



THE ROLE OF QUANTIFIER RESTRICTION IN ILLUSORY NPI LICENSING

Thesis submitted to the Faculty of Linguistics, Philology and Phonetics
at the University of Oxford for the Degree of Doctor of Philosophy

by

LUIS AUGUSTO HILDEBRANDT BELMONT

WOLFSON COLLEGE

UNIVERSITY OF OXFORD

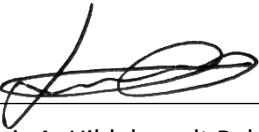
Michaelmas Term, 2024

Word Count: 63,770

I hereby certify that this thesis is the result of my own work except where otherwise indicated and due acknowledgement is given.

I also certify that this thesis is 63,770 words long excluding the bibliography.

SUBMITTED: 24/10/2024



Luis A. Hildebrandt Belmont 24/10/2024

THE ROLE OF QUANTIFIER RESTRICTION IN NPI LICENSING

LUIS AUGUSTO HILDEBRANDT BELMONT

Wolfson College

University of Oxford

Thesis submitted to the Faculty of Linguistics, Philology and Phonetics
at the University of Oxford for the Degree of Doctor of Philosophy

Michaelmas Term, 2024

Contents

Contents.....	iii
Abstract	v
Chapter 1: The Licensing of Negative Polarity Items.....	1
1.1 INTRODUCTION.....	1
1.2 MOTIVATION.....	3
1.3 BACKGROUND	4
1.3.1 <i>Negative Polarity Items (NPIs) Licensing</i>	4
1.3.2 <i>Upward and Downward Entailment</i>	5
1.3.3 <i>Cases of pragmatic licensing</i>	8
1.3.4 <i>Memory access-mechanisms</i>	12
1.3.5 <i>Types of Memory Mechanisms</i>	12
1.3.6 <i>Similarity-Based interference</i>	15
1.3.7 <i>Intervention</i>	24
1.4 THE CURRENT RESEARCH.....	29
1.4.1 <i>Theoretical questions</i>	31
1.4.2 <i>Pilot experiments</i>	33
Chapter 2: Illusory Licensing and the Universal Quantifier – Identifying the quantifier’s restriction clause	39
2.1 INTRODUCTION.....	39
2.2 BACKGROUND	40
2.2.1 <i>Illusory licensing</i>	40
2.2.2 <i>Quantifiers in NPI licensing</i>	42
2.3 EXPERIMENTAL FINDINGS.....	45
2.3.1 <i>Experiments 1 and 2: Illusion effects as a result of Quantifiers in NPI licensing</i>	45
2.3.2 <i>Experiment 1: Do illusions arise with the Universal Quantifier?</i>	45
2.3.3 <i>Experiment 2: Are illusions caused by general properties of quantifiers? The case of the existential quantifier some</i>	53
2.3.4 <i>Experiments 3 and 4: Illusory NPI licensing as a matter of ongoing restriction / scope calculations</i>	58
2.3.5 <i>Experiment 3: Identifying the universal quantifier’s restriction – Prenominal Modification</i> 59	
2.3.6 <i>Experiment 4: Identifying the universal quantifier’s restriction – Postnominal Modification</i>	64
2.4 GENERAL DISCUSSION.....	68
2.5 CONCLUSIONS	73
Chapter 3: Quantifiers in NPI Licensing – Self-Paced Reading Studies.....	75
3.1 INTRODUCTION.....	75
3.2 BACKGROUND	76
3.3 EXPERIMENTAL FINDINGS.....	87
3.3.1 <i>Experiment 5: Evidence of illusions in NPI licensing with online measures</i>	89
3.3.2 <i>Experiment 6</i>	96
3.4 GENERAL DISCUSSION.....	104
3.5 CONCLUSIONS	108
Chapter 4 - Identifying Restrictors Under Ambiguity	110

4.1	INTRODUCTION.....	110
4.2	BACKGROUND	111
4.2.1	<i>Garden Paths Sentences</i>	111
4.2.2	<i>Models of ambiguity resolution</i>	112
4.2.3	<i>Evidence from research</i>	118
4.3	EXPERIMENTAL FINDINGS	123
4.3.1	<i>Experiment 7: The role of NPI licensing on Garden Path sentences</i>	123
4.3.2	<i>Experiment 8: NPI presence manipulation on the Universal Quantifier</i>	141
4.3.3	<i>Experiment 9: NPI presence manipulation on the Determiner</i>	153
4.4	GENERAL DISCUSSION.....	162
4.4.1	<i>The garden path effect:</i>	164
4.4.2	<i>Effects within the relative clause</i>	166
4.5	CONCLUSIONS	170
Chapter 5: Summary of Results and General Discussion.....		172
5.1	SUMMARY	172
5.2	BRIEF REVIEW OF EXPERIMENTAL RESULTS	173
5.3	META STUDY OF EXPERIMENTS 1-4	178
5.4	IMPLICATIONS FOR NPI LICENSING.....	182
5.4.1	<i>Identifying local monotonicity as downward-entailing</i>	182
5.4.2	<i>Implications for illusory licensing</i>	185
5.4.3	<i>Implications for ambiguity resolution</i>	188
5.4.4	<i>Implications on memory mechanisms</i>	189
5.5	FUTURE RESEARCH:	193
5.6	CONCLUSION:.....	194
References		195

Abstract

Negative polarity items (NPIs) like *ever* have been of particular interest and much debate in psycholinguistics because they rely on semantic and pragmatic processes which are not involved in the processing of other dependencies. There are currently two competing accounts that explain how NPIs are processed. One theory considers that NPIs are licensed through memory retrieval mechanisms (Vasishth et al., 2008), a second theory considers that NPI licensing is primarily a matter of semantic and pragmatic properties of the sentence which contains it (Xiang, Dillon, & Philips, 2009; Xiang, Grove, & Giannakidou, 2013).

Given their specific semantic structure, universal quantifiers such as *every* are able to license NPIs within their restrictor and not within their scope. Therefore, they offer an interesting yet unexplored testing bed for both licensing hypotheses. We examined cases of illusory licensing to see whether clearly establishing a restrictor could reduce illusions in offline and online measures. We found that illusions arise with the universal quantifier, but that manipulating the restrictor can subvert illusions in offline ratings. We also examined whether licensed NPIs could reduce garden path effects by reducing structural ambiguity that produces them.

We found that licensed NPIs can help the parser choose the reduced relative clause interpretation at ambiguity regions and pre-empt the later appearance of garden path effects. We conclude that NPI licensing is driven primarily by semantic/pragmatic properties but memory retrieval plays a role. The universal quantifier provides the parser with information about an upcoming downward-entailing context and that the NPI itself provides information to the parser that the context where it is located is the appropriate one, regardless of whether it is located inside the appropriate one, which explains the rise of illusions. At this point the parser checks the availability of a context-providing element possibly through a memory retrieval mechanism.

Chapter 1: The Licensing of Negative Polarity Items

1.1 Introduction

The licensing of Negative Polarity Items (NPIs), words like ‘ever’ and ‘any’, is a kind of long-distance dependency that has been important to linguistic theory. Unlike other dependencies, NPI licensing has been argued to not rely solely on syntactic category and structural relationships for their distribution. Instead, they also rely on the semantic and pragmatic properties of whole propositions (Fauconnier, 1975; Ladusaw, 1979; Giannakidou, 1998, 2011). These properties of NPIs have been important for our understanding of the interaction of syntax, semantics, and pragmatics in linguistic theory. They also raise questions for psycholinguistic theory: How are the licensing conditions for NPIs realized during real-time sentence processing? In particular, NPI licensing likely relies on memory processes that retrieve previously encoded information important for the establishment of long-distance dependencies raising questions about how this information is retrieved from memory during sentence processing.

Classically, certain types of NPIs are thought to be licensed under the scope of a downward-entailing operator that c-commands them (Ladusaw, 1979). However, this explanation does not seem to account for the multitude of contexts in which NPIs appear, which has led linguists to try to find the correct theoretical characterization for NPI licensing. It is important to have a theory for NPI licensing as a starting point in order to establish a solid ground for this research. For this purpose, we will review the current theoretical studies and research on NPI licensing in this chapter. The theoretical approach that we take will have important consequences for raising the psycholinguistic questions that we will formulate to explain both the online processing behaviour of NPI licensing, as well as the mechanisms that drive this process.

Because long-distance dependencies (LDDs) involve a relationship between two linguistic elements that can be separated by several other elements, it has been argued that language processing relies on memory mechanisms that allow the parser to rapidly encode and navigate the linguistic structure, and retrieve previously encoded information (Phillips, Wagers, & Lau, 2011). Long-distance dependency research has focused on two main possible ways in which memory mechanisms may work: a serial-search mechanism, where information is accessed an item at a time; and a cue-based, content addressable retrieval mechanism, in which all items are accessed based on their content features at the same time. While some studies have revealed a *selective profile* in how memory is accessed that is dependent on each type of dependency, that is to say that different dependencies may use different ways of retrieving information from memory (Xiang et al., 2009; Phillips et al., 2011; Dillon et al., 2013; Parker, 2014), this has been a matter of recent debate and some studies consider that we should consider all these dependencies as functionally similar (Patil et al., 2016; Jäger et al., 2020).

Within this debate, some dependencies, like subject-verb agreement, have been shown to be prone to grammatical illusions, that is to say, unlicensed dependents are sometimes judged to be acceptable. Much of this research has focused on interference constructions, sentences in which a potential, but structurally inaccessible licenser is available, making these dependencies seem grammatical, even if momentarily (Bock & Miller, 1991; Vigliocco, & Nicol, 1998; Solomon & Pearlmutter, 2004; Phillips, Wagers, & Lau, 2011; Parker, 2014). Conversely, according to some researchers, reflexive anaphors, another type of dependency, seem to be strongly grammatically constrained and immune to this effect (Sturt, 2003; Kazanina, et al., 2006; Phillips, 2006; Kush, 2013), though more recent evidence (Parker and Phillips, 2017) seems to suggest that this may not be the case. NPIs have shown a mixed profile, with most studies finding evidence for grammatical illusions (Drenhaus, Saddy, & Frisch, 2005; Xiang, Dillon, & Phillips, 2009; Vasishth, et al., 2008), and others revealing selective immunity to these effects in some manipulations (Parker, 2014; Orth et al., 2021).

The study of how different types of dependencies are established has been crucial to understanding how language representations are hierarchically structured in our minds, and how these representations are navigated during online sentence processing. The licensing of NPIs has been an important part of these studies because, while similar to other dependencies in that they too require a licensor in a structurally licit position, the relationship between the NPI and its licensor seems to be different in that it does not solely rely on syntactic constraints.

1.2 Motivation

The aim of the current research is to fill in some of the gaps left from previous studies by focusing on the semantic and pragmatic aspects of Negative Polarity Licensing during real-time sentence processing using the universal quantifier *every*. While most research has concentrated on NPI licensing using negative quantifiers (such as *no*), which create downward-entailing (DE) contexts that license NPIs within their restrictor clause and their scope, universal quantifiers only produce DE contexts within their restrictor clause, but not within their nuclear scope. This unique property allows us to explore how semantic properties such as restriction and scope are identified, calculated, and integrated into the licensing process of NPIs, and can, therefore, provide us with an interesting testing bed for semantic and pragmatic processes involved in NPI licensing.

Because of this distinct entailing environment, universal quantifiers provide us with a controlled way to look at illusory licensing effects that the negative licensor cannot provide. For a universal quantifier to license an NPI, the NPI must necessarily be located within the quantifier's restrictor, so it is possible to "misplace" it inside the nuclear scope and observe whether illusory licensing effects occur, and whether they can be manipulated by tweaking the semantic context of the sentence.

In this sense, it allows us to test non-structural properties that have been argued to drive the licensing process, especially in the case of NPIs where the semantic and pragmatic properties of

the sentence seem to have an important role in the licensing process. This thesis will conduct a series of experiments examining the predictions from cases involving c-command, memory retrieval mechanisms, semantic restriction and scope, and pragmatic inferencing, trying to further address how NPIs are licensed in real time.

1.3 Background

1.3.1 Negative Polarity Items (NPIs) Licensing

In terms of their distribution, it has been argued that NPIs require to be under the scope of a negative licensor that c-commands them, hence their name (Fauconnier, 1975; Ladusaw, 1979; Giannakidou, 1998, 2011). As can be observed in (1b) compared to (1a), where the negative element is contained within an embedded sentence and does not c-command the NPI. Since the NPI is not under the scope of the licensor, the sentence is ungrammatical.

- (1) a. No teacher [with a beard] believes that the student will ever fail the exam.
b. *The teacher [with no beard] believes that the student will ever fail the exam.

While early accounts of NPI assumed that licensors had a specific feature in their grammar that allowed them to license NPIs as long as they were in a specific position relative the licensor (Klima, 1964), NPI licensing does not only require that the licensing element should c-command the NPI, but also a semantic requirement that the NPI must be in a negative logical environment, usually provided by the licensor. This can be observed in (2), where the NPI *ever* is only grammatical in a negative sentence (3a).

- (2) a. Charlie hasn't *ever* been abroad.
b. *Charlie has *ever* been abroad.

Nevertheless, NPIs have been observed in other contexts where the negative environment is not that explicit. For example, in the sentences shown in (3b-d) we can observe that, in comparison with (3a), NPI licensing has wandered away from the original notion of negativity. While both the quantifier 'few' and the verb 'doubts' convey negativity in their meanings, the conditional case (3d) may give us a stronger suggestion that it may not be negation at all that is licensing NPIs.

Moreover, the asymmetry observed in (4) where the universal quantifier ‘every’ admits the NPI within the subject (restrictor) (4a), but not within the verbal phrase (scope), gives us further evidence that the relation between the NPI and its licenser is more about meaning, rather than just a syntactic constraint, or feature as originally established by Klima (Jacobson, 2014).

- (3) a. No linguistic students have *ever* set foot on the moon.
b. Few linguistic students have *ever* set foot on the moon.
c. Sarah doubts that she will *ever* buy a new smartphone.
d. If Holly *ever* buys a new computer, she will be able to skype with Henry.

- (4) a. Every linguistic student who *ever* set foot on the moon was awarded a prize.
b. *Every linguistic student who set foot on the moon was *ever* awarded a prize.

Due to the nature of the previous examples, negation was not enough anymore to describe the context in which NPIs appear. It was with Fauconnier (1975) and Ladusaw (1979) that a new explanation arose for all the contexts above. In order for NPIs to be licensed, they have to be under the scope of a downward entailing (DE) licenser that c-commands them.

1.3.2 Upward and Downward Entailment

Upward entailing (UE) and downward entailing (DE) environments are defined as the relation in which a set can entail a superset or a subset, respectively. That is to say, a constituent is within an upward entailing environment if the proposition that contains it can entail another proposition in which a superset of that same constituent is also true (5). Both (5a) and (5b) present an ownership of an animal. Since a poodle is a breed of dog, the set of poodles also belong to the set of dogs because being a poodle entails being a dog. Considering this, (5c) shows the implication that if John owns a poodle, John owns a dog. Therefore, the environment of the sentence is upwards entailing if by substituting in a superset of the individual in the sentence preserves the truth of the original statement.

- (5) a. John owns a poodle.
b. John owns a dog.

c. John owns a poodle \rightarrow John owns a dog.
set superset

Certain operators, like negation, reverse this entailment relationship and create a downward entailing environment. As we defined above, since poodles are a breed of dog, we can say that poodles is a subset of the set dog. Therefore, not owning a dog, entails not owning any breed of dog. Shown in (6c), the sentence in (6a) entails the sentence in (6b). The environment of a sentence will be downward entailing if substituting with a subset of the constituent in the sentence preserves the true of the original statement.

- (6) a. John doesn't own a dog.
b. John doesn't own a poodle.
c. John doesn't own a dog \rightarrow John doesn't own a poodle.
set subset

The previous examples show two different directions of entailments. As the examples in (5) show, affirmative sentences are typically upward entailing environments, while negative sentences are typically downward entailing environments, as are the sentences presented in (3), presented again in (7), below.

In the case of (7a) and (7a'), semantics being a subset of linguistics, the fact that no linguistic student ever set foot on the moon implies that the students of semantics, also haven't ever set foot on the moon. The same applies for the examples in (7b) and (7b'), where both sentences maintain the same truth value. In (7c) and (7c'), as well, if Sarah doubts that she will ever buy a new smartphone, then the same holds for any kind or brand of smartphone.

Furthermore, the case of the conditional (7d), which clearly doesn't carry a negative connotation in its meaning, shows the same pattern as described in the other examples. Consider that the conditional is formed as [if S_1 then S_2], where the value of S is a set of possible worlds. Given an S_3 , which entails S_1 : "Holly buys a new laptop" \rightarrow "Holly buys a new computer". If S_3 entails S_1 , then the worlds in which S_3 is true is a subset of the worlds in which S_1 is true. As these sentences are now, they present a case of upwards entailing. Nonetheless, when we substitute the subset

sentence in the antecedent of the conditional: [If [Holly buys a new laptop] then [she will be able to skype with Henry]], truth is preserved, which shows that the antecedent of a conditional is a DE environment (Jacobson, 2014).

- (7) a. No linguistic students have *ever* set foot on the moon.
a'. No semantics students have *ever* set foot on the moon.
- b. Few linguistic students have *ever* set foot on the moon.
b'. Few semantics students have *ever* set foot on the moon.
- c. Sarah doubts that she will *ever* buy a new smartphone.
c'. Sarah doubts that she will *ever* buy a new iPhone.
- d. If Holly *ever* buys a new computer, she will be able to skype with Henry.
d'. If Holly *ever* buys a new laptop, she will be able to skype with Henry.

The same is true for the case of 'every'. While the scope (the VP) of 'every' is UE, the restriction (the NP) is, in fact, DE. Considering a sentence structure like [every N VP], it will hold true as long as the set N is a subset of the VP set (8a). If N (poodle) is a subset of VP (dog), it will also hold true that a subset of N (brown poodle) will also be a subset that same VP (8b). Therefore, substituting N with an even smaller subset will preserve the truth value of the original sentence. Because we can substitute with a subset of an individual in the restriction of every, we know it is a DE environment.

- (8) a. Every [poodle]_N [is a dog]_{VP}
b. Every [brown poodle]_N [is a dog]_{VP}

If we do the same with the example in (4a), presented below as (9b), we can observe that the same holds true when we compare it with (9a). If every linguistic student was awarded a prize, it is still true that the subset of those students who set foot on the moon were also awarded a prize. Therefore, the restriction of 'every' is a downward entailing environment.

- (9) a. Every linguistic student was awarded a prize.
b. Every linguistic student who ever set foot on the moon was awarded a prize.

As we can observe, all the licensing context that we have mentioned can be described to have the same characteristic: they are all downward entailing environments. This goes to show that direct semantic licensing is an important factor in the licensing of negative polarity items and that their distribution is not solely determined by syntactic constraints.

Furthermore, it has also been argued that c-command is not a necessary requirement for NPI licensing, but, since scope mirrors c-command, it could be an epiphenomenon of NPIs being required to be under the scope of their licensor (Xiang et al., 2009). This could explain cases in which scope does not align with c-command, but NPIs are still licensed, as in (10). This also shows another way in which NPIs are special compared to other dependencies, since even syntactic constraints such as c-command may not be necessary for NPI licensing.

- (10) Nobody's mother has ever complained about his grades.

1.3.3 Cases of pragmatic licensing

NPIs can also be found in other cases where negativity or downward entailment is not as clear as in the cases explained above, especially in cases of implicit negation, in which the source of the negative meaning is pragmatic, and not asserted by the meaning of the word. A good example of this are "emotive" predicates like *'be surprised'*, *'be amazed'*, *'be sorry'*, *'be irritated'*, *'be disappointed'*, *'regret'*, etc. (Xiang et al., 2015).

A good example of why emotives don't have a clear DE environment can be observed in the comparison between the examples in (11) and (12) where one instance of the emotive 'be surprised' can elicit a DE environment and license an NPI (11), while at the same time, the other example cannot elicit a DE environment in the sentence, but still license the NPI. The example in (11a) shows that an affirmative sentence with an emotive can license an NPI. Despite being an affirmative sentence, (11a) is able to produce a DE environment regardless of the lack of a

negative element or a negative meaning and thus is able to generate set to subset entailments like (11b) which also maintains the truth of the original statement: it's still true that John was surprised that Mary owned a poodle, as well as any other breed of dog.

- (11)a. John was surprised that Mary ever owned a dog.
- b. John was surprised that Mary ever owned a poodle.
- c. Implicature: *John expected that Mary wouldn't have ever owned a dog.*

Despite what was observed in (11a, 11b), the same does not occur for the examples in (12) even if the same emotive is used. As opposed to the previous examples, the sentences in (12) have an upward entailing environment and the direction of the entailments is opposite to the one observed in (11). We can see that if John was surprised that Mary found her poodle (set), it would still be true that he is surprised that she found her dog (superset). The same cannot be said of the inverse relationship, since John could have been surprised that she found her Pomeranian, but not necessarily her poodle.

- (12)a. John was surprised that Mary (ever) found her dog.
- b. John was surprised that Mary (ever) found her poodle.
- c. Implicature: *John expected that Mary wouldn't ever find her dog.*

For cases such as in (11), it has been argued that apart from semantic licensing, pragmatic processes can also play an important part in NPI licensing since speakers can draw negative inferences from positive statements which can be enough to license NPIs and even change the entailment context of the sentence (Linebarger, 1987; Giannakidou, 2006; Chierchia, 2006; Jacobson, 2014). This way, (11a) can give rise to the negative implication shown in (11c). Considering that we are surprised of a certain outcome in a situation, means that we were expecting something not to happen.

In the case of (12), though the sentence itself may not have a clear licensing environment like in (11), the implication that is drawn from the sentence through the meaning of the predicate 'be surprised' (12c), can give us the negative/DE environment that we needed to license the NPI.

This opens important questions about what the scope of this implicature may be, and whether it has any structural restriction whatsoever.

The same thing does not happen when we compare the sentence in (11a) with the sentence in (13), which uses the verb 'glad'. This comparison tells us that there is something about the meaning of the verb 'surprised' that triggers the rise of the negative implication. In other words, we can see in cases such as this that NPI licensing can also be affected by pragmatic elements, something that does not seem to happen with other dependencies.

(13)*John was glad that Mary ever owned a dog.

Furthermore, NPIs can also be licensed in some types of questions such as the ones shown in (14). While interrogative sentences have been classified together with negative sentences, entailing direction is difficult to define, and furthermore, even if we could define questions to be DE, not every question can license NPIs, only yes/no questions (14a,b) and wh-questions if they are plural, but not when they are singular (14c,d) (Guerzoni & Sharvit, 2007).

- (14) a. Has George ever been to Paris?
b. Did the student have any complaints?
c. Which students who have any interest in polarity attended the seminar?
d. *Which student who has any interest in polarity attended the seminar?

As observed in the previous examples, an NPI licensor can be drawn from different classes of linguistic elements, like quantifier determiners, verbs, conditionals (9); as well as pragmatic implicatures (11) and (13), and questions (14). Additionally, the class of elements that conform NPIs are not determined by syntactic category, since they can be adverbs (*ever*), determiners (*any*), as well as a class of expressions known as *minimizers*, that somehow describe minimal amounts and comprise verbal phrases like 'lift a finger' or noun phrases like 'a red cent' (15).

- (15)a. Jake didn't lift a finger to help at the party.
b. *Jake lifted a finger to help at the party.

Furthermore, there is a class of words known as ‘strong’ NPIs that includes expressions like ‘in weeks’, ‘in years’, etc. that can only be licensed in a subset of the contexts in which ‘weak’ NPIs, like ‘any’ or ‘ever’, can appear, most specifically, only in those that are very strongly negative (Zwarts, 1998; Giannakidou, 2011, Chierchia, 2014). We can observe this in the examples in (16). In contrast with the example in (16a), where “in weeks” is licensed, the examples in (16b, c, d) are not able to license the expression, while they could still license weak NPIs, as we have seen above. Additionally, strong NPIs need to be licensed within the same domain as their licensor, as can be observed in (17a), as opposed to weak NPIs, which can (17b). This also seems to enforce a restriction to pragmatic licensing, since an example as (17c) also doesn’t seem to be licensed.

- (16) a. No one saw John in weeks.
b. *Few students saw John in weeks.
c. *Only Mary saw John in weeks.
d. *Every student who saw John in weeks came to school late.

- (17) a. *John didn’t believe that Mary saw Jake in weeks.
b. John didn’t believe that Mary ever saw Jake.
c. *John was surprised that Mary saw Jake in weeks.

As we can see, NPIs have a variety of context in which they can be licensed, and the great variety of them can pose a difficulty in explaining how all of them are licensed. Likewise, the difference between strong and weak NPIs could demonstrate that different NPIs have different licensing requirements. For example, since strong NPIs need to be bound within the domain of their licensor, it is possible that a syntactic requirement like c-command could be much more involved in their licensing, and this could be reflected in the memory mechanisms that are involved in NPI licensing, since they have different predictions about how syntactic constraints are considered during memory retrieval, as we will see in the next section.

1.3.4 Memory access-mechanisms

Because linguistic dependencies establish a long-distance relationship between two linguistic elements, it has been argued that their processing must rely on memory mechanisms. These memory mechanisms are used by the parser to navigate the linguistic structure and retrieve the necessary information to establish the dependency (McElree, Foraker, & Dyer, 2003; Phillips, Wagers, & Lau, 2011). Current research has focused on two distinct memory mechanisms to explain how linguistic data is retrieved during processing: (i) Serial search and (ii) direct-access, content-addressable retrieval.

1.3.5 Types of Memory Mechanisms

1.3.5.1 Serial Search Mechanisms

The serial search mechanism executes pairwise evaluations of the contents of memory by comparing the desired information to each item, one at a time, accessing them by location, rather than by content (Gillund & Gibson, 2005). There are two possible ways for these mechanism to work. On one hand, a low-level serial search which operates in an ordered manner, but without making use of the higher structure of the input. Rather, it evaluates each item sequentially following back the sentence's word order attempting to match some criteria established by the task. This type of mechanism has been argued for certain types of relational information, such as establishing temporal and spatial order (McElree, 2001; McElree, 2006; Grodner & Gibson, 2005;).

On the other hand, a structurally driven serial search which follows the sentence's representational structure and restricting access to constituents in non-relevant structural positions along a path defined in terms of relational properties. This kind of serial search mechanism may use a node-by-node strategy to determine a search path and the syntactic relationships between nodes to determine where to look for the desired information (Berwick & Weinberg, 1984; McElree et al., 2003).

1.3.5.2 *Direct-Access, Content-Addressable Mechanisms*

The direct-access, content-addressable mechanism, in contrary, accesses all the items in memory at the same time (i.e., in parallel) using a set of cues to check for the necessary content cues. In this cases, content refers to inherent properties of an individual item such as *nominative* or *plural*, or *female*. In some cases, some theories have also included syntactic relations as part of the content cues needed to retrieve information. Most notably, Vasishth et al. (2008) assume the feature [+c-command] in addition to a [+negative] feature in order to explain how NPIs use content information to identify their licensors. Similarly, other authors have proposed a feature such as [+local] for reflexive dependencies (Kush, 2013; Parker and Phillips, 2017). Nonetheless, there has been debate in the literature as to whether relational features can be considered natural to a content-addressable retrieval system. (Gillund & Shiffrin, 1984; Lewis & Vasishth, 2005; Lewis et al., 2006; McElree, 2000; inter alia). This mechanism works through a function that checks the content features of the items in memory that match the probe initiated at the dependent element, thus increasing the probability of retrieving an item based on ‘probe-to-item’ strength that is described in terms of activation (Lewis & Vasishth, 2005; Phillips, Wagers, & Lau, 2011).

1.3.5.3 *Mechanisms Predictions*

Both of these mechanism represent different strategies for recovering information from memory, and both of them make contrasting predictions for the temporal dynamics and the accuracy of retrieval when there are items in structural irrelevant or relevant positions, and for cases where the target for retrieval does not share all of the necessary features required by the dependent element.

Because serial-search is constrained to follow the sentence structure, checking each element at a time, the speed in which the dependency is established depends on how much intervening material is found between the dependent element that initiates the process and its antecedent. In other words, the more linguistic elements the parser is required to navigate through to arrive

to the desired target, the longer it will take to retrieve the target (Ratcliff, 1978; McElree et al, 2003). Furthermore, because serial-search is structurally driven, elements in an illicit position should be considered inaccessible for retrieval and, therefore, should not affect the accuracy of the retrieval process.

In summary, serial-search processes predict variation in retrieval speed depending on how many elements are found between the dependency constituents, but accurate retrieval of the correct antecedent even if there are other elements that share features with the dependent element. We can see this in (18a) where the subject of the sentence is not considered as a possible antecedent, even though it contains the same cues as the anaphor. Furthermore, (18b) shows that this is not a string locality effect, since the anaphor requires its antecedent to be within the same domain and c-command it (Principle A, Chomsky, 1981).

- (18) a. Mary thinks that Jessica likes herself.
- b. The man with the boy likes himself.

On the other hand, because content-addressable mechanisms probe all items in memory at the same time, this mechanism predicts that the amount of material that separate the elements in the dependency should not affect the processing speed. Furthermore, that the probability of finding the correct item in memory is decreased based on how many other items share features with the retrieval probe. In other words, other items within the sentence are capable of affecting the processing accuracy based on how similar they are with the desired target, something that is known as *similarity-based interference*. This implies that partially-matched items could be retrieved erroneously, especially in cases like (19a) where there are no fully-matched licensors, or that a fully-matched item in an illicit position could be retrieved in cases like (19b), interfering with the retrieval process (McElree et al., 2003; Van Dyke & McElree, 2006).

- (19) a. *The key_[singular] to the cabinet_[singular] were_[plural] probably destroyed by the fire.
- b. *The key_[singular] to the cabinets_[plural] were_[plural] probably destroyed by the fire.

Furthermore, previous research has shown that the sensitivity of the memory mechanisms to structural properties of potential licensors varies with time, thus having more time to process increases accuracy (Vasishth et al., 2008; Xiang et al., 2009; Wagers et al., 2009). This is reflected by the difference between online and offline task where having more time could reflect an improvement over the signal-to-noise ratio (Lewis & Phillips, 2014). These accounts predict that the distance between the dependent element and its licensor should not impact illusory licensing, since it's assumed that components in the licensing process do not change over time.

1.3.6 Similarity-Based interference

Similarity-based interference in sentence processing is the disruption of the licensing process determined by the inclusion of an item ('distractor') which shares features with the targeted antecedent but is located in a structurally inappropriate position in the sentence.

Because both types of mechanisms make different predictions about how similar items affect retrieval accuracy, similarity-based interference has been argued to be a good way to test which mechanisms are involved during sentence processing. While content-addressable mechanisms predict that similar items will decrease the rate of accuracy, serial search mechanism predict that structurally irrelevant items will have no effect on accuracy, regardless of similarity.

The similarity-effects that can be observed are based on the comprehenders' interpretation of the sentences grammaticality, either by considering an ungrammatical sentence grammatical ("illusions of grammaticality"), or a grammatical sentence ungrammatical ("illusions of ungrammaticality"). When this happens we can consider that the memory mechanisms have failed to refer to the required syntactic constraints, either at encoding or retrieval processes, or both.

There are two types of similarity-based interference that have different time profiles and conclusions about the nature of memory access. On one hand, inhibitory interference which occurs in 'multi-matching contexts', and the target perfectly matches the retrieval probe, but

there is also a distractor that overlaps some or all features with the target. This can lead to slower retrieval latencies, thus disrupting the licensing process by providing an extra element that is initially considered by the parser (Nairne, 1988, 1990), Anderson, 1974; Anderson & Reder, 1999; Lewis and Vasishth, 2005).

On the other hand, facilitatory licensing occurs in contexts of 'partial match', where neither the target nor the distractor have all the necessary features content and structural features and required by the retrieval probe. If the parser retrieves the distractor as a possible antecedent, speakers may initially consider that the sentence is grammatical, giving rise to *illusions of grammaticality* (Parker et al., 2011). While inhibitory interference does not necessarily entail that the illicit material has been retrieved or considered, facilitatory interference does, and therefore gives clear conclusions about the nature of the memory mechanisms, which makes it a predictor of content-addressable retrieval mechanisms if illusions of grammaticality arise.

In general terms, a dependency that displays immunity to linguistic illusions shows that the grammatical representations of syntactic constraints are preserved in memory. But a selective implementation of syntactic constraints, that is, if only some dependencies show this immunity, may reveal an error from the memory mechanisms to implement these constraints during retrieval.

Research has found a profile of selective fallibility in the effects of interference in the processing of long-distance dependencies. This means that, while some dependencies seem to be prone to linguistic illusions (Dillon et al., 2013; Drenhaus et al., 2005; Vasishth et al., 2008; Wagers et al., 2009; Xiang, et al., 2009), others show immunity to these effects (Clifton et al., 1999; Dillon, 2011; Dillon et al., 2013; Dillon et al., 2014; Kush, 2013; Nicol & Sweeney, 1989; Sturt, 2003; Xiang et al., 2009). According to these results, different memory mechanisms should be involved in the resolution of dependencies during sentence processing.

1.3.6.1 Evidence from long-distance dependencies

Evidence from illusions: subject-verb agreement

Research has shown that subject-verb agreement is sensitive to feature-matching distractors in illicit structural positions. For example, Bock and Miller (1991) provided participants with fragments such as those in (20) in a sentence production task. They found an increasing use of plural verbs (e.g. were) in sentences where the object of the PP in the subject was plural, compared to when it was singular.

- (20) a. The key to the cabinet...
- b. The key to the cabinets...

Similar effects were reported by Pearlmutter et al. (1999) for cases of comprehension. They used self-paced reading and eye-tracking with sentences like the ones in (21). They found that the ungrammatical conditions produced slower reading times after the verb compared to grammatical cases, but this disruption effect was mitigated with a feature-matching distractor.

- (21) a. The key to the cabinet was rusty from years of disuse.
- b. The key to the cabinets was rusty from years of disuse.
- c. *The key to the cabinets were rusty from years of disuse.
- d. *The key to the cabinet were rusty from years of disuse.

Furthermore, Wagers et al. (2009) found results consistent with those of Perlmutter and colleagues, but, additionally, that the agreement attraction effect only facilitated the processing of the ungrammatical conditions. They also found that during judgement studies participants were more likely to accept ungrammatical sentences with feature-matching distractors, in comparison to when the distractor did not match the verb features, producing illusions of grammaticality.

These studies show that, in contexts of facilitatory interference, subject-verb agreement is sensitive to the effects of linguistic illusions, which is compatible with the use of direct-access, content-addressable mechanisms, as we have described above. Because the distractors are

accessed during retrieval, we must assume that all items are being accessed during this process and that content features drive dependency formation.

Evidence from immunity to illusions: reflexive anaphors

As opposed to the sensitivity to illusion effects found on subject-verb agreement, several studies have provided consistent evidence for reflexive anaphors' immunity to the effects of attraction and linguistic illusions. Among them, an eye-tracking study by Sturt (2003) where he manipulated the gender features of the reflexive pronoun to match or mismatch the target or the distractor. He used both a proactive interference paradigm, where the distractor precedes the target (22), and a retroactive paradigm, where the distractor follows the target (23).

(22)

- A. Jonathan was pretty worried at the City Hospital.
 - a. He remembered that the surgeon had pricked himself with a used syringe needle.
 - b. He remembered that the surgeon had pricked herself with a used syringe needle.

- B. Jennifer was pretty worried at the City Hospital.
 - a. She remembered that the surgeon had pricked himself with a used syringe needle.
 - b. She remembered that the surgeon had pricked herself with a used syringe needle.

(23)

- A. Jonathan was pretty worried at the City Hospital.
 - c. The surgeon who treated Jonathan had pricked himself with a used syringe needle.
 - d. The surgeon who treated Jonathan had pricked herself with a used syringe needle.

- B. Jennifer was pretty worried at the City Hospital.
 - a. The surgeon who treated Jennifer had pricked himself with a used syringe needle.
 - b. The surgeon who had treated Jennifer had pricked herself with a used syringe needle.

Both experiments examined whether the matching features of the distractor (Jonathan/Jennifer) affected processing. He found that reading times on the reflexive pronoun were only affected by

gender mismatch of the target antecedent, but not by mismatch of the distractor. A finding consistent with a serial-search memory access paradigm.

Similar results were found by Xiang et al. (2009) in an ERP study where they found that the P600 response that was elicited by an antecedent mismatching the reflexive pronoun was not reduced by a feature matching distractor in sentences like (24), based on those used by Sturt (2003).

- (24) a. The tough soldier that Fred treated in the military hospital introduced himself to all the nurses.
- b. The tough soldier that Katie treated in the military hospital introduced herself to all the nurses.
- c. The tough soldier that Fred treated in the military hospital introduced herself to all the nurses.

Likewise, Dillon et al., (2013), using eye-tracking, tested similar sentences in where they manipulated number feature, instead of the gender (25). They found that reading times at the reflexive were affected by the mismatch of the target, but not by the distractor.

- (25)a. The new executive who oversaw the middle manager apparently doubted himself on most major decisions.
- b. The new executive who oversaw the middle managers apparently doubted himself on most major decisions.
- c. *The new executive who oversaw the middle manager apparently doubted themselves on most major decisions.
- d. *The new executive who oversaw the middle managers apparently doubted themselves on most decisions.

The findings provided by these studies show that in the case of reflexive anaphors, memory access is driven by a high-level serial search mechanisms since dependency formation seems to be always syntactically constraint, regardless of feature manipulations, or the presence of structurally irrelevant distractors.

