)
Canonical		Example (191a)	Example (191b)

(190) Stripping

a. DP Remnant

i. C-Command:

Joe: <u>She</u> liked the manager who assigned the job *to Bill*.

Bill: No, *<u>Mary</u>*.

ii. Non-C-Command:

Joe: One of <u>her</u> friends liked the manager who assigned the job *to Bill*.

Bill: No, *<u>Mary</u>*.

b. PP Remnant

i. C-Command:

Joe: She liked the manager who assigned the job *to Bill*.

Bill: No, *to Mary*.

ii. Non-C-Command:

Joe: One of her friends liked the manager who assigned the job *to Bill*.

Bill: No, *to Mary*.

(191) Canonical

a. C-Command

Joe: She liked the manager who assigned the job *to Bill*.

Bill: No, she liked the manager who assigned the job *to Mary*.

b. Non-C-Command:

Joe: One of her friends liked the manager who assigned the job *to Bill*.

Bill: No, one of her friends liked the manager who assigned the job *to Mary*.

Forty eight lexicalizations of the six conditions were constructed. The gender of the names and pronouns used in the C-Command Status manipulation were balanced across the lexicalizations.

80 filler sentences were constructed to contain a range of grammatical constructions, yielding nearly a 2:1 filler to target item ratio. Each filler item contained two or more underlined phrases. 30 filler sentences contained underlined phrases, which, if understood to refer to the same person, would result in either a Binding Condition A violation or a Binding Condition C violation. The remaining 50 fillers contained underlined phrases, which, if understood to refer to the same person, would not result in any Binding Condition

violation. However, 10 of these 50 fillers contained some other defect that would result in reduced acceptability.

Six lists were constructed using a Latin square design, each containing seven distinct lexicalizations of each of the six conditions. Each list began with six practice items in a variety of grammatical configurations, each containing underlined words. One practice item would yield a Binding Condition A violation and another a Binding Condition B violation, if the underlined words in these items were understood to refer to the same person. The practice items did not vary in lexicalization or order between lists, and were not marked explicitly as practice items. In sum, each list contained 48 test items, 80 filler items, and 6 practice items, for a total of 134 items. Each of the six lists was pseudo-randomized to ensure that sequential items were not from the same condition.

The stimuli were presented to participants in the form of two-turn dialogues between 'Joe' and 'Bill'. Following the procedure of Kazanina et al. (2007), participants were instructed to read the presented dialogues and to rate how plausible it would be for the underlined phrases in the dialogue to refer to the same person, on a scale of 1 (implausible) to 7 (plausible). The entire dialogue and the rating scale appeared on the same screen. Complete instructions are included in the appendices.

3.3.4. Predictions

If the effect observed in Experiment 5 was due exclusively to the precedence relations between the pronoun and base position of the Focus, then we would expect no difference between the C-Command and Non-C-Command conditions, in the Stripping and Canonical conditions. However, if that result were due to the C-Command relationship between pronoun and base position of the Focus, then we would expect the C-Command conditions to be rated as less plausible than the Non-C-Command conditions. For the Canonical conditions, this is true regardless of what analysis of ellipsis is adopted.

For the Stripping conditions, the predictions of relative plausibility depend on the analysis of ellipsis adopted. Isomorphic structural analyses of ellipsis predict that the Stripping conditions should pattern with the Canonical conditions, as the ellipsis site is populated with structure that is isomorphic to that of the antecedent. This structural isomorphism would result in the same effect of C-Command Status both Stripping and Canonical conditions, as their structures would be isomorphic to each other.

For the non-isomorphic and non-structural analyses of ellipsis, the predictions vary with the hypothesized distribution of non-isomorphic construals, and the effect of C-Command Status in the Canonical conditions. If there is no effect of C-Command Status in the Canonical conditions, then no effect is predicted under any version of a non-isomorphic analysis of ellipsis. This follows because no C-Command Status effect would be observed even where the full structure is present, as in the Canonical conditions, and so the lack of a C-Command Status effect cannot be used to diagnose the presence of isomorphic or non-isomorphic structure within the ellipsis site.

However, if there were an effect of C-Command Status in the canonical conditions, then the non-isomorphic analyses would predict that there should be no effect of C-Command Status for either or both DP and PP Stripping conditions. It is conceivable that the availability of a non-isomorphic construal of the ellipsis site is conditioned on the type of Stripping remnant. For example, it could be that a non-isomorphic construal is possible only in case the remnant is a DP, and not if it is a PP. Let us call these *selectvely non-isomorphic analyses*. If so, we would expect no effect of C-Command Status in the DP Stripping conditions, as a non-isomorphic construal is possible in these cases, yielding insufficient structure within the ellipsis site to result in an effect of C-Command Status. Then, if there were an effect of C-Command Status in the Canonical conditions, we would still expect an effect of C-Command Status in the PP Stripping conditions, as, such selectively non-isomorphic analyses posit no non-isomorphic structure to be available in case of a PP remnant. Figure (3.6) illustrates the relevant contrasts that a selectively non-isomorphic analysis would predict.

Under a *consistently non-isomorphic analysis*, non-isomorphic structures are available for both PP and DP remnant Stripping structures. Consequently, if we were to see an effect of C-Command Status in the Canonical conditions, we would expect to see no such effect for either DP or PP Stripping conditions. Such Stripping conditions would permit a non-isomorphic construal, and so would lack the structure necessary to yield the effect of C-Command Status. Figure (3.7) illustrates the relevant contrasts that a consistently non-isomorphic analysis would predict.

3.3.5. Results

The data was analyzed with a logistic linear mixed-effects regression model (LMER; Baayen et al. 2008), using lme4 in R, with rating as the dependent variable. The fixed

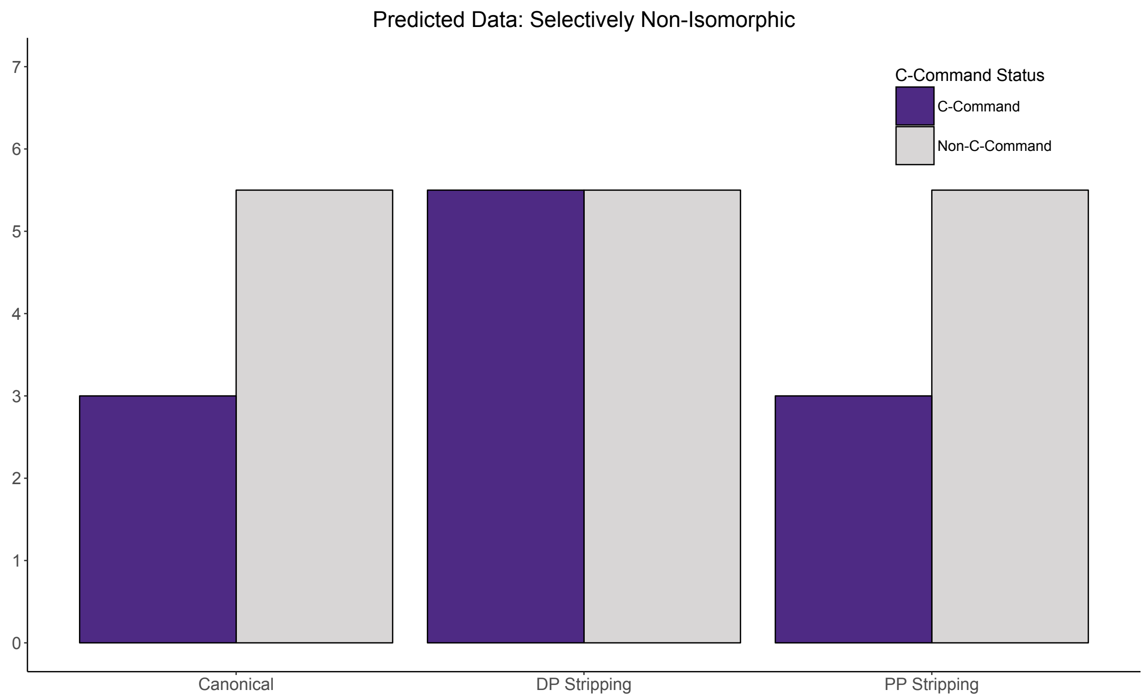

Figure 3.6. Experiment 6: Predicted Data, Limited Non-Isomorphic Analyses

effect of C-Command Status (C-Command, Non-C-Command) was contrast coded, while the three level factor of Construction Type (DP Stripping, PP Stripping, Canonical) was Helmert coded. The values used for the Helmert coding are given in table (**??**). The Helmert coding scheme allowed us to compare the Canonical condition against both Stripping conditions together and to compare just the DP Stripping and PP Stripping conditions. The maximal random effects structure that would converge was employed, which included random intercepts for Participant and Item, as well as random slopes by participant for Islandhood, Pronominal Status, and Construction Type, and the 2- and 3- way interactions. Model comparisons were performed to determine whether the inclusion of each of these fixed effects and their interactions made a significant contribution to the model.

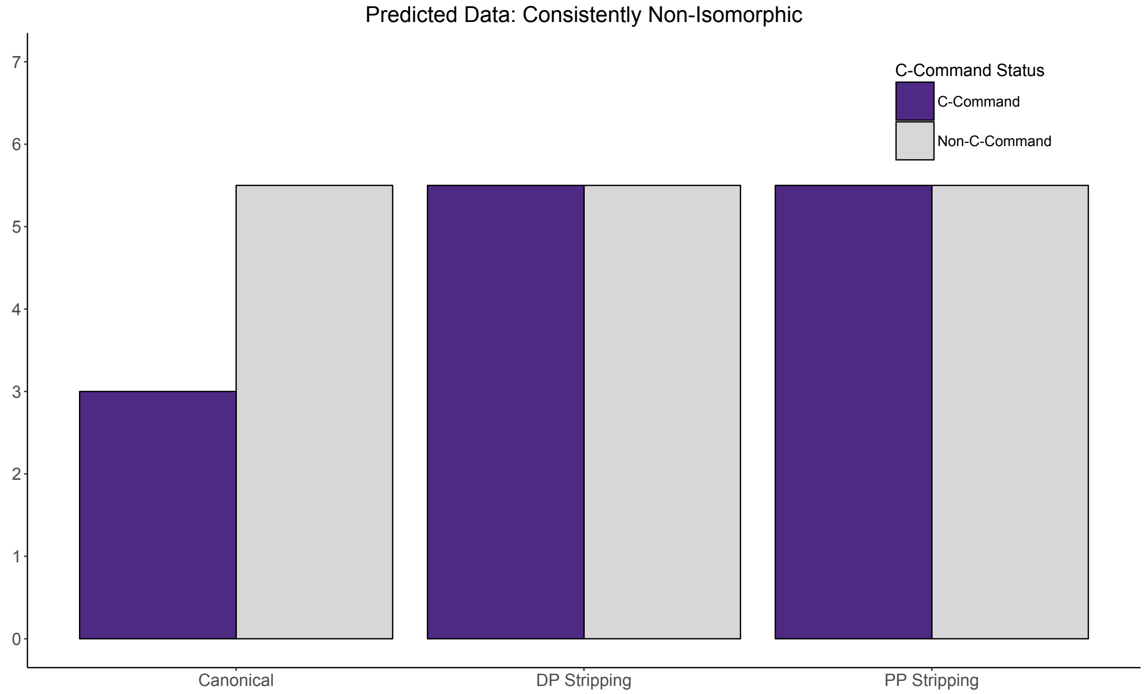

Figure 3.7. Experiment 6: Predicted Data, Full Non-Isomorphic Analyses

Table 3.7. Experiment 6: Helmert Coding Values

Construction Type	Stripping vs. Canonical	DP Stripping vs. PP Stripping
DP Stripping	-.25	.5
PP Stripping	-.25	-.5
Canonical	.5	0

Table 3.8. Experiment 6: Mean Plausibility Results, with standard errors in parentheses

Construction Type		Pronominal Status C-Command	Pronominal Status Non-C-Command
Stripping	DP Remnant	2.25 (0.10)	2.55 (0.10)
	PP Remnant	2.32 (0.10)	2.45 (0.10)
Canonical		3.70 (0.14)	3.98 (0.14)

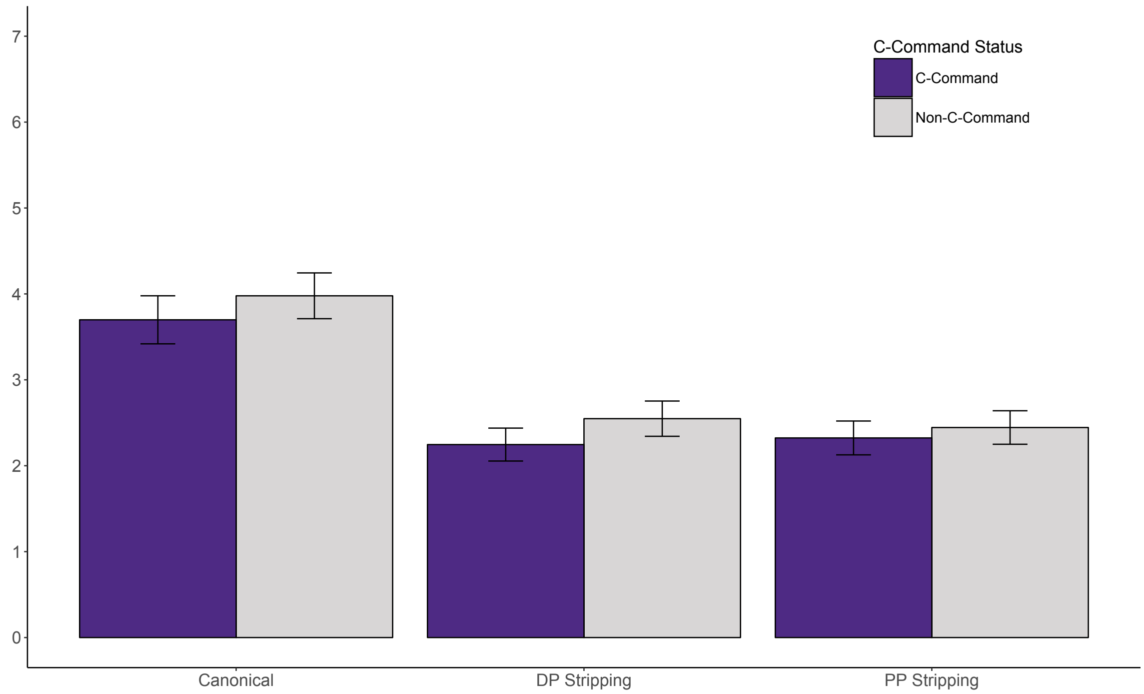

Figure 3.8. Experiment 6: Data

The results of these analyses revealed significant main effects of C-Command Status (β=0.23, SE β=0.08, χ^2(1)=7.66 p=0.006), and Canonical vs. Stripping Construction type (β=1.93, SE β=0.41, χ^2(1)=17.34, p<.001), but no effect of DP vs. PP Stripping Construction type (β=0.01, SE β=0.08, χ^2(1)=0.02 p=0.895). Additionally, neither the 2-way interaction between C-Command Status and Canonical vs. Stripping, nor that between C-Command Status and DP vs. PP Stripping were significant (χ^2(1)=1.37 p=0.242).

In summary, across the Canonical, DP Stripping and PP Stripping conditions, a coreference relation between a pronoun and a name was rated to be less plausible when the R-expression was C-commanded by the pronoun than when the pronoun preceded but did not C-command the name. The lack of interaction between C-Command Status and

Table 3.9. Experiment 6: Estimates of fixed effects

	Coefficient	StandardError	*t*	*p*
Intercept	2.87	0.21	13.89	*n/a*
C-Command Status (CS)	0.23	0.08	2.89	0.006
Canonical vs. Stripping Construction (CvS)	1.93	0.41	4.73	$<$0.001
DP Stripping vs. PP Stripping (DPvPP)	0.01	0.08	0.13	=0.895
CS:CvS	0.08	0.21	0.39	=0.697
CS:DPvPP	0.18	0.15	1.17	=0.242

construction type indicates that the magnitude of the effect of C-command relationship between name and pronoun did not differ between canonical, DP Stripping or PP Stripping continuations.

3.3.6. Discussion

Experiment 6 was designed to determine whether what we identified as a BCC effect in Experiment 5 could be entirely attributed to the precedence relations between the pronoun and the base position of the Focus. Experiment 5 confounded precedence and C-Command, such that it was implausible for pronouns which both preceded and C-commanded the correlate and the base position of the focus to refer to the same person. Experiment 6 eliminated this confound; in every condition the relevant pronoun preceded the correlate and focus, but only in the C-Command conditions did the pronoun C-command the correlate and focus.

The results of Experiment 6 indicate that the effects of Experiment 5 cannot be entirely attributed to the precedence relation between the pronoun and the name. Across Canonical, DP Stripping, and PP Stripping conditions, the plausibility of a pronoun and proper name was lower when the pronoun C-commanded the (base position of the) proper name, than when there was no such C-command relationship. Thus, the effect found in Experiment which we attributed to BCC cannot be completely explained in terms of precedence relations; the C-command relationship between a pronoun and a name influences whether those items can be construed as referring to the same person. Further, no differences in the magnitude of C-Command effect between the various construction types was found. These results confirm the predictions of the Isomorphic Structural analyses of ellipsis, and were not predicted by any non-structural or non-isomorphic analysis.

Isomorphic structural analyses of ellipsis posit structure within the ellipsis site that is isomorphic to that of the antecedent. Therefore, the prediction for the present experiment was that the Stripping conditions should pattern with the Canonical conditions. The Canonical conditions exhibited a C-Command Status effect, which was likewise found in both Stripping conditions. This suggests that for both DP and PP Stripping conditions, the ellipsis site is construed with material isomorphic to the antecedent.

The predictions of neither the limited non-isomorphic nor the full non-isomorphic analyses of ellipsis were born out. Given that an effect of C-Command Status was found in the Canonical conditions, the full non-isomorphic analyses would predict that no such effect of C-Command Status should be found for either of the Stripping conditions, as an alternative non-isomorphic construal is possible for such configurations, which would thereby yield no effect of C-Command Status. The limited non-isomorphic analyses predict that the DP and PP Stripping conditions should diverge with respect to the effect of C-Command Status; because non-isomorphic parses are only possible for DP Stripping under these limited non-isomorphic analyses, DP Stripping should yield of effect of C-Command Status, while PP Stripping should exhibit such an effect, in parallel to the Canonical conditions.

While the effect of C-Command Status was reliable across all three construction types, the magnitude of the effect was much smaller than that observed between Pronoun-Name and Name-Pronoun conditions in Experiment 5. One plausible way to interpret this relatively small effect is that that, while the C-command relationship between a pronoun and a name matters for establishing the plausibility of co-reference between these elements, that precedence relationship also play a significant role in these judgments. However, we'd like to highlight two additional factors that may have played a role in reducing the acceptability of the Non-Command conditions, and thereby reducing the size of the observed effect of C-Command Status.

The first factor is complexity. The pronouns in the C-Command conditions were bare nominative pronouns, while the pronouns in the Non-C-Command conditions were in genitive case and embedded within a complex DP. Thus the sentences in the Non-C-Command conditions were more complex than those in the C-Command conditions, in virtue of these more complex subjects. The paradigm in (192) largely eliminates any differences in complexity between an example in which the pronoun precedes and C-commands the proper name, as in (192a), and an example in which the pronoun precedes but does not C-Command the proper name, as in (192b). Thus, if the relatively small effect size were due to complexity, we would expect the difference between examples like those in (192) to be greater than seen in Experiment 6.

(192) a. While he$_i$ was at school, Mary met the girl who likes John$_i$?s friend.

b. While Mary was at school, he$_i$ met the girl who likes John$_i$?s friend.

Notice, however, that there is a dependency length asymmetry introduced, however. In (192a) the length between the pronoun and name is greater than in (192b). Longer *wh*-dependencies result in lowered acceptability (Sprouse et al., 2012). Let's assume that BCC, as a grammatical constraint, is also more difficult to evaluate over longer distances, such that longer distances between a pronoun and name result in a greater plausibility that the pronoun and name refer to the same person. The intuition, then, is that BCC becomes harder to evaluate at longer distances.

Consequently, comparing (192a) and (192b) provides a strong test for the hypothesis that the small effect size seen in Experiment 6 is due to complexity. Both examples in (192) are approximately equally complex, while the length between pronoun and name in (192a) is longer. If the small effect size in Experiment 6 were due to differences in complexity, we would expect the difference between examples like those in (192) to be greater. The size of this effect could be expected to be somewhat attenuated, however, by the additional dependency length in example (192a). However, if the small effect size in Experiment 6 were do to some other factor, we might expect that small effect to be entirely washed out by the dependency length difference in those in (192).