Evidence from NPIs: a selective profile

While both subject-verb agreement and anaphors present a clear profile of fallibility from illusion effects, NPIs seem to have a different profile. While speakers have demonstrated sensitivity to the presence and location of NPI licensors in offline tasks, studies have shown that illusions can arise during online speeded-tasks, and that unlicensed NPIs can be accepted with the presence of a distractor acting as an intrusive licensor, an observation that is consistent with direct-access memory mechanisms. Nonetheless, recent studies by Parker (2014) and Parker and Phillips (2016) show that the distance between the NPI and its licensor can have an effect on whether illusions appear, or not.

1.3.6.2 Sensitivity to illusions

Drenhaus et al. (2005) tested the German NPI *jemals* (identical to *ever* in English) by using speeded-acceptability judgments and ERPs in sentences such as (26). While (26a) contains a licensed NPI, the sentence in (26b) only contains a structurally-irrelevant licensor whose scope does not include the NPI. Their results showed an increase in the rates of acceptability in ungrammatical sentences that contained an intrusive licensor, in comparison to those without one (26c). Furthermore, the ERP data demonstrated that while both ungrammatical sentences (26b, c) elicited an N400 effect at the NPI, the amplitude of the N400 was significantly reduced in cases like (26b) compared to those like (26c).

(26)a. Klein Mann, der einen Bart hatte, var jemals glücklich.

‘No man who had a beard was ever happy’.

b. *Ein Mann, der keinen Bart hatte, var jemals glücklich.

*‘A man who had no beard was ever happy’.

c. *Ein Man, der einen Bart hatte, var jemals glücklich.

*‘A man who had a beard was ever happy’.

Similar results were found by Vasishth et al. (2008) in sentences similar to those in (26) for German. Their results showed facilitations effects in reading times at the NPI when an irrelevant licensor was available in comparison to cases without a licensor. These results were confirmed

for English by Xiang et al. using speeded-acceptability ratings (2006) and ERPs (2009) for sentences such as those in (27), where they compared two types of licensors: *no* and *very few*. They found increased rates of acceptance in the speeded-acceptability task for both licensor types in sentences which contained an intrusive licensor, compared to those without a licensor. For the ERPs, they found a reduced P600 effect in the irrelevant-licensor conditions compared to the no-licensor conditions. These results confirm that the presence of a distracting licensor can facilitate the processing of unlicensed NPIs during real-time processing.

- (27)a. {No/Very few} bills that the Democratic senators have supported will ever become law.
- b. *The bills that {no/very few} Democratic senators have supported will ever become law.
- c. *The bills that the Democratic senators have supported will ever become law.

These results are consistent with content-addressable mechanisms for NPIs. In this respect, Vasishth et al. (2008) considers the illusory licensing of NPIs as evidence for a general mechanism of dependency formation that is sensible to partial-match between memory encoding and cue retrieval. According to them, NPIs must initiate a retrieval of [+negative] and [+c-command] features, and that errors occurs during partial-match when the mechanisms recover an element that contains only the [+negative] feature. However, there is currently consensus that NPI licensing is determined by the interaction of the semantic and pragmatic features of the NPI and the entire context in which it is subscribed, rather than having an item-to-item structure like the one that drives subject-verb agreement (Chierchia, 2006; Giannakidou, 1998; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Ladusaw, 1992).

1.3.6.3 *Turning off illusory licensing in NPIs (Parker, 2014)*

Though NPIs have consistently been shown to be prone to linguistic illusions, Parker (2014) argues that these illusions can be “turned off” in specific configurations. In a series of experiments he demonstrated that the condition for illusions to appear in NPI licensing has a fleeting time profile that is dependent on where the NPI appears respect to its licensor and that

this is a consequence of changes that occur over time in the encoding or interpretation of the licensing process.

Furthermore, he considers that the viability of the parser to implement linguistic constraints is a consequence of general cognitive mechanisms that distinguish the encoding and retrieval of emerging syntactic and semantic representations.

Parker compared the NPIs *ever* and *any* to observe whether both behave similarly and found that, in terms of illusions, they did not. Furthermore, that intrusive licensing did not increase the acceptability ratings for *any*, as it did for *ever*. Parker proposes that this contrasting profile is conditioned by the different distance between *any* and *ever* to their licensor.

He argued that parsing operations associated with thematic binding at the main verb may reactivate the subject NP, hence a reactivated NP just before *any* could give the advantage of increased activation to the target over the distractor, assisting in avoiding illusory effects. Furthermore, the distance between the NPIs could also give rise to “anti-locality” effects, where the increased distance between *any* and the main subject could narrow down the number of possible continuations. In order to test for locality effects, Parker compared between pre-verbal *ever* and post-verbal *ever* in sentences like (28). This showed a contrasting profile where pre-verbal *ever* showed consistent effects of illusory licensing while post-verbal *ever* did not.

(28)

A. Pre-verb *ever*

- a. No journalists [that the editors recommended for the assignment] *ever* thought that the readers would understand the complicated situation.
- b. The Journalists [that no editors recommended for the assignment] *ever* thought that the readers would understand the complicated situation.
- c. The journalists [that the editors recommended for the assignment] *ever* though that the readers would understand the complicated situation.

B. Post-verb ever

- a. No journalists [that the editors recommended for the assignment] though that the readers would ever understand the complicated situation.
- b. The journalists [that no editors recommended for the assignment] though that the readers would ever understand the complicated situation.

According to this account, the main subject would be reactivated, and the post-verbal NPI could have access to a reactivated licenser. Interestingly, this reactivation of the subject could produce a recalculation of the semantic scope of the licenser at the verb, extending its reach and reaffirming the NPI under the scope of the appropriate licenser, or at least decreasing the possibility of retrieving the distractor.

In cases where the NPI occurs after the verb, recalculation of scope would happen after the licensing process has taken place. As observed before, *ever* tends to be attracted to an irrelevant licenser, possibly looking to be in the scope of its licenser.

It is possible that scope may be a more difficult thing for the parser to encode and retrieve, which would make NPIs more susceptible to illusions. Moreover, why this illusions disappear when an NPI is located after scope has been recalculated by a verb, since it would be more readily available.

So far, the interference profile previously observed on NPI licensing appears to be similar to agreement attraction in SVA. Negative licensers, containing some form of [+negative] lexical feature, trigger similarity-based interference effects with other structurally-illicit negative distractors. However, NPIs have numerous types of licensers, and not all of them are lexically negative. For example, the licenser *only* licenses NPIs through a “negative exceptive component”, where “[Only NP] VP” entails “[Nobody but NP] VP”. This way, non-negative licensers have been argued to license NPIs through a pragmatics mode which relies on negative inferences (Giannakidou, 1998, 2006).

Following this idea, Xiang et al. (2013) predict that the individual pragmatic skills of participants can determine how interference affects NPI licensing, but not other types of syntactically driven interference such as in SVA. According to them, interference is sometimes caused by the over-application of this pragmatic route and this will occasionally boost NPI acceptability even in the absence of a semantically-close-enough inference. They tested this hypothesis by determining each participant's pragmatics skills using the communicative skills (SC) scale in the Autistic-spectrum quotient (AQ; Baron-Cohen et al., 2001).

Their results confirmed that individual pragmatic skills had an effect on interference effects on NPI licensing showing that participants with better pragmatic skills had a higher rate of acceptance of interference conditions than those who had worse pragmatics skills. Furthermore, they observed that different licensors presented different interference profiles, with a higher rate of interference for *only* cases in comparison to *no* cases. Moreover, they observed that these effects were only present for the NPI cases, but not for SVA cases, and pragmatic skills had an effect on offline acceptability ratings, but not on online RTs.

This research provides further evidence that pragmatic inferencing has a major impact on the licensing of NPIs. Trying to understand how and when these pragmatic inferences are computed into the licensing of NPIs is a crucial question on how NPI licensing works and begs the question of whether pragmatic licensing is integrated through memory or whether there are other cognitive skills at work, such as reasoning or attention focus.

1.3.7 Intervention

Though most of the research has been done using similarity-based interference effects where the potential licensors vary in terms of their structural position and compete with each other, intervention effects rely on the inclusion of a quantified expression that disrupts the grammatical relation between an NPI and its licit licensor (Linebarger, 1980; Chierchia, 2006). Thus,

intervention can be used to examine the formation of LDDs in a sentence without disrupting the normal structural relationship of the licensing process.

Consider the example in (29). While (29a) is a perfectly grammatical sentence, (29b) seems degraded in comparison, even though *no one* and *not everyone* are equally DE (Chierchia, 2006). Linebarger argues that in order for the NPI to be licensed, it must be within the *immediate scope* of its licensor, what she calls the “Immediate Scope Constraint”, and that the licensing of NPIs is disrupted by the intervention of the universal quantifier. The general configuration of intervention (Linebarger, 1980; Krifka, 1995) is shown in (30).

- (29) a. No one had any objection.
- b. ?? Not everyone had any objection.

(30) * ... DE [...INTERVENER... [...NPI...]]...

Intervention gives us an interesting alternative to similarity-based interference in terms of how the structure of the sentence is being manipulated. Figure 1, below, shows the comparison between the sentence structures for *interference (A)* and *intervention (B)*, and *grammatical licensing (C)* in NPI licensing. In interference (A), we observe that the negative element *no* is found with an embedded sentence in the subject, as compared to the grammatical licensing case (C), and therefore is not able to either c-command or give scope to the NPI. In contrast, in the intervention case (B), the licensor is in a structurally licit position, the same as in (C), which both c-commands the NPI and, potentially gives scope to it, but, at the same time, there is a quantifier *every* between the NPI and its licensor, which also scopes over the NPI.

Chierchia explains that NPIs want to be in a DE environment, but that intermediate quantifiers embedded in a DE environment generate positive implicatures that can disrupt the DE character of the environment in which the NPI is found (31).

- (31)a. I doubt that John had any complaint.
- b. ?? I doubt that many/most of the students had any complaint.

c. I doubt that the student had any complaint.

A question that arises from the effects of intervention is how quantifiers affect the relationship between the NPI and its licenser, and what does that tell us about how NPIs are licensed and the memory processes involved. As we have observed above, NPIs can be licensed by pragmatic implicatures generated by the sentence context.

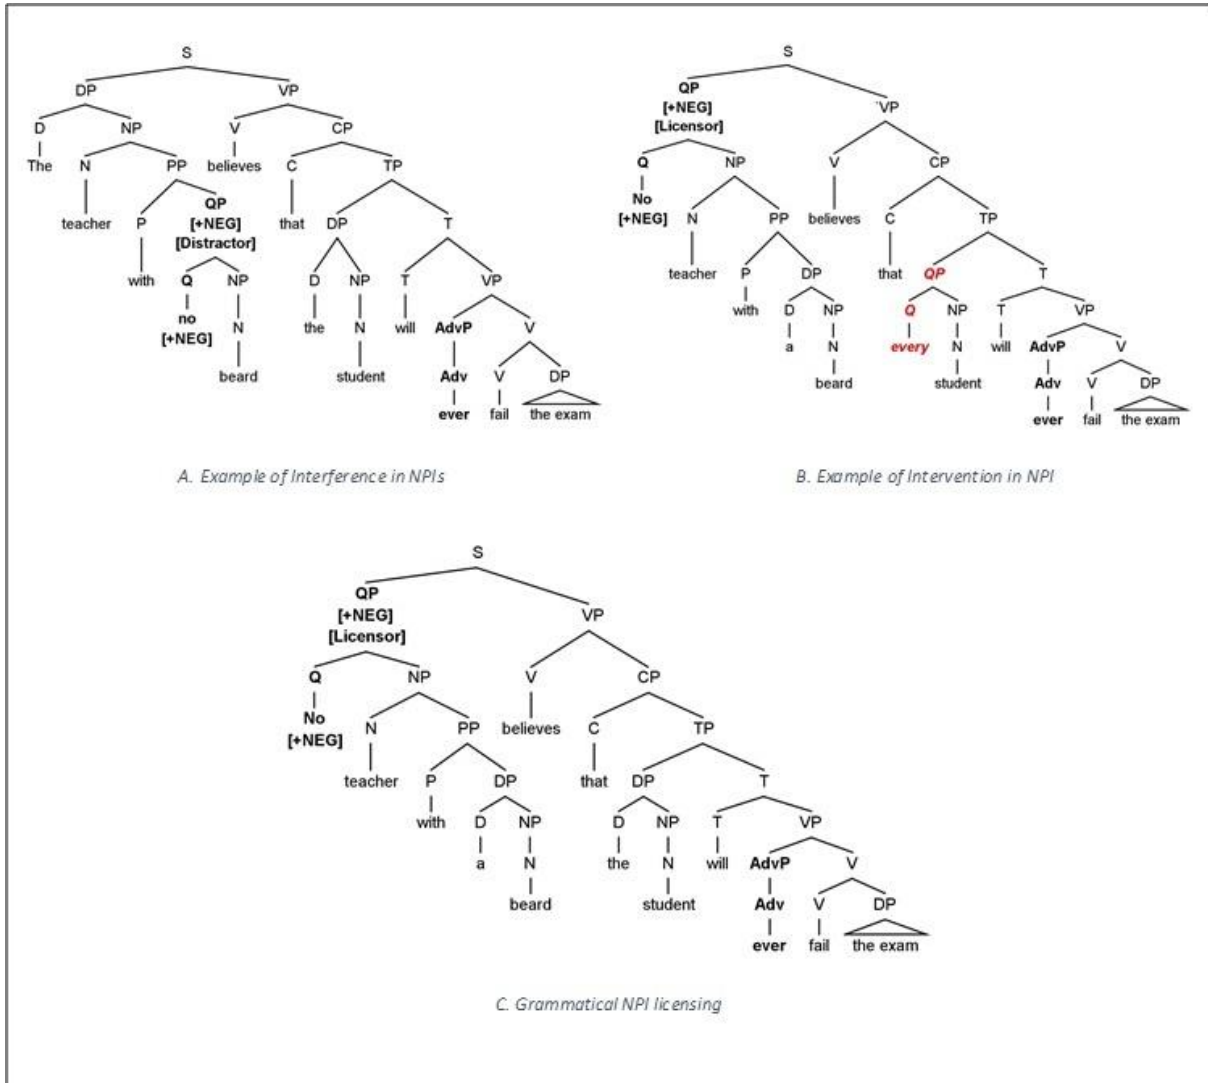


Figure 1: Interference (A) and Intervention (B) vs. Grammatical Licensing (C).

Interveners such as the universal quantifier (*every*) are a natural class of lexical items that trigger scalar implicatures (SIs) in DE environments (Chierchia, 2004). These SIs are optional inferences that arise in sentences that have two conditions, (i) the sentence can be argued to have a

minimally different alternative when replacing some lexical item with another in the same scale, for example, *some* and *every* (Horn, 1972, 1976; Gazdar, 1979; Katzir, 2007); and (ii) the sentence is consistent with such an alternative being false (Grice, 1975; Sauerland, 2004; van Rooij & Schulz, 2004; Chierchia et al., 2008; Franke, 2011, Bergen et al., 2014, inter alia). Following this conditions, the negation of the alternative can be added to the meaning of the sentence as a SI. This can be observed in the examples in (32).

- (32) Sentence: John didn't bring cookies for every colleague.
 Alternative: John didn't bring cookies for any colleague.
 Scalar Implicature: John brought cookies for some colleagues.

If the parser is able to access this information during NPI licensing, it is possible that having positive implicatures generated by an intervening quantifier could disrupt the parser's access to other negative implicatures. The "immediate scope constraint" suggested by Linebarger could provide an explanation to this, since the parser would have immediate access to the positive implicatures generated by the intervener, before being able to access the negative scope of the licenser.

Linebarger and Chierchia argue that the intervener might be affecting the DE environment of the sentence by either affecting the polarity, or by placing the NPI under the scope of the intervening quantifier, which begs the question of whether the NPI is also under the scope of the licenser at the same time. We observe this in terms of the inferences entailed by DE and UE environments in (33).

- (33) a. No student believes that the teacher owned a dog.
 b. No student believes that the teacher owned a poodle.

 c. No student believes that every teacher owned a dog.
 d. No student believes that every teacher owned a poodle.

As we discussed earlier, DE environments allow entailments where one can go from a set to a subset as seen in the relation between (33a) and (33b). Like so, the same relationship seems to be maintained between (33c) and (33d), where if no one in the group of students believes that every teacher owned a dog, then the same is true for any specific kind of dog. The environment remains DE even after the inclusion of the quantifier in (33c, d). If this is true, we could argue against just a change in polarity in the NPIs environment. Furthermore, if the scope of the intervening quantifier does not affect the downward entailing characteristic of the environment as argued by Linebarger and Chierchia, and the scope of the licenser should still reach an NPI in that position (34), then what drives the intervention effect?

(34) ?? No student believes that every teacher ever owned a dog.

If intervention does not seem to affect neither the structural constraints, nor the DE environment of licensing, another option could be that there is something in the meaning of *every* that is affecting the dependency. Furthermore, it could be possible that it is having a pragmatic effect, just by being there. It is then possible for the parser, if it is able to access pragmatic information during processing, to find itself with two contrasting sets of pragmatic inferences which could explain why processing this kind of sentences is so difficult.

A recent study by Denić, Chemla, and Tieu (2017) observed the effects of intervention on NPI licensing. They performed a series of experiments to assess whether intervention effects occur and whether they lead to categorical ungrammaticality, and whether there is evidence that intervention effects arise when SIs are derived. Using acceptability judgment tasks, they found that while intervention effects do arise in NPI licensing, having an intervener does not lead to categorical ungrammaticality as it was originally assumed. Across their study, they used acceptability judgments and picture selection tasks to test the link between SIs and intervention effects in sentences like (35).

(35) Monkey didn't give every lion juice.

Participants were asked to choose between one of two pictures, or both, that contained either some animals holding the item provided by the subject, or none of them holding any of it, in some tasks, or to decide whether a single picture was the appropriate description of the sentence. While in most cases they found no correlation between participants' ability to generate scalar implicatures and intervention effects observed in judgment tasks, in one experiment they found some evidence that when forced to assign meaning to NPI sentences with intervention, some participants used a "repair strategy" where they seemed to be able to block the derivation of a SI in order to complete the task.

Though this experiment was not able to yield consistent evidence of the relationship between SIs and intervention effects, this is the first quantitative study to explore the effects of intervention in NPI licensing and gives us a view that implicatures may somehow still be involved in the process, though how both phenomena are linked is still not clear and requires further study to understand both how interveners and pragmatic inferences affect NPI licensing.

Furthermore, if NPIs can be licensed through pragmatic information, studying the effects of intervention could not only shed light on what kind of information is required by NPIs to be licensed, but also on how this pragmatic information is retrieved from memory and used by the parser to establish the dependency between the NPI and its licenser.

1.4 The current research

As we have seen in the previous sections, NPI licensing does not seem to be as simple as a mechanical feature-checking process that other dependencies appear to use. While subject-verb agreement requires the parser to check for morphological features between the subject and the verb of the sentence, and anaphor binding further requires a syntactic constraint, NPIs seem to also involve certain semantic and pragmatic requirements that must be met. In this sense, we can say that NPI licensing is different than other long-distance dependencies.

Most research has attempted to observe the behaviour of NPIs in similarity-based interference contexts by either moving the licenser out of its structural position, or by adding a distractor in a structurally illicit position to disrupt the relationship between the NPI and its licenser. Furthermore, most studies have been conducted on a particular set of NPIs, such as *ever* and *any*, and a particular set of licensers, such as *no* and *not* which present a very strong form of negation, but not other NPIs or licensers which rely on other aspects of the licensing requirements such as DE and pragmatic licensing. In terms of the memory mechanisms involved, previous research has shown that NPIs are prone to linguistic illusions most of the time, but that, unlike other dependencies, this can be manipulated by the position of the NPI as shown by Parker (2014)¹. Moreover, researchers are not clear as to what kind of features are retrieved by the parser to fulfil the licensing process, nor how these are recovered from memory, or whether there are other cognitive abilities, such as reasoning or focal attention, involved during NPI licensing. Likewise, the difference between weak and strong NPIs show that there are differences in how and when they are licensed and could have consequences for the choice of memory mechanisms (serial-search vs. content-addressable) that are involved in the process, and we could wonder whether all NPIs always employ the same mechanism when retrieving their licenser.

The current research aims to observe how intervention effects disrupt the licensing context of NPIs in an attempt to fill in the gaps from previous research, and understand how memory retrieval is involved during licensing. With the exception of the recent study by Denić, et al. (2017), intervention effects have not been investigated in NPIs. This method can allow us to look into how the syntactic, semantic and pragmatic conditions of the licensing process interact during the licensing process, and observe whether intervention effects reflect the predictions of memory mechanisms. Furthermore, intervention can also shed light on whether NPI licensing is

¹ A follow up unpublished study by A. Alaberkyan studying NPI distance effects has found similar results to Parker (2014).

solely a matter of memory retrieval or whether there are other cognitive abilities involved in the process.

1.4.1 Theoretical questions

Question 1: Is NPI licensing a matter of memory retrieval? If so, is it serial-search or direct-access?

Because research into NPI licensing has revealed a selective profile of illusory licensing, the matter of understanding how memory is retrieved in the establishment of this dependency is still not clear. Furthermore, because NPIs seem to have semantic and pragmatic requirements, a method like similarity-based interference, which tests predictions based on syntactic constraints and content features, may not be enough to explore how these requirements are taken into consideration during the licensing process. Furthermore, because NPIs and their licensors are so varied, it is also possible that the memory mechanisms used for retrieval may be dependent on the NPI type. Therefore, in order to answer this question, it is necessary to test different types of NPIs using a method like intervention, which allows for the manipulation of the semantic context of the sentence, as well as the involvement of pragmatic inferences.

Question 2: If NPI licensing relies on memory access, what are the features required by NPIs during dependency formation to match their licensors?

For those long-distance dependencies that require memory retrieval, the parser uses a set of cues initiated by the dependent element to find a matching element in memory. Given that NPIs comprise such a varied set of elements, is it possible to assume that there is a single set of cues for NPI licensing? As we have observed from previous studies, both the licensor type (e.g., *no* vs. *only*, Xiang et al., 2013) and the NPI (e.g., *strong* vs. *weak*, Zwarts, 1998; Giannakidou, 2011, Chierchia, 2014; *any* vs. *ever*, Parker, 2014) can have an effect on the licensing process. This variation could indicate that it is not the case that all NPIs have the same feature requirements. While some studies have considered that NPIs could have a [+negative] feature or even a

structural [+c-command] feature (Vasishth et al., 2008), evidence from intervention effects, as well as illusory licensing with quantifiers other than a negative one, could point to other possible retrieval features, such as [+quantifier] since the universal quantifier is able to disrupt the relationship between the NPI and its licensor. If, for example, NPIs initiate a search for both [+quantifier] and [+negative], we could understand that the intervening quantifier scoping over the NPI could be producing a partial-match interference effect which could be powered by the immediate scope constraint proposed by (Linebarger, 1980), and that there could be an additional checking mechanism concerning the type of quantifier [+negative], which could suggest a two stage process. If a feature such as [+quantifier] is involved in the licensing process, can other quantifiers such as the universal quantifier or the existential quantifier give rise to illusions of grammaticality? Is illusory licensing available even if such a feature is not available?

Question 3: If not all the NPIs rely on memory access, what other mechanisms (semantic or pragmatic) are involved in NPI licensing, and how can these be manipulated?

As we have seen above, some NPIs seem to rely on pragmatic implicatures in order to be licensed (e.g. emotive predicates). A possibility is that NPIs can trigger the retrieval of an implied meaning. If this is the case, how are implicatures encoded and retrieved in memory? Xiang et al. (2013) show that pragmatic skills seem to have an effect in offline judgments but not a significant effect during online processing. Because of this, it has been argued that NPI licensing may not be solely a product of memory, but may involve other cognitive abilities such as pragmatic reasoning or focal attention. Another possibility, then, is that NPIs could be directly licensed by implied meanings, which could be kept in an active focus of attention. Furthermore, while it seems that NPIs can look for features, they also need to be within very specific semantic environments, i.e. being inside a DE environment. How does the parser recognise this environment? Is the semantic context readily available for the parser when the NPI is reached or

is semantic information of this kind accessed at a much later stage? Can these contexts be manipulated to modulate the appearance of illusions of grammaticality?

In order to answer these questions, we'll first examine the presence of intervention effects on NPI licensing, and whether they produce categorically ungrammatical sentences using judgment acceptability tasks. Furthermore, this tasks will be presented either in an offline setting where participants can review the experimental sentences while giving their judgements, or following an online task, where participants are no longer allowed to review the sentence and are forced to resort to memory to judge the sentences. This will allow us to compare whether intervention effects are pervasive in time. For the online tasks we will look at the self-paced reading (SPR) paradigm, which has been used to observe different aspects of long-distance dependencies licensing and should provide converging evidence about the NPI licensing process. SPR allows us to track processing difficulty over time with individual word resolution while also forcing participants to resort to memory retrieval at the dependent element.

1.4.2 Pilot experiments

1.4.2.1 *Design and Procedure*

In line with our research questions, we constructed 36 items with 3 variables in a 2 x 2 x 2 design. The variables varied NPI presence (NPI vs. No NPI), the presence of the universal quantifier *every* (Quantifier vs. Determiner), and whether the sentence had a configuration of intervention where the position occupied by the quantifier (or its alternating determiner) c-commanded the NPI (Intervener c-command "ICC" vs. No c-command "NoC"). The NoC conditions were achieved by placing the quantifier/determiner phrase within a relative clause. Examples of the experimental items are presented on Table 1.

Using these items we ran two studies. The first one consisted of an offline judgment task where participants were able to see the whole sentence and were asked to rate it on a scale of 1 (bad) to 7 (good) in order to assess whether participants were sensitive to intervention effects on the

intervention conditions (NPI x Quantifier x ICC). The second study, as a contrast, used an online self-paced reading task where participants were presented only a part of the sentence at a time (chunk-by-chunk) in a centred position in the screen that was replaced with the next part of the sentence at the press of a button (static-window display paradigm).

We predicted lower rating scores for the c-commanding (intervening) quantifier cases in sentences with NPis (intervention cases) in comparison to other cases, since the scope of the quantifier is expected to disrupt the licensing process.

Table 1: Sample items for Pilot Study 1

NPI	Quantifier	intervener c-command	Sentence
NPI	Quantifier	C-command	<i>No teacher</i> believes that every student will try to <u>ever</u> complain about the course.
NPI	Quantifier	No c-command	<i>No teacher</i> that every student believes will try to <u>ever</u> complain about the course.
NPI	Determiner	C-command	<i>No teacher</i> believes that the student will try to <u>ever</u> complain about the course.
NPI	Determiner	No c-command	<i>No teacher</i> that the student believes will try to <u>ever</u> complain about the course.
No NPI	Quantifier	C-command	<i>No teacher</i> believes that every student will try to complain about the course.
No NPI	Quantifier	No c-command	<i>No teacher</i> that every student believes will try to complain about the course.
No NPI	Determiner	C-command	<i>No teacher</i> believes that the student will try to complain about the course.
No NPI	Determiner	No c-command	<i>No teacher</i> that the student believes will try to complain about the course.

1.4.2.2 Results

While in experiment 1 we did not find a three-way interaction between NPI x Quantifier x C-command ($t = -1.811$), we did find a two-way interaction of Quantifier x C-command ($t = 3.995$). This results show that while there were no visible intervention effects in this experiment, there seems to be a trend in that direction. Furthermore, the two-way interaction seems to indicate that the position of the quantifier *every* (whether it's inside a relative clause or not) had an effect

on the overall rating of the sentence and suggest an interference effect with the NPI. Figure 2 shows the subtraction of the NPI cases minus the No NPI cases (NPI – No NPI) and seems to indicate that there could be a main effect.

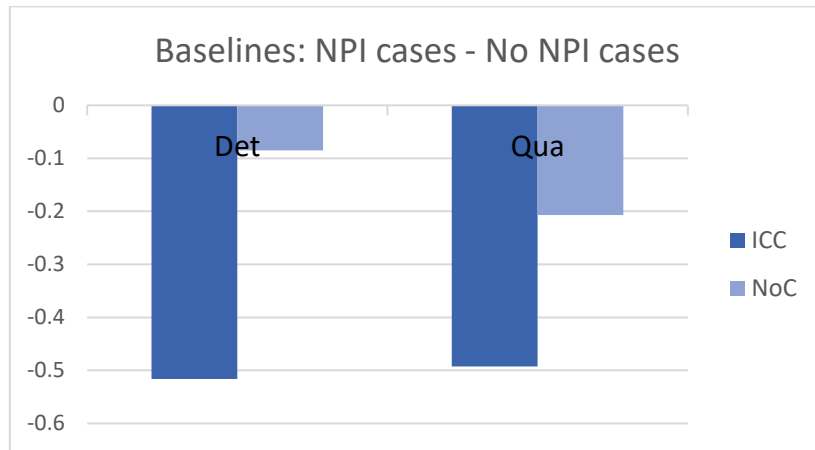


Figure 2: Pilot Study 1 Baselines (NPI - No NPI)

In pilot study 2, on the other hand, we found a three-way interaction of NPI x Quantifier x C-command ($t = -2.066$). Figure 3 shows the subtraction of the NPI cases minus the No NPI cases (NPI – No NPI). While the determiner cases show barely any variation, the quantifier cases show a significant difference between the intervener c-command (ICC) cases and the No-c-command cases (NoC), which show that, as we predicted, interveners disrupt the normal licensing of NPIs, therefore, in the intervention cases (NPI x Quantifier x ICC) the difference between having an NPI or not makes a big difference.

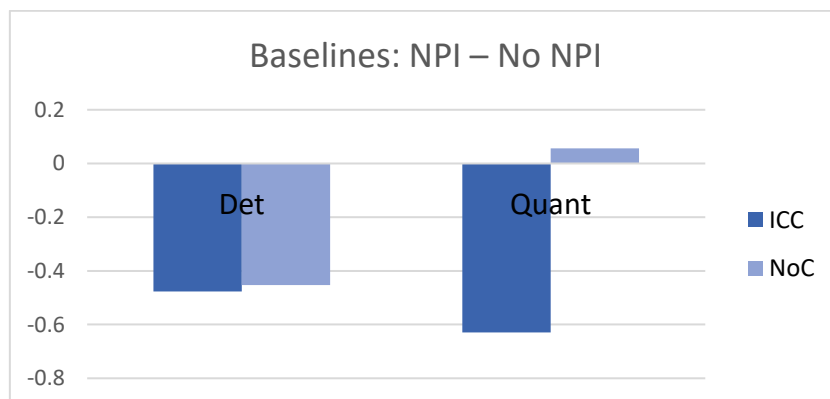


Figure 3: Pilot study 2 Baselines (NPI - No NPI)

Figures 4 and 5 below show the average rating for each condition for both pilot study 1 (Fig. 4), and pilot study 2 (Fig. 5). An interesting observation that we can make, is that in both studies, the non-intervention configuration conditions (NoC; represented by the lighter colour) seem to be degraded in comparison with the ICC conditions regardless of whether there is an NPI or not. This is unexpected as we predicted that the cases in which the quantifier c-commands should be rated worse than the rest. This could be due to a main effect from the relative clause in the No-command cases (NoC) which is consistent through all NoC conditions and through both experiments.




Figure 4: Experiment 1: Average ratings by condition




Figure 5: Experiment 2: Average ratings by condition

1.4.2.3 Discussion

The difference between both studies show that intervention effects are only evident in online measures, as opposed to offline ones, which is consistent with previous research that showed that the sensitivity to structural properties of potential licensors varies with time (Vasishth et al., 2008; Xiang et al., 2009; Wagers et al., 2009) in similarity-based interference. This seems to be the case as well with intervention effects.

Furthermore, follow-up experiments will require us to simplify the sentences used in these experiments to avoid the increase in complexity from a relative clause that has been observed in both experiments. Likewise, since the universal quantifier still seems to have some effect, even when it is not c-commanding the NPI (and therefore, should not scope over it). It seems that the presence of the universal quantifier *every* is affecting the licensing of the NPI even outside of an intervening structural location and is, actually, interfering with the NPI. If this is true, this results show that upon reaching the NPI, the parser initiates a search for a licensor and the universal quantifier is a prime candidate. If this is so, this opens an interesting line on inquiry about the nature of quantifiers and the licensing of NPIs that is worth exploring in more detail. Does this interference effect mean that universal quantifiers can give rise to illusions of grammaticality? Is so, does this mean that there could be a [+ quantifier] feature present and desirable by the NPI?

In terms of memory mechanisms, if the universal quantifier is able to interfere with the NPI despite its structural irrelevant position, this result is consistent with content-addressable memory retrieval, since information outside the structural requirements seems to be accessed, even though it should not be readily available.

1.4.2.4 Conclusion

The results of both pilot studies suggest that NPI licensing presents effects from both intervention effects, though only in online measures, and interference effects in both offline and online measures. Furthermore, in terms of memory retrieval, the interference effects found also suggest that a content addressable mechanism is involved in NPI licensing. The fact that these

interference effects are caused by the universal quantifier suggests that the *every*, and perhaps other quantifiers, may contain features desirable for the licensing process, provided it is features which are being retrieved. Alternatively, it could be some property inherent to the universal quantifier which rather than a more general property of quantifiers. In the following chapters we will look into how quantifiers are involved in the licensing of NPIs and whether they can elicit illusory licensing.

Chapter 2: Illusory Licensing and the Universal Quantifier – Identifying the quantifier’s restriction clause

2.1 Introduction

In order to construct the relationship between words and phrases, language comprehension relies on memory mechanisms. Linguistic relationships known as long-distance dependencies which span large amounts of intervening material seem to rely on such mechanisms to form associations between dependent elements (McElree, Foraker, & Dyer, 2003; Phillips, Wagers, & Lau, 2011).

Negative polarity items (NPIs) have been of particular interest to psycholinguistics because they rely on semantic and pragmatic processes which are not at play during the processing of other dependencies. Currently, research has shown that different dependencies have different profiles in terms of how information is retrieved from memory during sentence processing. For example, multiple studies have shown that dependencies such as subject-verb agreement (Bock & Miller, 1991; Vigliocco, & Nicol, 1998; Solomon & Pearlmuter, 2004; Phillips, Wagers, & Lau, 2011; Parker, 2014) and anaphora resolution (Sturt, 2003; Kazanina, et al., 2006; Phillips, 2006; Kush, 2013) have contrasting profiles in terms of how information is retrieved during their processing. While Subject-verb agreement, for example, seems to rely on a content-addressable memory mechanism and is predicted to be prone to linguistic illusions, a phenomenon where an ungrammatical sentence may be deemed grammatical by comprehenders or vice-versa, reflexive anaphora resolution seems to rely on a serial-search mechanism and is predicted to be immune to such illusions. Negative Polarity items seem to be less straight forward since it is not yet clear

how the semantic and pragmatic processes are integrated into the licensing process or whether their processing is driven by the same mechanisms as the other dependencies.

2.2 Background

2.2.1 Illusory licensing

Linguistic illusions are the cases when a subject temporarily accepts an ungrammatical sentence as correct or acceptable during time sensible measures but judge it as unacceptable given more time to reflect upon it. Illusion effects have been argued to show that subjects temporarily consider dependencies that are not structurally licensed (Perlmutter, et al., 1999; Vasishth et al., 2008; Wagers et al., 2009; Xiang et al., 2009). While most subjects are good at detecting erroneous sentences, in certain configurations, intrusive elements, or linguistic elements in structurally inaccessible locations can interfere with the processing of long-distance dependencies and give rise to illusory effects.

Previous studies have shown that NPIs are susceptible to illusions of grammaticality. In the examples in (36), taken from Parker & Phillips (2016), (36a) is a well-formed NPI dependency where negation scopes over the NPI forming a grammatical sentence. Meanwhile, the examples in (36b,c) are ungrammatical sentences where there licit licenser (36b), or where the negative element is not structurally higher than the NPI and does not scope over it (36c). Regardless of both sentences being ungrammatical, a sentence like (36c), which contains a negative element in a structurally irrelevant position, has been observed to be judged as acceptable in time constrained conditions and give rise to an illusion of grammaticality. These effects have been consistently reported across different measures such as eye-tracker (Vasishth et al., 2008), self-paced reading (Xiang et al., 2006, 2013), speeded acceptability judgments (Drenhaus et al., 2005, Parker & Phillips, 2016), and event-related potentials (Drenhaus et al., 2005, Xiang et al., 2009; Yanilmaz & Drury, 2014).

(36)a. **No** diplomats have *ever* supported a drone strike.

- b. *The diplomats have *ever* supported a drone strike.
- c. *The diplomats that no congressman could trust have *ever* supported a drone strike.

There are currently two competing accounts which explain the rise of illusions. The first one considers that illusions in NPI licensing reflect a partial match memory retrieval error (Vasishth et al., 2008). This account considers that NPI licensing is an item-to-item structural dependency that is formed by the retrieval of a negative licenser from memory by using a semantic feature [+negative] and a structural feature [+c-command]. A partially matched linguistic item, in the absence of a fully matched licenser, (36c), can be erroneously retrieved and create the illusion that the NPI is licensed and that the sentence is grammatical.

The second account considers that the illusions reflect an over-application of semantic and pragmatic processes involved in NPI licensing rather than the incorrect retrieval of partially matched information (Xiang et al., 2009). This account considers that in addition to direct licensing by negation, NPI licensing also occurs through semantic/pragmatic properties of entire propositions, as in the case of emotive predicates. Structural constraints such as c-command are assumed by current theories to be a by-product of these semantic/pragmatic mechanisms in NPI licensing.


These two accounts offer distinct views on how NPIs are licensed as well as how illusions arise during processing. An important question to ask based on these accounts is whether NPI licensing is mainly driven by a direct memory retrieval route, or whether pragmatic/semantic processes take precedence over a memory retrieval route. While most long-distance dependencies seem to rely on memory mechanisms in order to establish the dependency, NPI licensing is different and memory retrieval does not seem to be sufficient to explain all instances of licensing which may imply that either memory retrieval is an incomplete account of the licensing process, or that it may not be the driving mechanism behind NPI licensing. Other processes may occur simultaneously, or be applied under different circumstances depending on

the licensor. If this is the case, certain properties of licensors may continue to be at work during the licensing process and could be immediately checked upon reaching the NPI.

2.2.2 Quantifiers in NPI licensing

Quantifiers have been a great part of NPI licensing research. The most researched quantifier in the literature has been the negative quantifier *no*, along with other instances of explicit negation such as *not*, or implicit licensing with the focus operator *only*. Both explicit and implicit negation can offer evidence for a direct licensing route where the negative operator is retrieved through a memory mechanism, but a quantifier like *every*, which is not negative and does not have a negative meaning, can license NPIs within its restrictor, and allow us to look at some of the semantic/pragmatic processes involved in NPI licensing.

Semantically speaking, quantifiers occupy the operator position of a tripartite structure (Kamp, 1981; Heim, 1982, Von Stechow, 1994), and bind the other two parts of this structure which are the restrictor clause, and the nuclear scope of the quantifier as observed in (37) for the sentence *Every child smiled*.



The restrictor clause of this structure identifies the set to which the quantifier applies, thus restricting the entities that it describes. The restrictor will be delimited at the very least by the common noun in the quantified phrase but it may also be defined by preceding context, or in some cases by later context. The nuclear scope of the structure refers to the rest of the sentence and corresponds to what is being attributed to the quantified set. In this sense, both the restrictor clause and the nuclear scope can be thought of as identifying sets with the quantifier defining the relationship between those sets.

In NPI licensing, different quantifiers behave differently in terms of whether they can license NPIs that are found within their restrictor clauses or their scopes. For example, the negative quantifier *no* in (38) is able to license NPIs that are located in either its restrictor clause (38b) or its nuclear scope (38c), or even in both at the same time (38d) since both the restrictor clause and the nuclear scope of the negative quantifier are downward entailing.

- (38)a. **No** technician who has read the manual has been able to operate the machine.
- b. **No** technician who has *ever* read the manual has been able to operate the machine.
- c. **No** technician who has read the manual has *ever* been able to operate the machine.
- d. **No** technician who has *ever* read the manual has *ever* been able to operate the machine.

This is not the same case with the universal quantifier. As we can see in (39a), the sentence is an affirmative one and does not contain any negative context. Notwithstanding, it is able to license NPIs within its restrictor clause (39b) since the context in this case is downward entailing. Conversely, its nuclear scope is not downward entailing, and thus not able to license NPIs, as we can see in example (39c) which is ungrammatical.

- (39)a. **Every** technician who has read the manual has been able to operate the machine.
- b. **Every** technician who has *ever* read the manual has been able to operate the machine.
- c. ***Every** technician who has read the manual has *ever* been able to operate the machine.

In contrast with both the examples above, the existential quantifier *some* is not able to license NPIs in neither its restrictor clause (40b) or its nuclear scope (40c).

- (40)a. **Some** technician who has read the manual has been able to operate the machine.
- b. ***Some** technician who has *ever* read the manual has been able to operate the machine.
- c. ***Some** technician who has read the manual has *ever* been able to operate the machine.

Because quantifiers differ in whether they can license NPIs in their restrictor clause or their nuclear scope, looking at different quantifiers can shed light in how both accounts described before are applied in the licensing of NPIs. Since the negative quantifier can license NPIs in either its restrictor clause or its nuclear scope, the best possible manipulation comes from the universal quantifier since it can license an NPI in only one of the components. According to Parker and Phillips (2016) illusions are a general property of NPI licensing and the product of a faulty licensing mechanism. If this is true, we can expect illusion effects to appear when the NPI is located within the scope of the universal quantifier.

Furthermore, we should be able to manipulate the illusions depending on what is driving the licensing process. If licensing is driven mainly by semantic/pragmatic properties of the sentence, NPI illusions might be a consequence of the misidentification of the boundaries between the restrictor clause and the nuclear scope of the quantifier. If this is the case, manipulation of the restrictor clause in order to make its boundaries more clear should provide the parser with enough context to determine that the NPI is no longer within the restrictor clause but within the nuclear scope of the quantifier, thus reducing the illusory effect. If licensing is primarily driven by memory mechanisms, we should be able to observe illusory effects since the NPI would be able to access the licenser as part of noisy cue retrieval when the NPI is located outside of the restrictor clause, but the addition of an adjective or a relative clause to the restrictor clause should not greatly affect the appearance of illusions.

Experiment 1 of this study aims to find whether this illusions appear with *every*. Furthermore, in order to test whether this property is exclusive to the universal quantifier, or whether it stems from a general property of quantifiers, experiment 2 will look at whether illusions arise with a different quantifier, in this case the existential quantifier. If illusions were to appear despite *some* not being able to license NPIs, then it is possible that quantifiers in general contain a desirable feature that at least partially matches the retrieval cue initiated by an NPI.

Experiments 3 and 4 aim to observe whether manipulations to the restrictor clause can be used to modulate the illusory effect of the universal quantifier in NPI licensing by using prenominal and postnominal modification of the common noun in the restriction clause. We assume that a better defined restrictor clause will make the boundaries of the restrictor clause much more clear, and therefore evident that the NPI is located within the nuclear scope instead. If semantic and pragmatic processes drive the licensing mechanism, rather than memory, we should observe a reduction or disappearance of the illusory licensing effect.

2.3 Experimental Findings

2.3.1 Experiments 1 and 2: Illusion effects as a result of Quantifiers in NPI licensing

If NPI licensing is driven by memory mechanisms, one important question is what kind of features the parser uses upon reaching an NPI in order to retrieve its licenser. One of the main difficulties with defining this is that NPIs, as well as their licensers, can be very varied and so finding a definite set of ideal features that can be shared by all licensers seems like a monumental task.

In these first two experiments we focus on whether quantifiers can give rise to illusions of grammaticality in NPI licensing. On the first one, whether they appear with the universal quantifier *every*, and on the second experiment, whether they rise with the existential quantifier *some*. If both quantifiers are able to produce illusions, we can assume that there is a general property of quantifiers involved in NPI licensing. However, if only the quantifier *every* can drive illusory licensing, then something more particular about its properties must be considered.

2.3.2 Experiment 1: Do illusions arise with the Universal Quantifier?

The purpose of experiment 1 was to test whether the universal quantifier is able to generate illusions of grammaticality in NPI licensing. For this purpose, we constructed sentences where the NPI surfaces inside the scope of the universal quantifier, which is upward entailing and does not license NPIs, rather than inside the restrictor which is downward entailing and can license

NPIs. More specifically, we manipulated the presence of NPIs in the sentence, as well as the subject type to compare the results between a grammatical licensor (*no*), an unambiguously ungrammatical one (*the*), and the universal quantifier (*every*). Given that NPI licensing has been observed to be consistently prone to illusions in previous studies (Drenhaus et al., 2005; Vasishth et al., 2008; Xiang et al., 2006, 2009, 2013; Yanilmaz & Drury, 2014; Parker & Phillips, 2016), we predict that the presence of the universal quantifier is enough for illusory licensing when the NPI is not within the restrictor clause.

2.3.2.1 Participants

Thirty six native speakers of English were recruited through Amazon's Mechanical Turk (<https://www.mturk.com/>). All participants were informed about the experiment's main task (rating sentences) and provided consent. Through a small survey, participants were asked about their linguistic background to ensure that they were native speakers. They were compensated \$3.00. The experiment lasted approximately 30 minutes.

2.3.2.2 Materials

In order to test whether illusory licensing occurs with the universal quantifier *every* when an NPI surfaces outside of the restrictor of the quantifier, we designed a set of 48 sentences with 6 conditions organized in a 2 x 3 design where the variables were NPI Presence (No NPI vs. NPI) and Subject Type (Licensed No, Unlicensed The, Unlicensed Quantifier). Each item was formed of a subject noun phrase and a main clause predicate. For the NPI conditions, the NPI was found inside the predicate. The subject was formed with an Article or a Quantifier, which was manipulated according to Subject Type, followed by a main noun. This way, the factor Licensor Type comprised one grammatical condition which uses the licit licensor *no*, and two ungrammatical conditions. The first of the ungrammatical conditions uses the determiner *the* as an unambiguously ungrammatical licensor, while the second one is the universal quantifier *every* as an ambiguous one. This 3-way division allows us to not only test whether the universal quantifier can give rise to grammatical illusions but also to compare these results with both a

grammatical and an ungrammatical baseline. The subject noun was immediately followed by the auxiliary verb *has*, which unambiguously initiated the main clause predicate, after which the NPI *ever* was introduced for the NPI conditions. The rest of the predicate was formed by *been* plus a main verb followed by four words which formed a complement for the verb (either as a direct or indirect object or a different form of complement). NPI presence was used as a variable to test the effect of having an NPI on the sentence and to make sure that any differences found are caused by the interaction of an NPI, and not by intrinsic differences of each of the potential licensors, as well as reducing participants from developing expectations of finding NPIs whenever they came across our experimental conditions. All items consisted of 9 (No NPI conditions) or 10 words (NPI conditions). A sample of all conditions can be observed in Table 2.

Each participant was asked to judge a total of 144 sentences, 48 of which were experimental items. The other 96 items were used as fillers (48 were part of a second experiment involving picture noun phrases, 48 were constructed as pure fillers). These fillers were between 7 and 10 words in length. The items were divided into 6 lists using a Latin Square design.

Table 2: Sample items Experiment 1: Speeded Acceptability Judgments

NPI	Licensor	Sentence
NPI	Negation	No/researcher/has/ever/been/invited/to/the/annual/conference.
NPI	Universal	Every/researcher/has/ever/been/invited/to/the/annual/ conference.
NPI	Determiner	The/researcher/has/ever/been/invited/to/the/annual/conference.
No NPI	Negation	No/researcher/has/been/invited/to/the/annual/conference.
No NPI	Universal	Every/researcher/has/been/invited/to/the/annual/conference.
No NPI	Determiner	The/researcher/has/been/invited/to/the/annual/conference.

2.3.2.3 Procedure

The experiment was designed as a two-alternative forced choice (2AFC) speeded acceptability task, using an experimenter-paced, word-by-word, serial visual presentation paradigm. Words were presented in a 400ms interval, where each word appeared on the center of the screen for 300ms followed by a 100ms blank between each word. Immediately following each sentence, subjects were presented with the question “Was this a good sentence?” Participants were then

given 3 seconds to respond to the question with either “y” (‘yes’) or “n” (‘no’) by pressing the corresponding key on the keyboard. This time limit restricted the amount of time participants had to reflect on their initial acceptability intuitions (Drenhaus et al., 2005; Parker & Phillips, 2016; Wagers et al., 2009; Xiang et al., 2006). The experiment was programmed using Ibbex Farm (Alex Drummond, <http://spellout.net/ibexfarm/>). Before the presentation of the experimental items, participants were provided with a set of practice items to train them in the task ahead. Six items were used, out of which two were well formed sentences, two were “intermediate” anomalous sentences, and two were outright badly formed sentences.

2.3.2.4 Data Analysis

Judgement response data was analysed by using the `glmer` (generalized linear mixed-effects regression) function in the `lme4` package (Bates, Maechler, & Bolker, 2011) in R software (R Core Team, 2014) to perform logistic regression via mixed models with fixed factors for experimental manipulations and a specified random effect structure, including random intercepts and slopes for all fixed effects by both participants and items (Baayen, Davidson, & Bates, 2008; Barr et al., 2013). Critically, we look at whether sentences with the intervening universal quantifier ‘every’ are rated better than sentences with and intervening determiner ‘the’, and whether this difference is larger for sentences which contain an NPI versus those without one. NPI presence (NPI vs No NPI) was sum-coded (present = 1, absent = -1), and the three-level factor Licensor was Helmert coded (Vasishth & Broe, 2011; Venables & Ripley, 1999) into two orthogonal contrasts: grammaticality (‘no’ = 1, ‘every’ = -0.5, ‘the’ = -0.5) and illusion (‘no’ = 0, ‘every’ = 1, ‘the’ = -1). This way, relevant licensors contrasts were determined by the grammaticality of the licensors within the NPI cases (Negative vs Universal + Determiner cases), and illusory licensing within the NPI cases (Universal vs Determiner). The model reports a total of five fixed effects: NPI presence, Grammaticality, Illusion, and the interaction of NPI presence with both Grammaticality and Illusion. Additionally, we included planned comparisons analysis of with five fixed effects

which were Helmert coded to include the difference between all NPI conditions vs all No NPI conditions, and the comparisons for grammaticality and illusory licensing within both NPI and No NPI cases. Conditions were analysed using the same method as the model above. Following Baayen et al. (2008), in the case of convergence failure, the random effect structure was simplified. Effects were considered significant if the p values were $p \leq 0.05$ (Gelman & Hill, 2007). Extreme RTs were removed using model criticism (Baayen et al. 2008). A total of 1.90% data was lost during clean up.

Finally, we also looked at the difference of the means of the quantifier and the determiner in the NPI cases by subtracting the percentage of responses for the NPI cases of the determiner from the NPI cases of the universal quantifier, as expressed but the following formula (NPI+Universal – NPI+Definite).

2.3.2.5 Results

The results for experiment 1 can be observed in Figure 6, and tables 3 and 4.




Figure 6: Acceptability rating for experiment 1. Error bars show standard error.

Table 3: Experiment 1: Logistic Regression Mixed Effects Modelling

	Estimate	Std. Error	z value	p value	
(Intercept)	1.022	0.162	6.322	< 0.001	***
NPI Presence	-1.274	0.083	-15.295	< 0.001	***
Grammaticality (Licensed v. Unlic.)	-1.332	0.121	-10.960	< 0.001	***
Illusion (Universal v. Definite)	-0.047	0.096	-0.495	0.620	
NPI Presence x Grammaticality	-1.304	0.121	-10.745	< 0.001	***
NPI Presence x Illusion	-0.194	0.096	-2.028	0.042	*

First observation of the data through the analysis, shows that there is a main effect of NPI presence ($z = -8.630$; $p < 0.001$), which shows that having an NPI greatly affects how the experiment sentences were judged. This effect is driven by the universal and determiner conditions which are ungrammatical, compared to the negative condition which is the grammatical one. We also observed a main effect of Grammaticality ($z = 11.400$; $p < 0.001$), which shows that participants were able to distinguish grammatical and ungrammatical sentences consistently for the universal quantifier and the determiner cases. The analysis also revealed an interaction between the presence of an NPI and grammaticality ($z = 11.600$; $p < 0.001$), which is expected since only the NPI cases contained sentences that were ungrammatical (NPI + *the* or *every* as licensors). Furthermore, we found an interaction between NPI and Illusion ($z = 2.360$; $p = 0.019$), which shows that the difference between the universal quantifier and the determiner was bigger for sentences with and NPI in comparison to the sentences that did not have and NPI. In general, this shows that sentences that contained a universal quantifier and the NPI were rated significantly better than those that had the determiner with an NPI. This result could be evidence for illusory NPI licensing with the universal quantifier.

Table 4: Experiment 1: Planned comparisons

	Estimate	Std. Error	z value	p value	
(Intercept)	0.941	0.148	6.370	< 0.001	***
No NPI v. NPI	-1.250	0.145	-8.630	< 0.001	***
No NPI, Grammaticality (Licensed v. Unlicensed)	-0.027	0.171	-0.158	0.875	
No NPI, Illusion (Universal v. Definite)	-0.162	0.149	-1.090	0.277	
NPI, Grammaticality (Licensed v. Unlicensed)	2.900	0.186	15.600	< 0.001	***
NPI, Illusion (Universal v. Definite)	0.305	0.131	2.330	0.020	*

Planned comparisons (table 4) show an effect of NPI presence ($z = -8.630$, $p < 0.001$) and a significant effect of grammaticality in sentences which contained an NPI ($z = 15.600$; $p < 0.001$), and more importantly, that Illusion effects are driven by the presence of the universal quantifier ($z = 2.330$; $p = 0.020$). This difference confirms a contrast between both unlicensed conditions. No significant difference was found for grammaticality or illusion in the No NPI cases. The results obtained are consistent with the hypothesis that participants were more prone to believe that NPI sentences with a Universal Quantifier were grammatical, and shows consistent illusions of grammaticality for these cases.

Additionally, we looked at the differences between the means between the acceptability ratings of universal quantifier sentences in comparison to the mean acceptability of determiner cases. We observed that in the case of the No NPI cases, sentences with the universal quantifier were rated slightly worse than both negative and determiner conditions. This could be due to an increased difficulty of the universal quantifier in comparison to both other conditions. If this is the case, NPI cases could also be affected by this increase in difficulty and differences between the universal and the determiner conditions could differ. In the No NPI cases we observed that sentences with the universal quantifier were rated 2.91% worse than determiner cases, while in the NPI cases the universal quantifier was rated 7.67% better than the determiner ones. If this difference is taken as a baseline for particular difficulties of the universal quantifier 'every', we can observe a difference in the reduction of acceptability responses of 10.58% between the universal quantifier and the determiner cases as a cost of adding an NPI to those sentences. In this case, the drop in good judgments is bigger for the determiner cases compared to the universal quantifier. Regardless of this difference, the analysis demonstrated a significant difference between both conditions in the NPI cases in the paired comparisons and this difference was also evidenced by the NPI x Illusion effect observed in the logistic regression analysis.

2.3.2.6 Discussion

The results of Experiment 1 seem to show evidence that illusory NPI licensing could be occurring with the universal quantifier *every* when an NPI is located within the scope of the quantifier, rather than within the restriction. This result is consistent with a content addressable memory retrieval account in which a structurally inaccessible licenser can be incorrectly retrieved based on a noisy cue-based retrieval process (McElree, 2000, McElree et al., 2003). Memory theories establish that upon reaching the NPI, the parser initiates a search for a licenser based on a set of retrieval cues (Vasishth et al., 2008), and that the universal quantifier matches at least some of these cues in order for it to be retrieved. If this is the case, what could those cues be? A [+Negative] feature like the one proposed by Vasishth does not seem to be enough given all the cases in which NPIs can be licensed. Similarly, a [+Downward Entailing] cue seems problematic for the same reasons.

One possibility arising from Experiment 1 is that quantifiers in general bare retrieval cues that allow them to temporarily license unlicensed NPIs, leading to illusions of grammaticality. If this is possible, we could imagine that quantifiers contain at least some of the desirable features required by NPIs. We can propose a [+quantification] cue which could help NPIs search for their licensers if memory retrieval is still a possibility. A cue such as this one may also seem to be restrictive considering all the cases of NPI licensing, but if some general properties of quantifiers are able to partially license NPIs or are enough to give rise to illusions of grammaticality, it could at least lead us in the right direction. To test this hypothesis, we next examined whether this illusory licensing effect occurs with quantifiers more generally or if this is due to the potential ability of the universal quantifier to license NPIs. On the next experiment we turn to the existential quantifier. If it is able to give rise to illusions of grammaticality, we can expect a feature linked to general properties of quantifiers to be involved in the retrieval process.

2.3.3 Experiment 2: Are illusions caused by general properties of quantifiers? The case of the existential quantifier *some*

In this experiment we tested whether the illusory licensing effect found on experiment 1 can be found with the use of the existential quantifier *some*.

2.3.3.1 Participants

Thirty two native English speakers were recruited through Amazon’s Mechanical Turk. They were compensated \$3.00. The experiment lasted approximately 30 minutes.

2.3.3.2 Materials

For this experiment we used the same items as in experiment 1, but we replaced the universal quantifier “every” with the existential quantifier *some*. Sample items can be observed in Table 5. As before, the items are organized in a 2 (NPI vs. No NPI) x 3 (Negation *no*, Quantifier *some*, Determiner *the*) design. In this case, *some* is not able to license NPIs, but if the reason why illusions rose in the previous experiment is a memory issues based on cues related to a general property of quantifiers, we might find similar results if there is a partial match of cues between the universal quantifier and the existential quantifier.

Table 5: Sample items Experiment 2: Speeded Acceptability Judgments

NPI	Licensor	Item
NPI	Negation	No/researcher/has/ever/been/invited/to/the/annual/conference.
NPI	Existential	Some/researcher/has/ever/been/invited/to/the/annual/conference.
NPI	Determiner	The/researcher/has/ever/been/invited/to/the/annual/conference.
No NPI	Negation	No/researcher/has/been/invited/to/the/annual/conference.
No NPI	Existential	Some/researcher/has/been/invited/to/the/annual/conference.
No NPI	Determiner	The/researcher/has/been/invited/to/the/annual/conference.

2.3.3.3 Procedure

We followed the same procedures as in experiment 1. The experiment lasted approximately 30 minutes.

2.3.3.4 Data Analysis

Data was analysed utilizing the same methods that were used in experiment 1. Similarly to experiment 1 NPI Presence was sum coded (NPI = 1, No NPI = -1) and Licensor was Helmert

coded to include two orthogonal contrasts: grammaticality ('no' = 1, 'some' = -0.5, 'the' = -0.5), and illusion ('no' = 0, 'some' = 1, 'the' = -1). Paired comparisons were also included in a similar manner as with experiment 1, using the same methods. Based on model criticism for reaction times, 1.27% data was lost. No participants were excluded from this study.

2.3.3.5 Results




Figure 7: Acceptability rating for experiment 2. Error bars show standard error.

Table 6: Experiment 2 Linear Mixed Effects Modelling

	Estimate	Std. Error	z value	p value	
(Intercept)	0.923	0.211	4.376	< 0.001	***
NPI Presence	-0.754	0.110	-6.836	< 0.001	***
Grammaticality (Licensed v. Unlic.)	1.306	0.116	11.271	< 0.001	***
Illusion (Existential v. Definite)	-0.294	0.086	-3.436	< 0.001	***
NPI Presence x Grammaticality	1.239	0.116	10.702	< 0.001	***
NPI Presence x Illusion	0.271	0.08	3.151	0.002	**

Analysis using logistic regressions (Table 6) shows main effects of NPI presence ($z = -6.836$; $p < 0.001$), Grammaticality ($z = 11.271$, $p < 0.001$). As with experiment 1, the presence of an NPI greatly affects how experimental items were judged by participants, the same with grammaticality of sentences. Contrary to what we found in experiment 1, we now have a

main effect of Illusion ($z = -3.436$; $p < 0.001$) that was not present before. Observation of the graph on figure 7 though, shows that this effect is most probably driven by differences between the judgment of the existential quantifier and the determiner cases in the No NPI cases, rather than with the NPI cases. The analysis also revealed an interaction of NPI x Grammaticality ($z = 10.702$; $p < 0.001$) which was expected due to the ungrammatical cases are only found in the NPI part of the items, and an interaction between NPI x Illusion ($z = 3.151$; $p = 0.002$), which is driven by a bigger difference between the quantifier and the determiner in No NPI cases rather than the NPI cases.

Table 7: Experiment 2: Planned comparisons

	Estimate	Std. Error	z value	p value	
(Intercept)	0.923	0.211	4.380	< .001	***
No NPI v. NPI	-0.754	0.110	-6.840	< .001	***
No NPI, Grammaticality (Licensed v. Unlicensed)	0.067	0.144	0.465	0.642	
No NPI, Illusion (Existential v. Definite)	-0.565	0.126	-4.490	< .001	***
NPI, Grammaticality (Licensed v. Unlicensed)	2.550	0.182	14.000	< .001	***
NPI, Illusion (Existential v. Definite)	-0.024	0.116	-0.206	0.837	

Planned comparisons (Table 7) show an effect of NPI ($z = -6.840$; $p < 0.001$), as observed in the logistic regression model analysis. We also found a significant result in Grammaticality for NPI cases ($z = -14.000$; $p < 0.001$) due to the ungrammaticality of determiner and existential quantifier cases compared to the grammatical negative quantifier. Interestingly, there was a significant difference between the No NPI cases and Illusion (Existential vs. Definite; $z = -4.490$; $p < 0.001$). This difference, though, was not found for the NPI cases ($z = -0.206$; $p = 0.837$). While both conditions are equally grammatical on the No NPI cases, this could reflect a higher difficulty or higher processing demand when an existential quantifier is present, and explains that the NPI x Illusion interaction found on the analysis is driven by the No NPI cases rather than the NPI cases. The current results show, therefore, that in the case of the existential quantifier do not seem to trigger illusions with NPIs, but rather, that the existential quantifier is as easily identified as ungrammatical as the determiner cases.

When we looked at the intrusion effect in terms of the difference between the means, we observed that in the NPI cases the existential quantifier cases were rated marginally worse than the determiner cases by 0.89%. Meanwhile in the No NPI cases, the existential quantifier sentences were rated 15.48% worse than the determiner cases. The difference in rating between both conditions without an NPI seems to indicate that the presence of the existential quantifier itself has a negative impact in the perception of the sentence despite being equally grammatical as those with a negative quantifier or a determiner. On the other hand, the marginal difference in ratings between the determiner and the existential quantifier seem to indicate that people didn't have many problems recognising the existential quantifier as an inadequate licenser for an NPI at par with the unambiguous definite article. Due to this differences we can confirm that the interaction between NPI presence and "illusion" (Determiner vs Existential) found in the logistic regression analysis was driven by the differences in the No NPI cases.

2.3.3.6 Discussion

On their own, the results observed in experiment 2 indicate that there is no illusory licensing happening with the existential quantifier, while the results found in experiment 1 by themselves to show that illusory licensing with the universal quantifier is possible. If we compare the results of these individual experiments together, while experiment 1 shows that illusory licensing may occur with *every*, experiment 2 shows that this does not seem to be general effect of quantification. Quite the contrary, the results of experiment 2 show that NPI sentences with the existential quantifier as a potential licenser elicit similar, if not identical, results to those that contain a determiner. This also means that if memory mechanisms are involved in the licensing process, a feature [+quantification] cannot be involved since its existence fails to predict the patterns observed in experiment 2.

In terms of how the illusion is produced, while the presence of illusory licensing for the universal quantifier in a structural illicit position with respect to the NPI could be explained by a noisy content-addressable, cue-based memory process where at least some cues in *every* (but not

general to quantification) are retrieved, this is still not evidence that there are no other processes involved in NPI licensing. As discussed before, illusory licensing could also happen due to the application of semantic and pragmatic processes involved.

Our experiments have been able to show that illusory licensing with the universal quantifier exists, at least in some measure, but the nature of why or how it arises is still not yet clear. In order to get a better view of the processes involved more testing is necessary, this time manipulating semantic/pragmatic properties that may be related to NPI licensing. Fortunately, the universal quantifier *every* provides us with such a mechanism. The reason why illusory licensing arises may be due to semantic and pragmatic properties such as the calculation of the quantifier's restriction and scope. It is possible that at the time when the subjects encountered the NPI in the sentences from experiment 1, they still have not yet finished identifying the restriction of *every*, and therefore, didn't yet realized that the NPI was not within it since the restrictor was still undetermined.

In general, the results observed in experiment 1 are consistent with both the memory account (McElree, 2000; McElree et al., 2003, Vasisth et al., 2008), as well as through a different route which involves the semantic calculation of the restriction of *every*. In order to further test these theories, we must observe if manipulations that help the processor identify the restriction of a quantifier affect the rate of NPI licensing illusions. If the manipulations to the restrictor fail to affect the rate of illusory licensing and the results continue to be similar to experiment 1, then we could conclude that the illusory effect is driven by memory mechanisms and that content-addressable memory is accessing a structurally illicit licenser as part of noisy cue-based retrieval. However, if the rate of illusions is considerably reduced or disappears, we can conclude that the illusion is driven by identification of the quantifier's restriction. That is to say, that illusory licensing of NPIs is driven by the identification of the quantifier's restriction, and not by noisy memory retrieval.

2.3.4 Experiments 3 and 4: Illusory NPI licensing as a matter of ongoing restriction / scope calculations

As mentioned in the previous section, experiments 1 and 2 reveal that illusory licensing occurs with the universal quantifier in NPI licensing, at least in some measure, but that this illusion is due to a particular effect of the universal quantifier, and not due to properties of quantifiers generally. The current results agree with the theory that illusions are both a general property of NPI licensing and due to a faulty licensing mechanisms (Parker and Phillips, 2016). Nevertheless, these results do not show the underlying causes of the effect. While these results are compatible with a memory account of grammaticality illusions, they are also compatible with an account where illusions arise from the application of other semantic/pragmatic properties of NPI licensing (Xiang et al., 2009; 2013).

According to Frazier (1999), every overt NP in subject position will immediately initiate a search for a restriction clause. The subject of the sentence will always be either restricted by context, or will introduce a new entity into the discourse. Furthermore, she shows that restriction search may look for previous context if it is present, but also, that the search can continue into novel context that follows the main noun. It is possible that in experiment 1, due to the absence of a well-defined restrictor clause, as well as the lack of previous context, the search for the restrictor is still an ongoing process and that therefore, the mechanisms that define whether the restrictor is closed may still be “active” and have not yet fully interpreted a restriction for “Every N”. If this is the case, the reason why we find evidence of illusions in experiment 1 could be at least partially because the mechanisms involved in the licensing process could identify the NPI as being within the restrictor where it would be grammatical.

If illusory licensing is driven by a memory mechanism (exclusively), manipulation of the restriction with the addition of *should* not greatly weaken illusory licensing. Conversely, if it is driven by an active “restriction-checking” mechanism manipulating the limits of the quantifier’s restriction should effectively weaken the effect if these manipulations allow the parser to define

a limit to the restriction with more ease. Even more so, manipulation of the restrictor could be different between pre-nominal and post-nominal modification. While a prenominal modifier may define the restrictor early (similar as to how previous context would), a post-nominal modifier may keep it open until later, and therefore closer to the NPI, which to a certain extent, may still be considered erroneously to be within the restrictor. In order to test this hypothesis about the mechanisms behind universal quantifier produced illusions in NPI licensing, we manipulated the restriction of the universal quantifier with both types of nominal modification in two separate experiments.

2.3.5 Experiment 3: Identifying the universal quantifier's restriction – Prenominal Modification

In this experiment we increased the length of the quantifier's restriction by adding a pre-nominal modifier in the form of an adjective to the main noun of the subject. By increasing the elements within the restrictor we give participants a greater chance of identifying the boundaries between the restrictor and the scope of the quantifier, in this case, by quickly establishing a contrasting set. If the illusory effect is mainly driven by memory retrieval, then we should see no significant variation from the results found in experiment 1. Nevertheless, if the effect is driven by the correct online identification of the restriction, this modulation should show variant results.

2.3.5.1 *Participants*

We recruited thirty two native English speakers through Amazon's Mechanical Turk. Each participant was compensated with \$3.00. The experiment lasted approximately 30 minutes.

2.3.5.2 *Materials*

For this experiment, we used the same items from experiment 1 but with the addition of a pre-nominal adjective in order to manipulate the restriction of the universal quantifier. The subject of the sentence was formed by the potential licenser (*no*, for the licensed condition, and *every* or *the* for the unlicensed conditions), followed by an adjective and the subject main noun. The

main predicate phrase immediately followed the main noun in the same way that was described in experiment 1. Table 8 contains examples of all experimental conditions.

Table 8: Sample stimuli for experiment 3: Speeded Acceptability Judgments

NPI	Licensors	Item
NPI	Negation	No/young/researcher/has/ever/been/invited/to/the/annual/ conference.
NPI	Universal	Every/young/researcher/has/ever/been/invited/to/the/annual/ conference.
NPI	Determiner	The/young/researcher/has/ever/been/invited/to/the/annual/ conference.
No NPI	Negation	No/young/researcher/has/been/invited/to/the/annual/conference.
No NPI	Universal	Every/young/researcher/has/been/invited/to/the/annual/conference.
No NPI	Determiner	The/young/researcher/has/been/invited/to/the/annual/conference.

2.3.5.3 Procedure

This experiment followed the same procedure as experiments 1 and 2.

2.3.5.4 Data Analysis

The data was analysed with R utilizing the same methods described in experiment 1 and 2.

NPI presence was sum-coded (NPI = 1, No NPI = -1) and Licensors was Helmert coded into two orthogonal contrasts: grammaticality ('no' = 1, 'every + adj.' = -0.5, 'the' = -0.5), and Illusion ('no' = 0, 'every + adj.' = 1, 'the' = -1). Data was further cleaned using model criticism to remove extreme reaction time values with little data loss (2.72%).

2.3.5.5 Results




Figure 8: Acceptability rating for experiment 3. Error bars show standard error

Table 9: Experiment 3 Logistic Regression Mixed Effects Modelling

	Estimate	Std. Error	z value	p value	
(Intercept)	0.718	0.131	5.461	< 0.001	***
NPI Presence	-0.975	0.135	-7.209	< 0.001	***
Grammaticality (Licensed v. Unlic.)	1.042	0.103	10.050	< 0.001	***
Illusion (Universal v. Definite)	-0.011	0.086	-0.134	0.893	
NPI Presence x Grammaticality	0.949	0.103	9.204	< 0.001	***
NPI Presence x Illusion	0.173	0.085	2.026	0.043	*

Same as with the previous experiments, we started by performing logistic regression via mixed effects modelling which can be observed in table 9. The results obtained show an effect of NPI presence ($z = -7.209$; $p < 0.001$) as well as an effect of Grammaticality ($z = 10.050$; $p < 0.001$). The difference between the universal condition and the determiner condition (“illusion”) was not significant ($z = -0.134$; $p = 0.893$). Meanwhile, both the interaction of NPI x Grammaticality ($z = -8.886$; $p < 0.001$) and an interaction of NPI x Illusion ($z = -2.028$; $p = 0.043$) showed significant results. Significant grammaticality results are driven by the difference between the negative and the other two conditions in the NPI cases, where both the universal and the determiner are unlicensed. The presence of an interaction between the NPI and the Illusion effect could reflect that illusions of grammaticality were in some measure present in this experiment, like in experiment 1, which is consistent with a theory of memory driven licensing of NPIs. The lack of a general effect is probably due to the universal + adjective condition being rated slightly worse than the determiner in the No NPI cases, while being slightly better rated than the determiner in the NPI cases.

Table 10: Experiment 3 results by condition

	Estimate	Std. Error	z value	p value	
(Intercept)	0.717	0.131	5.461	< 0.001	***
No NPI v. NPI	-0.975	0.135	-7.209	< 0.001	***
No NPI, Grammaticality (Licensed v. Unlicensed)	0.093	0.146	0.638	0.523	
No NPI, Illusion (Universal v. Definite)	-0.184	0.123	-1.495	0.135	
NPI, Grammaticality (Licensed v. Unlicensed)	1.991	0.146	13.624	< 0.001	***
NPI, Illusion (Universal v. Definite)	0.161	0.118	1.363	0.173	

Following, we looked at the planned comparisons, show on table 10. These results show an effect of NPI presence ($z = -7.209$; $p < 0.001$), and an effect of Grammaticality within the

NPI conditions ($z = 13.624$, $p < 0.001$), which is expected since ungrammatical items were only used with NPI cases and only negative sentences were grammatical. Despite the results found in the logic regression model, significance was not found in the Illusion contrast in the NPI cases ($z = 1.363$; $p = 0.173$). This means that acceptability ratings for sentences that contained the universal quantifier with an adjective were not significantly different from the ratings of the determiner sentences with an NPI.

One thing to note is that the universal quantifier condition within No NPI cases is rated slightly worse than both the negative and the determiner conditions, an observation that was also noticed in experiment 1. It is possible that the universal quantifier presents intrinsic difficulties that are reflected during acceptability ratings as evidenced by the lowered rating in the No NPI cases. In order to assess the effect of the NPI itself, we compared the difference in the mean percentage of acceptability between the determiner and the quantifier cases in both the No NPI and NPI cases as a post-hoc analysis. In the No NPI cases, the quantifier was rated 4.83% worse than the determiner. Meanwhile, In the NPI cases, sentences with every + adjective were rated 5.7% better than those that contained the determiner. Comparing between the NPI and the No NPI cases, we can observe a 10.53% difference between the quantifier and the determiner as a cost of adding an NPI. While the results in the planned comparisons do not show a significant effect in the “illusion” comparison (Determiner vs Universal Quantifier) on the NPI cases, the interaction observed in the logistic regression analysis between NPI x Illusion is probably due to these differences in ratings between both the NPI and the No NPI cases. In general, the difference between mean ratings in the NPI cases in experiment 3 between the universal and the determiner is slightly smaller than the one found in experiment 1 (7.67%) and not big enough to reach significance, though the difference taking into account both the NPI and No NPI cases is roughly the same as in experiment 1 (10.58%). While we weren’t able to find a significant effect for illusory licensing, the differences between determiner and universal

quantifier + adjective cases still show a trend for slightly higher ratings compared to those that contained the definite article, a contrast with experiment 2 that showed no such difference. Whether these difference in acceptability ratings indicate that a baseline difficulty for the universal quantifier may also be taking effect in the NPI cases and that this could be a factor in ratings is not determined by these studies and would require further examination.