A second factor which might have contributed to the small effect size in Experiment 6 is Weak Cross Over (WCO). WCO typically describes a reduction in acceptability for examples like those in (193), in which a quantificational element moves, covertly or overtly, past an element that would be bound by that quantificational element (Chomsky, 1977; Lasnik and Stowell, 1991; Postal, 1993). Example (193a) illustrates this for QR, and example (193) for overt *wh*-movement.

(193) a. *His$_i$ brother likes every$_i$ student.

b. *Whose father does his$_i$ son despises?

Chomsky (1976) claimed that in-situ focused elements show WCO effects. Thus example (194b), in which *John* is focused, is reported as being less acceptable than example (194a), in which the verb is focused[9]. Chomsky argues that this is due to covert movement of the focus to the left periphery of the clause, yielding a WCO configuration.

(194) a. His$_i$ mother SHOT John.

b. ??His$_i$ mother shot JOHN.

If in-situ focused elements and remnants in Stripping configurations do involve a WCO configuration, then the relatively small effect size observed in Experiment 6 may be due to the resulting WCO effect. The plausibility of the Non-C-Command conditions would be reduced, as this reading of the examples would involve a WCO violation, thereby reducing the magnitude of the difference between the C-Command and Non-C-Command conditions.

[9]Note that the validity of this generalization is in question; there is experimental evidence that in-situ focus does not induce WCO effects (Moulton et al., 2017)

An effect of WCO could be controlled for by embedding the name within a DP, as in (195). In such a configuration, the name, contained within the remnant raised to the left periphery, would no longer C-command the pronoun, and so would not instantiate a WCO configuration.

(195) a. Joe: <u>She</u> liked the manager who assigned the job *to Bill's employee*.

b. Bill: No, <u>*Mary's employee*</u>.

3.4. Conclusion

This chapter presented the results of two plausibility judgment experiments, designed to probe the amount of structure contained within Stripping ellipsis sites. Experiment 5 showed that Stripping configurations are sensitive to Binding Condition C. For examples in which the remnant contains a name and the correlate is C-commanded by a pronoun, co-reference between the name and pronoun was judged to be as implausible as in comparable non-elliptical controls. Furthermore, this pattern held of both non-island and island violating Stripping. This sensitivity to BCC is predicted by isomorphic analyses of ellipsis, because it is in these analyses that the ellipsis site is populated with enough structure to recover the pronoun that leads to the BCC violation. Non-isomorphic and non-structural analyses predict no such sensitivity, as the ellipsis site in these analyses does not include enough material to induce a BCC violation.

Experiment 6 examined whether the BCC effect found in Experiment 5 could be exclusively attributed to the precedence relation between the pronoun and name, or whether the C-command relationship between these items also influenced the plausibility of co-reference between the pronoun and name. The results indicated that, when the precedence relation between pronoun and name was controlled for, the possibility of co-reference in conditions in which the pronoun C-commanded the base position of the name were still rated as less plausible than conditions without such a C-command relationship between these items.

Taken together, these results indicate that the ellipsis site in Stripping configurations is construed so as to contain syntactic structure isomorphic to that of the antecedent. Importantly, this holds true even in instances of island-violating Stripping, configurations which were taken to be prime candidates for non-isomorphic ellipsis site resolutions. As I will discuss next in Chapter 4, this means that the island insensitivity of Stripping cannot be attributed to a non-isomorphic construal of the ellipsis site.

CHAPTER 4

An Overt Move and Elide Approach to Stripping

4.1. Approaches to Ellipsis

In his survey of elliptical phenomena, Merchant 2016 breaks the analysis of ellipsis into three questions. The first, and the principal subject of this dissertation, is "In elliptical constructions, is there syntactic structure that is unpronounced?" The second question is the *identity* question: "What is the relationship between the understood material in ellipsis and its antecedent?" How an analysis would answer the structure question impacts what sorts of answers to the identity question are possible. Given that I am investigating various alternatives to the structure question, I largely set aside this question of identity. The third question is the *licensing* question: "What heads or positions or structures allow for 'ellipsis', and what are the locality conditions on the relation between these structures and ellipsis?" Although, as Merchant notes, this third question "has not attracted quite the attention" the others have received, I will also set this question aside (Merchant, 2016, pg. 3).

Approaches to ellipsis fall into two basic categories, with respect to what structure, if any, is contained within the ellipsis site. As the name implies, non-structural approaches posit no structure at all within the ellipsis site. The ellipsis remnants either exhibit the external syntax canonically associated with fragments of that category, or they project from their canonical category directly to the terminological equivalent of an S node. Thus, in a non-structural analysis, a contrastive Stripping example like that in (196b) would have a parse like that in (197a).

(196) a. James met the student who speaks *Thai*.

b. No, *Mandarin*.

Structural approaches, predictably, propose that some sort of structure is contained within the ellipsis site, though they vary as to how similar this structure is to that of non-elided structure and the antecedent structure. One prominent type of structural analysis posits that the ellipsis site contains syntactic structure more or less isomorphic to that of the antecedent, modulo the movement of the remnant from the constituent to be elided and, of course, ellipsis, or non-pronunciation, of that constituent. Such an analysis is depicted in (197b).

(197) a. No, [$_{\{S|DP\}}$ Mandarin]

b. No, [Mandarin$_i$ ~~James met the student who speaks t$_i$~~]

In this chapter, I survey the range of approaches to Ellipsis and its island insensitivity in light of the evidence from Chapters 2 and 3. As we will see, the insensitivity of Stripping to definite relative clauses, as presented in Chapter 2, along with the Condition C effects observed in Chapter 3, in conjunction, are problematic for all but the structural isomorphic theories of Ellipsis, a conclusion which implicates the existence of an amelioration effect of ellipsis on island violating movement.

I then examine in detail two approaches to these theories: a PF movement account, in which the remnant escapes the ellipsis site only at PF, and a narrow syntax movement account, in which the remnant moves overtly, prior to Spellout. For both conceptual and empirical reasons, including evidence from the comparison of remnant movement with other instances of overt and covert movement with respect to intermediate binding effects, I conclude that remnants move prior to spellout, but that, due to the nature of ellipsis, this overt movement is not always successive-cyclic.

Before we dive into a critical analysis of the available analyses of Ellipsis, let us recall briefly the results of Chapters 2 and 3. Chapter 2 explored the island sensitivity of a variety of Stripping configurations. Across a series of four experiments, we found that non-contrastive Stripping is completely insensitive to definite relative clause islands. Two varieties of contrastive Stripping, Corrective and Elaborative, were found to be partially sensitive, facts which I argued to arise as the result from a biased parser and subsequent reanalysis of the antecedent. Chapter 3 explored the interaction of Stripping, in island and non-island contexts, with potential Binding Condition C configurations. There we found Stripping to be sensitive to BCC, when the relevant R expression originated within a definite relative clause is or within a complement clause, and, crucially, when a pronoun C-commanded that originating position from outside the island.

Finally, I will reiterate here that I use the phrase *island* as shorthand in this chapter for one particular variety of island, that investigated throughout this dissertation, the definite relative clause island. There are a host of types of islands, and they have received a great many analyses throughout the literature. For a broad overview, see Den Dikken and

Szabolcsi (2002b). I have thus far focused on definite relative clause islands, and will continue to do so in this chapter, not because of any commitment to the idea that the analysis presented here of Stripping out of definite relative clause islands will generalize across to all islands. Rather, my intention here is to determine where there is at least one case of island insensitive ellipsis for which there is no viable analysis other than island repair, and to examine what properties such an island repair analysis must exhibit.

4.2. Non-Structural Approaches to Ellipsis

Non-Structural analyses of Ellipsis have no problem accounting for the island insensitivity of Stripping. As there is no structure in the ellipsis site, there is, correctly, predicted to be no effect of moving a remnant out of an island, as there is no island to move out of within the ellipsis site. The same lack of structural within the ellipsis site serves to make an incorrect prediction for the sensitivity of Stripping to Condition C. If there were no structure in the Stripping continuation in an example like (198), the fact that a co-indexed pronoun C-commanding the correlate should have no effect on the acceptability of the Stripping continuation. We saw in Chapter 3 that this prediction was not born out, that the co-indexation in examples like (198) was rated to be as implausible as in non-elliptical controls like (199). We therefore take the non-structural approaches to be empirically inadequate.

(198) a. She$_i$ appreciates [$_{Island}$ the student who respects Bill].

b. No, Mary$_i$.

(199) a. She$_i$ appreciates the student who respects Bill.

b. No, she$_i$ appreciates the student who respects Mary$_i$.

4.3. Structural Approaches to Ellipsis

Structural approaches vary vastly in terms of the amount of structure within the ellipsis site, how that structure is generated, and how the remnant escapes the elided material. Here in turn, I will discuss pronominal, non-isomorphic (evasion), LF-copying, and isomorphic approaches. Ultimately, I argue that it is the isomorphic approaches that best explain the island insensitivity and Binding Condition C sensitivity of Stripping.

4.3.1. Pronominal Approaches

The least structurally complex structural approach to ellipsis posit a null pronominal element within the ellipsis site above which the remnant has been base generated (Barker 2013; Chao 1988; Lobeck 1995, among others). The pronominal element is then assumed to be interpreted anaphorically to the antecedent phrase.

Such impoverished structure within the ellipsis site yields similar results for these pronominal accounts as we found for the non-structural accounts in explaining the island insensitivity of Stripping and sensitivity to Binding Condition C. As Barker (2013) points out, under such an analysis, the remnant hasn't raised out of an island, given that it hasn't moved at all and that there is no island contained within the ellipsis site. So long as the correlate itself can take matrix scope out of whatever element, island or otherwise, it is contained within, a Stripping continuation should be acceptable. Of course, the impoverished structure within the ellipsis site also means that three is no structure within the ellipsis site that would be expected to induce a BCC violation. If the structure within the ellipsis site is simply a pronominal, as in (200), the remnant wouldn't be, at any point in its derivation, C-commanded by a co-indexed pronoun. Consequently, and incorrectly, under a pronominal analysis of Stripping, we would not expect such configurations to yield BCC effects.

(200) a. She$_i$ appreciates [$_{Island}$ the student who respects Bill].

b. No, [Mary$_i$ e]

4.3.2. Non-Isomorphic Approaches

Non-Isomorphic approaches to ellipsis posit somewhat more structure within the ellipsis site, little enough that they need not appeal to island repair, but still not enough to account for the sensitivity to Binding Condition C. A non-isomorphic construal of an ellipsis site involves structure that is non-isomorphic to the antecedent site, such as the short source construal in (201b), or the cleft construal in (201c) (Merchant 2001, Fukaya 2007, Barros et al. 2013). In such instances, the remnant doesn't escape an island, since none is posited within the ellipsis site, but they also cannot account for the presence of BCC effects, as there is no pronoun within the ellipsis site which would C-command the base position of the R expression remnant either. Thus the bargain made by such approaches, avoiding the need to posit the notion of island repair by eliminating the islands within ellipsis sites, sabotages any possible explanation of BCC effects in these cases.

(201) a. She$_i$ appreciates [$_{Island}$ the student who respects Bill].

b. No, Mary$_i$ the student respects e.

c. No, Mary$_i$ it was e.

4.3.3. LF-Copying Approaches

Now, let us turn to analyses of Ellipsis in which the ellipsis site is populated by material that is more or less isomorphic to the antecedent. Such analyses come in two basic flavors: one in which the ellipsis site is syntactically identical to the antecedent having been copied into the ellipsis site at LF (Chung et al., 1995), and one in which the material within the ellipsis site is generated there, but subject to one or more of a variety of conditions specifying how and in which ways that material must be similar to that of the antecedent (Ross, 1967; Merchant, 2001).

Let us first consider LF copying accounts. Here, the remnant is merged high, above the ellipsis site, which remains empty, until, at LF, a suitable antecedent available in the context, is copied in. The analysis proposed in (Chung et al., 1995) was intended to apply to Sluicing, and would yield an LF along the lines illustrated in (202). The basic idea here is that indefinites are interpreted in such a way as to permit them to be bound by an ellipsis remnant, and that, because the remnant is base merged above the copied LF and not moved from within it, no island violation is expected.

(202) a. James appreciates [$_{Island}$ the student who respects someone].

b. Who [James appreciates [$_{Island}$ the student who respects someone]]?

Such an analysis seems more natural for remnants that have the inherent quantificational force that *wh*-phrases do. One could imagine extending such an analysis to instances of non-contrastive Stripping, if we assume that their interpretation as a focus yields the quantificational force necessary for the remnant, which itself is not necessarily quantificational, to bind the indefinite contained within the copied FL, as illustrated in (203).

(203) a. James appreciates [$_{Island}$ the student who respects someone].

b. Yeah, Mary$_F$ [James appreciates [$_{Island}$ the student who respects someone]].

Extending such an analysis to instances of contrastive Stripping is more problematic, as Romero (1998) and Merchant (2001) have observed in the context of contrastive Sluicing. Consider (204). Here the correlate, not being an indefinite at all, doesn't naturally provide the type of variable which could be bound by a quantifier, under the assumption that it is the indefiniteness of the correlate in examples like (202) or (203) which allows them to be bound.

(204) a. James appreciates [$_{Island}$ the student who respects Bill].

b. No, Mary$_F$ [James appreciates [$_{Island}$ the student who respects Bill]].

One could imagine a variant of the LF copying analysis, modified so as to provide a bindable variable even in the case of contrastive Ellipsis. Imagine that, within the antecedent phrase, contrastive correlates undergo LF movement to a position within the left periphery, such as a focus phrase. If we assume a trace theory of movement, this movement would leave a bindable trace within the TP that is to be copied into the ellipsis site. Then it seems plausible to assume that the remnant, base generated above the TP and endowed with the quantificational force of a focus, would be able to bind this trace, providing the correct interpretation. Because the intent behind this LF copying variant is to derive a bindable variable within the antecedent, the prominent alternative approach to focus scoping, Rooth's Alternative Semantics (Rooth, 1996), involves no movement of the focused element, thereby leaving no bindable trace, and so is unhelpful for the current task.

Such a variant would derive the island insensitivity of both contrastive and non-contrastive Stripping, but at a price. In the non-contrastive cases, we avoid any island insensitive movement. In the antecedent, the correlate scopes wide, being bound by an operator in the left periphery. The remnant is base generated in the left periphery, and so likewise is involved in no island violating movement. In the contrastive cases, however, we needed to assume that the correlate moved to the left periphery, so as to leave a bindable trace within the TP which would be copied into the ellipsis site. However, if the base position of the correlate were contained within an island, this would require island violating movement of the correlate. And so extending the LF copying analysis to both contrastive and non-contrastive Stripping can be done but only at the cost of admitting island violating movement into the grammar, thereby abandoning one of the main selling points of the LF copying analysis.

The above modified LF copying approach relies on a trace theory of movement to provide a bindable variable within the copied TP. It is not clear how such a story would work under a copy theory of movement (Chomsky, 1995). Let us assume that lower copies in a chain can be converted by the interpretive component to variables, for example as in (Fox, 1999, 2002). One might expect that the reused TP, copied into the ellipsis site, would then contain a bindable variable. However, this requires that two distinct lexical items, the Stripping remnant and the chain tail contained in the copied TP, be treated as members of a single chain, so that the lower copy of the remnant within the copied TP can be converted to a bindable variable. It is not clear that this should be possible.

Consider an example like (205). Let us assume that, in the continuation in (205b), *Bill* is base generated high, in the specifier of a focus phrase, which we assume to be the target of focus movement, whether at LF or overtly. If, as assumed above, it were possible for such distinct items as *Bill* and *Joe* to form a chain, then we would expect the undesirable result that such examples like these could be pronounceable as either of (206a) or (206b), depending on the context and other factors, all while having the just interpretation one would expect of the string in (206a).

(205) a. Mary liked Joe.

b. No [$_{FocP}$ Bill [$_{TP}$ Mary liked Joe]]

(206) a. No, Bill, Mary liked.

b. No, Mary liked Joe$_F$.

Kyle Johnson (p.c.) remarks that one might expect similar arguments to apply to an ellipsis approach to Stripping. After all, if the identity condition on ellipsis ensures that what is elided and the antecedent are identical in a particular manner, then wouldn't the elided TP in (207b) and the antecedent TP in (207a) fail to be identical, because *Joe* and the lower copy of *Bill* are not identical?

(207) a. Mary liked Joe$_F$.

b. No, Bill$_F$ ~~Mary liked Bill~~

One key difference between the LF copying approaches and the deletion approaches, however, is that the identity condition of the former is, in a sense, hardwired into the account. Under an LF copying approach, the identity condition on ellipsis is satisfied in virtue of the LF of the antecedent being copied into the ellipsis phrase. Given that the entire LF is copied, the correlate comes along for the ride, giving rise to the issues for trace conversion discussed above. The issue for the LF copying approaches, then, is not that the identity condition on ellipsis is violated, but rather that the correlate within the copied LF should not plausibly be able to form a chain with the remnant or participate in trace-conversion.

Deletion approaches, on the other hand, are not tied to a single approach to the identity condition on ellipsis. One could imagine an identity condition which ignores lower copies of focused elements, for example, along the lines of Merchant's e-GIVENness approach (Merchant, 2001). However such an identity condition is implemented, examples like (207) would not be problematic. Nor would the problem of trace conversion afflicting the LF copying approaches apply here, as the elided TP would contain a copy of the remnant, not the correlate.

Merchant (2001) noted a different problem with the LF copying analysis of non-contrastive Sluicing; such an analysis would predict that a *wh*-phrase base generated high in the left periphery should be able to bind an indefinite merged into what would otherwise be the base position of the *wh*-phrase, as in (208). Merchant seems to indicate that such examples, in the absence of Sluicing, ought to permit the pronunciation of both the indefinite and *Wh*-phrase. This, of course, depends on our assumptions about how well-formed chains are parsed at the PF interface. If we assume that, at least in English, only one link in a chain can be pronounced, and that chains such as these are in fact well formed, then such examples should only permit pronunciation of one of the two links. If this were the case, such examples would be as problematic as (206) above. Imagine there is a context which would license pronunciation of the lower link, yielding the pronounced string in (209a). Such a string presumably has the interpretation associated with the string in (209b). This is problematic because these two questions are associated with distinct congruent responses. (209a) can be naturally answered with a *Yes* or *No*, but responding to a constituent question like that in (209b) with *Yes* borders on farce.

(208) a. James likes someone.

b. Who does James like someone?

(209) a. Does James like someone?

b. Who does James like?

To summarize the discussion on LF Copying analyses and Stripping Island Insensitivity, it is the case that the data can be captured, but at high cost. First, the approach advocated by Chung et al. (1995) would need to be extended to account for not just the island insensitive contrastive Stripping cases, but for any sort of contrastive Stripping case, independent of islandhood. I argued the only way to do this, while yielding the requisite bindable variable within the reused TP, was for the antecedent to scope via movement at LF. This, in the island insensitive cases, of course requires that the theory admit island insensitive LF movement, abandoning a key advantage of such LF Copy theories, in their ability to do without such movement. Then I observed that the assumptions inherent in such an extension were incompatible with a copy theory of movement. The costs of maintaining an LF Copy theory of Stripping, in light of the island insensitivity data are prohibitively high. As it turns out, the Binding Condition C data, presented in Chapter 3, is even more challenging for this approach.

To start, observe that in-situ focused R-expressions, as in (210), produce BCC effects when C-commanded by a co-indexed pronoun. If we continue to assume that contrastively focused elements scope through LF movement, as per our modified LF copying analysis above, then the conclusion from this is that LF focus movement does not bleed BCC, that it reconstructs. Now, turn to the corresponding Stripping configuration, that in (211). Under our modified LF copying analysis, the LFs of the antecedent and Stripping phrases are roughly as in (212), prior to the conversion of the lower copies to variables.

(210) a. She$_i$ appreciates the student who respects Bill.

b. *No, She$_i$ appreciates the student who respects Mary$_{F/i}$.

(211) a. She$_i$ appreciates the student who respects Bill.

b. *No, Mary$_i$.

(212) a. Bill λx she$_i$ appreciates the student who respects Bill.

b. Mary$_i$ λx she$_i$ appreciates the student who respects Bill.

The time comes to consider the mechanisms behind reconstruction effects. Substantial evidence, much of it specifically referencing the BCC effects, has been established suggesting that reconstruction is the result of a syntactic mechanism (Chomsky, 1993; Fox, 1999, 2002; Romero, 1998; Sauerland, 1998). The common aspect to these syntactic approaches to reconstruction is the claim that reconstruction effects are the result of the representation of the 'reconstructed' lexical material at the position of the trace, or lower copy, of a movement chain. Under this view, reconstructed material behaves as if it were in the reconstructed position because, at least some part of the moved element *is* present in the lower position. Now, it is clear that, in our revised LF Copying analysis, there is no way we would expect such syntactic reconstruction to be possible. The tail of the chain <Mary,Bill> in (212b), prior to conversion to a variable, contains only the lexical material copied, along with the rest of the TP, from the antecedent. Thus reconstruction of *Mary* into a position where it would be C-commanded by the pronoun *she*, thereby inducing BCC effects, is impossible, and so no BCC effects are predicted. So, the LF copying accounts can account neither for the BBC insensitivity nor for the island insensitivity of Stripping.

4.3.4. Isomorphic Approaches

We have now reviewed Non-Structural, Pronominal, Evasion, and LF Copying analyses of Stripping, and found them each lacking. And so turn to an Isomorphic Ellipsis approach to these data (Lasnik, 1972; Tancredi, 1992; Chomsky and Lasnik, 1993; Fox, 1995; Merchant, 2001). In these approaches, the ellipsis site contains syntactic structure that is more or less isomorphic to the antecedent. Depending on one's theory of the identity condition on ellipsis, this could mean that the LF of the antecedent and ellipsis site are perfectly isomorphic (Sag, 1976; Williams, 1977; Fiengo and May, 1994), or that they are merely isomorphic enough to satisfy a semantic identity condition (Merchant, 2001), or some hybrid of the two (Kehler, 2002; Chung, 2013; Merchant, 2013). In either case, if an ellipsis site is “more or less isomorphic” to the antecedent will contain, in the case of island violating ellipsis, syntactic structure corresponding to both the content of the island and the material outside of the island, in a configuration matching that of the antecedent. Example (213), illustrates, with enough precision for the present purposes, an isomorphic approach to Stripping.

(213) a. James met the guy who likes Volkswagens.

b. No, Subarus$_i$ ~~James met [$_{Island}$ the guy who likes e$_i$]~~.

Under these ellipsis approaches, the syntactic structure in the ellipsis phrase is generated as would be generated any other syntactic structure, as opposed to e.g. the copying analysis of Chung et al. (1995). This structure is subject to an ellipsis mechanism, by which the elided material fails to be pronounced. A common way to think of this was discussed in Merchant (2001), who suggests that a special syntactic feature, the E-feature, which also happens to license ellipsis, sends an instruction to the PF interface to not spell out the material C-commanded by the head bearing the feature.

A final note on the representation in (213). I have represented the gap within the ellipsis site with *e*. I use this symbol to indicate the gap, without committing to a particular representation for the syntactic structure at that position. As we will see, various approaches to how the remnant comes to survive the ellipsis posit distinct content to this *e* character.

As is commonly mentioned in certain research circles, islands are constraints on overt movement, and so under an isomorphic ellipsis approach, if the remnant originates within and subsequently moves out of an island, ellipsis island insensitivity is not intuitively predicted. Clearly, for an Isomorphic Move and Elide approach to get off the ground, will be necessary to assume some sort of manner of repairing the resulting island violation. Before exploring island repair, let us consider an alternative [1].

4.4. Remnant Escape through Resumption

Setting aside for the moment the hypothesis that remnants move from within the ellipsis site, another way for ellipsis remnants to escape the ellipsis is to never have been merged within the ellipsis site at all. Suppose that ellipsis remnants are base generated above the ellipsis site, and, taking a nod from the LF Copying approach, that the remnant binds an element, of some sort or other, that is generated in what would, under other circumstances, have been the base position of the remnant.

Merchant (2001), in the context of Sluicing, gives three arguments against treating this null element as a type of resumption, and at least two of these apply to Stripping as well. The first argument relies upon the observation that not every type of Sluicing remnant can be resumed (Merchant, 2001, pg. 129). For example, *where* can serve as a remnant, but generally cannot participate in a resumptive structure, as illustrated in (214) However, the inability for *there* to participate in a resumptive structure seems to be limited to *wh*-resumption; in examples such as (215), from Prince (1990), *there* makes as fine a

[1] A third approach within the Move and Elide family is, of course, one we have already seen, and dismissed: Non-Isomorphic, Evasion approaches. Here I consider isomorphic approaches to ellipsis.

resumptive pronoun as pronouns generally make in English. This suggests that the difficulty Merchant (2001) noted for the resumption of *wh*-phrases by the pronouns *then, there* and *that* is due to some independent factor. The broader issue remains that if resumption were a strategy used in e.g. Stripping to evade the need for island violating movement, then some other strategy would still need to be provided by the grammar to explain island insensitive Sluicing with e.g. *where* remnants. And, if such a strategy were at play in these island insensitive Sluicing cases, some explanation would need to be given why it could not also apply to the Stripping cases.

(214) a. James wanted to camp somewhere, but I don't know where.

b. *Where$_i$ does he want to find a person [who camped (there$_i$)]?

(215) They were just towed across the Midway onto the bridle path, where they were just sitting there peacefully. (Prince, 1990, ex. 4 pg. 482)

The other two arguments Merchant gives against an explanation of the island insensitivity of Sluicing in terms of resumption hold fully for Stripping. Merchant notes that *wh*-resumption is impossible where a pronoun has pied-piped, as in (216), although pied-piping a pronoun is acceptable in instances of Sluicing, modulo the formality associated with it. Thus, given that Sluicing with a pied-piped preposition is no more island sensitive than P-stranded versions, (217), resumption cannot be a general explanation for Sluicing island insensitivity.

(216) a. (*For) which candidate$_2$ did they receive reports that more than 60 per cent of eligible voters were planning to vote for him$_2$?

b. More than 60 percent of eligible voters were planning to vote for one of the Red candidates, but I don't remember (for) which.

(217) James met with the prof who gave extra credit to a student, but I don't know to which student.

As attested by the island insensitivity of Stripping examples in Experiments 1 and 2, in which the remnant had pied-piped the preposition *to*, Stripping tolerates prepositional pied-piping. Stripping continuations like that in (218b) are acceptable with and without the preposition. However, pied-piping remains unacceptable in non-*wh* resumptives, as illustrated by the topicalization and it-cleft examples in (218c)-(218d).

(218) a. James met the student who gave the book to Bill.

b. No, (to) Mary.

c. No, ?(*to) Mary, James met the student who gave the book to her.

d. No, it was ?(*to) Mary that James met the student who gave the book to her.

The final argument Merchant makes against a resumptive explanation for the island insensitivity of Sluicing is that *wh*-resumptive structures don't generally show case connectivity. The *wh*-phrase in resumptive structure, like that in (219a), generally doesn't match the case required of the resumptive. In (219a), the *wh*-phrase must be *who*, not *whose*. In Sluicing, however, the case of the remnant must match that of the correlate. Thus, in (219b), the correlate must be *whose* and not *who*. Explanation of the island insensitivity of Sluicing in terms of the insensitivity of resumption incorrectly predicts that Sluicing should not show case matching effects.

(219) a. Who$_i$ did the police say that finding his$_i$ car took all morning?

b. The police said that finding someone's car took all morning, but I cannot remember {whose$_i$,*who$_i$}.

This same argument holds of Stripping, in which the case of the remnant must generally match that of the correlate (Merchant, 2004). Thus, the Stripping continuation in (220) requires the remnant to be in genitive case. Such genitive case in resumptive structures is impossible, even when the resumptive pronoun is genitive.

(220) a. James met the student who reviewed a student's book.

b. Yeah, Mary's.

c. *Yeah, Mary.

d. Yeah, {?Mary,*Mary's} James met the student who reviewed her book.

e. ?Yeah, it was {?Mary,*Mary's} that James met the student who reviewed her book.

4.5. Remnant Escape through Movement: Roll-Up Movement

If we cannot appeal to a resumptive strategy to explain how the remnant appears above the ellipsis site, we are left with an analysis in which the remnant moves to escape the ellipsis site. Here I will consider three alternative conceptions of this movement. The simplest approach is for the remnant to move, prior to spell out (Lasnik, 1999; Nakao, 2009; Merchant, 2004). In a second approach, the remnant moves at PF to escape ellipsis (Weir, 2014). I return to these approaches shortly.

A third approach involves roll-up movement, in which the remnant moves to the edge of some domain, followed by movement of that whole domain to some higher position. Roll-up ellipsis remnant movement therefore actually involves the movement of multiple phrases, and, as a consequence of this, the ellipsis of multiple constituents. The principal benefit of such an analysis is that it is no longer necessary to assume that ellipsis remnants can escape islands; they merely need to move to the edge of the island.

Roll-up movement was originally proposed independently of ellipsis as a mechanism by which foci, indefinites, and *wh*-phrases could scope out of islands, while avoiding the need to posit island violating movement. This general strategy originated with Nishigauchi (1990), and has since resurfaced in (Rooth, 1996; Krifka, 2006; Moritz and Valois, 1994; Wagner, 2006; Charlow, 2017). These versions vary in detail, all proposing some amount of pied-piping by an element that scopes at LF. The proposals presented in Rooth (1996) and Charlow (2017) are most useful for the present purposes in that the relevant element within the pied-piped phrase moves to the edge of the pied-piped phrase. This additional step of movement creates a constituent that could be targeted by an ellipsis mechanism.

Consider (221). A focused phrase, *Thai*, appears within an island, indicating that there are contextually salient alternatives to Thai such that, in those alternatives, James met the student who speaks the alternative. Thus the focus interpretation of *Thai* takes wide scope, out of the island. A roll-up movement analysis would derive such an interpretation from the movement of *Thai* to the edge of the island, followed by movement of the whole island to the matrix left periphery, as in (221b). Thus the focus never actually moves out of the island.

(221) a. James met [$_{Island}$ the student who speaks Thai$_F$]

b. [Thai$_F$ λx$_1$([the student who speaks x$_1$)]]$_F$ λx$_2$ ([James met x$_2$])

Notice though that to account for the behavior of island-insensitive stripping, we would have to assume that such examples involve not a single application of ellipsis but several. Let us take (222c) to be the representation of a Stripping configuration, at the level of representation that feeds ellipsis. The remnant, *Mandarin* had fronted within the definite relative clause island, and the whole island has moved to the edge of the matrix clause. This leave two independent constituents, *[the student who speaks e$_i$]* and *[James met e$_j$]*, which both would need to be each independently elided for a representation like (222c) to yield a string like (222b). Consequently, in a roll-up analysis of remnant movement, we would need to assume that the ellipsis process at work in island-violating Stripping is a series of ellipses, each targeting a single layer of roll-up.

(222) a. James met [$_{Island}$ the student who speaks Thai$_F$]

b. No, Mandarin$_F$.

c. [Mandarin$_F$ [the student who speaks e$_i$]]$_F$ James met e$_j$

To derive the string in (222b), two instances of ellipsis must occur. However, these two instances of ellipsis are not fully independent of each other. If the island is elided, the matrix clause must also be elided, but the inverse is not true. As illustration, consider (223)-(226). Let us take (223a) to be the antecedent, and (223b) to be the non-elliptical counterpart to the various Stripping configurations to be considered. In (224b), we have a typical Stripping configuration, and in (224b), the structural representation at the level at which ellipsis applies. The matrix clause and the relative clause each are subject to ellipsis. In (225a), we have an acceptable variant in which the island is pied-piped, but not elided. Here, the focused element in the island, *Mandarin*, remains in-situ within the island, while the island is fronted, permitting the matrix clause to be elided. Here we see that the matrix clause may be elided, while the island is not elided. But the unacceptability of (226a) demonstrates that the inverse is not the possible. To derive (226a) would require that the matrix clause remain unelided, while the island is elided. Apparently this is not possible.

(223) a. James met [$_{Island}$ the student who speaks Thai$_F$]

b. No, James met [$_{Island}$ the student who speaks Mandarin$_F$]

(224) a. No, Mandarin$_F$.

b. ... [Mandarin$_{F/i}$ ~~the student who speaks e$_i$~~]$_j$ ~~James met e$_j$~~

(225) a. No, the student who speaks Mandarin$_F$.

b. ... [the student who speaks Mandarin$_F$]$_j$ ~~James met e$_j$~~

(226) a. *No, Mandarin$_F$, James met.

b. ... [Mandarin$_{F/i}$ ~~the student who speaks e$_i$~~] James met e$_j$

If there were a constraint ruling out derivations like in (226b), it doesn't seem to be one that applies to various elliptical configurations freely. In particular, the interaction of NP ellipsis and Sluicing shows that configurations just like that in (226b) are possible. Consider (227)-(230). In (227) we have a Sluice antecedent, along with the unelided counterpart to various Sluicing variants. In (228), we have a Sluicing continuation, in which both Sluicing and NP ellipsis have independent taken place. (229) illustrates that Sluicing is possible independently of NP ellipsis, and (230) shows that NP ellipsis is possible independently of Sluicing. Thus the unacceptability of (226a) doesn't seem to stem from general constraints on how two instances of ellipsis interact in a single utterance, such as the admittedly implausible sounding: "if the linearly first ellipsis occurs, the second of two possible ellipses must also occur".

(227) a. James likes a particular type of gin ...

b. ... but I don't know which type of gin James likes.

(228) a. ... I don't know which.

b. ... I don't know [which ~~type of gin~~]$_i$ ~~James likes e$_i$~~

(229) a. ... I don't know which type of gin.

b. ... I don't know [which type of gin]$_i$ ~~James likes e$_i$~~

(230) a. ... I don't know which James likes.

b. ... I don't know [which ~~type of gin~~]$_i$ James likes e$_i$.

Returning to the Stripping examples in (224)-(226), it is important to note that they would all have the same LF representation under such a roll-up theory of focus scoping, that given in (231). The only difference between them is whether and where ellipsis, and the overt movement of the ellipsis remnant that goes along with the ellipsis, takes place. Consequently, we cannot rule out (226a) on the basis of its LF structure. Nor does it seem plausible to rule it out on the basis of its PF, given that the only difference between (226a) and (224a) is whether the matrix clause is elided, which must be independently possible, as in (225a). It seems then that we must stipulate a constraint against derivations like that in (226b).

(231) [Mandarin$_F$ λx$_1$([the student who speaks x$_1$)]]$_F$ λx$_2$ ([James met x$_2$])

Note also that a roll-up analysis of island violating Stripping predicts that fragment pied-piped Stripping continuations should receive the same interpretation as the non-pied-piped variants, given that they both have the same LFs. This prediction is not borne out. Consider the two Stripping continuations in (232). The non-pied-piped continuation in (232b) has an interpretation available to it that the pied-piped continuation in (232c) does not, namely that there is a single student who James met who speaks both Thai and Mandarin. This is unexpected if the two have the same LF structures.

(232) a. James met the student who speaks Thai$_F$.

b. Yeah, and Mandarin$_F$ too.

c. Yeah, and the student who speaks Mandarin$_F$ too.

The differences between a pied-piped and a non-pied-piped Stripping continuation can be further highlighted by adding the exhaustive focus sensitive operator *only* to the antecedent, as in (233). Here, the antecedent entails that there is no other language besides Thai such that James met the student who speaks it, and implicates[2] that James met the student who speaks Thai (Horn, 1969). Now, the non-pied-piped Stripping continuation in (233b) is acceptable, and here has the interpretation that the student who James met also speaks Mandarin. The other interpretation that would otherwise be available to such an

[2]See Roberts (2006) for a discussion as to whether this is a presupposition, implicature, or entailments.

example, and the only interpretation available to the pied-piped Stripping continuation in (233c), yields a contradiction. That is, the only reading available to (233c) is that there is another student, besides the one who speaks Thai, that James met, and this other student speaks Mandarin. This reading is contradictory to the implication of the antecedent, that there is only one student that James met. Again, if these two Stripping continuation had the same LF, we wouldn't expect such divergences in available readings.

(233) a. James only met the student who speaks Thai$_F$.

b. Yeah, and Mandarin$_F$ too.

c. Yeah, and the student who speaks Mandarin$_F$ too.

A further difficulty facing a roll-up analysis of ellipsis remnant movement concerns the notion that the remnants move to the very edge of the island, as in e.g. (Charlow, 2017). For Charlow's purposes, concerning the scope of indefinites and foci, as well as for ours, moving the remnant out of the ellipsis site(s), for a roll-up analysis to be successful, the focused element, or indefinite, must move to the very edge of the island. From the perspective of ellipsis remnant movement, only remnant movement to the edge of the island would yield a constituent that could be elided. In (234), if the remnant *Mandarin* were to land anywhere but in the edge of the DP, the ellipsis of no constituent would yield a Stripping continuation of just the remnant *Mandarin*.

(234) . . . [$_{DP}$ <Mandarin> [$_{DP}$ the student [$_{CP}$ <Mandarin> who [$_{vP}$ <Mandarin> [$_{vP}$ speaks <Mandarin>]]]]] . . .

This raises the issue of just how is it that extraction from islands is restricted in general, and how, in the case of roll-up movement, the remnant can move to the edge of the island but no further. A broadly accepted idea about why movement out of, at least some, islands is highly constrained is that the sort of movement that would be required to escape them is, in one of a variety of technical senses, too far. This idea has been cashed out in terms of Subjacency (Chomsky, 1973, 1977), Barriers Chomsky (1986), and, more recently, in terms of Phases (Chomsky, 1999). Underlying these various implementations are the notions that movement is required to be successive cyclic and that extraction from an island cannot happen successive cyclicly. Consider the illicit extraction of *which student* from the definite relative clause island in (235). Successive cyclic movement of *which student* through the edge of the relative clause is blocked by the presence of *who*, forcing extraction to skip this step in the succession of movement through the relevant intermediate positions.

(235) *Which student$_i$ did James meet [$_{Island}$ the student who speaks t$_i$]?

In the above depictions of roll-up remnant movement, I have portrayed the remnant as having moved to the left edge of the head of a definite relative clause. Such movement is predicted to be impossible, as movement through the edge of the relative clause is, by assumption, blocked, and acyclic movement to the edge of the relative clause head from within the relative clause is likewise blocked. Assuming that the remnant moves to the edge of the relative clause, which then is pied piped to the edge of the relative clause head, as in (236), fares no better, as the initial movement of the remnant to the edge of the relative clause is itself blocked[3]. If that initial movement were not blocked, as implicated in the fundamental concept behind the roll-up analysis, then there seems to be no reason why the remnant couldn't continue to move, escaping the island entirely.

(236) . . . [[Mandarin$_{F/i}$ ~~who speaks e$_i$~~]$_j$ ~~the student e$_j$~~]$_k$ ~~James met e$_k$~~

In summary, an analysis which attempts to avoid the need to posit island violating remnant movement through roll-up movement faces serious difficulties. Such an analysis would require assuming that a single Stripping derivation involves the ellipsis of multiple independent constituents, while also assuming the licensing of some of these instances of ellipsis to be dependent on the presence of others. Such assumptions must be stipulated, as they are not independently motivated by the interaction between other examples where multiple constituents are elided, such as the interaction between Sluicing and NP ellipsis. Furthermore, this type of analysis predicts island pied-piped Stripping remnants to have meanings identical to remnants that have not pied-piped an island, and for the island to take wide scope of certain other operators, like intensional verbs. This is not the case. Finally, for a roll-up analysis of island insensitive Stripping to even get off the ground requires that the remnant move to the edge of a pied-piped island. It is hard to square this assumption with the view that extraction from an island results from the inability to move through the edge of that island.

[3]In addition to apparently requiring CP ellipsis, which has been argued to not exist (Merchant, 2001, pg. 117).

4.6. Remnant Escape through Movement: “Simple” Movement

As we have seen so far in this chapter, analyses which propose no structure within the ellipsis site, or those which posit very limited structure, cannot account for the Binding Condition C effects reported in Chapter 3. Analyses which posit structure in the ellipsis site but which reject that the remnant has escaped the island, either by positing a resumptive binding relationship between the remnant and some element within the ellipsis site, or by positing multiple, but non-island violating, movements, are likewise untenable, for a host of reasons. So now I turn to analyses which propose that ellipsis remnants escapes the ellipsis site through movement, even when this movement is island violating and so would otherwise induce unacceptability.

Island insensitive Stripping configurations have a representation like that found in (237b). The ellipsis site is indicated with the ~~strikeout~~ diacritic, and the island, which I illustrate for legibility, is indicated by the labeled brackets: $[_{Island}$]. Likewise for legibility, I suppress any intermediate traces or copies from the representation. The syntactic material within the ellipsis site is isomorphic to that of the antecedent, and the remnant, which originates within the island, moves from within that island to a position above the ellipsis site.

(237) a. James met the student who speaks Thai$_F$.

b. No, Mandarin$_{F/i}$ ~~James met~~ ~~$[_{Island}$~~ ~~the student who speaks~~ ~~e_i]~~.