2.3.5.6 Discussion

Experiment 3 shows data which contrasts the results obtained in Experiment 1. The previous experiment was indicative that illusory licensing was a consequence of having a universal quantifier as a potential, but incorrect licenser to an NPI located within the scope of the quantifier, but it did not give us any information about the processes that drive the illusory licensing. This experiment, however, manipulated the restriction of the quantifier in order to make its boundaries more clear, allowing us to test whether this manipulation can counteract the illusory effects. In first instance, we see that with the addition of an adjective to the restrictor clause the difference between the acceptability ratings of the universal quantifier and the determiner did not yield a significant result. If the effect of the illusion was weakened by this manipulation, this could favour the theory that the illusion effect, at least in these cases, is primarily driven by semantic/pragmatic processes. If illusory licensing was a matter exclusively of cue-based memory retrieval we would have expected to find consistent evidence of illusory licensing regardless of the manipulations made to the restrictor since a cue-based mechanism would analyse all elements in memory simultaneously and would, therefore, be able to select a partially matched, or structurally inaccessible licenser regardless of the elements contained in the restrictor. We are still able to observe a trend towards an illusion effect, which means that it has not entirely disappeared, but is no longer strong enough to be significant. It is important to note that given that the restrictor of the quantifier can be further restricted by context, the

semantic/pragmatic processes of identifying a restrictor clause is fairly open-ended, and as such, some degree of illusory licensing can still be expected. Given this observation, manipulation of the restrictor's length could affect the amount of illusory licensing by providing increased context to it, as we will observe in Experiment 4.

These results provide some evidence that manipulating the restriction of the quantifier by adding a pre-nominal modifier in the form of an adjective, may have caused the illusory licensing of NPI to be weakened to the point that pairwise statistical analysis no longer shows a significant effect, at least when compared to the much more clear difference found in experiment 1. There is still a possibility that differences between the means could be obscured by the differences observed in the No NPI cases if an extra cost for the universal quantifier is also a factor that plays a part in participant judgements compared to the determiner. Regardless, experiment 1 still manages to show a significant illusion effect for NPI cases that does not seem to be present in experiment 3, despite a similar trend in better judgments for universal quantifier cases compared to the definite, which could lead us to speculate that it is possible that the differences between both experiments might be due to the presence of the adjective (though this speculation should be taken with a grain of salt since we didn't directly compare both conditions in experiments 3 and 4).

2.3.6 Experiment 4: Identifying the universal quantifier's restriction – Postnominal Modification

In order to corroborate the findings of experiment 3, we tested the previous hypothesis by further modifying the restrictor. In this case we added a post-nominal modifier to the subject of the sentence in the form of a relative clause rather than a pre-nominal one as we did in experiment 3. While this increases the length of the restrictor even more than in experiment 3, it also greatly increases the number of elements between the NPI and its potential licensor. This experiment allows us to further observe if the effects found in experiment 3 can be found in different kinds of modification.

2.3.6.1 Participants

We recruited thirty six native English speakers through Amazon’s Mechanical Turk. Participants. Each participant was compensated with \$3.00. The experiment lasted approximately 30 minutes.

2.3.6.2 Materials

For this experiment, we modified the items from experiment 1 by adding a relative clause following the main subject noun. The relative clause always started with *who was*, making sure that the relative clause was overtly established and to avoid the ambiguity of a covert relative clause. This was immediately followed by a verb and 3 words that formed a complement or complements to the verbal phrase. This was in most cases a direct object “who was read *a bedtime story*”, but it could also be a prepositional phrase “who was marched *through the trenches*”, or in one case, a direct object followed by a temporal adverb “who was cooked *dinner last night*”. In total, the inclusion of a relative clause increased the number of words between the Licensor and the NPI by 6 words relative to the items in experiment 1, while experiment 3 added one. Examples of all conditions can be found in Table 11.

Table 11: Sample items Experiment 4: Speeded Acceptability Judgments

NPI	Licensor	Item
NPI	Negation	No/researcher/who/was/awarded/a/scientific/prize/has/ever/been/invited/to/the/annual/conference.
NPI	Universal	Every/researcher/who/was/awarded/a/scientific/prize/has/ever/been/ invited/to/the/annual/conference.
NPI	Determiner	The/researcher/who/was/awarded/a/scientific/prize/has/ever/been/ invited/to/the/annual/conference.
No NPI	Negation	No/researcher/who/was/awarded/a/scientific/prize/has/been/invited/to/the/annual/conference.
No NPI	Universal	Some/researcher/who/was/awarded/a/scientific/prize/has/been/invited/to/the/annual/conference.
No NPI	Determiner	The/researcher/who/was/awarded/a/scientific/prize/has/been/invited/to/the/annual/conference.

2.3.6.3 Procedure

This experiment followed the same procedure described in experiment 1 and that was used for experiments 2 and 3 as well.

2.3.6.4 Data Analysis

The data was analysed using R with the same methods described in experiment 1. Data cleaned using model criticism resulted in a total of 1.76% data loss.

2.3.6.5 Results




Figure 9: Experiment 4 Acceptability Rating. Error bars indicate standard error.

Table 12: Experiment 4 Logistic Regression Mixed Effects Modelling

	Estimate	Std. Error	z value	p value	
(Intercept)	0.223	0.201	1.111	0.267	
NPI Presence	-0.716	0.104	-6.857	< 0.001	***
Grammaticality (Licensed v. Unlic.)	0.566	0.085	6.671	< 0.001	***
Illusion (Universal v. Definite)	0.058	0.073	0.781	0.435	
NPI Presence x Grammaticality	0.683	0.085	8.039	< 0.001	***
NPI Presence x Illusion	0.090	0.074	1.215	0.224	

Logistic regression results (table 12) show an effect of NPI ($z = -6.857$; $p < 0.001$), as well as an effect for our grammaticality condition ($z = 6.671$; $p < 0.001$). We also found the expected interaction between NPI and Grammaticality ($z = 8.039$; $p < 0.001$) since both universal and determiner cases are ungrammatical while the negative cases remain grammatical with an NPI. In this case, though, we did not find an interaction between NPI and Illusion effect ($z = 1.215$; $p = 0.224$) like we did in experiment 3.

Table 13: Experiment 4 Fixed Effects of Planned Comparisons

	Estimate	Std. Error	z value	p value	
(Intercept)	0.223	0.201	1.111	0.267	
No NPI v. NPI	-0.716	0.104	6.857	< 0.001	***
No NPI, Grammaticality (Licensed v. Unlicensed)	-0.117	0.113	-1.036	0.300	
No NPI, Illusion (Universal v. Definite)	-0.032	0.100	-0.322	0.747	
NPI, Grammaticality (Licensed v. Unlicensed)	1.248	0.126	9.877	< 0.001	***
NPI, Illusion (Universal v. Definite)	0.147	0.109	1.357	0.175	

Planned comparisons (Table 13) reflect the findings from the logistic regression model above: we found an effect of NPI ($z = 6.857$; $p < 0.001$) consistent with our previous experiments, as well as the expected effect of Grammaticality within the NPI conditions ($z = 9.877$; $p < 0.001$) due to only the negative conditions being grammatical while the determiner and the universal quantifier were not. More importantly, similarly to experiment 3, the results show no significant difference for the illusion effect within the NPI conditions ($z = 1.357$; $p = 0.175$) which shows that the acceptability rating of NPI sentences that contained the universal quantifier were not significantly different from those that had the definite article. These results, when compared to the effect observed in experiments 1 and 3, are consistent with our hypothesis that the illusion effect in NPI licensing with the universal quantifier could be weakened by the modulation of the restrictor clause by providing a previously determined contrasting set, this time by the inclusion of a relative clause, a result that is consistent with the theory that NPI illusory licensing could be driven by the application of semantic/pragmatic properties of the licensing process rather than by faulty memory retrieval mechanisms.

Once again, we looked at the differences in ratings between the determiner and the universal quantifier as a post-hoc test to observe the size of the effect. In this case, we found that in the No NPI cases, the universal quantifier was rated 1.27% worse than the determiner. While a smaller difference than in the other experiments, still consistent with what we observed in the previous experiments that contained the universal quantifier. Similarly, in the NPI cases, we found that the universal quantifier was rated 5.47% better than the determiner, continuing with

the trend that participants will rate the universal quantifier sentences slightly better than the determiner ones when it contains an NPI. Additionally, there are a couple of observations worth mentioning about the acceptability ratings in this experiment. Firstly, in experiment 4 the overall rating of sentences was lower than in all other experiments across all conditions, including the No NPI ones. This probably indicates that there is a great increase in difficulty of processing the relative clause. Secondly, though only by a small margin (2.96%), this is the only case in which the universal quantifier was rated better in the No NPI cases than the Negative.

2.3.6.6 Discussion

The results obtained from experiment 4 are consistent with the results found in experiment 3, which shows that the manipulation of the restrictor by a post-nominal modifier had a weakening effect over illusory NPI licensing with the universal quantifier. These results could also provide consistent evidence with the theory that illusory licensing is not driven by memory mechanisms, but rather by the application of semantic/pragmatic properties of NPI licensing. This effect has been observed consistently across different types of modification (pre-nominal and post-nominal) across both experiments 3 and 4 which shows that it is not an effect exclusive to specific types of nominal modification.

2.4 General Discussion

Through the current research, we looked into the effects of quantification in the rise of illusions of grammaticality in negative polarity items by testing two different accounts for illusory licensing in NPI licensing. While one account posits that illusions are caused by the incorrect retrieval of a structurally inaccessible licenser as part of noisy cue-based, content-addressable memory retrieval process, the other account suggested that the rise of illusions involves the application of semantic and pragmatic processes involved in the licensing of NPIs which do not include memory mechanisms (Xiang, Dillon, & Phillips, 2009; Xiang, Grove, & Giannakidou, 2013).

The experimental design we followed allowed us to observe, first, whether illusions could appear with the universal quantifier. Second, we compared these results with a different quantifier, the existential quantifier *some* which does not license NPis in order to observe whether the rise of illusions could be attributed to a general property of quantifiers, or if it is a property unique to the universal quantifier. Third, we compared the results of experiment 1 with two more experiments where we manipulated the restriction clause of the quantifier to make it more easily identifiable to the parser and to observe whether this manipulation reduced the rise of illusions. We divided this task into two experiments, one using a pre-nominal modifier (experiment 3) and another using a post-nominal modifier (experiment 4).

From experiment 1, we observe that given certain conditions, NPI sentences that contain the universal quantifier 'every' as a potential licensor are rated better than those that have the determiner 'the'. This provides evidence that illusory licensing can arise with the universal quantifier in NPI licensing when the NPI is outside the restrictor clause. The results of experiment 2 confirm that these illusions are not produced by the presence of quantifiers more generally but that the effect is specific to a property of the universal quantifier itself. Furthermore, if memory mechanisms are involved in the retrieval of the licensor, the results found in experiment 1 are consistent with content-addressable memory since the structurally inaccessible licensor was successfully retrieved.

Moreover, the reason why these illusions appear are because of something related to the universal quantifier *every*, and not as a requirement by the NPI for the parser to search for any general form of quantification. If cue-based search is involved in the retrieval of the licensor, [+quantification] is not a feature that is searched for. If this were the case, we could expect other quantifiers to be able to give rise to illusions, or at least show some difference in acceptability ratings, regardless of whether they naturally license NPis or not, because of partial feature match. It is still possible, if memory is involved, that a feature is retrieved that may more closely

resemble [+Downward Entailing] but this would be subject to the same difficulties as a [+Negative] feature in that it is not able to account for all possible cases of NPI licensing. Furthermore, given the semantic/pragmatic nature of the restriction clause and the nuclear scope of a quantifier as contexts, and the difficulty in clearly defining the limits of each of these components whose processing seems to be open-ended, it is difficult to easily attribute a [+DE] (or [+UE]) feature to specific words or phrases, in other words, when and where would these features be assigned? Where does the assigned feature end and how is this determined? Considering that upward and downward entailing are defined as contexts which allow certain types of entailment to arise, it would be too complicated to specify them as just a “feature”. We would have to explain that every single word or phrase contains these features, which in turn can be changed by single word with the opposite feature and affect the continuous context. So if we have an affirmative sentence, the phrase is considered to be UE, but then we encounter a word such as *doubt* that has a [+DE] feature which affects the following context and changes the “feature” of the ongoing phrase or sentence. Defining where this “feature” starts or ends would be too complicated and begs the question of what exactly is being retrieved by a memory mechanism. Furthermore, in the case of a quantifier such as *every* which introduces both a [DE] and a [UE] component, it is difficult, based on the quantifier itself to determine where and how these “features” are being assigned. In contrast, a quantifier such as *no* introduces two [DE] components (its restriction and its scope) so establishing that the end of a phrase resets the entailing context doesn’t seem to be a valid possibility. Given that these components are open-ended and that a calculation must be happening in order to determine the entailing context of the sentence or phrase, it seems to be more logical to assume that it is a semantic property which is not static or previously determined by the apparition of a single word. While a [+Negative] feature is difficult to explain for all cases, a [+DE] feature seems to be way too hard to implement in a simple system and would have to affect not just the licenser but the whole phrase. If this is the case, the retrieval of such a feature means that we have to retrieve whatever

onsets the context as a licenser without determining where it actually stops, or the whole phrase which contains the DE context, which at the same time would have to contain the NPI itself if we are required to determine the extent of the DE context. On one hand, it would be too costly to retrieve such a large portion of speech, and on the other hand, we would have to say that the licenser (the DE phrase) contains the licensee which seems, to say the least, improbable. It seems that, rather than a specific feature belonging to a specific word (a licenser in this case), we are looking at general semantic/pragmatic properties of whole constructions involving quantifiers as explained by the tripartite structure of quantifiers (Kamp, 1981; Heim, 1982, Von Stechow, 1994), and which are being constructed and calculated as the sentences unfold. Given this possibility, the next step would be to investigate the timing of the temporary illusory effect with online methods to observe when these properties are being integrated into the licensing process.

While the results of experiment 1 and 2 allowed us to see whether illusions arise with which quantifiers, they do not allow us to test both accounts of why these illusions arise. On the contrary, experiments 3 and 4 were designed specifically for the purpose of comparing both theories. We theorised that by increasing the number of constituents inside the restriction of *every* we would make the calculation of the delimitation of the restriction more clear than in experiment 1. If illusory licensing can be manipulated in this way, and the illusion weakens or disappears, the theory of memory as the main source for the illusions will also be weakened. Content-addressable memory as a mechanism for the retrieval of the licenser postulates that all elements in the preceding context of the triggering word are considered by the parser at the same time, and the most viable candidate is chosen among all the possibilities based on activation that is prompted by matching the desirable features of the retrieval cue. Based on this, we would expect the number of elements to not affect the viability of a potential licenser to be selected. Thus, if memory is involved in illusory licensing, we would expect the illusions to remain in experiments 3 and 4. Nevertheless, if the modulation of the restriction is enough for the illusions to weaken or disappear, we can no longer assume that a content-addressable

memory is involved, or at least that it is the main cause of illusory licensing. Conversely, if the modulation of the semantic context of the restriction, i.e. establishing a more clear boundary between the restriction and the scope of the universal quantifier is enough for illusions to no longer be present, then we can assume that illusory licensing is actually driven by semantic processes, rather than memory ones.

The results of experiments 3 and 4 show that illusory licensing is actually weakened by the presence of both pre-nominal and post-nominal modification. In both cases, the difference between the quantifier *every* and the determiner *the* conditions is no longer significant in comparison to the results observed in experiment 1 where the difference was indeed significant. We can also observe this by looking at the difference between the means across the experiments. Where experiment 1 showed a difference of almost 8%, experiments 3 and 4 show a difference between the quantifier and the determiner cases of about 5%, a more reduced difference. These results could provide support for an account of semantic properties driving the rise of illusions, rather than an account based on memory retrieval mechanisms. However, it is possible that memory retrieval might still play an important part on the licensing process as evidenced by the trend found by applying a base-factor analysis to observe the size of the effect (difference between the means). It is possible that for NPI licensing with the universal quantifier the licensing mechanisms involved resort to a different path of memory driven licensing in cases where the restriction of *every* is not well defined and another strategy must be used to try to retrieve a potential licensor for the NPI.

While some of the differences between the experiments could also be explained by the increased processing requirements of each modification, as observed by the general decrease in acceptability ratings in No NPI cases with a relative clause in experiment 4, there does seem to be a consistent (however small) observable difference in ratings between the universal quantifier and the determiner conditions in NPI sentences when compared to experiment 2, which

provides a contrast with the rest of the experiments in that it shows that a non-licensing quantifier ('some') is rated the same (if not slightly worse) than the more unambiguous definite article 'the'. Additionally, even though the percentage of positive ratings in experiment 1 for both the determiner and the quantifier conditions with NPIs are in general lower ('the' = 11.27%, 'every' = 18.95%) than in other experiments (experiment 2: 'the' = 25.35%, 'some' = 24.46; experiment 3: 'the' = 19.55%, 'every' + adj. = 25.13%; experiment 4: 'the' = 22.02%, 'every' + RC = 27.49%), the direct comparison between both conditions in experiment 1 is significantly different.

The results to each of these individual experiments could provide some evidence that illusory NPI licensing can be driven by semantic processes, and that when the extent of a universal quantifier's restrictor is uncertain, the parser may temporarily license a structurally unlicensed NPI. Nevertheless, further research is needed to corroborate the relationship between all four experiments. Future studies should consider gathering acceptability ratings of all universal quantifier conditions together to ensure that the results are comparable.

2.5 Conclusions

While experiment 1 showed that illusory licensing can occur with a structurally inaccessible universal quantifier in NPI licensing, experiment 2, showed that this effect is not due to a general property of quantifiers, but rather that the effect is caused by a specific property of the universal quantifier. This results discard the possibility that a feature such as [+quantifier] could be involved in the licensing process of NPIs, but rather that there is something more specific in the universal quantifier that licenses NPIs. In terms of memory retrieval mechanisms, if memory is involved, the results are consistent with a content-addressable mechanism rather than a serial search one.

Meanwhile, experiments 3 and 4 provided evidence to support the hypothesis that the illusory licensing could be driven mainly by semantic processes related to the correct identification of

the restriction of the universal quantifier, rather than noisy cue-based memory retrieval exclusively. In a sentence where the restriction of the quantifier is well delimited, illusions should not be expected to rise consistently, but uncertainty about the quantifier's restriction during processing may actually be able to temporarily license a structurally unlicensed NPI.

Chapter 3: Quantifiers in NPI Licensing – Self-Paced Reading Studies

3.1 Introduction

In the previous chapter, we explored the rise of illusions of grammaticality in NPI licensing with quantifiers as licensors using speeded acceptability judgments. We found that illusions arise in ungrammatical configurations that contain an NPI outside of the restrictor clause of a potential universal quantifier licensor. This effect was not found for the existential quantifier "some", thus discarding the hypothesis that illusory licensing in NPIs occurs due to general properties of quantifiers, but rather occurs exclusively with specific quantifiers. Furthermore, we observed that in sentences where the restriction of the universal quantifier was better defined by adding either a prenominal (adjective) or a postnominal (RC) modifier, the illusory licensing effect was reduced to the point where it was no longer statistically significant. Based on the observed data, we concluded that this was due to the restrictor clause providing the parser with a sufficient contrasting set, which therefore established that the boundaries of the restrictor clause were already determined. This made it more evident to the parser that upon reaching the NPI, it was already outside of the restrictor clause. Given that the parser has more clues about the limits of the restrictor clause, it becomes easier for the parser to realize that the NPI is not within a licit structural position and that the sentence is consequently not grammatical.

These findings provide us with evidence that the semantic and pragmatic properties of the sentence are not only involved during the processing of NPIs but that the licensing of NPIs is mainly driven by these semantic and pragmatic properties, rather than exclusively by memory mechanisms involved in long-distance dependency resolution (though not entirely discarding their involvement in the licensing process). Although offline methods such as judgment tasks allow us to observe people's perceptions of grammaticality, they only provide us with a limited

window into the processing involved in NPI licensing, since they only show us data collected after the whole sentence has been processed. Though it is possible for illusory effects to persist after initial processing, research has shown that illusions can be overcome given enough time to reflect upon the sentence. Therefore, even in speeded acceptability judgments, where people are given a few seconds to judge a sentence after having hastily read it, these methods cannot provide deep insight into the time course of the licensing process itself, the processing of illusions, or the mechanisms involved in rapidly overcoming these illusions.

In order for us to better understand the underlying mechanisms involved in the processing of NPIs and their illusory licensing, we turn to online methods that allow us to collect data during the ongoing processing of the sentence and, therefore, gain deeper and more fine-grained understanding of the illusory licensing. The current chapter will explore the effects discovered in the previous chapter using the self-paced reading task introduced by Just et al. (1980). This task allows us to collect reaction times of participants for each word they see, which offers a window into the incremental online processing of the sentence being read. We expect this method to provide a deeper and more fine-grained understanding of the processes involved in NPI licensing, as well as the time course of different licensing routes.

3.2 Background

Since its introduction, SPR has become widely used in psycholinguistics and it has been a reliable tool for researchers to better observe and understand phenomena in sentence processing in multiple areas of research such as subject-verb agreement, anaphora, and of course NPI licensing, among other long-distance dependencies.

For example, in subject-verb agreement, Pearlmutter et al. (1999) performed a series of experiments to test participants' sensitivity to agreement violations using self-paced reading (SPR) and eye-tracking methods using sentences such as the ones on (41) based on those used by Bock & Miller (1991).

(41)

[Grammatical, match]	The key to the <i>cabinet</i> _[s] <u>was</u> _[s] rusty from many years of disuse.
[Grammatical, mismatch]	The key to the <i>cabinets</i> _[pl] <u>was</u> _[s] rusty from many years of disuse.
[Ungrammatical, match]	The key to the <i>cabinet</i> _[s] <u>were</u> _[pl] rusty from many years of disuse.
[Ungrammatical, mismatch]	The key to the <i>cabinets</i> _[pl] <u>were</u> _[pl] rusty from many years of disuse.

In the sentences above, the main noun was always singular. The number of the noun inside the propositional phrase was manipulated to either match or mismatch the head noun, thus creating a local distractor between the head noun and the verb. The verb number was also manipulated to either maintain agreement with the head noun (singular – grammatical) or not (plural – ungrammatical). Pearlmutter and his colleagues wanted to observe whether comprehenders were sensitive to agreement violations during the course of normal reading, as well as the relative timing of this sensitivity, and whether the pattern of interference that had previously been observed in production could also be found in comprehension. The results of self-paced reading showed that participants were sensitive to both ungrammaticality and locally distracting plural (marked) nouns. They saw that sensitivity to head/local number-match appeared immediately at the verb where NP-match conditions were faster than NP-mismatch conditions. Sensitivity to grammaticality did not appear until the word following the verb, where reaction times were faster for grammatical conditions in comparison to ungrammatical ones. Furthermore, they noticed a specific interaction between both factors that was visible at the word following the verb where a mismatch in number between the head noun and the local noun increased the processing difficulty for the grammatical conditions at the verb which parallels the findings in production studies. However, for the ungrammatical cases the effect was inverted: a mismatch between both nouns actually eased the process. These results show that number mismatch between the two nouns increases the possibility of errors when calculating the number of the head noun in grammatical sentences, but fewer errors in ungrammatical ones. The effects found in SPR were corroborated by their eye-tracking study.

Similarly, Wagers et al. (2009) used a series of SPR experiments to study effects of agreement attraction in comprehension in order to understand the underlying processes that produce the effect.

Their study was designed to compare two types of theories that explain how agreement attractions occurs. On one hand, one type of account considers that agreement attraction happens due to the instability in the representation of features in the subject at encoding as a result of feature percolation (Eberhard et al., 2005; Franck et al., 2002; Nicol, Forster, & Veres, 1997; Vigliocco & Nicol, 1998). According to these accounts, the features of a given syntactic constituent can be transferred to the immediate dominating node. In this way, in a subject with a PP modifier such as (42), the plural number feature of the noun inside the PP percolate upwards, assigning the plural value to the higher phrasal projection. Given that the verb or VP is valued in agreement with the number of the subject phrase, when the PP-object's number percolates to the entire phrase the verb will be inappropriately valued as well. Agreement attraction is, in these kind of theories, a result of errors in the encoding of the number feature of the subject due to the transferral of the plural marked feature from the PP to the entire subject NP, rather than the singular non-marked feature of the head noun.

(42) The key to the cabinets [...]

On the other hand, the second type of theories propose that agreement attraction is primarily the result of errors during the construction of the dependency when accessing the subject's representation during planning or retrieval (Badecker & Lewis, 2007; Bock & Cutting, 1992; Solomon and Pearlmutter, 2004) According to Solomon and Pearlmutter (2004), in production attraction occurs because the features of both nouns are activated simultaneously during planning with the distractor noun interfering with the agreement processes before the representation of the subject is complete. Furthermore, the structural distance of the head noun and the attractor noun is also significant since structural distance often reflects conceptual

relatedness between the two nouns, and more tightly related nouns will induce more agreement errors during planning than less conceptually related ones: “the pizza with the yummy toppings” vs “the pizza with the tasty beverages”. In comprehension, the simultaneous activation of both nouns’ features might be realised at the moment the verb is encountered when the subject must be retrieved from memory to check for the verb’s number features. Crucially, according to these types of theories errors are not due to the faulty representation of the subject (as in the feature percolation), but rather involve errors in mapping (production) or retrieval (comprehension) when multiple elements are active at the same time.

(43) Wagers et al. (2009) : Example stimuli Experiment 1:

Singular Subject / Grammatical	The old key unsurprisingly was rusty from many years of disuse...
Plural Subject / Grammatical	The old keys unsurprisingly were rusty from many years of disuse...
Singular Subject / Ungrammatical	The old key unsurprisingly were rusty from many years of disuse...
Plural Subject / Ungrammatical	The old key unsurprisingly was rusty from many years of disuse...

In their first experiment they used SPR to determine the size and time course of the ungrammaticality effect in the baseline case, that is to say when the subject noun and the verb mismatch in number without the presence of an attractor NP in sentences such as (43). This experiment also allowed them to test whether the additional complexity of processing a plural noun would in itself have an effect on self-paced reading, as well as to contrast between the singular and plural subject-verb mismatch cases when there is no attractor noun. They found a significant main effect of number at the adverb region, where plural conditions had longer reading times, which shows that reading a plural noun does exact a significant reading time cost on the following region. Furthermore, they found a main effect of grammaticality (but not number or an interaction) at the critical verb region that extended into the following 2 or 3 spill-over regions. This shows that grammaticality effects can appear at the critical region itself during SPR and that the differential sensitivity of singular and plural subjects to attractors is not a

function of differential initial sensitivity to subject-verb agreement errors in the baseline no-attractor case.

(44) Wagers et al. (2009) : Example stimuli Experiments 2 and 3

Experiment 2: Singular subject items

Singular Attractor / Grammatical	The musician who the reviewer praises so highly will probably...
Plural Attractor / Grammatical	The musicians who the reviewer praises so highly will...
Singular Attractor / Ungrammatical	The musician who the reviewer praise so highly will...
Plural Attractor / Ungrammatical	The musicians who the reviewer praise so highly will..

Experiment 3: Plural subject items

Singular Attractor / Grammatical	The musician who the reviewers praise so highly will...
Plural Attractor / Grammatical	The musicians who the reviewers praise so highly will...
Singular Attractor / Ungrammatical	The musician who the reviewers praises so highly will...
Plural Attractor / Ungrammatical	The musicians who the reviewers praises so highly will...

Their second and third experiments were designed to test the online effects of attraction in object relative clause constructions in which the relative clause head noun is a potential attractor to the agreement between the relative clause subject and verb (44). This type of construction is interesting because the plural RC head does not intervene either linearly or hierarchically between the subject and the verb of the relative clause, which means that agreement attraction would present a challenge to theories of feature percolation. Experiment 2 tested only singular object cases where the attractor could be either singular or plural. In contrast, experiment 3 tested only plural object cases in the same environment to test for singular-plural asymmetry where strong attraction effects have been found for singular-subject/plural attractor pairs, but not for corresponding plural-subject/singular attractor pairs (Pearlmutter et al., 1999). Experiment 2 showed a main effect of attractor number, grammaticality and an interaction at the region following the critical verb. The significant grammaticality effect was found only when the relative clause head was singular with slower times for the ungrammatical condition (singular subject/plural verb). However, when the relative clause head was plural the ungrammatical singular subject/plural verb combination did not differ from the grammatical singular subject/singular verb condition. These results suggest that the head of the relative clause can act as a strong attractor for agreement and that facilitation happened when the attractor noun

matched the verb but the subject didn't. Experiment 3, yielded similar results to experiment 2 but only for singular subjects where plural attractor conditions had smaller grammaticality effects than singular attractors. But attractor effects did not appear when the subject of the relative clause was plural. In general, these experiments showed that the attractor had an effect on reading times only in ungrammatical sentences where the subject was singular and mismatched with the plural verb, and that there was no associated costs with having an attractor noun that mismatched the verb.

Given that Pearlmutter et al. (1999) reported effects of attraction on grammatical sentences, Wager et al. (2009) decided to test prepositional-modifier constructions such as the ones on (41) considering that a possible explanation was that these effects were dependent on the attractor-verb order of the PP modifier construction. Similarly to what they did on their previous experiment, they included an adverb intervening between the attractor noun and the verb as an index of the effects of noun number independently from the verb. Same as their previous experiments they found a main effect of attractor number that extended into the adverb region where the plural conditions were read slower than their corresponding singular conditions. Because these occur before the appearance of the verb, these costs are more likely associated to the additional complexity of reading plural nouns compared to singular ones. On the verb they found both a significant effect of grammaticality and an interaction on grammaticality and attractor number which was more pronounced on the following two regions. The addition of the intervening adverb showed that a clear attraction effect could only be seen on ungrammatical sentences and that significant differences of number found before the verb were caused by increased processing of plural nouns but the classic attraction effect was not found until the region after the verb. These results demonstrate a large facilitation effect for ungrammatical sentences in post-verbal regions with plural attractors without a corresponding slow-down for their grammatical analogues, counter to the predictions of a feature percolation account.

In general, the results obtained from their experiments were inconsistent with the first type of theories where attraction effects are driven by the erroneous representation of the subject's number. On object relative clause constructions an attractor noun was found to have strong attraction effects despite not intervening structurally between the subject and the verb. The same pattern was observed on prepositional modifier constructions where the attractor noun intervenes between the subject and the verb, but also that previously observed effects on grammatical sentences were more likely due to the increased processing demands of plural nouns that persists into later regions of the sentence compared to singular nouns. Their observations show that number marking on non-subject nouns had an impact on the processing of ungrammatical sentences but none on grammatical ones. In this manner, the results of their experiments taken together show that the representation of the subject number is not influenced by the mere presence of nouns with mismatching number features, but rather that the attraction effects observed at the regions past the verb reflect a retrieval mechanism that is active during the structure building process that requires to gain access to previously encoded information during comprehension which are consistent with a cue-based retrieval account in the manner in which they access the previously processed information. Self-paced reading was reliable in showing not only the effect of attraction effects on certain conditions but also the time-course of this effect. Furthermore, it allowed to observe that the increased cost of having a plural noun can extend further into later regions of the sentence and that intervening elements can help separate that effect from the classic attraction effect, and provide evidence into the underlying mechanisms involved in attraction phenomena.

Self-paced reading has also been used to study the licensing of negative polarity items. In a study by Xiang et al. (2013), self-paced reading was used together with judgment studies to compare how the individual pragmatic abilities of participants can modulate interference effects and the rise of illusory licensing in NPIs in comparison to agreement attraction in subject-verb agreement. Since errors in both dependencies seem very similar on the surface it can be

considered that they have the same source: similarity-based interference effects. However, the linguistic representations involved in NPI licensing are fundamentally distinct from those involved in subject-verb agreement and, therefore, errors from this dependency could potentially arise from different sources. Given that NPIs have been shown to be licensed at least partially by semantic and pragmatic properties of whole propositions (as discussed in Chapter 1 in addition to the retrieval of syntactic features from a previously encoded licensor, Xiang and her colleagues propose that there are two distinct routes through which NPIs are licensed. One where a previous licensor is retrieved from memory through the matching of retrieval cues and features encoded within the licensor, and a second that is driven by the application pragmatic inferences that are derived from the whole sentence. Furthermore, that in the case of the pragmatic route, differences in participants' individual pragmatic abilities can be used to modulate the rise of pragmatic driven illusory licensing.

In order to test whether individual pragmatic skills have an effect in the rise of illusory licensing, the authors compare two NPI licensors: the negative licensor *no* which is lexically encoded with an explicit [+negative] feature, and the licensor *only* which does not contain an explicit negative feature but rather licenses NPIs through the application of a pragmatic inference where "*Only NP VP*" implies "*Nobody but NP VP*" that contains the negative context that is required by the NPI. Based on the differences between both licensors, it is argued by Xiang and her colleagues that *only* should rely more heavily on the pragmatic route to license NPIs than the explicit licensor *no* since it can only license NPIs through the inference described above.

To account for each participant's individual pragmatic skills, they used the Communicative Subscale (CS) of the Autistic-Spectrum Quotient (AQ: Baron-Cohen, 2001) which has been shown to reflect speakers' pragmatic reasoning skills in language processing (Nieuwland et al., 2010; Sikos et al., 2013). The scores collected by the test represent the number of autistic traits presented by each participant, with fewer autistic traits (low CS score) representing better

pragmatic skills and more autistic traits (high CS score) representing worse pragmatic abilities in comparison. These scores were implemented into the experiment to compare whether groups with better pragmatic skills were more or less prone to interference effects during the licensing process. Then they performed a self-pace reading experiment together with a judgement task in which they compared people's sensitivity to interference effects in SVA in sentences such as the ones in (45) and NPI licensing in sentences such as (46).

- (45) a) The **receptionist** who the **boss** depends on never **fails** to do a stellar job.
b) The **receptionist** who the **bosses** depend on never **fail** to do a stellar job.
c) The **receptionist** who the **boss** depends on never **fail** to do a stellar job.
- (46) a/c) **No/Only** documentaries that the network TV stations have played during prime time have **ever** been controversial.
b/d) The documentaries that **no/only** network TV stations have played during prime time have **ever** been controversial.
e) The documentaries that the network TV stations have played during prime time have **ever** been controversial.

On the judgment studies, for the NPI cases, they found a clear interaction between interference effects and CS scores with a significant negative correlation between rating scores and participants' CS scores. Those with a high-CS score rejected both the interference and the plain unlicensed conditions, while those with a low-CS score (better pragmatic skills) were more likely to accept the interference NPI conditions. Furthermore, for the explicit negation condition *no*, illusory license was found only for participants with a low-CS score, and for the implicit negation, illusory licensing was found for those with low-CS scores and a weakened effect for those with a high-CS score ($p < 0.6$), which showed that the interference effect was indeed modulated by the pragmatic communicative skills across licensors. In the subject-verb agreement cases, they found a significant grammaticality effect and a significant interference effect where grammatical cases were judged better than both the interference and the plain ungrammatical ones, and the interference condition was judged better than the plain ungrammatical conditions. In contrast to the NPI cases, CS scores did not affect participants' judgments which shows, as expected, that

pragmatic licensing is exclusive to NPI licensing and that, therefore, the pragmatic ability of individual participants does not modulate interference effects in subject-verb agreement.

On the self-paced reading, NPI licensing was analysis was made for both licensors separately. With the explicit negative licensor *no*, at the critical NPI region the grammatical condition was read faster than both the plain unlicensed and the interference conditions, but the interference condition was also read faster than the plain unlicensed one, which suggested the presence of illusory licensing. The split group analysis between high and low CS score groups did not reveal a significant correlation between CS scores and interference effects. At the spill-over region (NPI+1) they found a significant effect of grammaticality, but no effect of interference, CS scores, or any interactions. The exploratory split-group analysis showed different trends between both groups where the high-CS group showed a significant effect of grammaticality but no interference effect, while the low-CS group showed no significant effect of grammaticality or interference.

On the implicit negation *only* cases, at the critical NPI region they found a significant effect of grammaticality where the licensed condition was read faster than the plain unlicensed condition and the interference condition with but without an interference effect or interaction. This result differed from the results from the explicit licensor at this region. In the split group analysis, the high-CS group presented a significant effect of grammaticality, but no effect of interference where the licensed condition was faster than the other two conditions, and no significant difference was found between both of them. Similar results were found for the low-CS group at the NPI region but with a numerical trend of interference that did not reach significance. At the spill-over region (NPI+1), they did not find significant effects of grammaticality or interferences, which was confirmed by mixed-effect models for both the low and high CS groups. In summary, for the licensor *only*, the self-paced reading the grand average data shows a pattern of

grammatical effect without an interference effect that holds for both the high and low CS groups, but with also a small trend of interference for the CS-low.

For the subject-verb agreement cases, the mixed effect model for self-paced reading showed a significant effect of both grammaticality and interference where the grammatical condition was read faster than both the plain unlicensed condition and the interference condition and, at the same time, the interference condition was read significantly faster than the plain unlicensed one. There was no interaction between interference and CS scores at this point. At the spill-over region (verb +1), they found a significant effect of grammaticality, but not an effect of interference. Similarly to the critical verb region, there was no difference between the CS groups on the split-group analysis.