Such an analysis warrants a few immediate comments. The first is that, as I take representations like (237b) to be grammatical, following the discussion of the data presented in Chapters 2 and 3, the movement of the remnant across an island boundary must not ultimately yield an ungrammatical structure. That is to say, whatever mechanism or characteristic that typically induces ungrammaticality in the sort of island violating movement seen in (237b) is suppressed in some way when ellipsis takes place.

This suppression of the ill effects of island violations in ellipsis contexts is therefore diagnostic of the sorts of mechanism that yields island effects. If island constraints were a derivational constraint, which is to say that the grammar would be simply incapable of generating structures with a movement dependency spanning an island boundary, then there would be no possible way for ellipsis to repair an island violation. The remnant would never be able to escape the island in the first place. Thus, that we know that remnants must escape the ellipsis site through movement and that they can do so across island boundaries tells us that island constraints cannot be derivational in nature. Likewise, island constraints cannot have a source in an LF or semantic representation. Island violating Stripping phrases are completely coherent and interpretable. If island constraints on overt movement were due to violations of these constraints creating malformed LF structures or semantic representations, then, Stripping phrases, which have full fledged syntactic representations, should be expected to reproduce those malformed LF or semantic structures.

The conclusion is that islands, at least the definite relative clause islands considered here, are PF representational constraints. The non-pronunciation of the violation of these PF constraints obviates that violation. Stripping island insensitivity is the result of island repair. The basic idea that islands are essentially PF constraints is well known in the literature on both island and ellipsis, going back to Ross (1969) through Chomsky (1970); Uriagereka (1999); Kennedy and Merchant (2000); Merchant (2001, 2004, 2008); Nunes and Uriagereka (2000); Fox and Lasnik (2003); Lasnik (2001); Bošković (2001), among many others.

In perhaps what is the best known research document on ellipsis and island insensitivity, Merchant (2001) weighs heavily on the ability of Sluicing to violate relative clauses, but ultimately concludes that relatives clauses are in fact a sort of island which is *not* the result of a PF representational constraint at all. Rather, Merchant would have RCs included in a class of propositional islands. Crucially, violations of RCs and other propositional islands cannot be, according to Merchant, repaired by ellipsis. Merchant instead proposes that the acceptability of relative clause violating Sluicing is due to a short source structure, of the sort discussed above in this chapter and extensively in Chapter 3, under the guise of non-isomorphism. This type of analysis is illustrated for Stripping in (238b), where the ellipsis site is construed as just a portion of the content from the relative clause and for which the subject is construed as an e-type pronoun. The evidence and argumentation presented in Chapter 3 weighs heavily against the adoption of such an analysis. Island-violating Stripping Exhibits Binding condition C effects in contexts like (239). Such BCC effect are predicted only if the ellipsis site were construed as containing syntactic material isomorphic to the antecedent, such that, within the ellipsis site, the base position of the remnant were C-commanded by a co-indexed pronoun.

(238) a. James met the student who speaks Thai$_F$.

b. No, Mandarin$_{F/i}$ ~~they$_j$ speaks e$_i$~~.

(239) a. She$_i$ met the student who photographed Bill$_F$.

b. *No, Mary$_{F/i}$

c. No, Mary$_{F/i}$ ~~they$_j$ photographed e~~.

d. No, Mary$_{F/i}$ ~~she$_i$ met the student who photographed e~~.

But what was the reason why Merchant (2001) adopted such a short source analysis in the first place? Merchant argues, effectively exclusively, that the source of island effects, whether PF or propositional, should track whether a null operator, such as that found in a comparative, can escape from within an island within an ellipsis site. Thus left branch extractions, (240a), are classified as PF islands and relative clauses, (240b), are classified as propositional islands.

(240) a. This professor speaks a rarer Balkan language than *Op*$_i$ that professor does ~~speak a t$_i$ Balkan language~~.

b. *The University of Chicago hired a professor who speaks more Balkan languages *Op*$_2$ that Northwestern did ~~hire a professor who speaks t$_2$~~.

I'd note that the acceptability of examples like (240b) increases greatly when placed in a context which would make the relevant comparison salient. For example, in the context given in (241a), my informants rate a similar example containing a relative clause, (241c), about as acceptable as the non-island control in (241d). Thus, assuming that this diagnostic is effective in distinguishing various types of sources for island constraints, if this is the sole diagnostic between the sources for island constraints, there seems to be no strong evidence that relative clauses cannot be repaired by ellipsis.

(241) a. Context: Mike and James made a bet about which of the two could interview a linguist who speaks a rarer language.

b. Who won?

c. Well, Mike interviewed a linguist who speaks a rarer language than James did.

d. Well, Mike's linguist spoke a rarer language than James's did.

Now, if we are to take the idea of island repair seriously, we need to have a clear idea of how repair is to be implemented by the grammar. I am aware of three approaches.

The first is represented by Merchant (2001), and proposes for each type of island considered a distinct type of phonological problem the ellipsis of which ameliorates the problem. For example, following the proposal of Kennedy and Merchant (2000), left branch extractions in English are impossible in overt movement because English lacks a suitable lexical items which would host the *wh+* features necessary for extraction from a left branch to occur. In elliptical configurations, a lexical item, by assumption unpronounceable in English occurs, the unpronounceability of which is irrelevant due to ellipsis of that element. Likewise, the constraint against extraction of a conjunct from a coordinate structure is proposed to be due to the Null Conjunct Constraint, which mitigates against null conjuncts. If the whole coordinate structure is not pronounced, presumably this constraint is irrelevant.

While there may be merit in supposing that there are various types of PF constraints that could be obviated through ellipsis, as a general approach, it leaves us identifying, and defending, piecemeal each and every phonological constraint. From a more pessimistic perspective, without a clear understanding about just what qualifies as a phonological constraint, a skeptic would have leave to render this position absurd by asking what ellipsis cannot repair (e.g. Marušič and Žaucer (2013)). If we are to make the claim that ellipsis can repair PF constraint violations, we need to be able to identify a broadly applicable class of PF constraint such that ellipsis could repair their violations.

A second type of approach to island repair goes back to Chomsky (1976). Here, the crossing of an island boundary places a new feature, e.g. a *, on the island node crossed, as in Chomsky (1970), or on each of the traces or copies left through successive cyclic movement of the remnant out of the ellipsis site, as in Merchant (2004). Only if each of these *s are eliminated through ellipsis will the resulting island violating structure remain grammatical. Assuming that the remnant movement moves cyclically through at least the *v*P and CP edges, leaving starred traces along the way, then ellipsis of the highest TP will eliminate each, results in a grammatical structure.

(242) a. James met the student who speaks Thai$_F$.

b. No, Mandarin$_F$ ~~James met~~ $\cancel{[_{Island}}$ ~~the student who speaks e]~~

c.

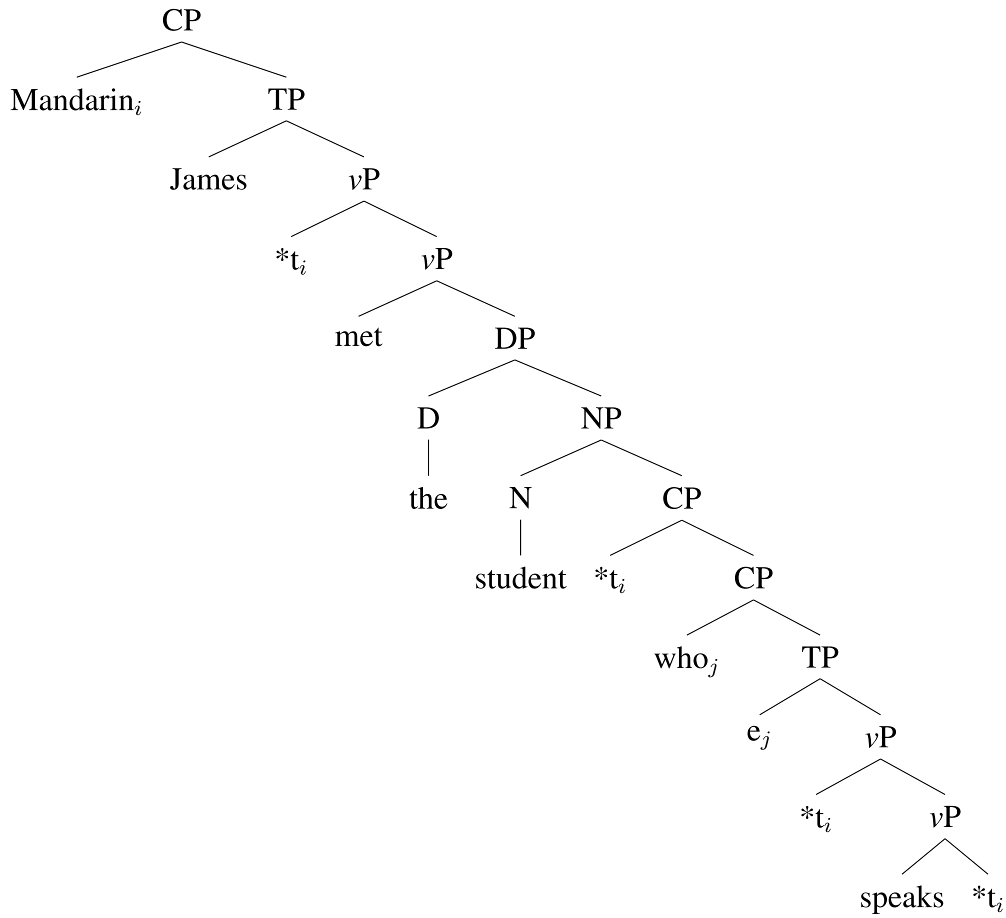

This approach does have a simplicity which is appealing, in that it gives us a clear element, presumably stable across island types, that must be eliminated under ellipsis for repair to be successful. Of course this simplicity is largely due to its stipulative nature. As discussed by Fox and Lasnik (2003), it doesn't follow from the nature of islands nor the nature of traces or copies that a diacritic should be added to traces when they cross an island boundary. Nor does this approach tell us anything about islands, or why deletion of island violating movement traces should be repair an island violation. All of this is stipulated on top of an existing theory of islands.

But it seems intuitive that the ability of ellipsis to repair island violations should indicate some deeper aspect of islandhood. A third approach to island repair, proposed in Fox and Pesetsky (2005), does just this. In their proposal, island effects are the result of the interaction between particular syntactic domains, *Spell-out domains*, and the linearization algorithm. The elements within a given Spell-out domain, which include CP, VP, and DP, are linearized when that domain is complete, yielding a list, the *Ordering Table*, of precedence statements for the elements within that domain. As successively larger structures are built, new precedence statements are added to the Ordering Table at the completion of each Spell-out domain. Crucially, for a syntactic object to be well formed, the statements on the Ordering Table must not be contradictory. As Fox and Pesetsky (2005) show, island violating movement, and other instances of acyclic movement, are ruled out under this analysis, as such movement, in general, yields linearization contradictions. If the elements which participate in these contradictory precedence statements are elided, however, the problematic contradictions are likewise eliminated.

Let us see how this system works in a simple example of *wh*-movement to exclude acyclic movement. First, I will assume here that, when linearizing a Spell-out domain, the highest copy of a moved element is the one that is relevant for the linearization algorithm[4].

[4](Fox and Pesetsky, 2005, pg. 40) assume a remerge theory of movement, rather than a copy theory. So a moved element is remerged into a higher position, only the highest of which is relevant for linearization purposes, by means of the definition of *Mother* that explicitly refers to "the most recent Merge of α". The authors suggest that one advantage of the adoption of remerge is that it allows them to avoid having to make statements about which copy to pronounce. However, the need to explicitly identify that the most recent application of merge determines which is a node's mother is, in effect, a statement about where to linearize an multiply merged element. A second advantage the authors provide is that remerge provides a more intuitive understanding of why linearization contradictions should be problematic. That is, "$\alpha < \beta$" and "$\beta > \alpha$" is a more clearly contradictory linearization than is "$\alpha < \beta$" and "$\beta > copy(\alpha)$". If the linearization algorithm is sensitive to the contrast between copies and originals, then point taken. But if one copy among many can be distinguished as the original, then it seems that the copying process is performing

Consider the linearization of (243), in which the *wh*-phrase has moved cyclicly to the CP domain. The derivation, and resulting precedence relations, at the point of the *v*P Spell-out domain, are given in (244), and for the CP Spell-out domain, in (245).

Upon the completion of the *v*P, that domain is linearized. *who* has moved to the edge of the *v*P and so it is this highest copy which is relevant for linearization. If we assume that asymmetric C-command maps onto linear precedence, this yields *who* preceding *Mary* and *Mary* preceding *call*[5]. The CP domain is linearized similarly, with two new precedence relations added to the Ordering Table: *who* precedes *did* and *did* precedes *Mary*. As the reader can verify this results in a total ordering, where each element is ordered with respect to the other elements and no contradiction results.

(243) Who did Mary call?

(244) *v*P spell-out domain

a. [$_{vP}$ who [$_{vP}$ Mary call who]]

b. who $>$ Mary, Mary $>$ call

(245) CP spell-out domain

a. [$_{CP}$ Who did Mary [$_{vP}$ who [$_{vP}$ Mary call who]]]

b. who $>$ did, did $>$ Mary, who $>$ Mary, Mary $>$ call

---

some additional annotation on one or more of the copes, which seems antithetical to the spirit of the copy theory of movement. In the body text, I will continue to assume a copy theory for expository ease.

[5](Fox and Pesetsky, 2005, pg. 40) assume syntactic geometry is linearized on the basis of various "Laws of Precedence", which, in English, include the ordering of specifiers before heads, and complements after. The particular formulation of the laws of precedence is largely orthogonal to the issues discussed here. As do Fox and Pesetsky, I only include the minimum number of precedence relations to show either a total ordering, or a contradiction.

Acyclic movement of a *wh*-phrase fails where cyclic movement was successful. Assume that, in the derivation of (243), the *wh*-phrase does not move through the edge of the *v*P in its way to the CP domain. In this case, *who* is linearized to succeed the other elements in the *v*P domain. Then, as illustrated in (247), the linearization of the CP domain adds new statements, to the effect that *who* is linearized to preceded not only *did* but also *Mary* and *call*. This results in a contradiction in precedence relations: during the spellout of the *v*P domain, *who* followed *Mary*, but during the spellout of the CP domain, *who* precedes *Mary*. Thus the derivation crashes.

(246) *v*P spell-out domain

a. [$_{vP}$ Mary call who]

b. Mary $>$ call, call $>$ who

(247) CP spell-out domain

a. [$_{CP}$ Who did Mary [$_{vP}$ Mary call who]]

b. who $>$ did, did $>$ Mary, Mary $>$ call, call $>$ who

Similar considerations of cyclicity derive island constraints in terms of failed, contradictory, linearizations, along with a few crucial additional assumptions. In particular, we need to assume that at least one of the relevant edge positions in an island are inaccessible landing sites. For example, in a definite relative clause, we would need to assume that edges of one or more of the following Spell-out domains contained within a relative clause cannot be an intermediate landing site: the relative clause *v*P, the relative clause CP, the edge of the relative clause head DP. Blocking intermediate movement though any one of these edges would suffice to block extraction from a relative clause. I will assume, as illustrated in (248), that intermediate movement through the relative clause CP and relative clause head DP edges are blocked, as these domains resemble domains from which extraction is independently constrained: complex noun phrase islands, as in (249a), and *wh*-islands, as in (249b).

(248) James met [$_X$ the student [$_X$ who speaks Thai]].

(249) a. *Which language does James like the book about?

b. *Which language did James wonder who speaks?

Now consider the derivation of *wh*-movement from within a definite relative clause island, as in (250). Within the relative clause *v*P, *which language* moves to the edge, yielding the precedence relations in (251b). Crucially, *language* precedes *who*. Already at the spell-out of the relative clause CP domain linearization problems arise. Here, the relative pronoun has moved to the edge of the domain while *which language* languishes at the *v*P edge, given that we have assume that relative clause CP edges cannot host intermediate movement landing sites. Linearization of the relative clause CP yields contradictory precedence relations: *who* must both precede and follow *which* and *language*. Further contradictory statements are added later in the derivation, when *which language* moves acyclicly to the matrix CP edge. There, *which* and *language* are linearized to precede the matrix subject *James*, which itself had been linearized to precede *which* and *language* when the matrix *v*P had been spelled out.

(250) Which language did James meet the student who speaks?

(251) Relative clause *v*P spell-out domain

a. [$_{vP}$ which language [$_{vP}$ who speaks which language]]

b. which $>$ language, language $>$ who, who $>$ speaks

(252) Relative clause CP spell-out domain

a. [$_{CP}$ who [$_{vP}$ which language [$_{vP}$ who speaks which language]]]

b. who $>$ which, which $>$ language, language $>$ who, who $>$ speaks

(253) Relative clause DP spell-out domain

a. [$_{DP}$ the student [$_{CP}$ who [$_{vP}$ which language [$_{vP}$ who speaks which language]]]]

b. the $>$ student, student $>$ who, who $>$ which, which $>$ language, language $>$ who, who $>$ speaks

(254) Matrix clause *v*P spell-out domain

a. [$_{vP}$ James meet [$_{DP}$ the student [$_{CP}$ who [$_{vP}$ which language [$_{vP}$ who speaks which language]]]]]

b. James $>$ meet, meet $>$ the, the $>$ student, student $>$ who, who $>$ which, which $>$ language, language $>$ who, who $>$ speaks

(255) Matrix clause CP spell-out domain

a. [$_{CP}$ Which language did James [$_{vP}$ James meet [$_{DP}$ the student [$_{CP}$ who [$_{vP}$ which language [$_{vP}$ who speaks which language]]]]]]

b. language $>$ did, did $>$ James, James $>$ meet, meet $>$ the, the $>$ student, student $>$ who, who $>$ which, which $>$ language, language $>$ who, who $>$ speaks

Now, if ellipsis of a constituent eliminates not only the pronunciation of the terminals within that constituent, but also eliminates the precedence statements concerning those terminals, as Fox and Pesetsky (2005) suggest, any eliminated contradictory statements would no longer produce ungrammaticality. Let us recast (250) as a Stripping continuation, as in (256). The ellipsis of the matrix TP eliminates all of the linearization statements concerning elements contained within the TP, leaving just *Mandarin*. No contradictions remain and so the island violation is repaired.

(256) a. James met the student who speaks Thai.

b. No, Mandarin ~~James met the student who speaks~~.

c. [$_{CP}$ Mandarin ~~[$_{TP}$ James [$_{vP}$ James met [$_{DP}$ the student [$_{CP}$ who [$_{vP}$ Mandarin [$_{vP}$ who speaks Mandarin]]]]]]~~]

d. Mandarin

4.7. Remnant Escape through Movement: PF vs. Overt Movement

Now, I have so far argued that the ellipsis remnants move to escape ellipsis, that they do so in an island violating manner, and illustrated a proposal about how ellipsis can come to repair such island violations. The final issue I'd like to address concerns the nature of that remnant movement. In one set of approaches, ellipsis remnants are pronounced high as a result of an interaction between ellipsis and what would otherwise be a covert movement chain Fox and Pesetsky (2005); Nakao (2009); Pesetsky (1997, 2000). The essence of these approaches is that focused elements would ordinarily raise to scope at a high position at LF. However, the deletion of the lower links in such a movement chain forces the highest link to be pronounced. How each of these approaches implements this notion varies substantially. Notice however, that if the head of a chain which spans the edge of an ellipsis site is necessarily pronounced, a clear prediction follows. Any type of covert movement, including QR and covert *wh*-movement, with landing site above an ellipsis site, should see the covertly moved element pronounced in the higher position, thereby escaping ellipsis as any ordinary remnant would. This prediction is not born out, however.

Consider the case of verb phrase ellipsis in (257). The VPE phrase is ambiguous between a wide-scope construal of the indefinite, above the negation, or a narrow scope interpretation. If the indefinite raises to take wide scope above the negation, then there is a landing site above negation, and consequently above the ellipsis site. The above accounts would predict a surface structure like that in (257c), where the unelided element in the QR movement chain is pronounced. We can construct a similar example with Sluicing. In (258), the antecedent to the Sluice is a multiple *wh*-question. If proposals that the in-situ *wh*-phrase moves at LF to the CP domain are correct Huang (1982), then we again have a situation in which there is a link in what would otherwise be a covert movement chain above the ellipsis site. Consequently, that higher link ought to be pronounced, and (258c) is incorrectly predicted to be acceptable.

(257) a. Mary called a senator …

b. but Bill didn't.[6] (not > a senator, a senator > not)

c. … a senator Bill didn't ~~call a senator~~

[6]Thanks to Bob Frank (p.c.) for suggesting this example.

(258) a. Which voter should have called which politician ...

b. ... and which voter shouldn't have?

c. ... which voter which politician shouldn't have ~~called which politician~~

Weir (2014) presents an alternative proposal about how remnants escape ellipsis, in which remnants escape the ellipsis at PF. Under Weir's proposal, this movement is of an ordinary syntactic sort, simply occurring at PF, as a last resort, because the focus feature on the remnant cannot be elided.

A final alternative is that the remnant movement is tied in some way to the licensing of ellipsis itself. One could pursue the idea, presented in Thoms (2010), that the movement of the remnant itself is what licenses ellipsis. Thoms suggests that non-A-movement can trigger ellipsis just in the event that the lower copy of a moved element is immediately deleted. The complement of the landing site of the remnant is then deleted as a last-resort measure to ensure linearizability. It is not clear on this account, however, why the entire complement would need to be deleted. Surely if the deletion of the entire complement would satisfy the needs of the linearization algorithm, then the deletion of just the lower copy of the moved element, or of some constituent intermediate between the moved element's lower copy and the entire complement, would too.

Another way to link remnant movement with ellipsis is to piggyback remnant movement on Merchant's [E](llipsis) feature (Merchant, 2001). In addition to instructing the PF interface not to pronounce the complement of the head bearing this feature, we could also assume that the E feature also induces any elements bearing a focus feature within the complement to move overtly to the specifier of the head bearing that feature[7] This is plausible, given that focused elements are generally known to be unelidable, but does admittedly have the feel of simply restating the data.

So, we have two plausible alternatives at hand for when in the derivation the remnants move to escape the ellipsis site: at PF or before Spell-out. Weir (2014) advances five arguments in favor of a PF movement analysis, but, as we will see, these arguments are not as convincing as they might initially seem. Notice, however, that for both alternatives, some notion of island repair will be necessary, as the remnant would still move, at some point in the derivation, from within an island domain.

Weir begins with a discussion on the acceptability of negative polarity items (NPIs) as Fragment Answers (Weir, 2014, pg. 167). As Weir points out, Merchant (2004) claims that NPIs that must be in the scope of their licenser are unacceptable as Fragment Answers, citing the contrast between examples like (259) and (260).

(259) a. Max didn't read anything.

b. *Anything, Max didn't read.

(260) a. What didn't Max read

b. *Anything.

Weir goes on to provide several examples which are claimed to be acceptable, following data reported by Valmala (2007) and Dikken et al. (2000) (Weir, 2014, pg.170) .

[7] I have not consider structural approaches to ellipsis in which a prosodic constituent, not a syntactic constituent, is elided, of the sort advocated for by Bruening (2015). Because the remnant in such approaches does not move to escape ellipsis, the same arguments surrounding a PF movement approach to remnant movement apply to such non-constituent deletion approaches.

Example (261) was reported in Dikken et al. (2000), who claimed the key to acceptability for such examples was that the context provides relevant alternatives to the NPI. I would agree that such examples are improved in context, but notice that the minimally different (262), understood in the same context, is much less acceptable. My informants and myself still find the acceptability of such examples nearly as unacceptable as the instance of NPI fronting illustrated in (259b), and much less acceptable than the minimally different (262) and than the in-situ NPI in (259a). If the NPI were licensed in (261), we would expect it to be as acceptable as (259a) and (262).

(261) a. John has returned with the shopping for the party. A and B know that he bought bread, cheese, olives, and juice, but suspect that he has forgotten something.

b. A: What *didn't* he buy?

c. B: Any wine. (Judgment due to Dikken et al. (2000))

(262) a. A: What didn't he buy?

b. B: Wine.

These examples have all centered on the NPI *any*. If we turn to some other of the many NPIs available to English (Zwarts, 1998), the weak NPIs, in (263)-(269), and the strong NPIs, in (270)-(276), all resist appearing as remnants in Stripping or Fragment Answer configurations.

Let us start with the weak NPIs *ever* and *at all*. Examples (263) and (264) illustrate the licensing of these items by the quantifier *few students*. The Stripping continuations in (265) are examples of Sprout-Stripping, in which there is no overt correlate. Yet, judged against a context like (265a), which makes alternatives to the NPIs available in the context, such NPIs make poor remnants, in comparison to non-NPI remnants in the same context, like (265e), or compared with licensed instances of these NPIs, as in (263a) and (264a). Likewise, including an instance of *not* preceding these NPI remnants renders them fully acceptable as in (266).

(263) a. Few students were ever willing to go to Minsk.

b. *Some students were ever willing to go to Minsk.

(264) a. Few students were willing to go to Minsk at all.

b. *Some students were willing to go to Minsk at all.

(265) a. Context: The students were required to travel at least one of three cities either earlier or later in their studies: Rome, Paris, or Minsk.

b. Few students were willing to go to Minsk.

c. *Yeah, ever.

d. *Yeah, at all.

e. Yeah, at any point during their studies.

(266) a. Few students were willing to go to Minsk.

b. Yeah, not ever.

c. Yeah, not at all.

The weak NPI *care to* returns the same results. Example (267) illustrates that we are dealing with an NPI, (268a) illustrates the inappropriateness of *care to* as a remnant, and (269) is a control, illustrating the general acceptability of pied-piping VP material, to ensure the unacceptability of (268b) is not due to this factor.

(267) a. Few students cared to do something important.

b. *Some students cared to do something important.

(268) a. Few students cared to do something important.

b. *Yeah, cared to vote.

(269) a. Few students wanted to eat well.

b. No, wanted to exercise.

Moving on to the strong NPIs *a finger* and *bat an eyelash*, we again see the unacceptability of NPIs as remnants. Examples (270) and (272) provide the baseline licensing of these NPIs, and (271) and (273) illustrate their unacceptability as Fragment Answers. Example (274) is an important baseline control for these examples, as it illustrates the general acceptability of a VP fragment answer to a negative island violating constituent question.

(270) a. No student lifted a finger.

b. *A student lifted a finger.

(271) a. (But) What did no student do in reaction to the janitor's request for help?

b. *Lift a finger (to help)

(272) a. No student batted an eyelash.

b. *A student batted an eyelash.

(273) a. (But) What did no student do (in reaction to the shocking news)?

b. *Bat an eyelash.

(274) a. (But) What did no student do (despite being given lots of free time between and after classes)?

b. Hand in the paper on time.

A final strong NPI is *any NP at all*, which illustrates the same pattern. Here I've used corrective Stripping to illustrate the unacceptability of this NPI as a remnant. Again, note that adding an explicit NPI licenser, the focus sensitive operator *not*, as in (276d), improves the acceptability of these items dramatically.

(275) a. Few students smoked cigarettes at all.

b. A student smoked cigarettes at all.

(276) a. Few students smoked any cigars at all.

b. *No, any cigarettes$_F$ at all.

c. No, cigarettes.

d. No, not any cigarettes$_F$ at all.

Notice that fronting of NPIs even in contexts where a contrast, even a correction, is explicitly made, is unacceptable. This is shown in (277b). On the other hand, the comparable in-situ focus configuration, (277c) is perfectly acceptable. These results directly follow from an analysis in which the remnant moves at surface structure to a position above the ellipsis site, as the NPI is no longer in a licensed position at the completion of the surface structure, (277d). Under a PF movement analysis, as Weir points out, NPIs are predicted to be as acceptable as the in-situ focus variant, in apparent contradiction to the judgments illustrated here.

(277) a. Few students smoked any cigars$_F$ at all.

b. *No, any cigarettes$_F$ at all, few students smoked.

c. No, few students smoked any cigarettes$_F$ at all.

d. *No, any cigarettes$_F$ at all ~~few students smoked e~~.

The next argument Weir makes concerns the binding of *the other* by the partitive *Each of NP* (Weir, 2014, pg. 171). Weir observes that *the other* can serve as a Fragment Answer, bound by *Each of NP*. This, he claims, is despite the inability for *the other* to be bound in topicalization, *wh*-question, and It-Cleft contexts, as illustrated by Weir's examples, repeated in (281). These data, Weir suggests, follow directly from a PF movement analysis of remnant movement. Only if the remnant moves just at PF would such binding relationship be licensed. An overt movement analysis, Weir counters, would predict Fragment Answers like in (278b) to be as unacceptable as the instances of overt movement in (281).

(278) a. Who does each of them hate?

b. The other.

(279) a. Each of them hates the other.

b. *The other, each of them hates.

c. ??Which of the others does each of them hate?

d. ??It is the other that each of them hates.

(280) a. Each of the boys hated someone.

b. Yeah, the other.

This argument rests on two key contrasts. The first is that between instances like those in (279b) through (279d), in which *the other* has raised overtly past *each of NP*, and cases of ellipsis in which *the other* serves as a remnant, e.g. the Fragment Answer in (278) and the Stripping continuation in (280). The second is the contrast between examples like (279a), in which *the other* remains in-situ, C-commanded by *Each of NP*, and the elliptical examples. Based on my informal investigations into the acceptability of these and similar examples, the relevant contrasts does not appear to be very clear.

For the examples in (281), a strong majority of these informants (n=11/15) reported contrasts more or less as depicted above. However, they also judged the relevant Fragment Answer continuation, in (278), and the Stripping continuation, in (280), to be roughly as degraded as the instances of overt *the other* movement in (279c) and (279d). That is, there is not a sharp contrast between all of the cases where *the other* has moved overtly and cases where it serves as a remnant.

One factor that most of the informants reported to make some of these examples difficult to judge was that *the other* was singular. Singular *the other* requires that *each of NP* refer to just two people, which is itself was reported to be quite awkward. Variants of the above examples with the plural *the others* reportedly are somewhat more acceptable and easier to judge.

Likewise embedding *the others* within, for example, a DP, as in (281), further improved the examples. Consider a context where there are several boys who each have several paintings. The relevant reading for the examples is one in which each of the boys admired the paintings of each of the other boys. Crucially, the acceptability of the Fragment Answer and Stripping continuations seems to track that of the It-Cleft and *wh*-question examples, while the Topicalization examples remain somewhat less acceptable. Obviously, the matter warrants further empirical study, but if these judgments hold, they are compatible with both an overt remnant movement account and a PF-only remnant movement account.

(281) a. Each boy admired the others' paintings.

b. ??The others' paintings, each boy admired.

c. ?It was the others' paintings that each boy admired.

d. ?Which of the others' paintings did each boy admire?

e. i. Which paintings did each boy admire?

ii. ?The others' paintings.

f. i. Each boy admired some paintings.

ii. ?Yeah, the others' paintings.

The next argument Weir raises involves Fragment Answers that are predicates (Weir, 2014, pg. 172). Weir notes, giving the example in (282), that predicate inversion is known to bleed inverse scope (Barss, 1986; Huang, 1993). Under a reading of (282a) where *every* scopes under *refuse*, John would be refusing to teach all of the students together, though he might be willing to teach some of them. For example, John considers a class of 50 students to be too large, although he would be happy teaching any given 25 of those 50 students at a time. Under a reading where *every* scopes over *refuse*, for each individual, John would be refusing to teach that individual. For example, a series of students approach John with a request for him to teach them, and for each student, John refused.

(282) a. John refused to teach every student. (refuse > every, every > refuse)

b. . . . and teach every student, John refused to. (refuse > every, *every > refuse)

The scope interaction with intensional verbs provides another example. In the canonical (283a), *three bakers* can scope over *want*, yielding an extensional reading: there are three bakers such that John wants to hire those three bakers. The predicate inverted (283b) lacks this extensional reading, instead yielding a reading in which the number of bakers that John wants to hire is three.[8]

(283) a. John wants to hire three bakers. (want > three bakers, three bakers > hire)

b. . . . and hire three bakers, John wants to. (want > three bakers, *three bakers > hire)

[8]Interestingly, an extensional reading does seem to be possible, both of the predicate inverted and elliptical examples, if *certain* is added to. Simple inverted predicate examples, like (i), remain distinctly odd under, but the corrective variant in (ii) seems possible, as do the FA and corrective Stripping versions. I will set aside why wide scope should be possible of these examples, as it seems to be possible of both predicate inverted and elliptical cases.

(i) a. John wants to hire a certain baker.

b. ??. . . and hire a certain baker, John wants to.

(ii) a. John wants her to hire a certain chef.

b. No, it's hire a certain *baker* that John wants her to.

(iii) a. What does James want her to do?

b. Hire a certain *baker*.

(iv) a. John wants her to hire a certain chef.

b. No, hire a certain *baker*.

Weir claims that predicate Fragment Answers, as in (284), admit both of the scope possibilities available to the non-inverted counterpart. He suggests that "[t]his asymmetry between predicate fronting and fragments answers casts doubt on the hypothesis that movement is involved in the derivation of [(284)]".

(284) a. What did John refuse to do?

b. Teach every student. (refuse > every, every > refuse)

Now, Weir notes that there is some speaker variation in these judgments. I would agree that the judgments are difficult, but my informants generally report that in such examples as (284b), the only reading available is that associated with the topicalized VP, that in (282b), the wide scope reading of *refuse*. I believe examples like (285), involving intensional verbs and numerals, are easier to judge. Here, my informants uniformly rejected the extensional reading. Likewise for the similar corrective Stripping example in (286).

(285) a. What does James want to do?

b. Hire three bakers. (want > three bakers, *three bakers > hire)

(286) a. James wants her to hire three chefs.

b. No, hire three bakers. (want > three bakers, *three bakers > hire)

Now, if predicate inversion bleeds inverse scope because of the LF structure of predicate inversion, then both PF and overt remnant movement analyses predict the above data, contrary to Weir's claims about PF movement. Whether the remnant moves overtly or at PF, in order to satisfy Parallelism, at LF, under both analyses, an example like (285b) must have an LF like that in (287). This is the structure for which the remnant scopes at the same position as the correlate. Thus, from the perspective of LF, it makes no difference if the remnant achieves this position prior to Spell-out, as in the overt remnant movement analysis, or at LF, as under a PF movement analysis.

(287) [Hire a baker]$_i$ James want to e$_i$

Weir reiterates Merchant (2004)'s observation that bare quantifiers, like *everyone*, make acceptable Fragment Answers, (288b), such bare quantifiers are generally unacceptable It-Cleft pivots, (289a), and when topicalized (289b) (Weir, 2014, pg. 173).

(288) a. Who did you interview?

b. Everyone/Someone.

(289) a. ??It was everyone/someone that they interviewed e.

b. ??Everyone/Someone, they interviewed e.

Merchant suggests that the ability for a bare quantifier to serve as a remnant, even when overt movement of such quantifiers is generally disallowed, could simply be the product of the repairing effects of ellipsis (Merchant, 2004, pg. 710). As discussed above, I think it is important to be cautious with proposing a multitude of repairing effects of ellipsis, at the pain of losing predictiveness.

I think we should also be cautious when using Topicalization and It-Clefts as our baseline of comparison, when clearly informational structural properties are at play. Both of these constructions demonstrate properties distinct from fragment ellipsis. For example, Prince suggests elements that have been topicalized in English must "represent an entity" (Prince, 1981, pg. 258). Prince gives example (290), in which the quantifier phrase *Few* *books* has been Topicalized, as an instance of Topicalization that is unacceptable due to the non-entity status of the topicalized element. Focus fronted elements, characterized by a lack of prosodic prominence anywhere but on the fronted element, on the other hand, have no such requirement. This is illustrated in (291b). It-Clefts encode an exhaustiveness requirement, which also does not hold of ellipsis remnants. This is to say that we should not expect ellipsis remnant movement to pattern exactly like either of these constructions, simply because they are not all the same construction.

(290) *Few books, I brought with me.

(291) a. James and Mary just bought a dog.

b. Few worries, it seemed to have.

A final argument Weir provides in favor of a PF remnant movement analysis is based on the claim that certain particles are immobile, such as *up* in (292) (Weir, 2014, pg. 174). If such elements are immobile, then certainly it is unexpected that they would move to serve as Fragment Answer under an overt remnant movement analysis, as in (293b). However, such particles do seem to be able to move, as in (294), and so this objection loses its force.

(292) a. He looked up.

b. *Up he looked.

(293) a. Did he intend look up?

b. No, down.

(294) a. He intended to look up and so up he looked.

b. He wanted to breath in deeply, and so in deeply he breathed.

c. He wanted to move the box inside, and so inside he moved the box. .

On the other hand, certain particles do not seem to be able to move at all, as in (295)[9]. These particles cannot serve as remnant, as in (296), which is expected if they cannot independently move at all. Notice that the unacceptability of such examples is not simply due to the repetition of the particle in the remnant; prepositional phrases make fine remnants, as in (297b).

(295) a. *John wanted to look up the answer, and so up the answer he looked.

b. *John wanted to write out the solution, and so out the solution he wrote.

(296) a. John wanted to look up the solution.

b. *No, up the theorem.

(297) a. James wanted to go to the mall.

b. No, to the bar.

I have argued above that the evidence Weir presents in favor of a PF theory of remnant movement is in fact either compatible with both an overt and PF conception of remnant movement, or favors the overt movement analysis. The arguments surrounding the behaviors of predicate inversion, bare quantifier remnants, particle remnants, and *the other* binding, are all predicted by both theories. The distribution of NPIs, however, strongly favors the overt remnant movement analysis. Further evidence points in the same direction.

[9]Thanks to Kyle Johnson for suggesting such examples

Consider first the evidence presented by Merchant (2004) concerning clausal remnants: such remnants require the presence of a complementizer, as illustrated in (298). The Fragment Answer in (298c) lacks a complementizer, and so it can only be understood as a contradictory statement, that one can be taller than one's height, instead of as a Fragment Answer to the antecedent question. Instead, when the complementizer is present, as in (298b), the Fragment Answer is indeed understood as an answer to the antecedent question. If the CP remnant moves overtly to escape the ellipsis site, then this pattern tracks the distribution of complementizers in in-situ and fronted CPs illustrated in (299). When a CP is fronted, the complementizer is obligatory.

(298) a. What does no one believe?

b. That I'm taller than I really am.

c. #I'm taller than I really am.

(299) a. No ones believes (that) I'm taller than I really am.

b. *(That) I'm taller than I really am, no one believes.

How would a PF theory of remnant movement account for such data? Notice that, in Weir (2014)'s analysis, the remnant movement is not feature driven, but rather driven by a last resort operation. Consequently, there is no indication on the basis of the lexical items selected for the derivation, or their feature structures to indicate that any movement would take place in the derivation. From the PF movement perspective, the presence of the complementizer in such examples should be as optional as it is in examples like (299a).

Nor does it seem that there is any plausible way to later exclude derivation in which a clausal remnant lacks the complementizer. Notice that in-situ CPs lacking complementizers can be focused, as illustrated by the full answer in (300). Thus, whatever LF mechanism provides the structure for examples like (300b) which is appropriately parallel to the antecedent, ought to also be possible for instances of the comparable Fragment Answer, e.g. (298c) in a PF theory of remnant movement, as they would have the same LF structures.

(300) a. What does James believe?

b. James believes [I'm taller than I really am]$_F$!

If we wish to exclude complementizer-less CP Fragment Answers like (299b) under a PF remnant movement analysis, we are left with the possibility that such examples are problematic along the PF branch. It is not clear in what way such examples could be construed as being PF defective. But notice that whatever sort of PF branch problem that might be proposed, it would have to be a PF problem that is unrepairable by ellipsis.

Now, let us turn to what would cause the ungrammaticality of the overt movement of a CP lacking a complementizer, as in (299b). Let us follow Stowell (1981) and assume that the problem with cases like (299b) is that the fronted CP contains an empty category, e.g. the null complementizer, which is no longer properly governed. In cases where the CP remains in-situ like (299a), that null complementizer is governed by the selected verb. If this is true of examples in which overt movement has taken place, what does this tell us able examples like (300b)? I will assume that whatever mechanism rules out (299b), such as the ECP, applies at LF. Then, if a CP with a null complementizer moves in both (299b) and (300b), why should only the former be ungrammatical?

Let us suppose, following Chomsky (1995), that the difference between overt movement and covert movement is that overt movement is the movement of categories, while convert movement is the movement of just formal features. Thus, the overt fronting of a CP would result in an LF, prior to trace conversion, along the lines of (301a). There, two copies of the CP appear in the LF, one in the overtly raised position, one in the base position[10]. In (301b), we have the representation of the covert movement of the CP to some left peripheral position, XP. Instead of the movement of the entire CP, just the formal features of that CP have raised, adjoined to the head of the XP.