From these results Xiang et al. (2013) present three main findings. First of all, subjects' individual pragmatic communicative skills had an effect on over NPI intervention only, but not an effect over subject-verb agreement interference. Subjects who had higher CS scores (i.e. worse general pragmatic communicative skills) were less prone to NPI interference effects, while those who had lower CS scores (i.e. better general pragmatic communicative skills) would accept the interference condition more often as grammatical. Second, the modulation of the interference effect by individuals' pragmatic skill, as indicated through the Communicative Subscale of the Autistic-Spectrum Quotient had its effect mostly on offline acceptability, but not on online reading times, though there was an observable trend of an effect in online RTs. Third, they observed a distinct profile between the NPI licensors *no* and *only*. While both licensors showed offline interference effects during the acceptability task, NPIs licensed under *only* did not show online interference. The difference between both licensors supports the perspective that multiple mechanisms are involved and acting in parallel to license NPIs. Taken together, these findings show that while NPIs and subject-verb agreement interference may look similar on the surface, they may arise from different sources. Crucially, Xiang and her colleagues propose that

the interference effect observed can be modulated by the over-application the pragmatic licensing route and use unwarranted negative inferences to license NPI, which results in illusory licensing. Furthermore, they explain that the lack of modulation by participants' communicative pragmatic skill suggests that NPI licensing may, at least partially, actually involve a syntactic matching process similar to subject-verb agreement and that the syntactic process may be applied at a different time course than the pragmatic route. This way, the effect of pragmatic interference may be visible in offline measures after participants have had more time to reflect on the sentence.

As the previous examples show, self-paced reading is a versatile and efficient method to record online responses from participants that can even be used in combination with offline methods (such as judgment studies), as well as with other online methods (such as eye-tracking). Self-paced reading has shown to be effective at measuring the onset, duration and relative time course of different effects in sentence processing, as well as the size of such effects in specific linguistic manipulations, and how this compares to other similar effects in other manipulations. In this manner, self-paced reading has been an invaluable instrument for the study of psycholinguistics and it's applications seems to be the next logical step into this investigation.

With the application of self-paced reading in this study, we expect to replicate the results of the experiments of the previous chapter in terms of the occurrence of illusory NPI licensing with the universal quantifier *every* using online methods, as well as the intervening effects of additional modification in the quantifier's restrictor clause that we observed in experiments 3 and 4.

3.3 Experimental findings

In experiment 5, if online methods are effective in showing the presence of illusory NPI licensing we expect to see a difference in reading times at the NPI region or immediately after, between sentences such as those in (47) depending on the type of licenser used. Based on the results of both *experiment 1* and *2* from Chapter 2, we expect to observe similar illusory effects where

there is no significant difference between the grammatical case with the licensor *no* and the universal quantifier *every* cases, while at the same time we should observe significant time differences between the grammatical *no* cases and the plain ungrammatical condition with *the*, where cases with *the* should be considerably slower at the NPI or in regions after. Likewise, we should observe a similar slowing down of reading times at the NPI region or immediately after with the existential quantifier, which is not able to license NPIs at all. If illusory effects were to arise with the existential quantifier, nevertheless, we should consider that there is something more general about quantifiers that licenses NPIs. If this is the case, the difference between online and offline methods could reveal a process that occurs at online stages but is quickly overcome once the sentence is finished. If no illusory effects arise with some, then we can be certain that illusory effect in NPI licensing with quantifiers is a more exclusive phenomenon related to universal quantifiers, rather than to quantifiers in general.

(47) (No/The/Every/Some) researcher has ever been invited to the annual conference that is hosted by the university.

In *experiment 6*, we use self-paced reading to replicate *experiments 3* and *4* from the previous chapter where we observed that the addition of modifiers within the restrictor clause of the universal quantifier *every* reduced the illusory licensing effect to the point where it was no longer significant in comparison to sentences without modification (Examples in Table 17, below) where the illusory effect was consistent with previous results. If modification is able to reduce the significant effect of illusory licensing in online methods, we expect to see comparable differences between the quantifier conditions in comparison to the grammatical negative condition in similar fashion to the results of *experiments 3* and *4*. Consequently, reading times should slow down significantly at the NPI region or regions following the NPI for the modified conditions relative to the negative condition, while the unmodified condition should not present significant differences in reading time speeds compared to the negative quantifier condition. Conversely, if the illusory effect remains undetected at these regions and there is no significant difference in

reading times between the modified conditions and the grammatical condition, we can conclude that the more clear identification of the boundaries of the restrictor clause thanks to increased semantic information, which we proposed as the reason why the illusory effect was reduced in chapter 2, is not immediately available at this point during online processing and that semantic and pragmatic information is probably integrated into the licensing process at a later time. Moreover, if we observe a difference between the modified conditions where one presents significant illusory effects and the other does not, then we can draw conclusions based on the time course of the differences.

3.3.1 Experiment 5: Evidence of illusions in NPI licensing with online measures

The aim of Experiment 5 was to replicate the results obtained in experiments 1 and 2 from the previous chapter using online measures instead of early offline measures to better understand the time course of illusory licensing effects in NPIs. For this purpose, we designed a self-paced reading task that included both the universal quantifier *every*, which was shown to produce illusory licensing in speeded acceptability measures (Exp. 1), and the existential quantifier *some* which didn't give rise to illusory licensing (Exp. 2).

While a speeded acceptability task like the one used on the previous chapter can shed light into the appearance of illusory licensing and its prevalence in time after the sentence processing has been finished, it does not tell us when the illusion appears during the sentence. While offline judgments can be useful, illusory licensing effects can be very short lived and participants, with enough time to think about the sentences, can ultimately overcome an effect that may have appeared during the sentence parsing. Online measures, on the other hand can provide us with information about effects that may rise before completion of the sentence and that may be transitory even before the completion of the sentence.

Despite having failed to find illusions of grammaticality with the existential quantifier in *experiment 2*, the reason to include *some* along with *every* in this experiment is to observe

whether any effects appear within the sentence that may be too short-lived to be observed in offline measures. Nevertheless, we expect to find a similar pattern to both experiments where the universal quantifier *every* produces illusions consistently, while the existential quantifier *some*, which does not license NPIs, does not.

3.3.1.1 Participants

We recruited 72 participants through Prolific (<https://www.prolific.co/>). All participants were informed about the main task of the experiment in which they would have to read through sentences one word at a time and answer comprehension questions about those sentences and asked to provide consent to participating in the experiment. Participants were asked about their linguistic background through a questionnaire included at the beginning of the experiment to ensure that they were native speakers of English. Each participant was paid £6.00 for the completion of the experiment which lasted approximately 30-45 minutes. 8 participants were excluded due to having a different native language.

3.3.1.2 Materials

Table 14: Sample stimuli for experiment 5

Subject Type	Sentence
Determiner	The researcher has ever been invited to the annual conference that is hosted by the university.
Existential	Some researcher has ever been invited to the annual conference that is hosted by the university.
Negation	No researcher has ever been invited to the annual conference that is hosted by the university.
Universal	Every researcher has ever been invited to the annual conference that is hosted by the university.

For this experiment we designed a set of 40 sentences with 4 conditions based on Subject Type depending on the word with which the sentence started and that served as a potential licenser for the NPI. In this manner, the conditions included one licensed condition with the negative quantifier *No*, and three unlicensed ones which included an unambiguous unlicensed condition *The*, and the two quantifier ones that we were testing *Every* and *Some*. Each sentence was formed of a noun phrase and a main clause predicate, similar to how sentences were constructed in the previous chapter. Each subject was formed of a determiner or a quantifier based on

Subject Type followed by the main noun. The predicate was formed by the auxiliary *has* which served to unambiguously indicate the start of the predicate, immediately followed by the NPI *ever*. For this experiment, in comparison to the previous ones, we did not include NO-NPI conditions therefore for every sentence the NPI was located at the 4th word in the string. The NPI was then followed by the word *been* and a main verb plus two additional phrases. All experimental items had between 13 to 18 words in total in order to ensure that spill-over regions after the NPI could be observed and to avoid end of sentence effects from affecting the reaction times near the NPI. A sample of items can be observed in Table 14.

Each participant saw a total of 130 sentences which comprised of 40 NPI experimental items and 90 fillers out of which 32 were part of a second experiment involving co-predication, another 40 were part of a third experiment involving garden path sentences and subject-verb agreement, and 18 were built as pure fillers to mask the NPI conditions.

3.3.1.3 Procedure

The experiment was designed using a self-paced, word-by-word, moving window serial visualisation paradigm where participants were originally shown a series of dashes which represented the words in the sentence. Upon a press of space bar, the first word appeared replacing the dashes corresponding to the first word and the program started recording the time between presses of the keyboard. Further presses of the space bar revealed the following word while masking the previously shown word. This way, a participant was only able ever able to see one word of the sentence at a time and was not be able to review previous content. Upon the end of each sentence, participants were presented with a comprehension question related to the sentence they had just read and respond to it with either ‘yes’ or ‘no’ by pressing the corresponding key on the keyboard (‘j’ for “yes”, and ‘f’ for “no”). Comprehension questions were used to make sure that participants were performing the task correctly and comprehending the sentences and not rushing through them. The experiment was coded in Ixweb (Alex Drummond, <http://spellout.net/ibexfarm/>). Before the experimental items were presented,

participants were given a set of practice items to train them on the task. A total of six practice sentences were used with sentence structures that did not fully resemble the experimental items to not give away the experiment's purpose.

3.3.1.4 *Data Analysis*

The data was analysed using the lme4 package (Bates, Maechler, & Bolker, 2011) for R Software (R Core Team, 2014) using linear mixed-model effects with fixed factors for experimental manipulations and a specified random effect structure, including random intercepts and slopes for all fixed effects by participant and by items (Baayen, Davidson, & Bates, 2008; Barr et al., 2013). In this experiment, we looked at whether the presence of an intervening universal quantifier ('every') or an intervening existential quantifier ('some') could elicit differences in reaction times compared to the definite article ('the') and the negative quantifier ('no') when located as potential licensors for an NPI, and give rise to illusions of grammaticality in online measures. Conditions for the variable Licensor were Helmert coded (Vasishth & Broe, 2011; Venables & Ripley, 1999) into three contrasts. "Grammaticality" was coded by comparing the negative quantifier to all other three ungrammatical conditions ('no' = 1, 'every' = -1/3, 'some' = -1/3, 'the' = -1/3). The other two contrasts compared each quantifier against the determiner condition. Determiner vs Universal ('every' = 1, 'the' = -1), and Determiner vs Existential ('some' = 1, 'the' = -1). Data was modelled across 7 regions from the Licensor up to three words after the NPI for plotting purposes. The statistical analysis below is reported for the NPI and spill-over region. Effects were considered to be significant if the p values were $p \leq 0.05$ (Gelman & Hill, 2007). Extreme RTs were removed by using model criticism (Baayen et al., 2008). A total of 0.17% data was lost during clean up.

3.3.1.5 *Results*

The results for experiment 5 can be observed in Figure 10 (SPR Graph), and tables 15 for the NPI region, and table 16 for the first spill over region.




Figure 10: Experiment 5 word-by-word reaction times (msec) with standard error.

Region 4: NPI region

Table 15: Region 4: NPI Region – Linear Mixed Effects

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	370.900	13.423	63	27.632	< 0.001 ***
Grammaticality: Neg vs (Uni + Exist + Det)	-11.675	5.204	2491	-2.243	0.025 *
Det vs Existential	11.466	4.906	2491	2.337	0.020 *
Det vs Universal	-10.214	4.904	2491	-2.083	0.037 *

The linear mixed effects analysis at the NPI region already shows significant differences across all conditions. The universal quantifier shows clear illusion effect ($t = -2.083$, $p = 0.037$) as early as the appearance of the NPI. While in the previous regions it shows no difference with the determiner, there is a clear acceleration at the NPI that is greater than the determiner condition. On the other hand, the existential quantifier shows a significant effect in the opposite direction ($t = 2.337$, $p = 0.020$), which shows that the existential quantifier was processed slower than the determiner and the universal quantifier at the NPI. We also found a significant effect of grammaticality ($t = -11.675$, $p < 0.001$). While this could show that participants were sensitive to grammaticality effects as early as the appearance of the NPI, the effect was present from the

previous region (Region 3 – Neg vs Determiner + Universal + Existential: $t = -2.005$, 0.045), probably driven by the spike in RTs from the existential quantifier at that stage.

Region 5: Spill-over region

Table 16: Region 5: Spill-over Region – Linear Mixed Effects

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	5.820	0.033	71	174.208	< 0.001 ***
Grammaticality: Neg vs (Uni + Exist + Det)	0.046	0.019	2800	2.394	0.017 *
Det vs Existential	0.041	0.019	2800	2.126	0.034 *
Det vs Universal	0.009	0.019	2800	0.464	0.642

At the spill-over region we found a significant effect of grammaticality ($t = -4.495$; $p < 0.001$). The difference between the existential quantifier and the determiner is no longer present at this region ($t = 0.380$, $p = 0.704$). This reduction in RTs could perhaps indicate that intrinsic difficulties specific to the existential quantifier at this stage are being overcome and that ungrammaticality of this condition is now similar to that of the determiner by itself. More importantly, we do find that the significant difference between the determiner and the universal quantifier persists at this region ($t = -2.302$, $p = 0.021$). Despite the universal quantifier being an inadequate licenser to the NPI like the other two unlicensed conditions, the reaction times seem to remain the same as in the previous region. This difference can be evidence of an illusion effect taking place at this region.

3.3.1.6 *Discussion*

Experiment 5 successfully replicates the results from experiments 1 and 2 with online measures. The results show that at least some measure of illusory NPI licensing occurs with the universal quantifier *every* and that it can be observed with online measures, same as was observed in judgment studies in experiment 1. The experiment also shows that illusory licensing cannot be observed in online measures with the existential quantifier *some*, similarly to what we observed with judgement studies in experiment 2.

Taking a better look at the course of the illusion effects, we find that at the NPI region there is a significant difference between the universal quantifier and the determiner, with the universal quantifier showing faster processing times compared to the determiner. Similarly, we see that there is a significant difference between the existential quantifier and the determiner, but on the opposite direction, with the existential quantifier having increased reaction times compared to the determiner. The significant grammaticality effect could show that at this point during the processing participants are already showing an established grammaticality effect between the negative condition and the other conditions, though it is possible that the effect is driven by the inflated reaction times of the existential quantifier. At the spill-over region we see that the difference between the universal quantifier and the determiner is maintained, while the difference between the existential quantifier and the determiner has been reduced and are no longer statistically different.

In terms of the existential quantifier, these results can help us discard the possibility of NPI illusory licensing since we failed to see any evidence of facilitatory effects happening at the NPI. Quite the contrary, it seems that the presence of the existential quantifier already increases the demands of processing at the NPI region. This is similar to what we found in chapter two, where we found that the existential quantifier had increased processing demands even in sentences without an NPI. This increase in processing seems to peak at the auxiliary verb region, though it is unclear whether this is a spill-over effect initiated at the noun or produced at the auxiliary itself, but even though it shows some recovery afterwards, the effect of inflated reaction times continues over to the NPI region.

In terms of the universal quantifier, the results of this experiment show that, compared to the determiner, the universal quantifier can produce some degree of facilitation at the NPI and spill-over region, which provide support for the hypothesis that at least some measure of NPI illusory licensing can occur with the universal quantifier 'every'. Potentially, the effects from having

'every' are enough to last until the end of the sentences, which could explain why we see some evidence of illusory licensing in offline measures. After the NPI and spill-over region we do see an increase in RTs that follows a similar pattern to the determiner, which means that increased processing times with the universal quantifier are still prevalent across most sentences and that the rise of the illusion, while significant, is not a consistent effect. While it is possible that some effects could be observed later on the sentence, the current experiment was not designed to test for late sentence effects.

Based on the time course data this experiment shows that the illusions that arise with the universal quantifier can start as early as the appearance of the NPI, while experiment 1 shows that these effects can last until after the sentence has finished processing, and appear in participants offline judgments. The current experiment was not only able to replicate the results found in experiment 1 regarding the rise of illusions, but also shows that these effects are visible with an online method such as Self-paced reading, and that illusory licensing can be identified precisely within the time course of the licensing process.

3.3.2 Experiment 6

The aim of experiment 6 was to observe whether the calculations of the restriction and scope of the universal quantifier *every* that we observed in experiments 3 and 4 from the previous chapter could be observed as well with online measures. And whether the addition of prenominal and postnominal material to the restrictor of the quantifier could have the same effect observed in offline measures of reducing the illusory licensing effect as was observed before. Furthermore, in the case of finding similar effects, whether we can identify the modulation of the effect within the time course of online sentence processing.

In offline measures we observed that there were no significant illusory effects in experiments 3 and 4, which added each a prenominal and a postnominal modifier respectively. Based on the evidence of illusory licensing found in experiment 1 with the unmodified universal quantifier, we

assumed that the addition of both prenominal material (adjective) and postnominal material (RC) was enough for illusions to be reduced and no longer appear to be significant and that this could be due to increased information that allows the parser to easily identify the restrictor clause of the quantifier and determine that the NPI was located outside of it.

The use of online methods, in this case, allows us to observe the time course of the sentence to see whether the calculations of restriction identification are applied immediately during processing at the NPI region or at a later time, if at all applied online during the licensing process. If the illusory licensing disappears in modified conditions compared to the unmodified condition during SPR at the NPI or spill-over regions we can conclude that restriction identification is readily available at the time the NPI is reached and is immediately applied during processing. On the contrary, if the illusion persists, we can conclude that this information is considered after the illusion has already risen and that its effect of resolving illusory licensing is applied a posteriori.

3.3.2.1 Participants

We recruited 72 native speakers of English through Prolific (<https://www.prolific.com/>). Each participant was remunerated with £6.00 for an experiment that lasted approximately 30-45 minutes. Same as before, we explained the main task of the experiment and collected data from their linguistic background through a survey included at the beginning of the experiment.

3.3.2.2 *Materials*

Table 17: Sample stimuli for experiment 6

Subject Type	Sentence
Determiner	The researcher has ever been invited to the annual conference that is hosted by the university.
Negative	No researcher has ever been invited to the annual conference that is hosted by the university.
Unmodified UQ	Every researcher has ever been invited to the annual conference that is hosted by the university
UQ + Adjective	Every young researcher has ever been invited to the annual conference that is hosted by the university.
UQ + RC	Every researcher who was awarded a scientific prize has ever been invited to the annual conference that is hosted by the university.

For this experiment we used the same materials as in experiment 5 without including the existential quantifier condition. Instead, we added two additional conditions that used both the universal quantifier as potential licenser for the NPI. These two additional conditions differed from the main universal quantifier one in that they included either an adjective as a prenominal modifier right before the head noun, and the second one added a postnominal modifier in the form of a relative clause immediately after the head noun. This way, the materials included five conditions based on potential licensors: Determiner *the*, Negative *no*, Quantifier (unmodified) *every*, Quantifier + Adj *every + Adj + N*, Quantifier + RC *every + N + RC*. Furthermore, same as in the previous experiment, we included a comprehension question to ensure that participants performed the task as intended and comprehended the sentences, as well as prevent rushing through the items. Example items are shown in Table (17).

3.3.2.3 *Procedure*

The experiment followed the same procedure as in experiment 5 above. We used a self-paced, word-by-word, moving window serial visualisation paradigm where participants were originally shown a series of dashes that corresponded to the words in the sentence. Upon the press of the space bar, the first word appeared replacing the dashes corresponding to the letters of the first word. Further presses of space bar revealed following words while masking each previous word. The program recorded the time between presses of the keyboard, measuring reaction times. At the end of each sentence participants were presented with a comprehension question related

to the sentence they had just read and asked to respond either 'yes' or 'no' by pressing a corresponding key on the keyboard. The experiment was coded in Ibox Farm (Alex Drummond, <http://spellout.net/ibexfarm/>).

3.3.2.4 *Data Analysis*

Data was analysed using R software with the same methods described in experiment 5. Conditions for the variable Licensor were Helmert coded (Vasishth & Broe, 2011; Venables & Ripley, 1999) into four contrasts. First, a grammaticality condition was by comparing the negative quantifier to all unlicensed conditions ('no' = 1, 'the' = -0.25, 'every' = -0.25, 'every + adjective' = -0.25, 'every + relative clause' = -0.25). The other contrasts compared each universal quantifier condition with the determiner: unmodified universal vs determiner ('every' = 1, 'the' = -1), universal + adjective vs determiner ('every + adjective' = 1, 'the' = -1), and universal + relative clause ('every + RC' = 1, 'the' = -1). Data was modelled across 5 regions, from a word before the NPI up to three words after the NPI. Two participants were removed due to not being a native speaker of English. A total of 0.84% data was lost during RT clean up.

3.3.2.5 *Results*

The results for experiment 6 can be observed in Figure (11) and Tables 18 (NPI region), 19 (Spill-over region), and table 20 (spill-over + 1 region).




Figure 11: Experiment 6 word-by-word reaction times (msec) with standard error.

Region 0: NPI

Table 18: Region 0: NPI Region – Linear Mixed Effects

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	358.464	12.403	68.982	28.902	< 0.001 ***
Grammaticality	-1.979	6.955	76.578	-0.285	0.777
Det vs Universal	-4.126	11.443	184.755	-0.361	0.719
Det vs Universal + Adjective	10.516	10.943	243.716	0.961	0.338
Det vs Universal + RC	1.949	13.582	90.733	0.144	0.886

Planned comparisons for the NPI region showed different results than the ones found on experiment 5. At this point, there were no significant differences between the determiner and the quantifier conditions: determiner vs unmodified universal ($t = -0.361$, $p = 0.719$), determiner vs universal + adjective ($t = 0.961$, $p = 0.338$), and determiner vs universal + relative clause ($t = 0.144$, $p = 0.886$). We also failed to find a significant effect of grammaticality ($t = -0.285$, $p = 0.777$). These results show that at this point participants have not yet reacted to the ungrammaticality of the NPI with each ungrammatical licensor condition. While this could show that universal conditions are producing illusions, the fact that there is also no difference with the

determiner condition shows that it is probably too soon for participants to react or slow down to the presence of the NPI.

Region 1: spill-over

Table 19: Region 1: Spill-over (NPI +1)

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	348.991	11.575	69.003	30.150	< 0.001 ***
Grammaticality	-19.032	4.902	256.616	-3.882	< 0.001 ***
Det vs Universal	-6.397	9.324	296.701	-0.686	0.493
Det vs Universal + Adjective	0.806	9.121	511.927	0.088	0.930
Det vs Universal + RC	-12.778	10.608	83.994	-1.205	0.232

At the spill-over region (region 1), in contrast with the NPI region, we already have a significant effect of grammaticality observed by the difference between the negative quantifier and all unlicensed conditions ($t = -3.882$, $p < 0.001$), this result shows that after the appearance of the NPI, all unlicensed conditions present a slowdown in comparison to the negative condition, which, on the contrary speeds up. In this experiment, compared to what we observed in experiment 5, though, we were not able to observe any significant difference between the determiner and any of the universal quantifier conditions. This lack of difference did not only extended between the determiner vs universal quantifier + adjective ($t = 0.088$, $p = 0.930$), and the determiner vs universal + relative clause ($t = 1.205$, $p = 0.232$), but also to the difference between the determiner and the unmodified universal quantifier ($t = -0.686$, $p = 0.493$) which had shown a significant difference in experiment 5 in favour of the universal quantifier.

Nevertheless, we do see a trend with the universal quantifier conditions to speed up between the appearance of the NPI and this region in a more similar fashion to the behaviour of the negative quantifier, while the determiner condition presents a slowdown compared to the previous region. While the difference is not enough to be significant, because all three universal quantifier conditions present this pattern at this stage, it is possible that the universal quantifier still has some positive effect in NPI, licensing, even if it is not consistent enough to present a

statistical difference, or that participants are prone to more easily and quickly identify the ungrammaticality of the determiner immediately after the appearance of the NPI.

Region 2: spill-over +1 (Main Verb)

Table 20: Region 2: Spill-over +1 (NPI+2) / Main verb

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	372.638	15.028	75.086	24.796	< 0.001 ***
Grammaticality	-20.205	6.442	3375.814	-3.136	0.002 **
Det vs Every	0.155	12.489	3376.604	0.012	0.990
Det vs Every + Adjective	1.350	12.459	3372.980	0.108	0.914
Det vs Every + RC	-17.925	12.465	3377.240	-1.438	0.151

In region 2, two words after the appearance of the NPI, which corresponds to the main verb of the sentence, we observe a general increase in reaction times for all conditions, probably due to the effect of processing the main verb. We still observe a significant main effect of grammaticality ($t = -3.136$, $p = 0.002$), but are still not able to observe any statistical differences in the reaction times between the determiner and the universal quantifier conditions: determiner vs unmodified universal ($t = 0.012$, $p = 0.990$); determiner vs universal + adjective ($t = 0.108$, $p = 0.914$); and determiner vs universal + relative clause ($t = -1.438$, $p = 0.151$). While at this region we still see the universal conditions with some intermediate reaction times between the negative and the determiner, the slight reduction of RTs for the universal conditions is mostly disappearing and are mostly gone by the next region, except for the relative clause conditions which seems to still be slightly faster than the other ones. Although close, it did not manage to reach a significant difference in NPI+3: determiner vs universal + RC ($t = -1.891$, $p = 0.060$).

3.3.2.6 *Discussion*

Experiment 6 managed to reproduce the results found on its analogous offline experiments (experiments 3 and 4 in Chapter 2), but not the ones found on experiment 1. While the offline results show a difference between the illusory licensing of the universal quantifier conditions based on whether modification (prenominal or postnominal) was introduced or not, experiment

6 did not demonstrate the same distinction. We saw no evidence of illusory licensing from any of the universal quantifier conditions regardless of modification in online measures for regions that contained NPI or immediately after. This means that experiment 6 also failed to replicate results found on experiment 5 where an unmodified quantifier did demonstrate a significant effect on reaction times at the NPI and the following region relative to the determiner. We did observe a pattern on the region following the NPI where the reactions times for all quantifier conditions sped up in contrast to the determiner, which slowed down in relation to the previous region, though this effect was too small to produce any statistical differences.

If restrictor modification can help overcome the illusory effect by clarifying the boundaries of the universal quantifier's restrictor clause, the current experiment was not able to reflect the difference between modified and unmodified universal quantifiers, at least not immediately after the appearance of the NPI, and as such, we were not able to determine whether increased semantic information to the restrictor clause was produced any differences in the appearance of illusory licensing, nor the timing of the integration of this type of information. In contrast to the results of the offline experiments, it seems that the effects of the semantic and pragmatic information that is provided by the increase in modification is not observable in online measures such as self-paced reading, at least not immediately. It is possible though that the integration of this information happens at a later stage, or that it's inclusion at this point is not immediately observable in these kind of structures.

Another thing to consider going forward is that the lack of an effect in experiment 6 could also be an issue of experimental design. While we included conditions that contained a relative clause and an adjective with the universal quantifier, we did not consistently include experimental sentences that contained the negative quantifier or the determiner with such modifications. If this is the case, it is possible that some participants could have been able develop some expectations about the sentences they were reading based on the initial word they saw and this

could have had an effect in expectations about an upcoming NPI. Future studies should consider these observations during item design to prevent, or discard this possibility.

3.4 General Discussion

The results observed in this chapter show that while illusory NPI licensing with the universal quantifier can be observed in online measures as evidenced by experiment 5, that effect does not seem to rise with enough consistency to be statistically significant in all cases, as was observed in experiment 6, at least at the NPI and immediately after regions. The results follow a pattern in terms of the different quantifiers that we included in these studies: the negative quantifier *no* which normally licenses NPIs in either its restrictor clause or its scope has faster reaction times in comparison to either the universal quantifier and the existential quantifier. The existential quantifier *some*, similar to the determiner *the*, which does not license NPIs either within its restrictor clause or its scope is very quickly perceived as ungrammatical and, therefore, shows an increase in reaction times relative to the grammatical condition. It is of note, though, that the increase in reaction times with the existential quantifier seems to begin shortly after its initial appearance and starts to accelerate before the appearance of the NPI. Finally, and most importantly, the universal quantifier *every*, in general shows a pattern of acceleration after the appearance of the NPI in comparison to the determiner in all cases, whether modified or not. Nonetheless, this pattern was only significant in one of the experiments and not on the other. Based on the current observations, we can determine that the time course and location of illusory NPI licensing phenomena can be observed for quantifiers with the use of online measures such as self-paced reading and that they can occur as early as the appearance of the NPI, but also that this type of illusory licensing does not occur with much reliability.

Given the results of experiment 5, we know where we should be finding the early signs of illusory licensing and how fast they are overcome on the sentence. Though some facilitatory effects can remain for the next couple of words after the appearance of the NPI, producing some

intermediate reaction times between the a grammatical and the ungrammatical conditions, experiment 6 showed that participants are able to reliably identify ungrammatical NPIs with the universal quantifier more often than not. The lack of a difference between modified and unmodified universal quantifier conditions at this stage make it difficult to determine whether modification of the restrictor clause has an immediate effect in how structurally irrelevant NPIs are perceived based on how much information is incorporated into the restrictor clause, or that, at the very least, the effect cannot be observed in an online measure such as self-paced reading at this particular stages. If modification of the restrictor clause has any effect on the occurrence of illusory NPI licensing, that information must be processed and incorporated at a later stage of the sentence, or after the whole sentence has been processed.

In experiment 3, the addition of a prenominal modifier in the form of an adjective was enough to reduce the appearance of illusory licensing to the point where it was no longer significant in comparison to the unmodified case. Similarly, in experiment 4 the addition of postnominal modification in the form of a relative clause was also enough to reduce the rise of illusory licensing to the point where it was no longer significant compared to the unmodified universal licensor case. Similar to the results found in Xiang et al. (2013), the semantic/pragmatic information does not seem to reflect a change in the illusory licensing on NPIs in online measures, even though we were able to observe an effect in offline measures. It seems that during online processing, not all information is readily available to be integrated into the licensing process, at least at this point. If our hypothesis that increased modification can help alleviate illusory licensing by determining the current phrasal context (in this case, defining the boundaries of the restrictor clause) to better mark the location of the NPI has any effect during sentence processing, these results point that the incorporation of the additional semantic and pragmatic information provided must be happening at a later stage than when the illusions typically first appear. It is still possible that the integration of this information is done at a later time during online processing, but this was not observable in this experiment.

At this point, we have two different experiments (experiment 1 and 5) that have shown that illusory licensing with the universal quantifier does occur relative to a less ambiguous ungrammatical licenser such as *the*, and three that have failed to show the same results (experiments 3, 4 and 6). While in chapter 2 we assumed that the lack of illusory licensing in experiments 3 and 4 were a direct consequence of increased information inside the restrictor clause compared to experiment 1 where the unmodified universal quantifier condition yielded significant results, these were isolated experiments and we didn't have a direct comparison across all universal conditions. Experiment 6, on the other hand, provided a direct comparison between all these conditions and showed no difference, at least at early stages of the NPI's appearance. If we had seen some difference between the unmodified condition and the modified ones, we could conclude that the results are reliably a consequence of this modification, but as it is, we are unable to provide evidence of this.

Experiment 6 leaves us with some more questions, at least regarding the effect that was observed for experiments 1 and 5 on the unmodified universal quantifier's restriction clause. While both experiments show an illusory effect on NPI licensing with the universal quantifier *every*, experiment 6 contradicts these observations. This could indicate that universal quantifier illusions are not too common and may appear only sporadically, and that in experiment 6 these were not consistent enough to produce visible statistical effects. The pattern observed in the spill-over region that shows a trend for a slight acceleration in reaction times for all universal quantifier conditions could support this hypothesis. While not enough illusions were raised for a significant difference at that specific time, there must have been at least some portion of cases in which the ungrammaticality of the NPI was either temporarily overlooked, or not immediately identified by some participants as soon as the NPI was seen, or that at least, the ungrammaticality of the determiner sentences was more easily identified than those that contained the universal quantifier. The reason why we can entertain looking at this trend is because it can be observed for all conditions that contained the universal quantifier in a very

similar fashion regardless of modification and, furthermore, because experiment 5 shows a similar pattern between the universal and the determiner immediately after the appearance of the NPI where the determiner reaction times increase while the universal quantifier does not. Of course, this pattern is not enough to provide necessary evidence and these observations should be taken with a grain of salt since the existential quantifier condition also slopes down after the appearance of the NPI. Given the extreme increase in reaction times observed for the existential quantifier before the appearance of the NPI, it is not evident whether this continued acceleration is a consequence of the NPI itself or a general acceleration process initiated earlier (though the results observed in experiment 2 could point at it not being a direct consequence of the NPI itself).

Considering that the restrictor clause can be further extended by modification, and even be, in some cases, translocated from the sentence, it is sensible to think that the integration of this information could happen at a later stage during the parsing, or after the sentence has been completed. If this is the case, determining the actual time moment in which this information is integrated may require further and better time-grained experimental methods. Perhaps one of the issues in this study is the type of sentences that we used since they do not allow us to observe how the appearance of an NPI in a specific configuration can affect the following context. We believe that following studies into the incorporation of semantic and pragmatic information of online sentence processing using NPIs should consider structures where the appearance of the NPI can influence how the rest of the sentence can be interpreted, with for example, garden path sentences where the location of an NPI can inform of the type of structure (i.e. a reduced relative clause) in which it is located. We will examine this possibility in the following chapter.

The difference found between the different kinds of licensors, namely between the universal quantifier and the existential quantifier, reinforces the conclusion from chapter 2 that illusory licensing is exclusive to the universal quantifiers and not a general property of quantifiers. At

least, in the measures used, there does not seem to be any reason to believe that quantification on itself carries any feature desirable for the licensing on NPIs and, consequently, we do not find any indication that [+quantification] is a possible retrievable feature used during online memory retrieval processes, except perhaps if we consider the small reduction in reaction times for the existential quantifier at the NPI respect to the increase observed for the determiner, but even so, this consideration seems to be very unlikely. An interesting finding of the study, is that there seems to be an increased difficulty in the processing of the existential quantifier relative to the other quantifiers which is specially visible immediately after the noun region of the subject before even the NPI is found. The increased difficulty is also reflected in experiment 2 in the previous chapter, where we saw evidence that the existential quantifier could have increased demands even when an NPI was not involved in the sentence, since the biggest difference in judgments were found in the No-NPI conditions. If the Existential Quantifier already increases the demands on the parser it could explain why the slowdown effect is so readily available for *some* in self-paced reading at the NPI conditions. This doesn't seem to be an effect only of grammaticality since it is not observed with the determiner condition which is equally ungrammatical given that neither of them license NPIs. Furthermore, this difference could also point to differences in the way in which quantifiers are processed in comparison to the determiner. Given the differences found between the different quantifiers at the NPI region, we can assume that the illusion effect is discarded very early with the existential quantifier in comparison to the other conditions. Perhaps quantifiers are initially processed 'more deeply' than the determiner and are therefore expected to show effects earlier in processing.

3.5 Conclusions

On this chapter we saw that illusory NPI licensing is observable in online measures for the universal quantifier, and that this is not a phenomenon exclusive to observations made in offline judgements studies, as evidenced by experiment 5. While online measures can provide us with some information about the location and the time course of the effect, in the current

observations we noticed the effect to be weak and inconsistent across experiments. Furthermore, the difference between the universal quantifier *every* and the existential quantifier *some* in online measures provides further evidence that the rise of illusions is not due to general properties of quantifiers, but that it is exclusive to the use universal quantifiers. Though evidence of illusory licensing with the universal quantifier was not consistent across all experiments, a pattern could be observed that could provide some indication that NPIs with structurally irrelevant universal quantifiers can sometimes be processed slightly faster than with a less ambiguous ungrammatical licenser. These results would require further study or replication in order to assess whether the results we've found constitute a type I error in experiment 5, or a type II error in the case of experiment 6. Additionally, the differences observed in offline measures between different levels of modification in conditions with the universal quantifier could not be observed in online measures in the regions immediately after the NPI since we failed to find a difference between the universal conditions, or any differences between them and the determiner condition. Further observation of semantic and pragmatic information in the online processing of NPIs would require more complex sentences, such as garden paths, where the processing of subsequent structure could be affected by the appearance of an NPI.

Chapter 4 - Identifying Restrictors Under Ambiguity

4.1 Introduction

In this chapter we will explore whether NPI licensing can affect how the sentence structures are formed and processed. Namely, we are interested in seeing whether the apparition of an NPI can provide contextual cues about the structure that holds it by potentially facilitating the determination of the sentence's monotonicity (i.e. its upward-entailingness or downward-entailingness). In order for the NPI to be licensed, it is required to be within an appropriate downward entailing (DE) context. If it appears outside of this environment, the NPI should be considered to be unlicensed. Yet, as we have seen in the previous chapters, illusions of grammaticality can arise when the NPI is outside the desired context and these illusions tend to persist through the online processing of the sentence. While we know that illusory licensing happens due to the proximity and availability of a previously encoded potential licenser (such as *every*) in both offline and online measures, and that modification of the licensing context (such as the presence of a universal quantifier's restrictor) can lead participant's appreciation of this illusion (in offline measures), we do not know for certain that the parser is actually integrating the semantic and pragmatic information during the NPI's licensing. We believe that it is possible to assess this based on whether the apparition of the NPI influences the way in which the sentential context is perceived and that NPIs, at least partially, reveal clues about their structural location by allowing the parser to interpret or identify the semantic context as downward-entailing; or at least consider that the NPI may be located within a DE-context. If this is true, then it is possible that NPI illusory licensing is at least partially driven by semantic and pragmatic assumptions related to its environment (as DE-contexts) that are processed at the time that the

NPI is encountered, and not only due to the interference of a nearby structurally irrelevant licenser.

For this purpose, we look at garden path sentences to see if the correct identification of a DE contexts by the NPI, within the restrictor of the universal quantifier *every*, allows the parser to correctly identify the structure in which the NPI is located and, therefore, chose the correct reduced relative clause structure (including the corresponding past participle rather than simple past verb form) to avoid being driven down the garden path that choosing the incorrect structural analysis would produce.