(301) a. [$_{XP}$ [$_{CP}$ That I'm taller than I really am] X^0 [$_{TP}$ no one believes [$_{CP}$ that I'm taller than I really am]]].

b. [$_{XP}$ [$_{X^0}$ FF_{CP} X^0] James believes [$_{CP}$ that I'm taller than I really am]]

Now let us turn to examples in which the CP, whether having moved overtly or covertly, lacks an overt complementizer. In the instance of overt movement of the CP, in (302a), we can see that the overtly moved CP contains an empty category which is not governed, which would thereby induce ungrammaticality. When the CP moves covertly, here understood to be the movement of formal features, we can see that the only instance of the null complementizer is in the base position of the CP, where that complementizer remains properly governed. Again, instead of a higher copy resulting from categorial covert movement, just the formal features of the CP have raised. Consequently, there is no ungoverned copy

[10]I ignore positions of intermediate movement here.

of the null complementizer in the XP domain. Such covert movement thereby remains grammatical.

(302) a. $[_{XP}$ $[_{CP}$ $[_{C}$ e] I'm taller than I really am] X^0 $[_{TP}$ no one believes $[_{CP}$ $[_{C}$ e] I'm taller than I really am]]].

b. $[_{XP}$ $[_{X^0}$ FF_{CP} $X^0]$ James believes $[_{CP}$ $[_{C}$ e] I'm taller than I really am]]

Thus, the distribution of complementizers in CP Fragment Answers, as observed by Merchant (2004), along with the optionality of null complementizers in in-situ CP focal structures, cannot be accounted for under a PF remnant movement analysis. There is simply no plausible way for such a theory to distinguish CP remnants that lack a complementizer from those that do not. On the other hand, an overt remnant movement analysis does predict the observed distribution. Overt movement of a CP, whether or not followed by ellipsis, requires an overt complementizer. I cashed out this correlation between movement and an obligatorily overt complementizer in terms of Chomsky (1995)'s distinction between overt and covert movement, in which the former involves categorial movement, and the latter simply formal feature movement, along with the idea that an empty complementizer must be licensed, e.g. by the ECP.

With this much in hand, I'd like to return the the issue of predicate inversion bleeding inverse scope. Recall the facts presented above. Predicate inversion bleeds inverse scope of a quantifier contained within the predicate. Thus, (303b) lacks an extensional reading of the numeral. I claimed that predicate Fragment Answers, like fronted predicates, likewise bleed inverse scope, as illustrated in (304).

(303) a. John wants to hire three bakers. (want $>$ three bakers, three bakers $>$ hire)

b. . . . and hire three bakers, John wants to. (want > three bakers, *three bakers > hire)

(304) a. What does James want to do?

b. Hire three bakers. (want > three bakers, *three bakers > hire)

I explained above that both a PF remnant movement analysis, and an overt movement analysis could account for this pattern. But notice that focused predicates do not bleed inverse scope, as illustrated in (305). Now, if the LF of examples containing focused predicates moved elements is parallel to that in Fragment Answers, as a PF remnant movement analysis would predict, then an extensional reading of the numeral should be impossible here.

(305) a. What does James want to do?

b. He wants to [hire three bakers]$_F$. (want > three bakers, three bakers > hire)

This contrast between focused in-situ predicates and predicate Fragment Answers can, however, be explained under an overt movement analysis of ellipsis remnants. Many overtly moved elements become "frozen" elements, islands for sub-extraction from within the moved element (Ross, 1967, 1974; Wexler and Culicover, 1977; Johnson, 1985). This is true of inverted predicates, as (306) illustrates (see also (Lasnik and Saito, 1990, pg. 101))[11].

[11]Many accounts, e.g. Johnson (1985), suggest that some version of Barriers, or its many spiritual descendents, is responsible for this effect. For the overt movement cases, these explanations would carry over to the linearization-oriented view of locality adopted here. However, the covert movement of interest here would not seem to be covered by such an analysis. As such, some other still mysterious mechanism must be at play, perhaps having to do with the clause-boundedness and constrained landing sites for QR, as proposed by Beghelli and Stowell (1997) and related works.

(306) a. *[Which three bakers]$_j$ do you think [hire e$_j$]$_i$ James wants to e$_i$?

b. *[Three bakers]$_j$ I think [hire e$_j$]$_i$ James wants to e$_i$.

If the lack of inverse scope from out of inverted predicates results from the overt movement of the predicate, the lack of freezing effects for covertly moved, e.g. in-situ focused, predicates, as seen in (305b), follows from the nature of covert movement. As above, if we assume that covert movement entails the movement of not constituents, but rather simply formal features, the inability for covert movement to freeze inverse scope follows. The moved formal features are structurally plain, without the internal structure from which the freezing effect would presumably result. Instead, the formal features of the focused *v*P, *hire three bakers*, and the numeral, here interpreted as taking wide scope in Beghelli and Stowell (1997)'s RefP, are free to each move independently, as illustrated in (307). Thus, while the PF remnant movement analyses predicts inverted predicates to feed inverse and canonical scope, it is the overt remnant movement analysis which correctly predicts frozen scope in these contexts.

(307) [$_{RefP}$ [$_{Ref^0}$ FF$_{DP}$ Ref^0] [$_{FocP}$ [$_{F^0}$ FF$_{vP}$ F^0] James want to [$_{vP}$ hire [$_{DP}$ three bakers]]]]

Finally, PF and overt movement analyses of remnant movement can also be distinguished in the context of intermediate reconstruction effects. Reflexives generally must be bound by a C-commanding element within the same minimal clause (Chomsky, 1981), accounting for the unacceptability of (308). However, overt movement of an element containing the reflexive is known to permit the reflexive to be bound outside the minimal clause the moved element originated in (Lebaux, 1990; Heycock, 1995; Fox, 2003).

(308) *John$_i$ thinks Sue is unlikely to admire [pictures of himself$_i$].

(309) a. [Which pictures of himself$_i$]$_k$ did John$_i$ think Sue admired t$_k$? (Heycock, 1995, pg. 548)

b. John and Bill know [[which picture of themselves] Mary bought t]. (Fox, 2003, pg. 22)

(310) a. [Pictures of himself$_i$]$_k$ John$_i$ thinks Mary is unlikely to admire t$_k$. (Heycock, 1995, pg. 548)

b. John and Bill know that [[the pictures of themselves] Mary bought t]. (Modeled on (Fox, 2003, pg. 22))

Under a copy theory of movement, these intermediate binding effects can be explained by the presence of copies left at the intermediate movement landing sites (Fox, 2003), as in (311). These higher intermediate copies yields the configuration necessary for *himself* to be bound by the matrix subject, *John*.

(311) [Which pictures of himself$_i$]$_k$ did John$_i$ [which pictures of himself$_i$]$_k$ think [which pictures of himself$_i$]$_k$ that Sue [which pictures of himself$_i$]$_k$ admired [which pictures of himself$_i$]$_k$?

Interestingly, if we adopt the contrast between overt movement and covert movement in terms of category and feature movement discussed above Chomsky (1995), then the appearance of intermediate reconstruction only in instances of overt movement can be explained. Overt movement involves the movement of an element through various intermediate positions, as explained above. Covert movement, on the other hand, involves the movement not of phrases, but of formal features. Such formal features do not contain the relevant structure which would license intermediate reconstruction, as illustrated in (312).

(312) *[$_{CP}$ [$_{C^0}$ FF$_{DP}$ C^0] John$_i$ thinks Sue admired [pictures of himself$_i$]].

As it turns out, we can see these same intermediate reconstruction effects in fragment Ellipsis. Example (313) illustrates this effect with a corrective Fragment Answer, (313) with elaborative Stripping, and (315) with corrective Stripping. Take (313) as an example. The reflexive in the Fragment Answer in (313b) contains a pronoun bound by an antecedent local to the base position of the Fragment Answer. My informants also judged it possible for the reflexive in Fragment Answers like (313c) to be bound by the matrix subject. Such examples were reported to be somewhat less acceptable than examples like (313b), but the contrast between the two was reported to be roughly comparable to that between the two binding options available to the examples in (309) and (310)[12]. Now, if the possibility of intermediate reconstruction is dependent on overt movement, these results clearly support an overt movement analysis of the remnant, over a PF movement analysis.

(313) a. Did John$_i$ say that the girl$_j$ admired pictures of Mary$_F$?

b. No, pictures of herself$_j$

c. No, pictures of himself$_i$

(314) a. John$_i$ thinks that the girl$_j$ admired pictures of someone.

b. Yeah, pictures of herself$_j$.

[12]The less than perfect status of intermediate reconstruction examples like (313b) might be the source of the "??" judgement reported for similar examples in (Merchant, 2004) The complete reconstruction example, (313c), is not the only relevant baseline, given that my informants judged the instances of intermediate reconstruction in overt movement cases, e.g. (310a), to be as acceptable as intermediate reconstruction in the ellipsis cases.

c. Yeah, pictures of himself$_i$.

(315) a. John$_i$ thinks that the girl$_j$ admired pictures of Mary$_F$.

b. No, (just) pictures of herself$_j$.

c. No, (just) pictures of himself$_i$.

The ability of the overt movement analysis to derive these intermediate reconstruction results relied not just on overt movement of the remnant out of the ellipsis site, but successive cyclic movement. This is somewhat surprising given the claims of Fox and Lasnik (2003) that the remnant in fragment ellipsis, and Sluicing in particular, always moves in one-fell-swoop. However, there is a configuration which indicates that one-fell-swoop remnant movement is possible for Fragment Answers and Stripping remnants: island violating movement.

Consider (316)-(318), the definite relative clause island parallels to the non-island examples above. Whereas we saw the possibility of intermediate reconstruction effects in those non-island contexts, that possibility seems entirely unavailable here. Again, while intermediate reconstruction effects in Stripping are not perfect, they are much worse in the island cases, roughly on par with the long distance reflexive binding cases in e.g. (308).

(316) a. Did John$_i$ meet the girl$_j$ who admired pictures of Mary$_F$?

b. No, pictures of herself$_j$

c. *No, pictures of himself$_i$

(317) a. John$_i$ met the girl$_j$ who admired pictures of someone.

b. Yeah, pictures of herself$_j$.

c. *Yeah, pictures of himself$_i$.

(318) a. John$_i$ met the girl$_j$ who admired pictures of Mary.

b. No, (just) pictures of herself$_j$.

c. *No, (just) pictures of himself$_i$.

The preceding discussion of intermediate reconstruction effects in Stripping has focused on *picture of NPs*, but bare reflexives appear to display the same paradigm as well. This indicates that the above effects are not simply due to particular qualities of such picture NPs.

(319) a. James likes someone.

b. Yeah, himself.

(320) a. James said that Mary likes someone.

b. Yeah, himself.

(321) a. James knows the girl that likes someone.

b. *Yeah, himself.

Above, I took the availability of intermediate reconstruction effects in non-island cases of Stripping and Fragment Answers to indicate that the remnants can move successive cyclicly. The absence of these effects when the remnant escapes an island suggests that, in these cases, the remnant cannot move successive cyclicly. Rather, in the context of island violating movement, remnant movement must proceed in one-fell-swoop. I would like to suggest that the inability for island-violating remnant movement to move successive cyclicly results from a sort of "improper movement" constraint, of the sort proposed by

(Fox and Lasnik, 2003, pg.151): "Metaphorically, once you get on the express train, you cannot switch to the local short of the destination." Movement out of an island involves at least some non-successive cyclic movement, as the moved element cannot, by hypothesis, move through the edge of the island. Consequently, if a movement chain cannot involve both successive cyclic and one-fell-swoop movement, then we would expect to see no intermediate reconstruction effects in island-violating Stripping. The remnant movement in such examples must involve some non-successive cyclic movement, and so no successive cyclic movement is possible, thereby excluding the possibility of intermediate reconstruction.

4.8. Conclusions

Let us take stock of the principal findings of this chapter. I argued that Stripping ellipsis sites must contain syntactic material that is isomorphic to the antecedent, even in cases of island-violating Stripping. The various non-structural accounts, as well as those accounts which posit some form of non-isomorphic or anaphoric construal of the ellipsis site, struggle to capture the combination of Stripping island insensitivity, as presented in Chapter 2, and sensitivity to Binding Condition C, as presented in Chapter 3.

I then argued that, among the isomorphic analyses applicable to Stripping, the only tenable manner in which the remnants could escape the ellipsis site was through overt movement. This is the case even in instances of island-violating movement. This in turn implicates some form of island repair, which I argued to be the result of the deletion of contradictory linearization statements, following Fox and Pesetsky (2005).

The comparison between an overt remnant movement analysis and the PF movement analysis proposed in Weir (2014) proved useful. It revealed that remnants can move successive cyclicly, as evidenced by the potential for intermediate reconstruction effects in Stripping, but that the remnants cannot do so out of islands. This I suggested was the result of the violation of a ban on non-successive cyclic movement followed by successive cyclic movement.

CHAPTER 5

Conclusion

In this dissertation, I investigated the island insensitivity of Stripping, how much syntactic material is contained within the ellipsis site, and how the remnants escape ellipsis. I focused primarily on three types of Stripping: elaborative, corrective, and non-contrastive. Examples of island violating versions of these configurations are given in (322b), (322c), and (323b).

(322) a. James met the guy who likes Lagavulin.

b. Yeah, and Ardbeg too.

c. No, Macallan.

(323) a. James met the guy who likes a type of Scotch.

b. Yeah, Ardbeg too.

The puzzling characteristics of the myriad of elliptical configurations are well known to researchers working on the topic. How is it that a fragment is interpreted as something more? How is it that this something more cannot be filled in by the confluence of imagination and context? Do the mechanisms by which this something more is filled in conform to what we know about the broader syntactic, or semantics, or pragmatic systems?

I have focused on this last question in this dissertation, with a series of acceptability and plausibility judgement experiments, and a critical review of the literature. In the next section, I put the major results of Chapters 2-4 together, telling a story of Stripping as overt remnant movement, followed by ellipsis of an isomorphic syntactic structure, with a dash of island repair thrown in to taste.

5.1. Dissertation Review

5.1.1. Chapter 2 Review

The first step is, as ever, to establish the generalizations. Are examples of island violating Stripping like these, setting aside the whatever possible interpretation these might receive, are *strings* like these acceptable? A look to the literature, as reported in Chapter 2, revealed a deep literature with no clear consensus on whether fragments like Stripping and Fragment Answers were sensitive to islands or not[1]. Authors disagree as to whether such examples were ever acceptable, whether the remnant being a subject mattered, or whether the type of contrast between correlate and remnant influenced the acceptability of these island violating ellipses. It is difficult to see how such contradictory judgments could be used to reasonably formulate a theory of ellipsis.

(324) a. James met the guy who likes Lagavulin.

b. Yeah, and Ardbeg too.

c. No, Macallan.

(325) a. James met the guy who likes a type of Scotch.

b. Yeah, Ardbeg too.

[1] See Table (2.2) for a summary of the judgments reported for English in the literature

In light of these complicated judgments, I conducted a series of four acceptability judgement tasks, as reported in Chapter 2, concerning the acceptability of examples like (324b), (324c), and (325b). These were 2x2x2 experiments, which means that there are three factors, each with two levels. One factor was the comparison between Stripping examples against It-Clefts. In this way, It-Clefts, which are known to be island sensitive, served as an baseline for a possible island effect. A second factor was the comparison between island conditions, involving definite relative clause islands, as above, against simple complement clauses.

These two factors were crucial to be able to determine whether Stripping was indeed island sensitive. If I had used just It-Clefts as a baseline, the cleft items would have been longer than the Stripping items, a possible confound. Likewise using only the non-island items as a baseline would have meant that only the island conditions contained an island, another possible confound. By using these two factors in conjunction, we would be able to determine the independent effects of length and islandhood, and crucially test for an interaction between the two. An interaction would indicate that the effect of islandhood were different in the Stripping conditions than in the It-Cleft conditions, and so would mean that Stripping were either completely insensitive to island effects, or only partially sensitive to them.

The final factor was a comparison between two of the three contrast types illustrated above. Experiments 1 and 3 compared the elaborative items against non-contrastive items. Experiments 2 and 4 compared elaborative items against corrective items. Thus in each of the four experiments the elaborative items were tested, effectively replicating the test for the island sensitivity of elaborative Stripping four times over.

Finally, the four experiments differed in the grammatical category of the Stripping remnant and It-Cleft pivots. In Experiments 1 and 2, the pivots and remnants were PP, and in 3 and 4, they were DPs. PPs are sometimes said to resist extraction from within islands to a greater degree than DPs, see Den Dikken and Szabolcsi (2002a) for a recent discussion. It has also been said that the relative ambiguity of a DP remnant could reduce the acceptability of such examples in island violating contexts Valmala (2007), and it has been said that the potential ambiguity of a DP remnant could *increase* the acceptability of such island violating examples Winkler (2013). DP remnants could, focus marking aside, correspond to a number of elements in the antededent, including the matrix subject, the relative clause head, or the intended correlate itself, the DP containted within a PP object. Varying the remnant type allowed us to examine these hypotheses.

During analysis, it was discovered that, in the non-island factor, exactly half of the the matrix clause verbs were factive verbs. This is a problem, because factive verbs have been claimed to be themselves islands for extraction from their clausal complements (Ross 1967, pg. 449; Erteschik-Shir and Lappin 1979 Melvold 1991; Krifka 2006). Thus, three sets of analyses were performed. In one, the factive and non-factive items were compared together, in another the non-factive items from the non-island conditions were compared against the island conditions, and in a third, the factive items were compared with the non-factive items.

The comparison between factive and non-factive items revealed that, in the It-Cleft conditions, extraction from within the clausal complement of factive verb yielded lower acceptability than extraction from the complement of a non-factive verb. This held of all four experiments, across all contrastivity types. While we cannot know from the design whether this reduction is the result of extraction from within a factive domain, or whether the mere precence of a factive verb itself reduced the acceptability, it is clear that factivity was indeed a confound. Interestingly, the Stripping conditions showed no such factivity effect. Whatever the source of this effect, ellipsis eliminates it. I will return to the issue of factive islands below.

Because factivity was a true confound, let us focus on the results from the comparison between island conditions and non-factive conditions. Here we saw an interesting and reliable pattern across the four experiments. It-Clefts, across contrastivity types and pivot types, revealed an island effect; the island conditions were worse than the non-island (non-factive) conditions. This, so far, was as expected.

None of the Stripping conditions showed an effect of Islandhood as large as that of the It-Cleft conditions[2]. This was true across both PP and DP remnant types. However, the magnitude of the effect differed. Non-Contrastive Stripping consistently showed no island effect whatsoever. The two corrective Stripping types, corrective and elaborative, consistently showed a partial island effect. This effect was quite small, however, with

[2]With the exception of the elaborative Stripping conditions in experiment 1, where the effect was comparable with that of the elaborative It-Cleft conditions. However, as discussed in the general discussion of Chapter 2, given that the other experiments, which all featured elaborative Stripping, showed a smaller effect for the Stripping Island conditions than the the It-Cleft counterparts, I attributed these results of Experiment 1 to a type 2 error.

even the island violating contrastive Stripping items rated much higher than the non-island It-Cleft items. Thus, while the items were not "unacceptable", a small sensitivity remnant.

What to make of the partial island sensitivity of contrastive Stripping? In Chapter 2, I argued that this effect was unlikely to arise as a result of a violation of a Parallelism constraint on ellipsis. Simply put, whatever mechanism by which the remnant in an island violating Stripping example escapes ellipsis should be able to apply equally in the antecedent to derive the scope of the correlate. Thus no Parallelism violation ought to arise.

Instead, I argued that the partial sensitivity of contrastive Stripping to islands, as compared to the complete insensitivity of non-contrastive Stripping was due to a bias in the parse of an ambiguous antecedent. In the antecedents of contrastive Stripping examples, just what element is taken to be the focus is ambiguous. It could be just the emphasized word itself, or it could be the whole island containing the focus. If there is a bias towards projection of the focus to the whole island, so that island-violating wide scope can be avoided, then the preferred parse, with projected focus, would yield a mismatch in focal structures in the antecedent and the island-violating Stripping continuation. Reanalysis of the antecedent, to yield parallel focal structures, might then yield a reduction in acceptability for these contrastive island-violating Stripping cases. The non-contrastive island-violating Stripping cases would be subject to no such reduction, as the indefinite correlates in these cases do not project their indefiniteness, and so there is no comparable ambiguity. Likewise for the non-island contrastive Stripping instances; there is no parse which would require an island-violating wide focus scope, and so the bias of the parser is irrelevant here. Thus, the partial island sensitivity of the contrastive Stripping cases may simply be due to a biased parser and the resulting reanalysis.

This analysis makes the prediction that varying the discourse context to disambiguate the focal structure of the antecedent should increase the acceptability of the island violating contrastive Stripping items to match that of the non-contrastive counterparts. Awaiting the experimental testing of such a prediction, I tentatively assumed that, from the perspective of the grammar, both contrastive and non-contrastive Stripping are island insensitive.