4.2 Background

4.2.1 Garden Paths Sentences

Garden path sentences refer to a type of sentences that initially seem to have a straightforward and simple structure but, owing to their ambiguity, suddenly turn out to be more complex than originally assumed. This causes comprehenders who read or listened to such sentences to have difficulties understanding them. In terms of language comprehension, this sudden change in the interpretation of the sentence requires the parser to reanalyse the initial structure to make sense of it. The name comes from the idea that the sentence leads comprehenders down a clear path of analysis that suddenly dead ends and requires them to take a new unexpected direction. This type of sentences has been very important for psycholinguistic research to explore the cognitive processes involved in sentence parsing and how people comprehend language, especially in terms of ambiguity resolution. Consider the classic garden path sentence (Bever, 1970) in (48).

(48) The horse raced past the barn fell.

The sentence above is initially parsed as a main clause where the verb *raced*, which is morphologically ambiguous, is interpreted in the simple past tense, rather than in the past participle. Once the parser reaches the verb *fell*, this initial interpretation can no longer be sustained since the verb *fell* requires a subject and the main clause interpretation of the sentence

doesn't allow this to happen. Instead, in order to make sense of the sentence, the parser is forced to reanalyse the structure and reinterpret *raced* as a subordinate past participle verb of a reduced relative clause, rather than as a simple past main verb, allowing the whole string to be interpreted as the subject, rather than a sentence on its own. Sentences like these are shown to be difficult due to an overwhelming preference to analyse *raced* as a main verb (Townsend & Bever, 2001). Psycholinguistics has taken a great interest in such sentences because the existence of garden path effects, as well as other anomalies, is evidence that the parser makes rapid decisions during processing about which alternative, among different grammatical representations, must be adopted when the linguistic input is ambiguous (Ferreira & Clifton, 1986; Frazier & Rayner, 1982; MacDonald & Christiansen (2002); Traxler, 2019). There are four main models that explain the difficulty presented by garden path sentences: the Garden path model, the constraint satisfaction model, the referential model, and more recently, the Race-Based Model.

4.2.2 Models of ambiguity resolution

4.2.2.1 *The garden path model*

The garden path model is a syntax-first model that considers that the parser is a serial processing device and that the initial analysis of ambiguity is based solely on structural properties of the linguistic input (Clifton & Ferreira, 1989; Ferreira & Clifton, 1986; Frazier, 1979; Frazier & Rayner, 1982; Rayner et al, 1983). This model proposes that during sentence processing, the parser initially chooses the most structurally simple analysis with the fewest non-terminal nodes through a mechanism called Minimal Attachment. Because of this, the parser prefers to interpret the verb *raced* as a main verb since the alternative form, the past participle, that requires to be within a relative clause, implies an additional non-terminal node. This leads to the incorrect initial interpretation and, therefore, the main verb *fell* cannot be readily incorporated into the analysis.

However, the model suggests that information available at subsequent processing stages, such as a thematic processor (Ferreira and Clifton, 1986), can assist with reanalysis when the initial interpretation is incompatible with the rest of the sentence. This second stage processor supplies alternative thematic representation that can help the parser reanalyse the structure. Similarly, Mitchel, Corley, and Graham (1992) suggest that the parser may start to revise an unresolved ambiguity even before disambiguation. According to them, the effects of minimal attachment persists only briefly. This allows the parser to begin processing semantic and pragmatic information that may later be used for reanalysis at a later stage if needed. Nevertheless, these non-structural sources of information can only confirm or disconfirm the parser's initial decision and are not integrated at the same time as structural information. Other research has focused on the costs of reanalysis and diagnostic tools for recovering from unresolved ambiguities (Frazier, 1994; Frazier & Clifton, 1996).

4.2.2.2 The constraint-based theory

The constraint-satisfaction model (Tanenhaus & Trueswell, 1995) is an interactive and parallel alternative model of ambiguity resolution that considers that the parser's online decisions on sentence comprehension are affected by the interaction of various types of structural and non-structural factors which includes syntactic, semantic and pragmatic information, and world knowledge, all of which have some role in evaluating the alternative structural representations of ambiguous sentences (Boland, Tanenhaus, & Garnsey, 1990; Taraban & McClelland, 1990; Juliano & Tanenhaus, 1993, 1994; MacDonald, 1994; MacDonald, Pearlmutter, & Seidenberg, 1994; Merlo, 1994; Tabossi, Spivey-Knowlton, McRae, & Tanenhaus, 1994; Trueswell & Tanenhaus, 1994; Trueswell, Tanenhaus, & Garnsey, 1994 ; Spivey-Knowlton & Sedivy, 1995).

Constraint-based models are parallel and interactive, so the parser is supposed to build multiple analyses and rank them according to constraints using multiple sources of information. Garden path effects are explained by these models, assuming that the parser is initially biased towards

simple and frequent structures, which can lead to an incorrect interpretation of the sentence. For example, in (48) above, the parser initially generates a structure where *raced* is the main verb since this is a more frequent structure than where *raced* is a past participle. This initial interpretation leads to the garden path effect. However the parser is sensitive to a variety of constraints that can help avoid garden path effects. These constraints include syntactic constraints, which are rules that dictate how words can be combined into phrases and sentences, as well as semantic and pragmatic constraints, which involve the meaning and context of the sentence. For example, if a sentence contains a verb that is not consistent with the expected meaning of the sentence, the parser will be more likely to revise its initial interpretation. Additionally, the theory also incorporates the idea that constraints are probabilistic, like the statistical likelihood of certain syntactic or semantic patterns which can help the parser choose the most likely interpretation of a sentence based on frequency.

One of the strengths of the constraint-based theory is that it can account for a wide range of garden path effects, including those that involve ambiguity, syntactic complexity, and lexical ambiguity. Overall, the constraint-based theory of garden paths proposes that sentence comprehension is a dynamic process that involves the interaction of multiple sources of information, and that the parser is able to use a variety of constraints to help avoid garden path effects.

4.2.2.3 The referential theory

The referential theory considers that the referential properties of sentences and the complexity of alternative discourse representations plays a key role in resolving structural ambiguities in sentences. Parsing preferences that were previously attributed to structural properties of sentences are instead seen as consequences of semantic and referential principles. Similarly to the constraint-based theory, this theory is considered to compute multiple partial structural

analysis of an ambiguous phrase, which means that the referential theory is considered a parallel processing mechanism.

The referential theory prioritizes principles of discourse over a priori plausibility, while the constraint satisfaction model emphasizes plausibility over discourse principles. Both theories suggest that parsing decisions are influenced by general knowledge about the world or the universe of discourse. Crain and Steedman (1985) argue that specific knowledge should take precedence over general knowledge in case of conflict. They propose the *principle of parsimony* which establishes that “[i]f there is a reading that carries fewer unsatisfied but consistent presuppositions than any other, then that reading will be adopted and the presuppositions in question will be incorporated into the perceiver’s mental model” (p. 333).

4.2.2.4 *The Unrestricted Race Model*

The unrestricted base model is a more recent theoretical framework that combines some of the aspects of both syntax-first models like the garden path model, and more interactive constraint-based models. The unrestricted race model (Traxler, Pickering, & Clifton, 1998; van Gompel, et al., 2005; van Gompel, Pickering, & Traxler, 2001) is a theoretical framework that proposes that language processing involves a competition among different possible syntactic structures. According to this model, when someone encounters an ambiguous sentence, multiple alternative syntactic structures are initially activated and compete with one another for processing priority. The different syntactic structures "race" against one another, with the winner being the structure that is most consistent with the context of the sentence and the listener's or reader's expectations. The Race-Based Model proposes that this competition among different syntactic structures occurs at different levels of processing, from the initial processing of individual words to the construction of complex sentence structures. In this model, the different syntactic structures are represented as activation levels within a network, and the competition among them is determined by the relative strength of their activation levels.

Similar to constraint-based models, the unrestricted race model proposes simultaneous integration of several constraints from different sources such as statistical, semantic, and contextual sources. However, according to this model, ambiguity resolution is based on the immediate probabilistic selection among weighted alternatives of an ambiguity. Therefore, as in syntax-first models like the garden path model, it hypothesises that there must be a separate reanalysis mechanism that is responsible for the garden path effect when the initial structural choice is proven to be the inappropriate one. Nevertheless, unlike the garden path model, it also predicts that where there is equal competition between structures, the garden path will only sometimes be elicited and thus both alternatives can occur withing the same experimental settings.

The Race-Based Model in psycholinguistics offers a flexible and comprehensive framework for understanding language processing. By allowing for parallel processing of multiple alternative syntactic structures and integrating various sources of information, including statistical, semantic, and contextual cues.

4.2.2.5 Retrieval-based models

Retrieval-based models propose that comprehension difficulties in garden path sentences arise when the parser encounters a word or structure that does not match its expectations based on previous processing. These models suggest that the parser relies on memory retrieval processes to access previously encountered linguistic information and resolve ambiguities. One influential work in this area is "An activation-based model of sentence processing as skilled memory retrieval" by Lewis and Vasishth (2005). In this paper, the authors introduced the concept of retrieval-based parsing, which suggests that language comprehension involves the retrieval of previously encountered linguistic information from memory. According to this model, garden path effects occur when the parser retrieves an incorrect interpretation due to similarities between the current input and previously encountered structures, leading to initial

misinterpretations that must be revised. Studies using retrieval-based models often investigate how memory retrieval processes operate during sentence processing, such as the activation of competing linguistic representations and the role of working memory in resolving ambiguities (Van Dyke, 2011).

Retrieval-based models have the advantage of explicitly considering the role of memory retrieval processes in language comprehension, which provides a mechanism for resolving ambiguities based on previously encountered linguistic information. Furthermore, these models can also account for a variety of empirical findings relating to sentence processing that not only include garden path sentences, but also other phenomena such as structural priming and syntactic ambiguity resolution. Additionally, retrieval-based models can account for individual differences in language processing ability by including parameters related to working memory capacity or load.

4.2.2.6 Surprisal-based models

Levy (2008) proposed a model of syntactic comprehension based on the principle of surprisal, which measures the unexpectedness of linguistic input given preceding context. This work laid the foundation for surprisal-based models of language processing. In this manner, surprisal-based models focus on the degree of unexpectedness or *surprise* associated with different linguistic structures. These models calculate the surprisal of each word or structure in a sentence based on its probability given the preceding context. According to surprisal-based models, garden path effects occur when the parser encounters a low-probability structure that is unexpected given the preceding context, which leads to processing difficulties and potential misinterpretations of the sentence. Studies using surprisal-based models often investigate how the surprisal of different linguistic structures influences processing difficulty, as well as how factors such as syntactic complexity and lexical ambiguity affect surprisal calculations (Levy, 2013).

The surprisal-based model quantifies the unexpectedness of linguistic input based on the probability of the linguistic event given the preceding context. This probability can be estimated by the integration of probabilistic information through the use of statistical language models trained on large corpora of text data, and then assign a surprise value to each linguistic element in specific contexts, which can range from high expectations (low surprise value) or highly unexpected (high surprise value). Surprisal-based models have been shown to be consistent with neuroscientific evidence, such as event-related brain potentials (ERPs), providing support for their validity in explaining language processing phenomena.

4.2.3 Evidence from research

Ni, Crain and Shakweiler (1996) used garden path sentences to test the predictions of the referential theory. They performed 4 experiments in which they tested main verb clause and reduced relative clause ambiguities, similar to what we will do in this chapter, and ambiguous attachment sites for prepositional phrases. For their two garden path experiments, they compared the focus operator *only* and the definite determiner *the* in order to observe whether the parser's decision for a preferred structure could be manipulated by the referential properties of the focus operator.

The authors explain that the semantic function of the focus operator *only* is to indicate that the focus element (the element that *only* modifies) is contrasted with a set of alternatives. For example, in (49), the focus element *Willoughby's coffee* is being contrasted with the coffee from other shops in New Haven. The use of *only*, though, does not guarantee that this comparison has actually been made across all possible options, but rather presupposes that it has been made. This presupposition is triggered by the presence of the focus operator *only*.

(49) In New Haven, *only Willoughby's coffee* is really good.

Semantically speaking, the focus operator has three components: background information (what we say about the focused element), the element in focus (directly modified by the operator),

and a contrast set that is not mentioned explicitly, but represents the presupposed list of alternatives that are being compared to the focus element. For a sentence with *only* to be true it is necessary that the background information applies to the focused element (*Willoughby's coffee is really good*), but not to the contrasting set (*coffee from other places in New Haven is not [really] good*). Given that the appearance of the word *only* in the sentence predisposes the parser to expect a contrasting set, the referential theory predicts that garden path effects will not rise in sentences like (50), despite being the same as the sentence in (51) from the first noun onwards as explained below.

(50) Only businessmen loaned money at low interest were told to record their expenses.

(51) The businessmen loaned money at low interests were told to record their expenses.

The focus operator *only* gives an instruction to the parser to look for the contrasting set that is being compared to the focus element. If this set has not been previously established (in preceding context), the parser has two options to define it: it will either construct the contrasting set from scratch by accommodating a new set of entities (*those who are not businessmen*) in comparison with *businessmen*, or it can interpret the following ambiguous phrase as a reduced relative clause and partition the set of businessmen, who are already incorporated in the sentence as part of the mental model, into a subset of businessmen with a certain specific quality (in this case, *those who were loaned money at a low interest*) and contrast them to a subset that does not contain the same specific quality (e.g. *those who were loaned money at a high interest rate*). Given the principle of parsimony which suggests that the preferred interpretation is the one that requires the fewest additional assumptions, the parser is more likely to adopt the analysis that includes the restrictor on “businessmen”. Interpretation of the phrase as a reduced relative clause results in fewer commitments and presuppositions compared to having to construct a contrasting set from scratch.

A further prediction of the referential theory put to the test by Ni, Crain and Shankweiler is that if the contrast set is actually established before the point of ambiguity, sentences with *only* will elicit garden paths (52).

(52) Only *wealthy* businessmen loaned money at a low interest were told to record their expenses.

In this case, the adjective *wealthy* already fulfils the requirement for a contrast set (*wealthy businessmen vs businessmen who are not wealthy*) and, therefore, the parser will prefer the main clause interpretation since it the one that requires fewer presuppositions. This interpretation will result in a garden path effect at the verb *were* as the point of disambiguation.

In their garden path experiments 1 and 2, they tested sentences such as the ones in (53) plus a 4 unambiguous controls sentences using two online techniques, self-paced reading and eye tracking (but without the adjectives) respectively.

(53) Experimental sentences for experiments 1 and 2 (Ni, Crain, Shankweiler, 1996)

Ambiguous test sentences

The-amb	The businessmen loaned money at a low interest were told to record their expenses.
Only-amb	Only businessmen loaned money at a low interest were told to record their expenses.
The-adj-amb	The wealthy businessmen loaned money at a low interest were told to record their expenses.
Only-adj-amb	Only wealthy businessmen loaned money at a low interest were told to record their expenses.

Unambiguous test sentences

The-unamb	The vans stolen from the parking lot were found in a back alley.
Only-unamb	Only vans stolen from the parking lot were found in a back alley.
The-adj-unamb	The new vans stolen from the parking lot were found in a back alley.
Only-adj-unamb	Only new vans stolen from the parking lot were found in a back alley.

The results of these experiments found support for the predictions of the referential theory in both studies. In self-paced reading they found a decrease in both reaction times and errors in the only-amb condition at disambiguation in comparison to the the-amb condition. In eye-tracking they found that first time reading times were also shorter at disambiguation for the

only-amb condition but much longer for the-amb with a clear garden path for the latter condition but no clear garden path effect for the former one. It is possible that the requirement of the focus operator for a contrast set can direct the parser into preferring the reduced relative clause alternative so the appearance of the main verb came to no surprise, thus preventing the appearance of a garden path effect. Alternatively, it is also possible that providing a richer description of the referred set can lead to easier accommodation of the relevant contrast set.

When adjectives were included, the only-adj-amb condition did not present any difference to the the-amb condition (sans adjective), revealing that the appearance of the adjective before the focused noun fulfilled enough requirements for the parser to promptly establish a contrasting set and, therefore, prefer a main verb analysis at ambiguity, hence the garden path effect was not averted by the appearance of *only*.

In comparison to unambiguous conditions, the only-amb condition presented more errors and slightly higher reaction times. According to their research, these differences can be explained as a consequence of parallel processing and that the results show that, despite the focus operator averting the garden path effect at the ambiguous region, at least briefly, more than one analysis was considered. These results were not consistent with a serial approach such as the garden path model, but are consistent with the referential model showing both that referential effects occur online in ambiguity resolution, and that the referential information of *only* heavily influenced the parser's initial analysis of ambiguity.

Clifton and Frazier (2010), in a more recent study, use NPIs to try to determine when the monotonicity of a context is identified by the parser. By using the domain-widener NPI *ever*, they analysed sentences in which the NPI initially appears to be unlicensed but is later shown to be licensed. These experiments were designed to test whether computations of monotonicity (i.e. the downward entailingness or upward entailingness) of a context are done locally and on-line,

or in the global context of the whole sentence and whether a temporary failure of NPI licensing results in the decrease in acceptability ratings of globally acceptable sentences.

According to the global context approach NPIs are semantically licensed by DE contexts and monotonicity computations that come into play only within the global context of the whole sentence, and therefore, a temporary misanalysis shouldn't influence the processing of NPIs. Meanwhile, the local approach predicts that difficulties in the processing of NPIs should arise if an NPI appears to be initially unlicensed, regardless of whether it is later shown to be licensed.

In order to test these two approaches, Clifton and Frazier performed two written acceptability judgment tasks to see whether the initial incorrect (non-DE) analysis of NPIs inside a reduced relative clause (RRC) resulted in decreased acceptability of sentences that contained them.

In their first experiment, they used sentences such as the ones in (54), where they compared sentences with a reduced relative clause against sentences that contained a full relative clause, in which an NPI may be present or not. They predicted that the [RRC +ever] cases should be considered slightly odd due to the temporary misanalysis of *ever* inside the apparent non-DE context and therefore show a reduction in acceptability ratings compared to the other sentences.

- (54) Experimental sentences for experiment 1 (Clifton and Frazier, 2010)
- [RRC -ever] A man arrested in this country won't break the law here again.
 - [RRC +ever] A man ever arrested in this country won't break the law here again.
 - [FRC -ever] A man who was arrested in this country won't break the law here again.
 - [FRC + ever] A man who was ever arrested in this country won't break the law here again.

The results from these experiment showed that reduced relative clauses were rated slightly worse than full relative clauses, and that sentences that contained *ever* were also rated slightly worse than those that did not contain the NPI. Crucially, though, they found an interaction between clause type and NPI presence where, as predicted, [RRC +ever] sentences had the lowest rating among all conditions.

In their second experiment, they tested the same sentences, but using negation (*No*) in the head of the relatives (55). The prediction in this case was that if the difficulty observed in experiment 1 was due to the NPI appearing to be initially unlicensed, it should not be present since the presence of the negative head would create an unambiguous DE context for the NPI regardless of relative clause type.

- (55) Experimental sentences for experiment 2 (Clifton and Frazier, 2010)
- [RRC -ever] No man arrested in this country won't break the law here again.
- [RRC +ever] No man ever arrested in this country won't break the law here again.
- [FRC -ever] No man who was arrested in this country won't break the law here again.
- [FRC + ever] No man who was ever arrested in this country won't break the law here again.

In this case, the results showed that no effect approached significance, but found only a slight cost of having an NPI in general. The interaction found in experiment 1 was not found in experiment 2 where the presence of *No* immediately licenses the NPI.

The results from these two experiments confirm the prediction that monotonicity of the sentence is computed locally since a temporary ungrammatical analysis of the NPI was enough to lower the judgment ratings even when the sentence was ultimately interpreted as grammatical in a global interpretation of the sentence. Additionally, experiment 2 proves that the reduction in acceptability judgments was indeed due to an initial misanalysis of the licensing of the NPI, which enforces the view that the language processor is able to determine the monotonicity of the current sentence on a word-by-word basis, rather than only globally when the whole sentence is finalised.

4.3 Experimental findings

4.3.1 Experiment 7: The role of NPI licensing on Garden Path sentences

In order to observe whether the inclusion of an NPI can help reduce garden path effects, we created garden path sentences that contain a reduced relative clause (RRC), where a past participle verb is typically initially parsed as the main verb of the sentence rather than a

subordinate verb and included an NPI inside the relative clause preceding the verb. The experiment included full relative clauses which do not cause garden path effects since the presence of a relative pronoun and a copula heading the phrase eliminates the ambiguity that is observed in reduced relative clauses which do not contain these elements and, thus, overtly present the relative clause.

Furthermore, we manipulated whether the NPI was licensed by starting the sentence with either a universal quantifier (*every*), which licenses NPIs within the relative clause (quantifier restriction), or a definite determiner (*the*) which doesn't license NPIs at all. Presence of the NPI is crucial for the current analysis based equally on whether it is licensed (from the beginning) or not, and whether it is located within an full overt relative clause or within a reduced relative clause. For this experiment, NPIs were always included in the stimuli so it is important to establish the distinction between the licensing status (which concerns the determiner used), but is also dependent on whether comprehenders identify its structural location. An NPI inside a full relative clause will be more easily identified as licensed compared to when it is found inside a reduced relative clause where its structural location may be ambiguous, unless the NPI itself is enough to clarify the structural context.

As we have seen in previous chapters, NPIs are licensed inside the restrictor of the universal quantifier, but not inside its scope, and the chance of successfully determining the boundaries between the restrictor clause and the nuclear scope could increase when more information is provided inside the restrictor, making it easier to determine a contrasting set for the noun. While evidence from experiment 6 did not reveal any visible effects in self-paced reading immediately after the appearance of the NPI between the unmodified and the modified conditions or the determiner, it may be possible to observe effects of the NPI when the structural context of the sentence is at least partly dependent on the NPI itself, whether it is the current structural

context, such as being able to identify that the NPI is located inside a (reduced) relative clause, or the ongoing structure, such as in the case of the garden path effect.

In a full relative clause, the relative pronoun and the copula clearly informs comprehenders of the upcoming structure, but in a reduced relative clause where this elements do not precede the NPI this information is not readily available to indicate that the NPI is located within a relative clause. Crucially, in order for the NPI to be grammatical at this point it must meet both requirements: a proper licenser and being within the appropriate structural context. If the correct structural location is not immediately clear to the readers, the sudden appearance of an NPI may seem weird but, in order for them to recuperate the licensing of the NPI, comprehenders may use the NPI itself as a sign as to what the present structure is and, therefore, consider a different interpretation of the tense of the upcoming verb (i.e. the past participle rather than the simple past). If the presence of the NPI can help with identifying the structural context of its location, we should expect readers to favour the past participle interpretation of the subordinate verb over the main verb reading and thus alleviate the garden path effect expected to appear at the disambiguation regions.

For the determiner cases this might be different given that the NPI is not initially licensed and its appearance could, as a result, be more confusing than in the universal quantifier cases. In full relative clauses the apparition of the NPI should not contribute any additional information about the structure of the subject since the preceding elements already introduce the RC. As such, encountering the NPI at this point would probably cause more confusion than in the quantifier condition. We anticipate that these types of sentences will be less prone to a reduction in the garden path effect expected to appear after the disambiguation regions.

There are two main areas of ROIs in the experiment. On one hand, the subject of the sentence is where we would observe whether the presence of the NPI, either licensed or not, has an effect in comprehenders identification of the reduced relative clause early on. While we can expect to

see a difference between licensed and unlicensed NPIs (effects of Determiner), we can also see if there are any differences between the RRC and the FRC conditions relative the determiners that are being used. As such, we would expect to see that the condition that has the universal quantifier and the full relative clause (UNI.FRC) should be faster and easier to process than the other three conditions, at least after *ever* makes an appearance. Furthermore, in the reduced relative clauses, if *ever* does indeed help comprehenders to consider the following content as a relative clause, we should see at this point differences between both RRC conditions, where a licensed NPI should facilitate the processing of the subordinate verb and succeeding spill-over regions. If the NPI does not have any immediate effect yet at this point on determining or giving clues about the nature of the structure, that is to say that comprehenders are not taking the RRC hypothesis as a possibility, we should expect to see no significant difference between both RRC cases since both would be considered equally bad.

4.3.1.1 Participants

For the current experiment, 72 participants were recruited through the Prolific website (<https://www.prolific.co/>). The experiment was advertised as a reading task and were informed of the general procedure of the experiment. They were asked to read through sentences one word at a time and asked to answer to comprehension questions after some of the sentences. All participants were asked for their consent to be part of the experiment and were required to be Native Speakers of English. The experiment included a questionnaire about their linguistic background in order to ensure that all participants had English as their first language and to account for other languages they may know. Each participant was offered £6.00 for the completion of the experiment which should last approximately 30-45 minutes.

4.3.1.2 Materials

We designed a set of 44 experimental items that included 4 different conditions based on ‘Determiner’ depending on the starting word of the sentence and that served as a potential licenser for the NPI, and ‘Structure’ depending on the type of relative clause that was used within the subject of the sentence. Based on ‘Determiner’ each sentence either started with the definite determiner *the* (DET) or the universal quantifier *every* (UNI). Based on ‘Structure’, each

subject contained either a full relative clause (FRC) or a reduced relative clause (RRC). Each sentence included a subject with a relative clause followed by a main clause predicate. Subjects were formed by a potential licenser (*Every/The*) followed by a main noun and a relative clause. In the case of FRCs, the noun was immediately followed by *who was*, which overtly indicated the beginning of the relative clause and was followed by the NPI *ever* and a verb plus three more words. RRCs also included an NPI and a verb followed by three words, but without the use of *who was*, therefore not giving any overt indication that what followed after the noun was a relative clause. After the relative clause, the predicate started with an auxiliary verb (*did, has, or was*), followed by a negation *not* and a main verb plus three to seven words in order to ensure that spill-over regions after disambiguation and the negation could be observed, and to avoid sentence final effects. The negation was added to the predicate in order to serve as a potential secondary licenser for the NPI and to make sentences that started with the definite article (in which the NPI is not initially licensed) grammatical. An example of all conditions can be observed in Table 21. Each participant saw a total of 120 sentences that included 44 experimental items and 76 fillers.

Table 21: Experiment 7 – sample stimuli

Subject	RC	Sentence
Determiner	FRC	The soldier who was ever marched into a battle was not scarred by the experience.
Determiner	RRC	The soldier ever marched into a battle was not scarred by the experience.
Universal	FRC	Every soldier who was ever marched into a battle was not scarred by the experience.
Universal	RRC	Evert soldier ever marched into a battle was not scarred by the experience.

4.3.1.3 Procedure

Similarly to the experiments in chapter 3, experiment 7 was designed to use a self-paced, word-by-word, moving window serial visualisation paradigm in which participants were shown a series of dashes that corresponded to the words in the sentence, including spaces. Upon pressing the space bar key, the first word was revealed, replacing the dashes that masked it, allowing participants to start reading and the program to start recording reaction times between key presses. Further hits of the space bar revealed the following word while masking the previous

one with dashes once more. In this manner, a participant was only able to read one word of the sentence at a time and was not able to review previous content. At the end of the sentence participants were either shown a comprehension question (50% of the stimuli) or presented with a new sentence to read. Participants were prompted to respond each question with either 'yes' or 'no' by pressing the corresponding key on the keyboard ('j' for "yes", and 'f' for "no"). Comprehension questions were used mainly to ensure that participants were paying attention to the sentences and not simply rushing through them. Before the task was initiated, participants were trained into the task through a set of practice items that did not fully resemble the experimental items in order to ensure that participants knew how to perform the task correctly while not giving away the experiment's purpose. The experiment was coded using PennController for Ilex Farm (<https://farm.pcihex.net/>).

4.3.1.4 Data Analysis

The experimental data was analysed on R software (R Core Team, 2014) using the lme4 package (Bates, Maechler, & Bolker, 2011) for linear mixed-model effects with fixed factors for experimental manipulations and a specified random effect structure, including random intercepts and slopes for all fixed effects by participants and by items (Bayeen, Davidson, & Bates, 2008; Barr et al., 2013). In this experiment, we look at whether the presence of a licensed NPI inside a reduced relative clause can help the parser prevent the appearance of a garden path effect at the point of disambiguation by providing with clues about the context in which the NPI should be found. Contrast coding for both experimental conditions were organised as follows: Determiner ('the' = 1, 'every' = -1) and Structure, which refers to the relative clause type (FRC = 1, RRC = -1). Data was modelled across 15 regions from the Licensor (sentence onset) up to three words after the Main Verb for plotting purposes. The statistical analysis was divided into two separate sections, the first one inside the relative clause and the second one at the point of disambiguation of the garden path sentence. For the relative clause, we analysed the NPI as the critical word plus two following words, while in the garden path

disambiguation section we looked at the first disambiguation word plus two following regions corresponding to the negation and the main verb of the sentence. Effects were considered significant if the p values were $p \leq 0.05$ (Gelman & Hill, 2007). Additionally, paired comparisons were calculated using emmeans package. Extreme RTs were removed by using model criticism (Baayen et al., 2008). No participants were removed due to extreme reading times or inaccurate answering of the comprehension questions. A total of 0.26% data was lost during clean-up of extreme individual RTs.

4.3.1.5 Results




Figure 102: Experiment 7 Reading Times by Region

Figure 12. shows the general plot for all 4 experimental conditions observed in experiment 7. The graph shows each region in a word-by-word manner and the corresponding mean reaction times in each region for each condition. Paired comparisons were calculated for each region by both Determiner and Structure to better observe the differences between conditions. In this experiment we separate the regions into two main areas: primarily the garden path effect which corresponds to the disambiguation regions 'Disamb' (region 5) followed by *not* and the main verb (MV), and secondly, the area corresponding to the relative clause, starting with critical word (CW) *ever* (region 0). Below, we go through the analysis of all regions from *ever*, region 0, through MV (region 7).

Regions 0-2: NPI licensing effects inside the relative clause

Region 0: Critical Word: NPI ever

Table 22: Region 0 (NPI) - Linear mixed effects modelling

Fixed effects:	Estimate	Std. Error	df	t value	p value	
(Intercept)	392.762	16.128	74.344	24.352	< 0.001	***
Determiner	1.460	3.006	2931.634	0.486	0.627	
Structure	22.481	4.933	63.240	4.557	< 0.001	***
Det x Str	-3.072	3.006	2931.640	-1.022	0.307	

Table 22 shows the linear mixed effect results of region 0 that corresponds to the NPI *ever*. At this region, we found a significant effect of Structure ($t = 4.557$ $p < 0.001$), but no effect from determiner, nor an interaction. In both cases, the FRC are read faster than the RRC. It appears that comprehenders are surprised to find an NPI within a main clause when it is located within the reduced relative clause compared to when it is inside an overt relative clause. Furthermore, examination of the graph, shows that in general, the presence of the relative pronoun and the copula lowers the reading times considerably right before region 0 and that reading times at the time of apparition of the NPI after the copula even drops slightly while in the RRC cases, going directly into the NPI causes some delays.

Considering the SPR graph, we can also observe that at the time the NPI appears inside an RRC, it is slightly slower on the universal quantifier condition in comparison to the determiner case. It could be possible that at the time the NPI appears, the parser may take some time to consider the quantifier as a possible licenser, or that some processes are being considered to resolve the structure as a RRC. Effects from these possibilities can be observed in the following regions.

Region 1: CW +1 (Subordinate Verb) and Region 2 (CW +2)

Table 23: Region 1 CW +1 (Subordinate Verb) – Linear mixed effect modelling

Fixed effects:	Estimate	Std. Error	df	t value	p value	
(Intercept)	422.644	18.311	69.991	23.081	< 0.001	***
Determiner	13.819	3.795	2972.410	3.641	< 0.001	***
Structure	14.485	4.389	69.436	3.300	0.002	**
Det x Str	1.982	3.795	2972.536	0.522	0.601	

Results for region 1 can be found in Table 23. In this region we found main effects of both Determiner ($t = 3.641$, $p = < 0.001$) and Structure ($t = 3.300$, $p = 0.002$). Since the subordinate verb is the source of the ambiguity for the reduced relative clause (RRC), it is at this point that the parser may consider the possible interpretations of the structure. Given that this ambiguity does not exist for the FRC, increased times for the RRC represent the increased processing times required by the parser and this difference is represented statistically. The significant difference in determiner is driven by the difference observed between the universal quantifier and the definite article within both types of structure where the definite article conditions show inflated reading times relative to the universal conditions.

What is interesting at this point is that the NPI can only be grammatical if it is identified as being inside the restrictor of the quantifier, which requires the parser to interpret the verb in this region as participle and not as the main verb of the sentence. This distinction is facilitated in the FRC cases by making overt the relative clause ('who was') which explains why the increase in reaction times for the universal quantifier is reduced compared to the increase in times for the definite article at this stage. Nevertheless, this distinction is not presently available for the RRC, yet we see a very similar pattern between the universal and the definite article conditions despite that the ambiguity should be similar. It is possible that the NPI from the previous region has provided some facilitation to the parser in determining the form of the verb and thus we see a reduction in the time the parser requires to consider possible structures. If this is the case, could it be possible that the parser was able to, at least temporarily, license the NPI on the previous region based solely on the universal quantifier, but not on the determiner cases where we see an increase in reaction times, or alternatively, that the NPI itself was able to inform the parser that the more likely structure is a (reduced) relative clause? While there is an increase in RTs in the quantifier cases, this may reflect processes relevant to licensing of the NPI itself or recognition of the RRC thanks to the licensed NPI. What is most clear is that the increase in RTs for the quantifier cases is not as big as the increase for the determiner cases. This is most

noticeable in the RRC cases where the increase observed in the determiner case is quite noticeable, while the quantifier cases remain relatively flat with respect to the previous region.

Table 24: Paired Comparisons Region 1:

Region 1: CR+1 (Sub Verb) by Structure						Region 1: CW+1 (Sub. Verb) by Determiner					
	Estimate	SE	df	z ratio	p value		Estimate	SE	df	z ratio	p value
<u>FRC</u>						<u>The:</u>					
DET - UNI	23.7	10.7	Inf	2.205	0.028	FRC - RRC	-32.9	11.6	Inf	-2.837	0.005
<u>RRC:</u>						<u>Every:</u>					
DET - UNI	31.6	10.7	Inf	2.944	0.003	FRC - RRC	-25.0	11.6	Inf	-2.156	0.031

In order to compare the size of the effects between the definite article and the universal quantifier inside both structures separately, we looked at paired comparisons (table 24) in this region. Inside the full relative clause (FRC) we see a significant difference between the definite and the quantifier ($z = 2.205$, $p = 0.028$). We also see a significant difference between both conditions inside the reduced relative clause ($z = 2.944$, $p = 0.003$). Moreover, the estimated difference in RTs between both differences is very similar, being slightly greater for the RRC (estimate DET – UNI = 31.6, SE = 10.7) compared to the FRC (estimate DET - UNI = 23.7, SE = 10.7). While we didn't find an interaction structure x determiner, the difference between the RRC conditions could point to the parser identifying the NPI as being licensed, or that the presence of the NPI could have helped the parser determine the correct form of the verb. Unfortunately, these results are not strongly suggested by the main model.

Table 25: Region 2 CW+2 – Linear mixed effect modelling

Fixed effects:	Estimate	Std. Error	df	t value	p value
(Intercept)	413.685	14.611	71.699	28.313	< 0.001 ***
Determiner	13.538	3.820	63.569	3.544	< 0.001 ***
Structure	11.2636	4.428	46.438	2.544	0.014 *
Det x Str	-0.350	4.146	56.065	-0.084	0.933

These results remain consistent through Region 2 (CW+2) where all conditions follow the same pattern and relatively flat with respect to region 1. We found a main effect of determiner ($t = 3.544$, $p < 0.001$), and a main effect of Structure ($t = 2.544$, $p = 0.014$), but no interaction between

the two. Similar to region 1, paired comparison revealed a similar pattern for the differences between the universal quantifier and the definite article within the RRC (DET – UNI: estimate = 27.8, SE = 11.0, df = inf, z = 2.533, p = 0.011) and the FRC (DET – UNI: estimate = 26.4, SE = 11.6, df = inf, z = 2.275, p = 0.023).

Regions 5 – 7 Garden path effects

Regions 5 and 6 correspond to the disambiguation and negation *not* where we would expect to find garden path effects. Table 26 shows the results for region 5, and Table 27 shows the results for region 6.

Regions 5 (disambiguation)

Table 26: Region 5 - Linear mixed effects modelling

Fixed Effects	Estimate	Std. Error	df	t value	p value
(Intercept)	422.756	18.315	74.718	23.083	< 0.001 ***
Determiner	6.412	3.513	2996.945	1.825	0.068 .
Structure	-0.276	3.518	3009.338	-0.079	0.937
Det x Str	0.937	3.513	2996.673	0.267	0.790

In region 5 (disambiguation), results show a trend on Determiner that did not reach significance (t = 1.825; p = 0.068) which shows that at this time we see a slight difference between both determiner types. We found no effect of structure or interaction. Both universal conditions are slightly faster than the definite conditions. The lack of a difference between both Universal conditions can point to the lack of a garden path effect at this region between these conditions, but the determiner cases don't seem to display a difference between themselves either so it is difficult to point out exactly what is happening. Since the full relative clause with the determiner should have clarified that the verb was inside a relative clause, we can't expect this condition to reflect a garden path, but a slight increased difficulty with the definite article with respect to the universal quantifier.