5.1.2. Chapter 3 Review

Chapter 3 pushes the story forward. The question of what sort of structure, if any, is contained within that ellipsis site has long been at the front of the discussion on ellipsis. The fact that Stripping, of both contrastive and non-contrastive types, is island insensitive complicates the issue, in a satisfying way. If there is no, or little, syntactic structure within the ellipsis site, then the island insensitivity would seem to follow directly. If there were enough structure within the ellipsis site such that the island were reconstructed there, this would raise very serious questions about the nature of ellipsis, and of islandhood itself.

In some ways, these are questions about where to place the complexity of natural language within the theoretical architecture. If ellipsis sites were radically different from their antecedents in syntactic structure, this would require a complex identity condition, such that, for example, an elided It-Cleft could be anteceded by a syntactically and lexically item rich representation, even containing a definite relative clause. This complex identity condition would be rewarded with a simpler notion of islandhood; because, under such theories, no island is contained within the ellipsis site of island violating ellipsis, the notion of island can be maintained without exceptions or repair. Remnants of island violating ellipsis in such theories do not move out of islands, and so there is no need to posit island repair.

Conversely, an approach positing complete isomorphic structure within the ellipsis site would require a relatively simpler theory of identity. The ellipsis site and antecedents are, by hypothesis, isomorphic to each other, which would presumably require less identity condition gymnastics. The converse is that the notion of islandhood would seem to become more complex. If the ellipsis site contains an island, then the fact that the remnant escapes the island requires that the dependency between the remnant and a base position within an island cross that island boundary. Islands, as constraints on locality, find exception under ellipsis.

The experiment reported in Chapter 3 turns what might seem like an unmanageable, unarbitrable discussion on theoretical elegance into a testable empirical question. In this joint work with Tim Hunter and Masaya Yoshida, we sought to test whether Stripping configurations were judged to behave as if their ellipsis sites were populated with full syntactic structure, isomorphic to the antecedent.

To do so, we created a plausibility judgment experiment manipulating the presence or absence of a Binding Condition C violation. Consider the examples in(326). Take (326a) to be the antecedent, and either (326b) or (326c) to be the continuation. The continuation in (326b) clearly instantiates a BCC violation. But how would Stripping continuations like (326c) be judged?

(326) a. She$_i$ met the judge who gave an award to Bill.

b. No, she$_i$ met the judge who gave an award to Mary$_i$.

c. No, to Mary$_i$.

(327) a. Mary$_i$ met the judge who gave an award to Bill.

b. No, Mary$_i$ met the judge who gave an award to her$_i$.

c. No, to her$_i$.

(328) a. She$_i$ likes that the judge gave an award to Bill.

b. No, she$_i$ likes that the judge gave an award to Mary$_i$.

c. No, to Mary$_i$.

(329) a. Mary$_i$ likes that the judge gave an award to Bill.

b. No, Mary$_i$ likes that the judge gave an award to her$_i$.

c. No, to her$_i$.

The experimental design included two baselines. One is the C-command relationship between the co-referential R-expressions and pronouns. In (326) and (328), the pronoun C-commands the R-expression, and in (327) and (329) the inverse. The second baseline for the island conditions in (326) and (327), are the non-island conditions in (328) and (329).

We found striking results in which the Stripping conditions patterned extremely similarly to the non-elided canonically ordered controls, for both island and non-island conditions. That is, for both canonical and Stripping items the *pronoun-name* conditions were worse than the *name-pronoun* conditions, with no differences between island and non-island conditions. Stripping patterned almost exactly as if the ellipsis site were filled with the structure of the antecedent[3].

The fact that we observed a BCC effect at all suggests that there is enough syntactic structure within the ellipsis site to resolve Binding Condition C. Further, the fact that this effect was observed equally in non-island and island conditions, suggests that, in both cases, the ellipsis site is resolved with structure isomorphic to the antecedent. Notice that if the ellipsis site were resolved as just the content of the relative clause, the ellipsis site would not contain enough material to resolve the BCC violation, as the pronoun sits outside the relative clause.

5.1.3. Chapter 4 Review

Chapter 4 brings these two sets of results together, in an evaluation of the major approaches to the resolution of the ellipsis site. A host of approaches naturally capture the island insensitivity, but fail to capture the BCC effects. These include the non-structural approaches, pronominal approaches, and non-isomorphic structural approaches. LF copying accounts run into trouble with the island insensitivity of contrastive Stripping, an account of which would still require some sort of island repair mechanism, thereby ruining the appeal of these approaches.

[3] I say almost, because the rating of the *name-pronoun* Stripping conditions were somewhat lower than that of the canonical conditions. This is likely simply due to a resistance to contrasting a proper name with a pronoun.

No, it is the oft adopted, and perhaps even more often critiqued, isomorphic structural analysis that I argue to best explain these results. That the BCC effects are captured is plain to see. The island insensitivity of Stripping is somewhat more complicated to explain, but, as I argue in Chapter 4, understanding this complexity in terms of movement and ellipsis yields not stipulation, but a greater understanding of islands and ellipsis are, such that island violations are repaired under ellipsis.

Of course, one could imagine that the ellipsis site were populated by an isomorphic syntactic structure, but that the remnant achieves its escape from the ellipsis site in a manner which doesn't require island violating movement. I discuss two such approaches, a resumption approach and a roll-up movement approach, along the lines of (Charlow, 2017). Neither approach, I argue, yields particularly strong insight into the workings of ellipsis or islands, and each face various empirical issues as well.

Instead, I outline an approach to the island insensitivity of ellipsis based in which the remnants do indeed move from within the ellipsis site, even if this entails island violating movement. The key is reframing the source of island violating movement as a particular type of representational problem. Movement out of islands yields unlinearizable structures, as proposed by Fox and Pesetsky (2005). Ellipsis of particular categories eliminates any contradictory precedence relations associated with movement out of those categories. Thus if movement out of an island would ordinarily yield an unlinearizable structure, eliding that structure renders the resulting configuration linearizable, and so grammatical.

Note that this approach is informative about the nature of islands; it tells us that islands are the result not just of a representational constraint, but one of a particular kind.

Alternative representational approaches to island repair, such as the grab bag of PF repairs Merchant (2001) proposes or the starred approach of Merchant (2004) and others, make weak predictions about what sorts of violations should be repairable, and so are less testable, and do little to tell us why it is that the island violations should be unacceptable in the first place.

Finally, I confront the question of how to characterize the Stripping remnant movement that I have argued does indeed take place. I compare the PF remnant movement analysis of Weir (2014) with an overt remnant movement analysis. I show that the arguments presented by Weir can either be successfully analyzed in effectively the same manner under both approaches, or that, upon closer inspection, the facts suggest that the remnants must move overtly, before spell-out.

Perhaps the most interesting data from this discussion concerns intermediate reconstruction effects. I reported judgments in which a remnant containing a reflexive appear to be able to be bound by an element no local to the base position of the remnant, as in (330). Such follows if the remnants move overtly, as covert focus movement, or scoping, does not yield this same effect. But these results also only follow if the remnant has moved successive cyclicly. Otherwise, there would be no intermediate position from which the reflexive within the remnant could be bound. Things get really interesting when we consider the same paradigm nested within island violating contexts, as in (331). Here, it seems that the intermediate reconstruction effects disappear, suggesting that the remnant does not move successive cyclicly when it moves to escape an island. I suggest that this distribution follows from a constraint, proposed by Fox and Lasnik (2003), against one-fell-swoop movement followed by successive cyclic movement.

(330) a. Did John$_i$ say that the girl$_j$ admired pictures of Mary$_F$?

b. No, pictures of herself$_j$

c. No, pictures of himself$_i$

(331) a. Did John$_i$ meet the girl$_j$ who admired pictures of Mary$_F$?

b. No, pictures of herself$_j$

c. *No, pictures of himself$_i$

5.2. Looking Ahead

The sorts of larger scale empirical investigations I have undertaken in this dissertation require that a substantial tradeoff be made. The conception, design, implementation, analysis, and reporting of judgment experiments takes a tremendous time, energy, and page investment. In the case of the studies reported here, such an investment was absolutely crucial in order to tame, to the degree that I have been successful, the wild and contradictory empirical landscape found in the literature on ellipsis island sensitivity.

Another salient effect of this tradeoff is that there are a substantial number of immediately clear ways to continue this research. Here, I will highlight what I take to be the most critical of these.

5.2.1. Island (In)Sensitivity and the Menagerie of Contrastive Ellipses

Contrastive ellipsis comes in many shapes and sizes, including not just the Stripping cases examined in detail here, but also Fragment Answers, which I have grouped with Stripping during the analysis of Chapter 4, PseudoGapping, Contrastive Sluicing, discussed in Merchant (2008), and Gapping, among others.

To the degree that the present research is correct in its analysis that contrastive Stripping is island insensitive, this predicts, all things being equal, that these other contrastive ellipsis configurations should also be island insensitive. This prediction is not obviously correct. Merchant, for example, claims that contrastive Sluicing is indeed island sensitive.

As for the contrastive Sluicing cases, Merchant does not give us many examples to work with. However, I suspect that contrastive Sluicing is as island insensitive as contrastive Stripping is. Consider Merchant's example (52b), repeated here as (332) (Merchant, 2008, pg. 148). I agree that this example is unacceptable, but not due because contrastive Sluicing is island sensitive.

(332) *The radio played a song that Ringo$_F$ wrote, but I don't know who else. (Merchant's judgment)

First, I'll note it is difficult to construe indefinites under ellipsis as referring to distinct entities. Thus the person who likes gin in (333b) is also the person who likes whiskey in (333a). Consequently the interpretation of (332) is something along the lines of (334a). This is a distinctly odd interpretation, because of the real world knowledge that if a song, α for example, was written by a single person, no one else could have written it.

(333) a. A girl likes whiskey.

b. Yeah, and gin too.

(334) a. The radio played a song α that Ringo$_F$ wrote but I don't know who else the radio played α such that e wrote α.

b. ...but I don't know who else the radio played a song that Ringo wrote that song with e.

Notice that the very sensible interpretation in (334b) is unavailable to (332), as this latter interpretation is contingent on the presence of an additional item: *with*. Under such an interpretation, Merchant's example actually constitutes not contrastive Sluicing, but Sprout-Sluicing, in which the correlate is implicit. The actual correlate under this interpretation is something like *with person*. But not only are instances of Sprout-Sluicing known to be island sensitive (Chung et al., 1995), they also are subject to the constraint that no sprouted material can be elided Chung (2013). This is illustrated in (335). When the prepositional phrase *with* is pied-piped, the example is acceptable. In turn, adding that prepositional phrase to Merchant's original example seems to improve it tremendously, as in (336).

(335) a. Ringo$_F$ wrote this song, ...

b. *...but I don't know who else ~~Ringo wrote that song with e~~.

c. ...but I don't know with who else ~~Ringo wrote that song e~~.

(336) The radio played a song that Ringo$_F$ wrote, but I don't know with who else.

I would like to also suggest that ensuring that these examples are judged in an appropriate context improves their acceptability. Consider (337a), the other example of the island sensitivity of contrastive Sluicing that Merchant gives (Merchant, 2008, 148). Judged in the context of (337b), the example seems perfectly acceptable.

(337) a. *Abby wants to hire someone who speaks Greek$_F$ but I'm not sure what other languages.

b. Context: Abby runs an internationally distributed artisanal distillery. Anyone who she would hire must speak multiple languages. Given that she is looking to expand into Greece …

This brief discussion in no way does the breadth of contrastive ellipsis justice. But the careful testing of prediction this present research makes, that this breadth of contrastive ellipsis should be island insensitive, all things being equal, requires us to be sure that, in fact all things are equal among the configurations considered, and that the judgements carefully control for confounding factors and the effects of context. In this way, we can determine whether this not obviously correct prediction is obviously not correct, or simply correct.

5.2.2. Other Islands

Another clear next step is to examine the island insensitivity of Stripping not just in the context of definite relative clause islands, but the full range of islands. In the account proposed in Chapter 4, I linked the ability for ellipsis to repair violations of locality constraints to the deletion of contradictory linearization constraints. Thus, the general prediction is that, wherever a locality constraint can be explained in terms of contradictory linearization constraints, ellipsis ought to be able to repair those violations. Locality constraints not attributable to linearization constraints are not predicted by this approach to be repairable.

However, we have already seen one instance of in which a known locality constraint, which can be explained in terms of derivational constraints, cannot readily be explained in terms of a representational constraint like linearizability. This was the case of scope freezing in predicate inversion, discussed in Chapter 4. Freezing effects on overt movement have been attributed to Barriers (Johnson, 1985), in that any overt movement that would be required to derive the illicit cases would be movement too far. In a linearization framing of these facts, these too far movements would result in linearization contradictions. But since the movement involved in scope freezing in predicate inversion would be covert, we wouldn't expect any linearization failure to arise. Thus, we would expect that such cases could be explained in alternative terms. In footnote (11), I suggest that restrictions on possible QR landing sites is to blame (see Szabolcsi (2010) for discussion).

Similar considerations apply to the case of factive islands. We saw in Chapter 2 that extraction from within the clausal complements of factive verbs resulted in a reduced acceptability in It-Clefts and that Stripping was unaffected by this reduction. Furthermore, in Chapter 3, we saw that Stripping from within the complements of factive verbs was just as sensitive to BCC as was Stripping from the complements of non-factives. Together, these results suggest that ellipsis can indeed repair violations of factive islands, if factive verbs are indeed islands for extraction of arguments.

However, judgments in the literature vary, from fully unacceptable (338a), to merely "slightly degraded", (338d) (Basse, 2008, pg. 1). The results of Experiments 1-4 consistently revealed that the factive It-Cleft conditions were less acceptable than the non-factive It-Cleft conditions, although the average difference between the factive and non-factive It-Clefts across the four experiments was smaller, at .85 points difference, than the average difference between non-factive and definite relative clause island It-Clefts, at 1.38 points. Where these same sources report judgments concerning extraction of subjects or adjuncts from within factive complements, they are uniformly reported to be fully unacceptable.

(338) a. Who does Mary {think / *know} that Sue saw yesterday? (Krifka, 2006, pg. 20)

b. ??Who did she realize that he saw? (Ambridge and Goldberg, 2008, pg. 360)

c. ??What did you notice that Mary had fixed? (Haegeman and Ürögdi, 2010, pg. 120)

d. ?What do you regret that John stole? (Basse, 2008, pg. 1)

Moving forward, it seems the next step should be to establish whether the penalty for extracting from within a factive complement can be characterized as an island in the terms discussed in Sprouse et al. (2012), as simply the sum of the effects of the presence of an island and of a long distance dependency. Or, is it that extraction of an argument, or indeed a subject or adjunct, induces some reduction in acceptability greater than the sum of these two individual effects. Only if factives prove to be islands for extraction, in the relevant sense, would their behavior in elliptical contexts bear on their nature. If the reduced acceptability of extraction from within a factive were merely the result of the sum of the effects of the presence of a factive verb and of a long distance dependency, then factives would prove not to be islands in the grammatically interesting sense, being reducible to independent factors. However, if they were not reducible to independent factors, the present account predicts that ellipsis should only be able to repair the violation of a factive island if factive islands were representational in nature, rather than, for example, semantic in nature, as argued by Szabolcsi and Zwarts (1997)

Another locality constraint that has been argued to result from a factor other than linearization constraints are those that have been attributed to the ECP. Lasnik and Park observe that extraction of an adjunct from a DP is unacceptable, and that Sprout Sluicing does not ameliorate such a violation (Lasnik and Park, 2003, pg. 653). See also Nakao (2009) for further discussion.

The same seems to hold of Stripping, as in (340), which only seems to have the reading in which the reading was done on the relevant shelf, rather than the intended reading in which the book Mary read was on the relevant shelf. If this judgment holds, it suggest that the constraint at play is not merely a linearization constraint.

(339) a. *Which shelf did Mary read [books e]?

b. *Mary read a book, but I don't know on which shelf.

(340) a. Mary read [a book on one of these shelves].

b. ?*Yeah, on that shelf.

Finally, the present account predicts that, if ellipsis is able to repair a violation of given locality constraint, then that locality constraint should be due to linearization constraints.

Consequently, overt string vacuous movement out of such islands should remain acceptable. Kyle Johnson suggested (p.c.) that this prediction might be able to be tested using the coordinate structure constraint (CSC). Stripping appears to repair violations of the CSC, as in (341). This predicts, all things being equal, that string vacuous movement out of a coordinate structure should be acceptable. However, examples like (342a), where *whose cat* has plausibly vacuously raised from the coordinate structure to the edge of the embedded clause, remain unacceptable. Unfortunately, such examples can be independently ruled out by appeal to a ban on the coordination of unlikes, given that (342b) is acceptable, and so such examples remain inconclusive.

(341) a. I heard that Alexi and someone saw the movie last week.

b. Yeah, James.

(342) a. *I wonder whose cat and Mary's dog are best friends. (Under non-echoic interpretation)

b. I wonder whose cat and whose dog are best friends.

5.2.3. Other Langauges

This research has entirely focused on English. In part, this was the result of constraints on designing and executing a series of experiments in a predominately English speaking context. There is, however, a rich literature on the island sensitivity of ellipsis from a broad range of languages (Fukaya and Hoji 1999; Fortin 2007; Fukaya 2007; Merchant 2001, 2004; Ince 2009; Temmerman 2013; Ince 2012, among many others). In the interest of space, I will not attempt a summary of the results of these cross-linguistic studies.

However I think it should be clear from the multitude of often contradictory claims of the island sensitivity of Stripping and Fragment Answers in the literature on English, that the judgements of these types of examples is exceedingly difficult. In discussions of ellipsis island insensitivity outside of English, the contradictory claims made of English ellipsis island (in)sensitivity are not widely cited. Often the only relevant work cited is Merchant (2004), or others reporting similar judgments. It is at least possible that this in turn, depending on the degree to which one expects a level of universality in surface level syntactic phenomena, e.g. the expression of island constraints, could unconsciously bias researchers towards expecting or accepting judgements in a given language that mirror those reported in Merchant (2004).

5.2.4. Focus Scoping and Movement

A final observation concerns the implications that the present analysis have on the analysis of the movement or scoping of the correlate in such elliptical phenomena, and the manner in which foci and other elements scope more generally. The general prediction made is that, if Parallelism holds, whatever structures result from the movement of the remnant to escape ellipsis, the mechanism by which the correlate scopes should produce a parallel structure. I argued that remnants, in both contrastive and non-contrastive Stripping move to escape ellipsis. In non-island configurations, I argued that this movement was successive cyclic, and in island contexts in one-fell-swoop. Given that there are analyses of the scope of indefinites and foci which posit that these elements scope through some sort of non-movement operator-variable relationship, e.g. Rooth (1992); Reinhart (1998), and alternative analyses which posit that these elements scope through movement, e.g. Von Stechow (1991) it seems plausible to permit the correlate for island Stripping configurations to scope via an operator-variable relationship and the correlate for island stripping configurations to scope via movement. I leave the details and inevitable implications to further research.

References

Abels, Klaus. 2011. Dont fix that island! it aint broke. In *Islands in contemporary linguistic theory conference, vitoria-gasteiz*.

Agüero-Bautista, Calixto. 2007. Diagnosing cyclicity in sluicing. *Linguistic Inquiry* 38 (3): 413–443.

Ambridge, Ben, and Adele E Goldberg. 2008. The island status of clausal complements: Evidence in favor of an information structure explanation. *Cognitive Linguistics* 19 (3): 357–389.

Baayen, R Harald, Douglas J Davidson, and Douglas M Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of memory and language* 59 (4): 390–412.

Barker, Chris. 2013. Scopability and sluicing. *Linguistics and Philosophy* 36 (3): 187–223.

Barros, Matthew. 2012a. A non-repair approach to island sensitivity in contrastive tp ellipsis. In *Proceedings from the annual meeting of the chicago linguistic society*, Vol. 48, 61–75. Chicago Linguistic Society. Chicago Linguistic Society.

Barros, Matthew. 2012b. Sluiced fragment answers: another puzzle involving islands and ellipsis. *Snippets*.

Barros, Matthew. 2014. Sluicing and identity in ellipsis. PhD diss, Rutgers The State University of New Jersey-New Brunswick.

Barros, Matthew, Patrick D Elliott, and Gary Thoms. 2013. More variation in island repair: Clausal vs. non-clausal islands. In *The proceedings of the 49th chicago linguistics society*, Vol. 49.

Barss, Andrew. 1986. Chains and anaphoric dependence: On reconstruction and its implications. PhD diss, Massachusetts Institute of Technology.

Basse, Galen. 2008. Factive complements as defective phases. In *Proceedings of the 27th west coast conference on formal linguistics*, 54–62.

Beaver, David, and Brady Clark. 2003. Always and only: Why not all focus-sensitive operators are alike. *Natural language semantics* 11 (4): 323–362.

Beghelli, Filippo, and Tim Stowell. 1997. Distributivity and negation: The syntax of each and every. In *Ways of scope taking*, 71–107. Springer.

Bošković, Željko. 2001. *On the nature of the syntax-phonology interface: Cliticization and related phenomena*, Vol. 60. Elsevier.

Bowers, John. 1998. On pseudo-gapping. *Ms., Cornell University*.

Bresnan, Joan W. 1973. Syntax of the comparative clause construction in english. *Linguistic inquiry* 4 (3): 275–343.

Bruening, Benjamin. 2015. Non-constituent coordination: Prosody, not movement. *University of Pennsylvania Working Papers in Linguistics* 21 (1): 5.

Casielles, Eugenia. 2006. Big questions, small answers. In *The syntax of nonsententials*, 117–145. John Benjamins Publishing.

Chao, Wynn. 1988. On ellipsis. PhD diss, University of Massachusetts Amherst.

Charlow, Simon. 2017. The scope of alternatives: Indefiniteness and islands. lingbuzz/003302.

Chomsky, Noam. 1970. Some empirical issues in the theory of transformational grammar. In *Goals of linguistic theory*, ed. Paul Stanley Peters, 63–130. Englewood Cliffs, N.J: Prentice-Hall..

Chomsky, Noam. 1973. Conditions on transformations. In *A festschrift for morris halle*, 232–286. New York: Holt, Rinehart and Winston.

Chomsky, Noam. 1976. Constraints on rules of grammar. *Linguistic Analysis* 2: 303–351.

Chomsky, Noam. 1977. On wh-movement. In *Formal syntax*, 71–132. New York: Academic Press.

Chomsky, Noam. 1981. Lectures on government and binding, foris, dordrecht. *ChomskyLectures on Government and Binding1981*.

Chomsky, Noam. 1986. *Barriers*, Vol. 13. MIT press.

Chomsky, Noam. 1993. *Lectures on government and binding: The pisa lectures*. Walter de Gruyter.

Chomsky, Noam. 1995. The minimalist program. *Cambridge et al*.

Chomsky, Noam. 1999. *Derivation by phase*. *Mit occasional papers in linguistics*. MIT, Department of Linguistics.

Chomsky, Noam, and Howard Lasnik. 1993. The theory of principles and parameters. *Syntax: An international handbook of contemporary research* 1: 506–569.

Chung, Sandra. 2013. Syntactic identity in sluicing: How much and why. *Linguistic Inquiry* 44 (1): 1–44.

Chung, Sandra, William A Ladusaw, and James McCloskey. 1995. Sluicing and logical form. *Natural language semantics* 3 (3): 239–282.

Cinque, Guglielmo. 1990. *Types of ā-dependencies*. MIT press.

Culicover, Peter W, and Ray Jackendoff. 2005. *Simpler syntax*. Oxford University Press Oxford.

De Cuba, Carlos Francisco. 2007. On (non) factivity, clausal complementation and the cp-field. PhD diss, State University of New York at Stony Brook.

Den Dikken, M, and A Szabolcsi. 2002a. Islands. *The Second State of the Article Book. Berlin: Mouton de Gruyter*.

Den Dikken, M, and A Szabolcsi. 2002b. Islands. *The Second State of the Article Book. Berlin: Mouton de Gruyter*.

Dikken, Marcel den, André Meinunger, and Chris Wilder. 2000. Pseudoclefts and ellipsis. *Studia linguistica* 54 (1): 41–89.

Diogo, A de A, and Masaya Yoshida. 2007. A problem for the preposition stranding generalization. *Linguistic Inquiry* 38 (2): 349–362.

Drubig, Hans Bernhard. 1994. *Island constraints and the syntactic nature of focus and association with focus*. Universitäten Stuttgart und Tübingen in Kooperation mit der IBM Deutschland GmbH.

Erteschik-Shir, Nomi, and Shalom Lappin. 1979. Dominance and the functional explanation of island phenomena. *Theoretical linguistics* 6 (1-3): 41–86.

Fiengo, Robert, and Robert May. 1994. *Indices and identity*. The MIT Press.

Fortin, Catherine Rose. 2007. Indonesian sluicing and verb phrase ellipsis: Description and explanation in a minimalist framework. PhD diss, University of Michigan.

Fox, Danny. 1995. Economy and scope. *Natural language semantics* 3 (3): 283–341.

Fox, Danny. 1999. Reconstruction, binding theory, and the interpretation of chains. *Linguistic Inquiry* 30 (2): 157–196.

Fox, Danny. 2002. Antecedent-contained deletion and the copy theory of movement. *Linguistic Inquiry* 33 (1): 63–96.

Fox, Danny. 2003. On logical form. *Minimalist syntax*.

Fox, Danny, and Howard Lasnik. 2003. Successive-cyclic movement and island repair: The difference between sluicing and vp-ellipsis. *Linguistic inquiry* 34 (1): 143–154.

Fox, Danny, and David Pesetsky. 2005. Cyclic linearization of syntactic structure. *Theoretical linguistics* 31 (1-2): 1–45.

Frazier, Michael, David Potter, and Masaya Yoshida. 2012. Pseudo noun phrase coordination. In *Proceedings of wccfl*, Vol. 30, 142–152.

Fukaya, Teruhiko. 2007. Sluicing and stripping in japanese and some implications. PhD diss, University of Southern California.

Fukaya, Teruhiko, and Hajime Hoji. 1999. Stripping and sluicing in japanese and some implications. In *Proceedings of wccfl*, Vol. 18, 145–158.

Ginzburg, Jonathan, and Ivan Sag. 2000. *Interrogative investigations*. Stanford: CSLI publications.

Griffiths, James, and Anikó Lipták. 2014. Contrast and island sensitivity in clausal ellipsis.

Syntax 17 (3): 189–234.

Haegeman, Liliane, and Barbara Ürögdi. 2010. Referential cps and dps: An operator movement account. *Theoretical linguistics* 36 (2-3): 111–152.

Heycock, C. 1995. Asymmetries in reconstruction. *Linguistic Inquiry*.

Hooper, Joan B, and Sandra A Thompson. 1973. On the applicability of root transformations. *Linguistic inquiry* 4 (4): 465–497.

Horn, Larry. 1969. A presuppositional approach to only andeven. In *5th regional meeting, chicago linguistic society*, 98–107.

Horn, Larry. 1981. Exhaustiveness and the semantics of clefts. In *Proceedings of nels*, Vol. 11, 125–142.

Huang, C-T James. 1982. Logical relations in chinese and the theory of grammar. PhD diss, MIT.

Huang, C-T James. 1993. Reconstruction and the structure of vp: Some theoretical consequences. *Linguistic Inquiry*.

Hunter, Tim, and Masaya Yoshida. 2016. A restriction on vehicle change and its interaction with movement. *Linguistic Inquiry*.

Hunter, Tim, David Potter, and Masaya Yoshida. Condition c reconstruction, clausal ellipsis and island repair. Under Review.

Ince, Atakan. 2009. *Dimensions of ellipsis: Investigations in turkish*. University of Maryland, College Park.

Ince, Atakan. 2012. Sluicing in turkish. In *Sluicing: Cross-linguistic perspectives*, eds. Jason Merchant and Andrew Simpson. Oxford University Press.

Johnson, Kyle. 2004. Introduction to transformational grammar.

Johnson, Kyle Brian. 1985. A case for movement. PhD diss, Massachusetts Institute of Technology.

Karttunen, Lauri. 1971. Implicative verbs. *Language*.

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 (3): 384–409.

Kehler, Andrew. 2002. *Coherence in discourse*. CSLI: Stanford.

Kennedy, Christopher, and Jason Merchant. 2000. Attributive comparative deletion. *Natural Language & Linguistic Theory* 18 (1): 89–146.

Kiss, Katalin É. 1998. Identificational focus versus information focus. *Language*.

Kluender, Robert, and Marta Kutas. 1993. Subjacency as a processing phenomenon. *Language and cognitive processes* 8 (4): 573–633.

Krifka, Manfred. 2006. Association with focus phrases* manfred krifka. *The architecture of focus* 82: 105.

Lappin, Shalom. 1996. The interpretation of ellipsis. In *The handbook of contemporary semantic theory*, 145–176. Blackwell.

Lappin, Shalom. 2005. A sequenced model of anaphora and ellipsis resolution. *Anaphora Processing: Linguistic, Cognitive, and Computational Modelling. Amsterdam: John Benjamins*.

Lasnik, Howard. 1972. Analyses of negation in english. PhD diss, Massachusetts Institute of Technology.

Lasnik, Howard. 1989. On the necessity of binding conditions. In *Principles and parameters in comparative grammar*. MIT Press.

Lasnik, Howard. 1999. On feature strength: Three minimalist approaches to overt movement. *Linguistic inquiry* 30 (2): 197–217.

Lasnik, Howard. 2001. When can you save a structure by destroying it? In *Proceedings-nels*, Vol. 31, 301–320.

Lasnik, Howard, and Myung-Kwan Park. 2003. The epp and the subject condition under sluicing. *Linguistic Inquiry* 34 (4): 649–660.

Lasnik, Howard, and Mamoru Saito. 1990. On the nature of proper government. In *Essays on restrictiveness and learnability*, 198–255. Springer.

Lasnik, Howard, and Tim Stowell. 1991. Weakest crossover. *Linguistic inquiry* 22 (4): 687–720.

Lebaux, David. 1990. Relative clauses, licensing, and the nature of the derivation. In *Proceedings of nels 20*, ed. Carter J, Vol. 20, 318–332. Amherst:GLSA.

Lobeck, Anne C. 1995. *Ellipsis: Functional heads, licensing, and identification*. Oxford University Press on Demand.

Marušič, Franc, and Rok Žaucer. 2013. A note on sluicing and island repair. In *Annual workshop on formal approaches to slavic linguistics, the third indiana meeting 2012*, 176–189.

Melvold, Janis. 1991. Factivity and definiteness. *MIT working papers in linguistics* 15: 97–117.

Merchant, J. 2001. *The syntax of silence: Sluicing, islands, and the theory of ellipsis*.

Oxford University Press, USA.

Merchant, Jason. 2004. Fragments and ellipsis. *Linguistics and philosophy* 27 (6): 661–738.

Merchant, Jason. 2008. Variable island repair under ellipsis. *Topics in ellipsis*.

Merchant, Jason. 2010. Three types of ellipsis. *Context-Dependence, Perspective and Relativity* 6: 141.

Merchant, Jason. 2013. Voice and ellipsis. *Linguistic Inquiry* 44 (1): 77–108.

Merchant, Jason. 2016. Ellipsis: A survey of analytical approaches. *Ms., University of Chicago*.

Morgan, Jerry. 1973. Sentence fragments and the notion sentence. *Issues in linguistics: Papers in honor of Henry and Renée Kahane*.

Morgan, Jerry. 1989. Sentence fragments revisited. *CLS* 25: 228–241.

Moritz, Luc, and Daniel Valois. 1994. Pied-piping and specifier-head agreement. *Linguistic inquiry*.

Moulton, Keir, Queenie Chan, Tanie Cheng, Chung-hye Han, Kyeong-min Kim, and Sophie Nickel-Thompson. 2017. Focus on cataphora: Experiments in context. Ms..

Nakao, Chizuru. 2009. Island repair and non-repair by PF strategies. PhD diss, University of Maryland.

Nishigauchi, Taisuke. 1990. *Quantification in the theory of grammar*. Vol. 37 of *Studies in linguistics and philosophy*. Springer.

Nunes, Jairo, and Juan Uriagereka. 2000. Cyclicity and extraction domains. *Syntax* 3 (1): 20–43.

Pesetsky, David Michael. 1997. Optimality theory and syntax: Movement and pronunciation. *Optimality theory: An overview* 134: 170.

Pesetsky, David Michael. 2000. Phrasal movement and its kin MIT press.

Postal, Paul M. 1993. Remarks on weak crossover effects. *Linguistic Inquiry* 24 (3): 539–556.

Potter, David, Michael Frazier, and Masaya Yoshida. 2017. A two-source hypothesis for gapping. *Natural Language & Linguistic Theory*.

Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford University Press on Demand.

Prince, Ellen F. 1981. Topicalization, focus-movement, and yiddish-movement: A pragmatic differentiation. In *Annual meeting of the berkeley linguistics society*, Vol. 7, 249–264.

Prince, Ellen F. 1990. Syntax and discourse: A look at resumptive pronouns. In *Annual meeting of the berkeley linguistics society*, Vol. 16, 482–497.

Reinhart, Tanya. 1991. Elliptic conjunctions-non-quantificational lf. In *The chomskyan turn*, ed. Aka Kasher, 360–384. Blackwell.

Reinhart, Tanya. 1998. Wh-in-situ in the framework of the minimalist program. *Natural language semantics* 6 (1): 29–56.

Roberts, Craige. 2006. Only, presupposition and implicature. *Ms., The Ohio State University*.

Rodrigues, Cilene, Andrew Nevins, and Luis Vicente. 2009. Cleaving the interactions between sluicing and p-stranding. *Romance languages and linguistic theory 2006*.

Romero, Maribel. 1998. Focus and reconstruction effects in wh-phrases. PhD diss, University of Massachusetts Amherst.

Rooth, Mats. 1992. A theory of focus interpretation. *Natural language semantics* 1 (1): 75–116.

Rooth, Mats. 1996. Focus. In *The handbook of contemporary semantic theory*, ed. Shalom Lappin, 271–297. Oxford: Blackwell.

Ross, John R. 1974. Three batons for cognitive psychology..

Ross, John Robert. 1967. Constraints on variables in syntax. PhD diss, Massachusetts Institute of Technology.

Ross, John Robert. 1969. Guess who. In *Fifth regional meeting of the chicago linguistic society*, Vol. 252286.

Sag, Ivan A. 1976. Deletion and logical form. PhD diss, Massachusetts Institute of Technology.

Sauerland, Uli. 1996. Guess how?.

Sauerland, Uli. 1998. The meaning of chains. PhD diss, Massachusetts Institute of Technology.

Selkirk, Elisabeth. 1996. Sentence prosody: Intonation, stress, and phrasing. *The handbook of phonological theory* 1: 550–569.

Sportiche, Dominique, Hilda Koopman, and Edward Stabler. 2013. *An introduction to syntactic analysis and theory*. John Wiley & Sons.

Sprouse, Jon, Matt Wagers, and Colin Phillips. 2012. A test of the relation between working-memory capacity and syntactic island effects. *Language* 88 (1): 82–123.

Stainton, Robert. 2006. Neither fragments nor ellipsis. *The syntax of nonsententials*.

Stowell, Timothy Angus. 1981. Origins of phrase structure. PhD diss, Massachusetts Institute of Technology.

Szabolcsi, Anna. 2010. *Quantification*. Cambridge University Press.

Szabolcsi, Anna, and Frans Zwarts. 1997. Weak islands and an algebraic semantics for scope taking. In *Ways of scope taking*, 217–262. Springer.

Tancredi, Christopher Damian. 1992. Deletion, deaccenting, and presupposition. PhD diss, Massachusetts Institute of Technology.

Temmerman, Tanja. 2013. The syntax of dutch embedded fragment answers: On the pf-theory of islands and the wh/sluicing correlation. *Natural Language & Linguistic Theory* 31 (1): 235–285.

Thoms, Gary. 2010. 'verb floating'and vp-ellipsis: towards a movement account of ellipsis licensing. *Linguistic variation yearbook* 10 (1): 252–297.

Uriagereka, Juan. 1999. Multiple spell-out. *Current Studies in Linguistics* 32: 251–282.

Valmala, Vidal. 2007. The syntax of little things. In *17th colloquium on generative grammar, girona*.

van Craenenbroeck, Jeroen. 2010. Invisible last resort: A note on clefts as the underlying source for sluicing. *Lingua* 120 (7): 1714–1726.

Van Craenenbroeck, Jeroen. 2010. Invisible last resort: A note on clefts as the underlying source for sluicing. *Lingua* 120 (7): 1714–1726.

Van Craenenbroeck, Jeroen, and Anikó Lipták. 2006. The crosslinguistic syntax of sluicing: Evidence from hungarian relatives. *Syntax* 9 (3): 248–274.

Van Craenenbroeck, Jeroen, and Anikó Lipták. 2013. What sluicing can do, what it can't, and in which language: on the cross-linguistic syntax of ellipsis. In *Diagnosing syntax*, eds. Lisa Lai-Shen Cheng and Norbert Corver, Vol. 46, 502–536. Oxford University Press.

Vicente, Luis. 2008. Syntactic isomorphism and non-isomorphism under ellipsis. *Ms., University of California, Santa Cruz.*

Von Fintel, Kai. 1994. Restrictions on quantifier domains. PhD diss, University of Massachusetts at Amherst.

Von Stechow, Arnim. 1991. Focusing and backgrounding operators. *Discourse particles* 6: 37–84.

Wagner, Michael. 2006. Association by movement: evidence from npi-licensing. *Natural Language Semantics* 14 (4): 297–324.

Wasow, Thomas. 1972. Anaphoric relations in english. PhD diss, Massachusetts Institute of Technology Cambridge.

Weir, Andrew. 2014. Fragments and clausal ellipsis. PhD diss, University of Massachusetts Amherst.

Wellwood, Alexis, Roumyana Pancheva, Valentine Hacquard, and Colin Phillips. 2017. The anatomy of a comparative illusion. Ms..

Wexler, Kenneth, and Peter W Culicover. 1977. Some syntactic implications of a theory of language learnability. *Formal syntax*.

Williams, Edwin S. 1977. Discourse and logical form. *Linguistic Inquiry*.

Winkler, Susanne. 2013. Syntactic diagnostics for island sensitivity of contrastive focus in ellipsis. In *Diagnosing syntax*, 463–484. Oxford University Press Oxford.

Zwarts, Frans. 1998. Three types of polarity. In *Plurality and quantification*, 177–238. Springer.

APPENDIX A

Participant Experiments 1-5

A.1. Instructions Experiments 1-4

“In this survey, you’ll be reading short dialogues and rating how natural one of the phrases in each of these dialogues sounds.

In each dialogue, one phrase will be underlined. We would like you to rate how natural the underlined phrase sounds, on a scale of one to seven. It is important to consider the naturalness of the underlined phrase in the context of the whole dialogue, so please read the entire dialogue carefully.

Some of these underlined phrases will seem natural. A natural phrase is one that you might say yourself or one that you wouldn’t be surprised to hear someone else say, even if you wouldn’t say it yourself. If the phrase sounds natural, you should give it a high rating. Some of the other underlined phrases will seem less natural. These less natural phrases would be surprising to hear, they might sound like a mistake, or they might just sound odd. You should give these unnatural phrases a low rating. Please use the whole scale; if a phrase sounds completely natural or completely unnatural, don’t be afraid to assign it a seven or a one.

Even though these dialogues are written, try not to think of them as formal writing. When you are reading them, you could imagine that the scene is being spoken by regular people. (You could even try saying them aloud to yourself.) In other words, don't worry about what you learned about proper language in school. Just use your intuition about what a natural English phrase would sound like."

A.2. Instructions Experiment 5

"In this survey, you'll be reading short dialogues and rating how plausible it would be for the underlined phrases in the dialogue to refer to the same person.

In each dialogue, there will be two or more phrases that are underlined. We would like you to rate how plausible it would be for these words to refer to the same person, on a scale from one to seven. It is important to consider this plausibility in the context of both sentences in the dialogue, so please read the entire dialogue carefully.

In some of the dialogues, you might not be surprised to discover that the underlined words refer to the same person. For example, it would be very plausible for 'John' and 'he' to refer to the same person in this sentence: '<u>John</u> thought <u>he</u> was happy'. Although the underlined words could also refer to different people, because it is plausible for them to refer to the same person, you should give these examples like this a high rating. In other dialogues, it will be less plausible, or even impossible, for the underlined phrases to refer to the same person. In the next example, it is very implausible that 'her' and 'Mary' would refer to the same person: '<u>She</u> told <u>Mary</u> the secret.' You should give these cases a low rating. Please use the whole scale; depending on how plausible it is that the underlined phrases refer to the same person, don't be afraid to give a rating of seven (very plausible) or a one (very implausible).

Even though these dialogues are written, try not to think of them as formal writing. When you are reading them, you could imagine that the scene is being spoken by regular people. (You could even try saying them aloud to yourself.) In other words, don't worry about what you learned about proper language in school. Just use your intuition about how likely it is that the underlined words refer to the same person."