Regions 6 (negation)

Table 27: Region 6 – Linear mixed effects modelling

Fixed Effects	Estimate	Std. Error	df	t value	p value
(Intercept)	384.568	13.370	73.371	28.764	<2e-16 ***
Determiner	7.149	2.705	108.472	2.642	0.009 **
Structure	7.397	3.009	93.123	2.458	0.016 *
Det x Str	4.399	2.840	107.092	1.549	0.124

In region 6 (negation), we can observe a main effect of determiner ($t = 2.642$, $p = 0.009$), as well as a main effect of Structure ($t = 2.458$, $p = 0.016$). The effect on structure shows that in general, FRC are faster than RRCs. The effect of determiner, on the other hand shows that the universal conditions were processed faster compared to the definite articles and the trend observed on the previous region has consolidated at this point. Despite not finding any interaction, though, graphic data seems to suggest that reaction times for the universal quantifier + RRC condition is very close to the processing times of the definite article condition in the FRC. These results could demonstrate a garden path effect for reduced relative clauses at the negation. Since we did not find an interaction between structure and determiner, we are unable to determine whether the NPI had an effect on the garden path effect for the universal condition despite it showing lower reaction times compared to the determiner with the RRC that closely match those of the determiner in the full relative clause.

Given the relative difference displayed on figure 12 we decided to look at the paired comparisons to see if we can gain any insights from the observed conditions.

Table 28: Paired Comparisons Region 6: based on Raw RTs

Region 6: Negation by Structure						Region 6: Negation by Determiner					
	Estimate	SE	df	z ratio	p value		Estimate	SE	df	z ratio	p value
FRC:						The:					
DET - UNI	5.40	7.61	Inf	0.710	0.478	FRC - RRC	-23.70	9.70	Inf	-2.444	0.015
RRC:						Every:					
DET - UNI	23.20	9.20	Inf	2.522	0.012	FRC - RRC	-5.89	7.26	Inf	-0.811	0.417

Paired comparisons on table 28 show no significant difference between both FRC conditions (DET-UNI: $z = 0.710$, $p = 0.478$), but a significant difference between both RRC conditions (DET-UNI: $z = 2.522$, $p = 0.012$). At the same time, we can see a significant difference between both definite article conditions (FRC-RRC: $z = -2.444$, $p = 0.015$) and no significant difference between both universal quantifiers conditions (FRC-RRC: $z = -0.811$, $p = 0.417$). Again, while we can see a significant difference between the universal and the determiner at this stage in the RRC, we can not overrule the presence of a garden effect for the universal quantifier due to the lack of a significant interaction between structure and determiner. While we could speculate that this difference could be due to facilitation by the NPI, these results are not conclusive.

Region 7: Main Verb

Table 29: Region 7 – Linear Mixed Effect Modelling

Fixed effects	Estimate	Std. Error	df	t value	p value
(Intercept)	387.540	15.576	74.634	24.881	< 0.001 ***
Determiner	1.286	3.092	102.966	0.416	0.678
Structure	6.113	3.143	85.933	1.945	0.055 .
Det x Str	-2.396	3.300	79.612	-0.726	0.470

The results of for the main verb can be observed in table 29. At the main verb of the sentence, we only found a trend of Structure that did not reached significance ($t = 1.945$, $p = 0.055$), but no longer an effect of determiner or any interaction. While observations from the previous region could lead us to speculate that there could be an effect of the NPI on the reduction of the garden path effect, results at this region no longer display that possibility. If the garden path effect extend to this region, the results show that both RRC display garden path effects and any reduction between the definite article and the quantifier conditions (if any) has been overcome.

4.3.1.6 Discussion

Looking at the data analysis, we can observe noticeable differences between the experimental conditions in terms of their reading times across each region. Self-paced reading analysis shows different timing at several ROIs that correspond in structure and the determiners used in each

sentences. These differences include marked differences between relative clause types across several regions, as well as differences based on the determiners, but we didn't find any interaction after the NPI or at or after the garden-path disambiguation regions. In order to better understand what may be going on we will discuss that results of the experiment in terms of both sections of the sentence.

The garden path effect

While the previous regions may indicate the consideration of the relative clause in earlier processing, the garden path effect is expected to appear in the regions from the disambiguation onwards. In other words, the appearance of the auxiliary verb should be able to shed light for comprehenders into realising that a main verb hypothesis immediately after the NPI was the wrong one and would be forced into reanalysing the structure they have formed. Therefore, we should expect to see differences in reading times between the full relative clause and the reduced relative clause conditions.

At the disambiguation region, we don't yet see a big difference between all conditions though determiner cases seem to be slightly slower than the universal ones. At the disambiguation region we saw that both universal quantifier conditions were read faster than both definite article conditions. Nevertheless, we found no effect from structure, which we should expect as a sign of the garden path effect. This probably means it was, at this point, too early to shown any garden-path effects at all. At the negation and the main verb regions we found clear effects for relative clause type, where both FRC were read faster than their RRC counterparts. Additionally, we found an effect from determiner which showed that each universal quantifier condition was read faster than their definite article counterparts, but we did not find any interaction between these factors, which does not allow us to draw any clear conclusions about the effect of a licensed NPI on the garden path effect. While looking more closely at the individual differences between the RRC cases showed that the universal quantifier was significantly slower than the

definite, we cannot guarantee that this is due to the NPI providing clues about the structure that are enough to prevent the garden-path effect at this point.

While we could *guess* that the reduction observed between both determiners at the negation could be due a licensed NPI driving the parser towards preference of an interpretation of the relative clause's verb, the difference does not seem to warrant that conclusion. At least the licensed NPI itself would not be able to produce an effect as strong as the presence of the overt relative clause, and the parser will continue to consider all possible interpretations of the verb. If this reduction in reading times at the negation is indeed an influence of the NPI, the fact that we observe it over the reduced relative clause with the universal quantifier, and not with the definite article, shows that for an NPI to produce a positive effect, it *must be licensed*, but also, that these effects are not strong enough to tip the scales towards the consideration of a single structure and are probably also short lived. Furthermore, even if some reduction of the effect is taking place at the negation region, evidence from the Main Verb region, shows a more clear preference for relative clause type as a main driver of the garden-path effect. One possibility for this is that the garden path effect for the universal quantifier + RRC condition was delayed up until the Main Verb appeared and it is only now that the parser is realising that the structure needs to be reanalysed. This possibility seems unlikely considering the steep increase in speed that happens at the disambiguation region for all the conditions, except for the determiner + RRC condition. It seems that disambiguation is actually confirming a presupposed structure at this point, rather than delaying it. Notwithstanding, if the garden path effect of the universal + RRC condition is being delayed, it could be possible that the negative element in region 6 is responsible for the increase of speed of this condition, perhaps serving as a secondary, translocated, maybe interfering, NPI licenser which gives the potentially unlicensed NPI a boost. While this could help explain why the UNI + RRC and the DET + FRC receive a boost in reading times, it doesn't fully explain why this wouldn't happen for the DET + RRC case if it helps with an unlicensed NPI (DET) and a potentially structurally unlicensed case (RRC) as well. Perhaps having

both factors increases the difficulty too much to be easily recovered. A third possible hypothesis is that, as previous studies have shown, despite reanalysis processes occurring and that readers are able to reconsider the original wrong structure, the original assumed structure is still kept in memory and may still be the active therefore interfering with the reanalysis process (Inoue & Fodor, 1995; Christianson et al., 2001; Slattery et al., 2013, Quian et al., 2018; Huang & Ferreira, 2021). It is therefore possible, that if the garden path was actually reduced to some point due to some facilitation from the NPI, consideration of multiple structures still drives the appearance of the garden-path effect and is more strongly being reflected at this point in the sentence. Perhaps this merits taking a closer look in future studies.

In general terms, though, it seems that processes involved in the processing of NPIs could continue to affect the upcoming structure by determining the licensing environment and integrating information at later times than originally observed with the SPR measure and the types of sentences that we used in the previous chapter. In order to determine whether the NPI was actually able to help resolve the structural ambiguity of the initial sentence structure, it is necessary to take a look inside the relative clause and see if there are any significant differences between licensed and unlicensed NPIs. Furthermore, this could also give evidence as to whether NPI processed are performed continuously throughout the entire sentence or if the effects observed at the disambiguation areas are due to a retrieval that happens locally. Comparison between sentences that do not contain an NPI could also help us determine if our speculations about the reduced reading times for the universal condition found at the negation with the reduced relative clause can be attributed to the presence of a licensed NPI or not.

Effects of the NPI in terms of the relative clause

At the beginning of the sentence we can already observe differences between conditions where the presence of a licensed NPI is giving some signs of facilitation in reduced relative clauses. Comprehenders seem to be surprised at the appearance of *ever* inside the reduced relative

clause, but more willing to accept it after a full relative clause regardless of the determiner at the beginning of the sentence. In all conditions after the NPI, at the subordinate verb, we see an increase in reading times for all conditions, but mostly evident in cases in which the NPI is not licensed. This could indicate that a correct licensing environment has not been identified. This can be attributed to the lack of a potentially valid licenser in sentences that start with the determiner *the*.

For the overt relative clauses, the universal case sees a mild increase which is probably caused by retrieval licensing effects, which at this point seem to be the first path the parser takes when encountering an NPI. This effect is increased even further for the determiner cases in which the NPI is not licensed in both types of reduced relative clauses. Most notably, in the case of the reduced relative clauses, we see that there is already a big significant difference between both determiner types.

At the subordinate verb, the determiner case has had a much steeper increase in comparison to the universal case. This is our first indication in the sentence that a licensed NPI already has an effect in the parsing of the sentence and that comprehenders may be showing facilitation effects from the NPI. Given that there is no indication that structurally there is a relative clause, we should expect both NPIs to be considered anomalous in both RRC cases since at this point comprehenders should expect to be within at the main clause verb. As mentioned earlier in this chapter, for the NPI to be grammatical there are two requirements that must be met: that it has a readily available licenser, in this case the universal quantifier, and to be within the restriction of the universal quantifier (in a structurally appropriate location). Identification or at least consideration of the structural context in this case is necessary to explain the difference between both reduced relative clause cases, otherwise we should also see at least some increase in reading times in the universal + RRC conditions as well as with the determiner + RRC, even if one is much bigger than the other. Despite not having found an interaction at this or the following

region, we do see a very similar pattern in the difference between reaction times in the reduced relative clauses that closely resembles that of the full relative clauses. While we do not see a reduction in reaction times for the quantifier after the NPI, the pattern seems to show that for sentences that start with the universal quantifier, the parser requires less time to process. This pattern is similar to the one found in the previous chapter where, though the NPI was not licensed, when a universal quantifier is involved (regardless of modification) the reaction times recorded are generally faster than if the licenser is the definite article. This pattern shows that in some cases, the parser was able to, at least temporarily, license the NPI, or that the universal quantifier facilitated processing of the NPI. Additionally, given lower times at the subordinate verb, it is possible that the positive identification of the NPI could have allowed the parser to quickly consider the verb as a past participle and spend less time to make this decision. Results observed at the disambiguation of the garden-path, nevertheless, show that this consideration might not have completely overridden the main verb interpretation, at least not as well as an over relative clause would.

By the time we reach the end of the relative clause, participants seem to have already reached certain decisions regarding some of these sentences. In general, the UNI + FRC, which is the baseline condition for this experiment, is accepted as licensed and easy to comprehend, while the DET.FRC case seems to be unlicensed beyond hope. For both cases of the FRC we can only observe NPI effects since there is no other possible grammatical anomaly or ambiguity. For the RRC cases, there is at this point no difference between them and they are somewhere between the overt relative cases. It is possible that readers are not yet decided whether these sentences are ungrammatical yet and are waiting to see what comes next. The fact that neither has been abandoned yet may indicate that a relative clause hypothesis is being considered while waiting for further information to disambiguate the context or resolve it into an irrecoverable sentence.

Despite the findings of this experiment, it is still difficult to see whether the results are due to the sole presence of the NPI. While we do see some evidence that reading times for the quantifier condition compared to the definite article within the RRC cases after disambiguation are faster (though not conclusive due to the lack of an interaction), we do not see a clear explanation of what happens within the first part of the sentences and how much is due to the NPI itself. Because of this, we decided to complement the current results with two more experiments where we look independently at both the universal quantifier (experiment 8) and the determiner (experiment 9) while modulating the presence of the NPI and the type of relative clause. We believe this will give us a more fine grained look as to what processes are involved, as well as a better explanation of how the NPI actually behaves in these cases.

4.3.2 Experiment 8: NPI presence manipulation on the Universal Quantifier

Experiment 7 shows important differences between conditions. Namely, we see some differences between the universal quantifier conditions and the definite articles, as well as some structural differences at both the relative clause regions and the disambiguation regions. Though we see a reduction in reaction times at the negation region for the universal quantifier, it is not clear whether this is actually an effect of the NPI itself. This could as well be a difficulty intrinsically determined by having both an unlicensed NPI as well as a reduced relative clause.

Additionally, we observe some signs of facilitation at the NPI region and following regions for the universal-RRC condition compared to the determiner-RRC condition which shows a large peak in RTs after the appearance of the NPI. We attribute these differences mainly to the presence of a licensed NPI and not to the presence of the NPI itself. In other words, it seems that for the NPI to give adequate clues about its environment, the licensing element that determines what this structural environment is (i.e. the universal quantifier determines that the relative clause is a downward entailing environment) must be present as well. Nevertheless, it is difficult to determine whether this effect is due the NPI being licensed by the parser, or whether this is due to the NPI itself providing facilitation to the parser to the nature of the subordinate verb.

Furthermore, there is a possibility that the difference in reaction times after disambiguation between both conditions could be produced due to the presence of the negation which could act as a potential secondary licenser to the NPI in the universal-RRC condition thus increasing speed of processing. The presence of the negation could also have an effect on the determiner-FRC condition as a form of a late licenser for the NPI, making the sentence potentially grammatical. In this experiment we approach these problems by looking at similar sentences where we modulate both the NPI presence and the type of relative clause, but leave the universal quantifier *every* constant throughout all conditions.

4.3.2.1 Participants

As in the previous experiment, 72 participants were recruited through the Prolific (<https://www.prolific.co/>) website. The task was advertised as a reading task and participants were informed of the general procedure of the experiment, using the same instructions as the previous SPR experiments. Prolific filters were used to ensure that all participants were native speakers of English and the experiment included a questionnaire about their linguistic background. Each participant was offered £6.75 for their participation on the task which lasted approximately between 30-45 minutes.

4.3.2.2 Materials

Table 30: Experiment 8 example stimuli

NPI	RC	Sentence
NO	FRC	Every researcher who was awarded a scientific prize was not invited to the conference.
NO	RRC	Every researcher awarded a scientific prize was not invited to the conference.
YES	FRC	Every researcher who was ever awarded a scientific prize was not invited to the conference.
YES	RRC	Every researcher ever awarded a scientific prize was not invited to the conference.

This experiment used the same items as in the previous one, but this were modified to suit the purposes of the study. In this case, we did not include a Determiner Type variation since all sentences started with the universal quantifier *every*. Instead, we manipulated the presence of the NPI by either including *ever* inside the relative clause before the subordinate verb (NPI =

YES), or not (NPI = NO). Same as before, the conditions were divided depending on whether they included a full relative clause (FRC), or a reduced relative clause (RRC), which typically produces garden paths when the subordinate verb tense is ambiguous (simple past = past participle). The rest of the structure was the same as in experiment 7. Examples of the stimuli used can be observed in Table 30. Each participant saw a total of 120 sentences that included 44 experimental items and 76 fillers.

The No-NPI + FRC should be the simplest sentence from our conditions so this is considered our baseline for comparison. Nevertheless, the NPI + FRC condition should also be as easy to process as the baseline but with the added difficulty of having an NPI. While both RRC conditions are grammatical, the No-NPI is our classic garden path sentence. It should be easily read throughout the entire RC since the Main Verb interpretation of the sentence is a completely valid one at this point, but should present difficulties at the disambiguation, since at this point reanalysis of the sentence should happen.

Finally, it is the NPI version of the RRC that we are most interested in. If the NPI doesn't facilitate the identification of the structural context (i.e. assuming there has to be a relative clause) we should expect a different pattern in the relative clause areas where an assumption of it being unlicensed due to the NPI being outside of a licensing position should cause delays in processing. Nevertheless, an assumption that the NPI is in fact licensed should also imply that it is located within a structurally relevant domain, in other words, inside a relative clause. If this is the case, we should see facilitation happening both inside the relative clause (compared to the No-NPI case at least) and at the disambiguation regions since the interpretation of the first verb as past participle should override the need to reanalyse at the disambiguation areas.

4.3.2.3 Procedure

The procedure used in experiment 8 was the same one used in experiment 7.

4.3.2.4 Data Analysis

Data was analysed using the same methods as in experiment 7. In this experiment, we once again looked at whether the presence of a licensed NPI inside a reduced relative clause can help the parser prevent the appearance of a garden path effect at the point of disambiguation, but in this case instead of including two types of licensors (a definite article and the universal quantifier *every*), we used only the universal quantifier condition. Instead, we manipulated the appearance of the NPI within either a full relative clause or a reduced relative clause in order to ascertain that the presence of the NPI has an effect on the later reduction (or not) of the garden path effect at the point of disambiguation. Contrast coding for both experimental conditions were coded as follows: NPI presence (no = -1, yes = 1), and Structure (FRC = -1, RRC = 1). Similar to experiment 7, data was modelled across 15 regions from the licensor (sentence onset) up to three words after the Main Verb for plotting purposes. Statistical analysis was divided into two separate sections for observation of the effects of both NPI presence and disambiguation of the garden path effect. Two participants were removed from the experiment since they did not achieve the minimum score in the comprehension questions (75%) leaving us with 70 participants for the study. No other participant was removed due to extreme mean reading times, and only 0.46% of the data was lost after the removal of individual extreme RTs.

4.3.2.5 Results



Figure 13: Experiment 8 Reading Times by Region

The plot corresponding to the self-paced reading results of all 4 experimental conditions can be observed in Figure 13, above. The graph shows each region in a word-by-word manner on the x-axis and the corresponding mean reactions for each region in the y-axis. Paired comparisons were calculated for each region by both NPI presence and Structure to better observe the differences between conditions. Same as in experiment 7, we will observe the data based on two main areas of interest, one corresponding to the Garden Path Effect after disambiguation to see whether NPI presence modulates the appearance (or reduction) of the effect, and then the inside of the Relative Clause to observe how NPI presence intervenes in helping (or not) to determine the structural context during the early parts of the sentence (whether participants consider a RC or not after the NPI).

Regions 0-4: NPI modulation effect inside the relative clause

Region 0: NPI ever

Given that we have modulated the appearance of the NPI ever in this experiment, only half of the trials included it. For the conditions that included the NPI, we found a significant main effect of Structure (Estimate = 16.323, Std Error = 3.423, df = 1462.982, t = 4.768, p < 0.001). This effect

is both expected and consistent with the previous experiment since licensing of the NPI is dependent on recognition of the relative clause that contains it.

Region 1: Critical word +1 (Subordinate Verb)

Table 31: Region 1: CW+1 – mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	388.070	17.182	70.871	22.585	< 0.001 ***
NPI	-3.257	2.889	2920.210	-1.128	0.260
Structure	19.997	3.991	43.897	5.011	< 0.001 ***
NPI x Structure	-8.859	2.889	2920.985	-3.067	0.002 **

The results for region 1 which corresponds to the subordinate verb (i.e. the verb of the relative clause) can be observed in Table 31. We found a main effect of Structure ($t = 5.011$, $p < 0.001$) which shows that RRC conditions in general were read much slower than FRC in general regardless of NPI presence. More importantly, we found an interaction between NPI and Structure ($t = -3.067$, $p = 0.002$). This interaction shows that the RRC without an NPI is statistically different than all other three conditions, and that, not having an NPI inside the relative clause has a direct impact in the processing of the sentence at this region. While an NPI inside the overt relative clause increases processing times at the subordinate verb compared to processing times at this region in FRC without the NPI, we see the opposite effect inside the reduced relative clause where the NPI condition clearly maintains lower reaction times at the verb than a reduced relative clause that does not contain an NPI. In order to better understand this interaction, we looked at paired comparisons between the experimental conditions.

Table 32: Paired Comparisons Region 1 (Raw RTs)

Region 1: CW+1 (Sub. Verb) by NPI presence						Region 1: CW+1 (Sub. Verb) by Structure					
	Estimate	SE	df	z ratio	p value		Estimate	SE	df	z ratio	p value
No:						FRC:					
FRC - RRC	-57.7	9.85	Inf	-5.857	< 0.001	NO - YES	-11.2	8.13	Inf	-1.378	0.168
Yes:						RRC:					
FRC - RRC	-22.3	9.85	Inf	-2.261	0.024	NO - YES	24.2	8.21	Inf	2.953	0.003

Paired comparisons at the subordinate verb (table 32) show a significant difference between both RRC conditions (NO-YES: $z = 2.953$, $p = 0.003$), which demonstrates that the NPI presence

allows for faster reading times at the subordinate verb compared to sentences of these type that did not have an NPI. The lack of difference between both FRC (NO-YES: $z = -1.378$, $p = 0.168$) conditions at the verb shows that while the NPI may present increased reading times compared to not having an NPI, this increment does not negatively affect processing at the verb, but also that it does not exercise the same facilitation that we see in the reduced relative clause. Furthermore, as expected, significant differences observed in both NPI and No NPI conditions between FRC and RRC evidence that reduced relative clauses are generally more difficult to process than full relative clauses. If increased reading times at the subordinate verb indicate longer consideration by the parser of possible verb interpretations and syntactic structures, the inclusion of the NPI clearly reduces the time it takes the parser to make these decisions by either removing ambiguity or by driving the parser to prefer one interpretation over the other.

From these results, we could assume that, on one hand, presence of the relative pronoun and the copula, which make the relative clause overt, present a much stronger facilitation effect in terms of recognising the structural position of the subordinate verb and its interpretation than the NPI does especially when both are present. On the other hand, that in order for an NPI to provide significant facilitation at the subordinate verb, it must be inside the reduced relative clause (and potentially, licensed). Additionally, we can also assume that the increased reading times shown by the NPI in the full relative clause probably represent additional requirements of processing related to NPI licensing, but that these do not significantly affect processing at the verb. The pattern found for the NPI cases is consistent with the results of experiment 7.

In this case, these effects did not extend into the following region where we only found a significant effect of structure (Estimate = 12.406, Std. Error = 2.325, $t = 5.335$, $p < 0.001$) caused by increased reading times for the RRC conditions compared to the FRCs. This probably shows that the NPI effects on the RRC are limited to the reduction of times on the processing of the subordinate verb, and may be short-lived.

Regions 5 – 7: Garden path effects

Regions 5 through 7 correspond to the areas where we expect to find the garden path effect. Both FRCs regardless of the NPI presence are well-formed, unambiguous sentences that should not lead us down the garden path. These conditions should serve as baseline for their RRC counterparts in determining the appearance of the garden path effect. Because the No-NPI + RRC sentence is a classical garden path effect sentence we should expect to find the effect when compared to its FRC counterpart. A similar comparison between the NPI + RRC and the NPI + FRC should be possible in order for us to determine whether garden path effects are reduced with the early introduction of the NPI *ever*.

Region 5: Disambiguation

Table 33: Region 5: Disambiguation – mixed effects modelling

Fixed Effects	Estimate	Std. Error	df	t value	p value
(Intercept)	375.910	13.964	75.011	26.919	< 0.001 ***
NPI	-3.136	2.519	2992.595	-1.245	0.213
Structure	6.347	2.533	2997.781	2.506	0.012 *
NPI x Structure	-5.110	2.519	2992.593	-2.028	0.043 *

In region 5 corresponding to the disambiguation region (table 33), we found, in this case a main effect of Structure ($t = 2.125$, $p = 0.034$), but no main effect of NPI presence. More importantly, we found an interaction of NPI x Structure ($t = -2.028$, 0.043) that shows that reaction times were modulated by both the presence of the NPI, as well as the type of relative clause which provides support to the hypothesis that the NPI help mitigate the garden path effect. In order to better understand the effects we found, and corroborate that effectively the NPI had a positive effect on the processing of the RRC condition, we looked at the paired conditions.

Table 34: Paired Comparisons Region 5

Region 5: Disambiguation by NPI presence						Region 5: Disambiguation by Structure					
	Estimate	SE	df	z ratio	p value		Estimate	SE	df	z ratio	p value
No:						FRC:					
FRC - RRC	-22.91	7.15	Inf	-3.206	0.001	NO - YES	-3.95	7.13	Inf	-0.554	0.580
Yes:						RRC:					
FRC - RRC	-2.48	7.14	Inf	-0.347	0.729	NO - YES	16.49	7.13	Inf	2.314	0.021

Paired comparisons for the disambiguation region (table 34) showed no difference between both full relative clauses, a result that is expected since both are grammatical sentences that should not elicit a garden-path. Meanwhile, we did find a significant difference between both No-NPI conditions ($z = -3.206$, $p = 0.001$), where the RRC condition was significantly slower at disambiguation compared to the FRC, a result consistent with the expected garden-path effect. More importantly, we found that both RRCs conditions were significantly different ($z = 2.314$, $p = 0.021$) where the condition that contained the NPI was processed faster than the one that did not have an NPI. Additionally, both NPI conditions did not show any statistical difference between them regardless of the structure though we would expect the RRC to produce a garden path effect.

These results confirm the interaction that we saw on the mixed effects model showing that while the reduced relative clause without an NPI clearly displays a garden path effect, the presence of the NPI was able to produce a significant reduction of the effect. Because we found a main effect of structure, it seems that despite this mitigation of the garden path effect by the (licensed) NPI, reduced relative clauses still increase processing at this region compared to sentences with a full relative clause. This probably indicates that while the parser was able to receive facilitation by the NPI in deciding on the nature of the subordinate verb and locate it within a relative clause, the parser still engages in disambiguation processes at this point, even if just confirming its previous decision.

Region 6: Negation

Table 35: Region 6: Negation – mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	362.748	13.330	73.502	27.212	< 0.001 ***
NPI	0.405	2.607	70.987	0.155	0.877
Structure	5.113	2.279	2995.031	2.244	0.025 *
NPI x Str	-3.619	2.271	2993.182	-1.593	0.111

In region 6 (table 35) we continue to observe a significant effect of structure ($t = 2.244$, $p = 0.025$). This time, we didn't find an interaction. While the NPI conditions still seem to remain in an intermediate state, after the disambiguation region, reaction times for the reduced relative clause without an NPI have dropped drastically after the disambiguation region showing that, probably, repair has already happened in the previous region. Paired comparisons only showed a significant difference between the No NPI conditions (FRC-RRC: estimate = -17.46 , std. error = 6.39 , $z = -2.734$, $p = 0.006$), but no longer a difference between both RRC regardless of NPI presence.

The analysis did not find any further main effects of either NPI or structure, or any interactions in the following region (main verb). It is possible that at this point, all processes related to the appearance of garden path effects may have already been resolved.

4.3.2.6 Discussion

Effects of the NPI in terms of the relative clause

The results obtained for experiment 8 demonstrate that, at least in the case of licensed NPIs, our previous hypothesis that NPIs facilitate structural detection / verb tense selection inside the relative clause seem to be correct. Throughout the entirety relative clause we saw that all conditions were processed with certain ease though the overt relative clause conditions (i.e. the baseline and the baseline + NPI) were the faster ones. This was of course an expected result since the header elements of the relative (the relative pronoun *who* and the copula *was*) clearly

established the start of the RC which in turn facilitates the selection of the verb form that follows regardless of the presence of an NPI.

For the reduced relative clauses the situation was much more interesting. We predicted that the No-NPI condition should be easy to process since the main verb interpretation at this point is totally valid. Nevertheless, we saw that this was the slowest condition of the experiment at the verb and this accounts to the main difference found between both RRCs throughout this segment of the sentence. This shows that in sentences with the universal quantifier *every*, the inclusion of an NPI as the head of a reduced relative clause does not only seem to help with determining the structural domain in which the NPI is located, but it also seems to facilitate the processing of the verb in contrast with not having an NPI. This could explain why for the reduced relative clause after the NPI we see barely any variation in reading times in the regions following the NPI. We discuss the possible explanations in the general discussion.

The garden path effect

In terms of the garden path effects in experiment 8, we see, for starters, that the universal quantifier version of the classic garden path sentence (No-NPI + RRC), as expected, does elicit a garden path effect at disambiguation, and maintains a significant difference compared to the baseline (No-NPI + FRC) at the following region. This sets up an important comparison criterion for the appearance, or not, of the garden path effect in the critical condition (NPI + RRC): we have confirmation that garden path effects are prevalent in sentences that contain the universal quantifier *every*, and that, any effects related to the diminishing of the effect should be attributed to the presence of the NPI itself.

Given that we found an interaction between the relative clause type and the NPI, and that the reduced relative clause with the NPI was significantly different from its analogous RRC without an NPI and not statistically different from its baseline (NPI + FRC), we can conclude that the experiment yielded no garden path effects for the reduced relative clause condition that

contained an NPI, or at least that the NPI was able to reduce the effect significantly. This results are consistent with the hypothesis that the presence of the NPI facilitates the identification of the reduced relative clause earlier during processing, and that, furthermore, allows the parser to swiftly and correctly make a choice at the subordinate verb that tips the scales towards a higher ranking on the past participle interpretation, rather than as a main verb (simple past tense) of the sentence. In general, the inclusion of a licensed NPI seems to provide the parser with enough information early on the sentence to mitigate garden path effects that may otherwise appear during sentence processing, and that the effect of the NPI can have lasting effects that are visible at later regions of the sentence.

While these results provide evidence on licensed NPIs, we still do not know whether the NPI presence itself, regardless of it being within a licensing context, is enough to determine the clausal structure of its location, much less whether its presence is enough for the parser to prefer a past participle interpretation of the subordinate verb, or whether these are dependent on it being licensed. For this reason we must look at what happens when the NPI is encountered in the RRC of a non-licensing environment (in sentences that start with the definite determiner *the*).

This experiment provides evidence to support the hypothesis that not only the universal quantifier was necessary for the facilitation effects found in the RRC when the NPI is located within it, but also shows that the NPI is capable to provide clues necessary for the identification of the structure type that eventually override the appearance of the garden path effect. Since we did not see any difficulties in the licensing of the NPI at early stages, we can conclude that the parser was also able to successfully determine that the NPI was located inside a valid downward-entailing environment, or that it was in some measure able to accept it as licensed. In terms of how the parser manages this, we believe that both the universal quantifier and the NPI are important factors. On one hand, the universal quantifier *every* provides the parser with

information about an upcoming downward entailing environment (the restriction), and the NPI provides clues to the parser that at the point it appears the context should be downward-entailing, and that that information might be integrated at this moment.

4.3.3 Experiment 9: NPI presence manipulation on the Determiner

In order to more thoroughly observe the effects found in experiment 7, we decided to also include a variation of experiment 8 in which the sentences start with the definite determiner *the* instead of the universal quantifier *every*. While in experiment 8 all sentences were grammatical due to *every* licensing the NPI, in experiment 9 both of the NPI conditions are, at least initially, ungrammatical. This means that upon encountering the NPI *ever* in the sentence, there is no previously established licenser for it, and that participants will have to wait until the appearance of *not* after the disambiguation region in order to find a plausible grammatical licenser. In experiment 7 we were able to observe that the presence of a licensed NPI seems to facilitate the interpretation of the RRC and reduce garden path effects. Experiment 8 also seems to show some facilitation effects for the NPI + RRC condition at the disambiguation region. In this case, we observe whether the presence of the NPI by itself is enough to produce a facilitation effect, or whether the NPI must be previously licensed in order to do so.

4.3.3.1 Participants

Once again, we recruited 72 participants through the Prolific webpage (<https://www.prolific.co/>). The task was advertised as a reading task and participants were given instructions on the experimental procedure using the same instructions as in the previous experiments. Prolific filters were used to ensure that all participants were native speakers of English and we included a questionnaire in the experiment about their linguistic background. Each participant was paid £6.75 for their participation in the study which lasted approximately 30-45 minutes.

4.3.3.2 Materials

The materials used in experiment 9 were the same materials that we used in experiment 8. The only difference was that in this case all experimental items started with the determiner *the* rather than the universal quantifier *every*. Same as before, the conditions were formed by manipulation of the sentence structure by having either a full relative clause (FRC) presented by a relative pronoun *who* and the copula *was*, or a reduced relative clause (RRC) that was introduced without these two elements, making the following verb ambiguous. We also manipulating the presence of NPI by either placing *ever* before the subordinate verb (NPI = Yes), or not (NPI = No). An example of the stimuli can be observed in Table 36 Each participant saw a total of 120 sentences that included 44 experimental items and 76 fillers.

Table 36: Experiment 9 example stimuli

NPI	RC Type	Sentence
NO	FRC	The researcher who was awarded a scientific prize was not invited to the conference.
NO	RRC	The researcher awarded a scientific prize was not invited to the conference.
YES	FRC	The researcher who was ever awarded a scientific prize was not invited to the conference.
YES	RRC	The researcher ever awarded a scientific prize was not invited to the conference.

In this case, the No-NPI + FRC sentence is a well-formed, grammatical sentence that should not lead down the garden path, while the No-NPI + RRC sentence is a classical garden path effect sentence. Both of these sentences should serve as a baseline for determining the effects the NPI has over the garden path effect of the subsequent regions.

4.3.3.3 Procedure

We used the same procedure in experiment 9 that we used in both experiment 7 and 8.

4.3.3.4 Data Analysis

Data was analysed using the same methods used in both experiment 7 and 8. Contrasts coding was done following the same method described in experiment 8. No participants were removed for either inaccurate answers or extreme reading times, and only 0.61% data was lost after removal of extreme RTs.

4.3.3.5 Results




Figure 14: Experiment 8 Reading Times by Region

Figure 14 shows the plot corresponding to the self-paced reading results for all 4 experimental conditions of experiment 9. The graph shows each region in a word-by-word manner on the x-axis and their corresponding mean reading times for each region in the y-axis. For each region we calculated the linear mixed effects models and the paired comparisons separated by both Structure and NPI presence to better observe the differences between each condition. Following the same criteria that we used for experiments 7 and 8, we will look at the data based on two main areas. First, we will look at the area corresponding to the inside of the relative clause to see if there are any immediate effects of the NPI in its following regions that may give clues to whether participants take consideration of the context in which the NPI is located even if the NPI is not immediately licensed, and second, we look at the garden path effects after disambiguation to see whether the presence of an NPI had any reducing effects on the garden path.

Regions 0-4: NPI modulation effect inside the relative clause

Region 0: NPI ever

Because we modulated the appearance of *ever*, only half of the items contained an NPI. For the NPI conditions we found a significant main effect of Structure (Estimate = 20.141, Std. Error = 5.278, $t = 3.816$, $p < 0.001$). This effect is consistent with previous experiments.

Region 1: Critical Word +1 (Subordinate Verb)

Table 37: Region 1: CW+1 – mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	444.423	21.434	73.773	20.734	< 0.001 ***
NPI	23.213	6.324	41.977	3.670	0.001 ***
Structure	6.453	5.284	146.825	1.221	0.224
NPI x Str	0.251	4.664	2947.601	0.054	0.957

Linear mixed effects results (table 37) show a main effect of NPI ($t = 3.670$, $p = 0.001$) which demonstrates that conditions with an NPI were read much slower than regions that didn't contain an NPI regardless of relative clause type. This is expected since in these cases none of the NPIs are licensed under any condition, and the addition of an unlicensed NPI seems to have a great effect on the processing of the sentence at the subordinate verb. We did not find any other effects at this point. The lack of a difference between both No-NPI cases is expected since the parser has no information about the verb's structure and so far both sentences can be considered grammatical with a main verb interpretation. In general, this region shows the massive cost in processing of an unlicensed NPI.

Region 2: CW +2

Table 38: Region 2: CW+2 - mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	422.382	15.654	73.944	26.982	< 0.001 ***
NPI	17.847	5.353	52.898	3.334	0.002 **
Structure	4.596	3.587	2951.659	1.281	0.200
NPI x Str	-4.024	3.587	2952.427	-1.122	0.262

Results for region 2, the word following the subordinate verb (table 38), continue showing a strong main effect of NPI presence ($t = 3.334, p = 0.002$) where both NPI conditions are generally slower than the No-NPI conditions. No other effects are present. The same significant effect for NPI presence continues over the course of the entire relative clause: Region 3 ($t = 3.603, p < 0.001$) and Region 4 ($t = 2.459, p = 0.016$). This shows that the cost of having an unlicensed NPI endures throughout the processing of the sentence regardless of the type of relative clause. Contrary to the results of the previous experiment, the NPI shows no facilitations effects at any point and neither does it seem to be a recovery at any point so far.

Garden Path: Regions 5 (Disambiguation), 6 (Negation), and 7 (Main Verb)

In these regions, the contrast that we are looking for are whether the NPI conditions differ from both No-NPI conditions. These two conditions serve as a baseline to both ends of the spectrum: the No-NPI + FRC conditions is a perfectly formed, grammatical sentence that does not lead down a garden path; the No-NPI + RRC condition, on the other hand, represents a perfectly formed, grammatical sentence that is a classic example of a garden path sentence.

Region 5: Disambiguation

Table 39: Region 5: Disambiguation - mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	427.106	18.920	73.890	22.574	< 0.001 ***
NPI	13.278	5.551	60.907	2.392	0.020 *
Structure	14.720	5.463	60.779	2.695	0.009 **
NPI x Str	-5.762	5.269	52.135	-1.093	0.279

Linear mixed effects modelling for the disambiguation (table 39) shows both a main effect of NPI presence ($t = 2.393, p = 0.020$) and a main effect of Structure ($t = 2.695, p = 0.009$). This shows that both factors had an effect on processing times by the moment in participants reached disambiguation. Regardless, we did not find an interaction between both variables. These results show that both NPI conditions were processed much slower than the baseline, similar to what has been happening throughout the sentence since the appearance of the NPI. At the same time,

the reduced relative clause without an NPI is also significantly slower than the baseline. These results show that all three conditions presented great increase in reaction times relative to the No-NPI x FRC condition. These results show that both RRC conditions present a garden path effect. Increased reading times for the full relative clause with the NPI cannot be attributed to the garden path effect, but clearly the NPI by itself is enough to cause similar difficulties in processing, even at this distance from the NPI a result that was not found in experiment 8 where the NPI was initially licensed. Interestingly, this shows that despite seeing fleeting NPI effects a couple of regions after its appearance in previous research, its presence is still affecting the processing demands of the sentence several words away. Additionally, given that the high RTs of the NPI+FRC cannot be attributed to a garden path effect it is difficult to determine whether the increased demands of the NPI+RRC condition are entirely due to a garden path effect or to general difficulty of the NPI at this point. Regardless, given that the general RTs of the RRC case are slightly higher than the FRC, we can imagine that the increase in reading times is due to penalisation from both effects.

Results found in this region seem to be consistent with the data observed in experiment 7 since, as we saw, the determiner + reduced relative clause condition (along with the determiner + full relative clause condition) were the slowest at this region.

Region 6: Negation

Table 40: Region 6: Negation - mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	393.562	14.231	70.951	27.655	< 0.001 ***
NPI	3.933	3.728	71.172	1.055	0.295
Structure	14.526	3.133	71.248	4.637	< 0.001 ***
NPI x Str	-9.285	2.652	2932.692	-3.501	< 0.001 ***

Linear mixed effects results (table 40) for this region display a main effect of Structure ($t = 4.637$, $p < 0.001$) which shows that RRCs are processed slower than FRCs. In this case, there was no main effect of NPI presence. In addition, we now see a clear interaction of both NPI and Structure

($t = -3.501, p < 0.001$) which seems to be driven by the difference between the baseline condition (No-NPI + FRC) with the other three conditions. Simply put, the baseline condition is much easier to process than all three other conditions. The interaction is probably due to the fact that at this point both NPI conditions have fallen below the No-NPI + RRC condition.

Table 41: Region 6 Paired Comparisons

Region 6: Negation by NPI Presence						Region 6: Negation by Structure					
	Estimate	SE	df	z ratio	p value		Estimate	SE	df	z ratio	p value
No:						FRC:					
FRC - RRC	-47.6	8.20	Inf	-5.805	< 0.001	NO - YES	-26.4	9.15	Inf	-2.888	0.004
Yes:						RRC:					
FRC - RRC	-10.5	8.21	Inf	-1.276	0.202	NO - YES	10.7	9.15	Inf	1.170	0.242

Planned comparisons (table 41) shows a significant difference between both No NPI conditions (FRC – RRC: $z = -5.805, p < 0.001$) where the reduced relative clause is slower than the full relative clause, but we see no significant difference between both NPI conditions. Similarly, we see a significant difference between both FRC (No-NPI – NPI: $z = -2.888, p = 0.004$) where the condition with the NPI is significantly slower than the one without one. These results show that while at this region the effects from NPI have improved a little, they are still much more demanding than the baseline, and that the garden path effect is still strongly sustained.

Region 7: Main Verb

Table 42: Region 7: MV - mixed effects modelling

	Estimate	Std. Error	df	t value	p value
(Intercept)	389.177	16.294	74.270	23.885	< 0.001 ***
NPI	2.476	3.307	55.521	0.749	0.457
Structure	13.685	2.775	3046.444	4.932	< 0.001 ***
NPI x Str	-1.593	2.772	3037.425	-0.575	0.566

By the time we reach the main verb (table 64), the interaction is no longer present and we only see a main effect of structure ($t = 4.932, p < 0.001$). The current results seem to show that though the NPI presence effect no longer seems to be sustained, the garden-path effect is still prevalent. This effect is no longer present in any of the regions following the main verb. It is probable that the main verb of the sentence serves as a threshold for identification of garden

path effects and at this point the parser has been able to resolve any structural ambiguity. It is possible that unlicensed NPI effects also have a similar threshold and are no longer considered at or after the main verb of the sentence.

4.3.3.6 Discussion

The results of experiment 9 reveal several key elements of the NPI effects that carry over through the entire sentence, in this case, when it is not licensed. We have separated the discussion into two sections, same as with experiment 8. We first look at the effects inside the relative clause, and then we discuss the repercussions this has over the garden path effect.

Effects of the NPI in terms of the relative clause

At this area of the sentence we did not find any signs of facilitation from the NPI throughout the relative clause. Quite the contrary, both of the NPI versions of the conditions were in general the slowest overall regardless of the structure after the appearance of the NPI. Crucially, this shows the massive cost on reading times of having an unlicensed NPI throughout the entire phrase, which is even greater in a reduced relative clause. These results demonstrate that the NPI did not contribute any facilitation effects and that the parser even had to spend more resources trying to make sense of the subordinate verb itself, assuming that the peak observed in reaction times is caused mainly due to morphological and structural decisions of parser at the verb. The difference between RTs in regions 1 and 2 seem to be caused by the parser having reached a decision about the verb's morphology since we see this reduction across all conditions right after the verb. Meanwhile, the fact that all NPI conditions continue to have increased reaction times compared to the No-NPI cases probably points to sustained effects of the NPI that persist throughout the sentence until the disambiguation areas. Once verb decision effects are resolved, the parser continues to search for a solution to the unlicensed NPI and it is possible that the parser is still considering the ongoing context to integrate into the licensing process. Given that there is still no potential licenser, the parser may still keep an open search for

upcoming information and this is why we do not see any difference between both NPI conditions until after disambiguation regions. Whether the disappearance of the NPI effect at this point points towards a potential resolution of the NPI or a threshold for the process is not entirely clear.

The baseline condition (No-NPI + FRC) was in general read faster than all other conditions and even the classic garden path sentence (No-NPI + RRC) was, at the beginning of the phrase read slightly slower though not significantly so, which indicates that while the overt relative clause may allow the parser to make a more informed decision at the subordinate verb, both interpretations of the verb are considered grammatical and, therefore, the parser does not find a general increase in difficulty processing these sentences until disambiguation is reached. The decrease in speed for the baseline could be an effect of having the relative pronoun and the copula in preceding positions, since these are short words length effects may be a factor to consider as accelerators. Notwithstanding, there is also a possibility that upon encountering the ambiguous verb, the parser initially takes slightly more time to consider the alternatives when deciding how to generate the syntactic structure at the verb.

The garden path effect

The first thing to note at the disambiguation area is that the classic garden path sentence presents the expected garden path effect in comparison to the baseline sentence. In this experiment, though, the classic garden path sentence (No-NPI + RRC) is not the slowest, but the NPI + RRC one, with the NPI + FRC being very close to the classic garden path sentence. Considering that a FRC should not produce a garden path effect, it is not possible to attribute this effect to a garden path, but rather to the parser still trying to make sense of the NPI. Since encountering the NPI the general reading times of each region have been high for both the FRC and the RRC conditions and it is not until after this region that we observe any steep decrease of reading times. It is, perhaps, at these disambiguation areas that other types of information (such

as semantic and pragmatic information) start being consolidated and why we see that difficulties from the unlicensed the NPI are carried over up to this point and only just now start to decrease while structural effects start taking over. Whether the steep decrease in reading time that we see at the negation region is due to a potential licensing of the NPI by the negative element (which could function as potential licenser) or whether after the appearance of the auxiliary verb the parser just gives up on licensing the NPI still requires further research.

Since reaching the disambiguation region, past the negative region and into the main verb, we have seen a steady trend in which reading times for the NPI conditions have become faster. Nevertheless, the rate of acceleration has been different for both types of structure. While in the FRC the reading times have reached similar speeds as the baseline, in the reduced relative clause this did not happen, and, instead, we find a persistent garden-path effect similar to the one in the classic garden path sentence. This means that, if the NPI was able to be somewhat resolved after the appearance of the disambiguation or the negation this advantage was only readily available when *ever* was located inside the full relative clause and not inside the RRC. Conversely, if the licensing of the NPI was simply abandoned after a threshold determined between the auxiliary and the main verb of the sentence, the difficulties observed are solely determined by the restructuring demands of the garden-path effect.

Compared to the results found in experiment 8, it seems that, as we predicted before, the early licensing of the NPI seems to be crucial for it to contribute to the reduction of the garden path effect at or after the disambiguation region.

4.4 General Discussion

In the current chapter we explored the hypothesis that NPIs can help predict the upcoming structure of the sentence by facilitating the correct identification of the clause that contains it as a relative clause and, thus, later on help prevent the appearance of garden path effects. We

assumed that the NPI, in order to be able to reduce the garden path effect, must comply with its licensing requirements. Crucially, that the NPI must be able to find an appropriate and valid licenser, in this case the universal quantifier *every*, preceding it, and also that the NPI must be identified to be within the appropriate structure for *every* to properly license it, in other words, within the restrictor clause of the universal quantifier which is a DE environment.

This is not entirely different from what is proposed by Clifton and Frazier (2010) in the sense that the parser must be able to correctly identify the monotonicity of the context in which the NPI is located (DE) locally. We do however look at this issue from a different perspective. In their research they used temporarily unlicensed NPIs that were globally licensed (by the negation *won't* inside the main clause) to show that local computations that determined that the NPI was unlicensed had an effect on judgment ratings despite global licensing (experiment 1). In this sense, reduced ratings are attributed to the effect of a necessary reanalysis of the NPI after the negation inside the main clause is found, which affects the NPI licensing globally. Given that computations are done locally at the NPI which determines it to be unlicensed, the appearance of the negation requires the parser to reanalyse the NPI licensing again producing a garden path type effect, thus reconsidering the original analysis. Furthermore, in their second experiment, they looked at NPI licensing by using a negative quantifier to show that the effect found on their first experiment disappears when the NPI is immediately licensed. The sentences that we used in our research are very similar to the ones used by Clifton and Frazier with some important differences. For one, instead of relying on judgment ratings of seemingly unlicensed NPIs that are licensed globally, our experiments relied on the correct identification of a licensed NPI to determine the type of clause that the NPI is located in and thus the correct identification of the verb form inside the RRC. In this way, our measure, rather than reduced acceptability ratings are seen as a reduction of the garden path effect that should be expected at the disambiguation regions of the main clause if the parser was not able to correctly identify the correct form of the verb (or if, for that matter, considered the NPI to be outside of its intended structural and

semantic context). Additionally, the use of the universal quantifier *every* instead of the negative quantifier *no* in our experiments adds a dimension that is not entirely apparent in Clifton and Frazier's experiments. While the negative quantifier licenses NPIs inside both its restrictor and its scope, universal quantifiers can only license NPIs inside their restrictor and not inside their scope. This difference between both studies means that in order for our experiments to show the desired effect we were looking for, the parser must not just assume that the universal quantifier is able to license the NPI but must correctly identify the correct location of the NPI locally. That is to say that the NPI is located inside the correct clausal location and thus, that the monotonicity of the context is correctly identified as DE.

In our first experiment we manipulated the potential licenser for the NPI between a valid licenser (the universal quantifier *every*) and the definite determiner *the*; and we manipulated the way in which the restrictor of the quantifier was presented as either an overt relative clause or a reduced one which would elicit the desired garden path effects at the disambiguation regions of the sentence. We left the presence of the NPI constant across all conditions. In order to more precisely observe the effects of the presence of the NPI in these constructions we also included two variations of the first experiment where we manipulated the NPI presence and the form of the relative clause while leaving the potential licenser constant. One experiment with the universal quantifier *every* and the other with the definite determiner *the*. Then we analysed these experiments based on two main areas of interest: the garden path effect and the effects found within the relative clause where the NPI was found. Here we explore our findings across all three experiments by looking at the effects found under these two main areas.

4.4.1 The garden path effect:

Our main area of interest in this chapter was to observe whether the appearance of an NPI early in the sentence could have an effect dissipating garden-path effects. Across our experiments we found confirmation that the NPI can modulate the appearance of the garden path effect, but that in order for this to happen, the NPI must be licensed early in the sentence, and more

importantly, that the parser must be able to promptly recognise this by determining that the NPI is within its desired DE context, locally.

In general, our experiments show that at and after disambiguation sentences that had a reduced relative clause that contained an NPI were able to significantly reduce garden path effects when they were licensed by the universal quantifier *every*, but not when the NPI was unlicensed (preceded by the definite article *the*). Since experiment 8 allowed us to prove that garden path effects occur naturally in sentences that have the universal quantifier without an NPI, we can be sure that the reduction of the garden path effect observed at the disambiguation areas in experiment 8 were in fact a product of the presence of the NPI. Additionally, the results observed from the definite determiner from both experiment 7 and 9 show that the presence of an unlicensed NPI not only does not reduce the garden path effect but is also detrimental for the processing of the whole sentence, including the disambiguation areas regardless of the structure of the relative clause that contains the NPI. This allowed us to prove that the garden path reducing effects found with the universal quantifier are not due exclusively to the presence or interference of the NPI itself, but by the presence of a *licensed* NPI.

For cases with the definite article, experiment 9 also shows that after the negation region the full relative clause with an unlicensed NPI presents a greater decrease in reaction times when compared to both the reduced relative clause with and without an NPI, the it is no longer significantly different from the baseline. While we may suggest that the negation may work as a late secondary licenser for the NPI or have an interference effect, we cannot be sure. If the detrimental effects of the NPI still have an impact in the sentence at this point, this could explain why reaction times for both NPI cases with the determiner drop under the classic garden path sentence, but not why it would continue to have an effect for the NPI + FRC over the following region, but not for the NPI + RRC which actually shows an increase. While a possibility could be that the parser takes temporarily consideration of the negation, thus the RRC drops slightly, it

promptly abandons this possibility, it seems doubtful. It is more likely that this is a spurious effect. A more likely possibility is that NPI effects at this point no longer have a strong enough effect to greatly affect processing, and that any increased processing requirements are entirely due to the garden path effect and both RRCs are now being processed with similar difficulty.

4.4.2 Effects within the relative clause

By looking at the results from all experiments, we believe that our hypothesis that a licensed NPI facilitates the correct identification of a reduced relative clause seems to be correct. Furthermore, we find evidence that when the sentences contains a valid licenser (such as the universal quantifier *every*) and an NPI the parser has no difficulties with either identifying the NPI as being inside a reduced relative clause, nor identifying the correct form of the verb. This is evidenced by the mitigation of the garden path effect later in the sentence. Crucially, we believe that the correct identification of the restrictor clause of the universal quantifier means that the parser was able to correctly assume that the NPI was located inside its required semantic context (DE) and that the appearance of the NPI may be at least partially responsible for this for this computation to be done locally. While the universal quantifier provides information about the upcoming contexts, the NPI, which requires to be licensed under a downward entailing context, provides evidence that the context should be the desired one. Considering that the parser will try license the NPI and that it already has a DE context providing element, it could assume that the NPI must be licensed within its required context, thus applying monotonic computations locally and determine that the context satisfies the licensing requirements.

Furthermore, we believe that the NPI may also facilitate certain processes which are solely related to the subordinate verb in the relative clause and which are probably involved in the selection of the proper verb tense (past participle vs simple past forms) and therefore, the parsers interpretation of the upcoming structure. What we observe at the verb area, across both NPI modulation experiments (and in some measure in the original experiment) is that most conditions present large peaks at the subordinate verb region (CW + 1) and that these peaks

precipitate at the next couple of regions. We believe that these represent the parser's verb selection mechanisms which explains why they are most pronounced on more difficult or ambiguous conditions, compared to less ambiguous conditions. Since it seems that any peak effects that appear at the verb in the relative clauses are then resolved by the following word and the slower sentences (regardless of the origin of the difficulty) continue through the sentence with similar reading times, we can assume that this shows that the parser has encountered the verb and made an interpretation. This could explain why we see very little variation in reading times in the reduced relative clause + NPI with the universal quantifier. While the appearance of the NPI can produce an increase in processing times related to the licensing process, as evidenced by most NPI conditions across all three experiments (with much higher peaks for ungrammatical NPIs), in both experiment 7 and 8, we see that the universal quantifier + RRC conditions will only increase slightly when there is an NPI compared to when there is no NPI (experiment 8) or when the NPI is licensed (experiment 7).

There is a couple of potential explanations for this. On one hand, it is possible that the presence of the universal quantifier *every* at the beginning of the sentence predicts the upcoming appearance of a contrasting set and that a restriction of the universal quantifier phrase should be present in the sentence, especially since we don't have other elements that further constrict the presupposed quantified set (Ni, Crain, & Shankweiler, 1996 explain this for the operator *only*). This explanation is consistent with the referential theory in that, similar to how *only* requires a contrasting set based on its referential properties, the universal quantifier also presupposes the existence of a contrasting set based on its referential properties. If some sort of restriction is expected due to the universal quantifier at the beginning of the sentence we could see that the reading of the verb is much easily interpreted as past participle so that it would fit with a restriction. If this was solely a prediction done based on the universal quantifier, though, we should expect to see at least some similar effects happening in the No-NPI case given that a restriction does not necessarily require an NPI. Perhaps we could consider adding a

condition with an adjective in future studies to see if we can get similar results as those of Ni, Crain, and Shankweiler (1996) where the adjective was enough to act as a contrasting set and allow the reduced relative clause to be interpreted in an ambiguous manner.

A second possibility is that this difference in reading times is due to considerations done at the verb in order to select an interpretation of its tense. Given that all the verbs that we used here were ambiguous, it is possible that the parser recognises these alternatives and takes some time in order to make a decision. This approach is consistent with the constraint satisfaction model that states that the online decisions made by the parser are affected by a range of different non-structural sources of information (Trueswell, Tanenhaus, & Kello, 1993; Tanenhaus & Trueswell, 1995), as well as with the referential theory. Studies into this theory have shown that lexically-based factors such as verb frequency, argument structure, and conceptual-semantic information are important factors taken into account when evaluating alternative structural representations in real time (Boland, Tanenhaus, & Garnsey, 1990, Juliano & Tanenhaus, 1993, 1994; MacDonald, 1994; Vasishth, Drenhaus, & Srinivasan, 2010; Sturt, Haan, & Steinhauer, 2014; Wagers & Phillips, 2014). Given the differences in RTs between both NPI and No-NPI conditions we know that the universal quantifier itself is not the reason why we see facilitation of one alternative over the other, but is an important factor that allows the parser to license the NPI. This would also explain why the difference in reading times is so short lived and why by the time we reach the word following the verb the RTs between both RRCs have become virtually the same: the processes initiated at the subordinate verb are primarily the parser deciding on the upcoming structure with the information available to it. Considering that reading times after the verb remain consistent between both reduced relative clauses regardless of an NPI all the way up until the disambiguation, this explanation works out. Furthermore, this also explains why we don't see the same increase in the NPI condition. The appearance of the NPI facilitates identification of the reduced relative clause (i.e. starting boundaries of the restrictor clause) by the parser, and in turn, this allows the parser to use the NPIs structural information to swiftly make a decision

regarding the nature of the verb, tipping the scales towards a correct interpretation of the verb as past participle.

Importantly, for this to happen, the parser must assume that the NPI is licensed the moment that it appears and apply the correct monotonicity computations locally. While it could simply do this by the retrieval of the licenser (independently of recognising the structural position of the NPI), this seems unlikely because in that case, the structural information necessary for the parser to make an informed decision about the subordinate verb form would not necessarily be available, and we would not be able to see facilitation or reduction of the garden path effect at disambiguation. Since we do see that the (licensed) NPI mitigates the appearance of the GP effect, we have to assume that the parser is able to identify the correct structure where the NPI is located at the same time it licenses it, and therefore use this information to decide on the correct form of the verb. Consequently, the NPI cannot be licensed by the parser by simply checking for the appropriate licenser, but must also assume that it is located within its appropriate structural licensing context, and perhaps also the semantic context (downward entailing) of this structure which must be done by interpreting the monotonicity of the context locally.

In this sense, the function of the overt relative clause is now provided by the NPI instead and the parser doesn't need to spend extra processing time selecting an interpretation for the verb: while it may still consider the main verb interpretation, thanks to the NPI, the subordinate verb interpretation is most likely the highest ranked interpretation. While the No-NPI conditions still requires the parser to deal with the ambiguity of the verb, in the NPI condition this ambiguity has already been resolved by the preceding NPI.

In experiment 9, compared to experiment 8, we did not find any signs of facilitation from the NPI through the relative clause. Quite the opposite the results suggest massive processing costs reflected on reading times from having an unlicensed NPI. Additionally, it is also an important

distinction compared to the universal quantifier cases that this effect is even greater inside the reduced relative clause. With the universal quantifier we saw that the inclusion of the NPI reduced the processing times of the RRC at the verb compared to when the RRC didn't contain an NPI. In this case we are observing the opposite where the NPI is actually making things even slower at the verb than in the classic garden path effect where we previously saw a large peak in processing.

4.5 Conclusions

The experiments of this chapter provide evidence that the inclusion of a licensed NPI in a garden path sentence can mitigate the garden path effect at the point of disambiguation to the point that the effect can disappear from the sentence. This effect of the NPI seems to be due to effects that happen at the beginning of the sentence when the structural context is analysed at the subordinate verb. Our experiments suggest that an NPI with a universal quantifier (*every*) as a licenser can help provide information that not only allows the parser to correctly assess the current structure in which the NPI is located, but also that this structural information must be available immediately at the moment that the parser licenses the NPI in order to correctly choose the verb tense at the subordinate verb and reduce the time it takes to assess the verbal form in ambiguous verbs. An unlicensed NPI, on the other hand, does not only not facilitate the identification of the structural context or the verbal form, but also has detrimental effects that can be seen through the entire sentence past the points of garden path disambiguation. Taken together these results show that the parser continues to attempt to integrate information to the licensing of NPIs in much later regions that we had previously observed in the previous chapter with self-paced reading.

Finally, in our view, our experiments build on the research of Clifton and Frazier (2010) and confirm that monotonicity is, indeed, computed locally. What is more difficult to determine, though, is whether monotonicity is actually computed in a word-by-word manner constantly, or

whether the parser takes cues from specific words inside the sentence to check the appropriate context. One thing that seems evident from these experiments, though, is that if the monotonicity of the context was not correctly computed at the appearance of the NPI, we would not be able to find any evidence that the garden path effect of the sentence was reduced at the point of disambiguation.

Chapter 5: Summary of Results and General Discussion

5.1 Summary

In this dissertation we argue that universal quantifiers such as *every* provide us with a very interesting, yet unexplored vehicle to test how semantic and pragmatic processes are involved in the licensing of negative polarity items in addition to the involvement of memory retrieval mechanism, compared to the more explored negative quantifiers.

While negative quantifiers also provide semantic contexts that license NPIs, they supply a structure where both their restrictor clause and their nuclear scope are downward-entailing, the type of semantic context which licenses NPIs, and thus allow NPIs to appear in either of these structures. Because of this, it is difficult to observe whether the licensing process occurs due to the application of these semantic contexts or through a more direct licensing route based on memory mechanisms where the parser retrieves a [+negative] feature from the licensor upon reaching the NPI.

On the other hand, universal quantifiers do not contain neither lexically or semantically encoded negation and can only license NPIs exclusively through a downward-entailing context provided within their restrictor clause, but not within their nuclear scope. Taken this into account, we believe that universal quantifiers can not only provide evidence for a memory retrieval account, but are also more suitable to explore how the different semantic contexts affect NPI licensing and the rise of illusions, and how the parser integrates semantic and pragmatic information in sentence processing.

Throughout our research we proposed that with the use of universal quantifiers we can see whether illusory licensing occurs with *every*, similarly to how it has been observed with the

negative quantifier *no*, and whether this provides evidence of the memory mechanisms involved in the licensing process.

Furthermore, we propose that the specificity of the restrictor clause can be manipulated by increasing modification to the noun, making the quantified set much more specific, and therefore, making the boundaries of the restrictor clause much more clear. We believe that if the parser is able to correctly identify the universal quantifier's structural context based on semantic information, we can manipulate the appearance of illusion, similar to how Parker and Phillips (2016) did with distance effects, but without having to move the NPI a long distance away from its licenser. If this is true, then the selective profile of NPI illusions could be explained by the prompt integration, or not, of semantic and pragmatic information during sentence processing.

More importantly, we propose that the presence of an NPI helps the parser identify the monotonicity of the context in which it appears, and therefore allows the parser to override preference bias in ambiguous sentences such as garden path sentences. Rather than the parser checking the semantic context based on clues from external elements, or attempting to retrieve a specific feature to license the NPI, it is the NPI itself that provides the semantic information about the context at the time it appears, and the parser then checks for the presence of a context-giving element through a memory retrieval, or similar process.

5.2 Brief review of experimental results

In Chapter 2 we took a look into the effects of quantification in the rise of illusions in NPI licensing by testing for two accounts of illusory licensing using offline speeded acceptability judgments. The first account suggests that NPI licensing is a matter of memory retrieval mechanisms, and that illusory licensing occurs due to a noisy cue-based, content-addressable memory retrieval process. The second account considers that the rise of illusory licensing is driven by the

application of semantic and pragmatic processes which are involved in the licensing of NPIs and do not include memory mechanisms.

In order to test these accounts, we looked at when illusions arise when quantifiers are involved. First we found that NPI illusions of grammaticality arise with the universal quantifier *every* located in a structurally irrelevant position. Additionally, we compared this to the existential quantifier *some* which does not license NPIs to determine whether NPI illusory licensing is due to a property of quantification in general, or specific to certain quantifiers. Given the evident involvement of memory mechanisms in NPI licensing, identifying probable features that could be searched for during the retrieval process is important for any theory that attempts to explain what information is used in the licensing process, and how this can explain illusory licensing from a memory point of view. If non-licensing quantifiers can give rise to illusory licensing effects we can expect a feature such as [+quantification] to be possibly involved. We did not find the same effects of illusory licensing with the existential quantifier.

These results show that illusory licensing is not due to a general property of quantification, but due to a sole property of specific quantifiers, in this case, of the universal quantifier *every*. Additionally, given that retrieval of an irrelevant licensor was achieved, we can conclude that if memory retrieval drives the licensing process, retrieval must be content-addressable, cue-based retrieval mechanisms, otherwise the irrelevant structural positions would not be considered. Furthermore, these results also show that if memory retrieval is involved, a feature such as [+quantification] is not a possibility for retrieval. Also, in two separate experiments we found that modulation of the restriction of the universal quantifier by addition of prenominal or postnominal modifiers may cause a reduction of illusory licensing to the point where the appearance of illusions is no longer significant. We believe that these modifiers could increase the chance of the parser being able to identify the boundaries of the restrictor clause by setting up a more salient contrasting set which precludes the parser from assuming that the NPI is

located within the universal quantifier's restriction. Taken together, these results are consistent with the account that NPI licensing is driven by the application of semantic properties of the sentence, though it is possible that the parser may still resort to memory retrieval when the semantic or pragmatic information of the context is not clear or immediately accessible to integrate into the licensing process.

In Chapter 3, we tried to replicate the results from Chapter 2 using online self-paced reading to see whether the rise of illusory licensing can be detected with online measures in an attempt to determine the time course and location of the effect. Furthermore, we wanted to see whether the application of semantic and pragmatic information provided by the modulation of the restrictor by the addition of prenominal and postnominal modification could also be observed with self-paced reading.

On one hand, we found that illusory licensing can be observed in SPR for the universal quantifier *every*, but not for the existential quantifier *some*, consistent with the observation of experiments 1 and 2 from chapter 2. We observed that these illusions we're observable since the appearance of the NPI and the spill-over region. Given that the effects are still observable in offline measures (experiment 1) we believe that the effects should persist until the sentence has finished processing. On the other hand, experiment 6 was not able to fully replicate the same results found on chapter 2. While we didn't find evidence of illusory licensing in SPR for both modified conditions, similar to what we saw in experiments 3 and 4 where we proposed that defining a contrast set to modulate the restrictor clause with the addition of prenominal and postnominal modification had reduced the effects of illusory licensing in offline methods, we also failed to find evidence of illusory licensing for the unmodified condition, more specifically, we didn't find any significant difference between all universal quantifier conditions. The results of experiment 6 regarding the unmodified condition contrast drastically not only with what we observed in experiments 1 with acceptability judgements, but also what we found in experiment 5 using self-

paced reading, the same method used in experiment 6. A possibility for this is that the effects of modulation of the restrictor clause are not visible at this point in the sentence and may come in play later during processing, or after processing of the whole sentence has finalized, which could be a possible explanation as to why we saw a difference in offline measures.

An alternative that we should consider is that since the results that we got on chapter 2 were product of four independent experiments, they might not be entirely comparable. If that is the case, the difference between modulated and unmodulated universal quantifiers might not be as big as we initially thought. While the universal quantifier does seem to elicit at least some measure of illusory NPI licensing, it might not have a strong enough effect that illusions are prevalent or that they appear in all cases but only sporadically under certain conditions. This could explain why we saw evidence of this in experiments 1 and 5 where the effect just managed to be prevalent enough to appear, while in the other experiments the effect was just not strong enough. In order to get a better idea of what is happening, we decided to run a meta study to compare all four experiments from chapter 2 which will be presented and discussed in the next section.

Future research should consider examining eye movements as a potential way to better examine the time course of these effects. A problem with this approach, though, is that this does not explain why we didn't find an effect for the unmodified universal quantifier as we did in experiment 5. A possible explanation could have something to do with the experimental design itself. Since most of the sentences contained the universal quantifier and neither determiner and negative sentences contained the same variations in modification, participants could have developed predictions sentences they were reading.

In chapter 4 we looked at whether the use of an NPI within a classic garden path sentence can help reduce the garden path effect usually observed at disambiguation, or make it disappear entirely by potentially facilitating the parser's identification of the monotonicity (the direction of the entailing context) of the reduced relative clause that contains it. If this happens, the parser should be able to identify, not only the structure location of the NPI but also have immediate access to this information to correctly interpret the subordinate verb as a past participle, rather than as the main verb of the sentence, thus not being taken by surprise when it reaches the real main verb.

For this purpose, we inspected the results of three self-paced reading experiments in which we examined how the presence or absence of an NPI can affect the ambiguity of the initial reduced relative clause structure where the parser is usually driven down the garden path, in compared to an unambiguous overt relative clause. Furthermore, by comparing between the definite article *the* and the universal quantifier *every*, we observed whether the NPI requires to have a valid licenser preceding it, or whether its presence is enough to affect the parser's decision. This also allowed us to compare between effects that are produced by the NPI itself and those produced by the universal quantifier, and those that require a combination of the two.

The results of the experiments revealed that, first of all, the presence of an NPI embedded inside a reduced relative clause can pre-empt the appearance of garden path effects, but also that in order for the NPI to have this effect, it must be licensed. Second of all, when we looked at the reduced relative clause we saw that in garden path sentences (RRC) that contained the universal quantifier *every* as the head of the sentence, the NPI had a facilitation effect not only on the reading times at the disambiguation regions, but also at the subordinate verb compared to ambiguous sentences without an NPI despite the fact that those that didn't have an NPI should have been easily interpreted at that stage as normal main verb clauses.

These results provide evidence that at the moment in which the parser encounters the NPI, the parser can recognise the contextual environment of the NPI and identify that the NPI is located within a reduced relative. We believe that it is possible that the licensed NPI provides the semantic information at this stage of processing that allows the parser to recognise the downward entailing context of the sentence.

5.3 Meta Study of Experiments 1-4

As mentioned in the previous section, experiment 6 contradicted some of our results from chapter 2, as well as those of experiment 5 from Chapter 3. To recapitulate, experiment 6 failed to show a significant difference between the determiner and the unmodified universal quantifier, as well as any differences between the unmodified and modified universal quantifiers. Because the results from Chapter 2 were based on four individual experiments, we decided to run a meta study to compare all of these experiments and confirm whether there are, or not, any differences between their results.

The study was performed using a 3 x 4 design that used the variables Licensor (Negation, Determiner, and Quantifier) and Experiment, which corresponded to each experiment and the quantifiers used in each of them (Experiment 1: *every*, Experiment 2: *some*, Experiment 3: *Every + Adjective*, Experiment *Every + RC*). For this comparison, we left out the No-NPI cases and we compared only those that contained an NPI. The study was analysed using the same methods that were used and described originally in chapter 2. Similarly, the three-level factor Licensor was Helmert coded into two contrasts as Grammaticality (Negation = 1, Determiner = -0.5, Quantifier = -0.5), and illusion (Negation = 0, Determiner = -1, Quantifier = 1). The four-level factor Experiment was Helmert coded as Quantifier Type (Experiment 2 = 1, Experiments 1, 3 and 4 = -1/3) which refers to the difference between *some* vs *every*, Modification (Experiment 1 = 1, Experiment 2 = 0, Experiments 3 and 4 = -0.5) which refers to the difference between the unmodified universal quantifier vs the modified universal quantifiers, and modifier type

(Experiments 1 and 2 = 0, Experiment 3 = 1, Experiment 4 = -1), which refers to the difference between *every + adjective* vs *every + RC*. Results can be observed in Figure 15 and table 43.




Figure 15: Meta study: Acceptability ratings. Error bars show standard error.

Table 43: Meta Study (Experiments 1-4) – Logistic Regression Mixed Effects Modelling

Fixed effects:	Estimate	Std. Error	z value	p value
(Intercept)	-0.237	0.047	-5.019	< 0.001 ***
Grammaticality: Neg vs (Det + Quant)	1.850	0.072	25.559	< 0.001 ***
Illusion: Det vs Quant	0.126	0.054	2.345	0.019 *
Quantifier Type: Existential vs Universal	0.244	0.087	2.801	0.005 **
Modification: Unmodified vs (Adj. + RC)	0.071	0.091	0.779	0.436
Mod. Type: Adj. vs RC	0.092	0.067	1.369	0.171
Grammaticality x Quantifier Type	0.274	0.134	2.046	0.041 *
Illusion x Quantifier Type	-0.153	0.092	-1.652	0.098 .
Grammaticality x Modification	0.753	0.124	6.050	< 0.001 ***
Illusion x Modification	0.086	0.092	0.937	0.349
Grammaticality x Mod. Type	0.469	0.087	5.396	< 0.001 ***
Illusion x Mod. Type	-0.007	0.073	-0.101	0.920

Briefly, these results show a very strong main effect of Grammaticality ($z = 25.559$, $p < 0.001$), a main effect of Illusion ($z = 2.345$, $p = 0.019$), and a main effect of Quantifier Type ($z = 2.801$, $p = 0.005$). We didn't find other main effects. We also found interactions between Grammaticality and Quantifier Type ($z = 2.046$, $p = 0.041$), Grammaticality and Modification ($z = 6.050$, $p < 0.001$),

and Grammaticality and Modifier Type ($z = 5.396$, $p < 0.001$). Crucially, we didn't find an interaction between Illusion and Modification ($z = 0.937$, $p = 0.349$). Finally, the Illusion x Quantifier Type nearly showed an interaction but it did not reach significance.

The lack of an interaction between Illusion and Modification means that the meta study was not able to confirm that the significant difference between the universal and the determiner from experiment 1 was statistically different from the insignificant effect from the modified conditions. While we can assume that this shows that no difference exists between the results of the experiments, there is a possibility that the differences are too small to be captured by the meta study, and we could speculate on some of the reasons for this.

As we can easily observe from Figure 15, the ratings for the determiner and the quantifier in experiment 1 are slightly lower in general than those found in all other experiments, while the negative was rated better more consistently in experiment 1 compared to experiments 3 and 4. Given that all negative conditions were perfectly grammatical, we should expect them all to be similarly graded. Nevertheless, the adjective modifier shows a slight decrease, while the RC conditions shows a larger decrease, in comparison. What this shows is that the additional modification did have an effect on the general rating of the conditions in both experiment 3 and even more so in experiment 4. These are the reasons why all the interactions that we found had something to do with grammaticality.

This could, of course, be a matter of word count, but in this case, higher word count also correlates with increased modification so it is difficult to say if it's just one or the other. Fortunately, chapter 4 also shows that increased word count in the case of an overt relative clause may actually decrease reading times until the appearance of the NPI which shows at least some facilitation on processing at early stages, and only when the NPI was unlicensed did reading times increase significantly, showing the opposite. Despite the slight increase in ratings compared to experiment 1, both the universal and the determiner in the RC and adjective cases

were still rated mostly negatively and no significant difference was observed between them individually, and though the negative conditions were slightly worse than in other experiments, they were still significantly better rated than the other two, which correlates with the main effect of Grammaticality we saw before. This probably means that word count on itself cannot be the only reason why we see these results, but that it was also necessary that the NPI was promptly identified as being unlicensed. While the universal quantifier in general may still produce some increase in ratings in general, it is possible that in some part the increase in modulation, which also increases complexity, can help reduce the possibility that the NPI is considered temporarily licensed. This could explain why despite not seeing a significant interaction between Illusion and Modification, experiment 1 still managed to produce a significant Illusion effect while experiments 3 and 4 didn't. After all the main effect of Illusion does show that the quantifier conditions are in general significantly better than the determiners. If the universal quantifier can cause some illusions to appear and modulation can decrease this possibility, it does make sense to assume that we can see some general difference between the universal quantifier and the determiner but that modulation probably cannot completely override it. Interestingly, both times that we have seen evidence of illusory license (Experiments 1 and 5) the structures that elicit the illusion are the simplest ones without any additional modification.

Additionally, the existential quantifier could not be contributing to this main effect of Illusion, since we do see a main effect of Licensor Type that shows that the results for the quantifier *some* were significantly different than the experiments that had the universal quantifier *every* in general which is probably driven by the noticeable lack of difference between both unlicensed conditions. Clearly, in this case there is no chance of illusion so the effect is definitely driven by the universal quantifier experiments. Given this, the fact that we didn't find a significant interaction between the Illusion effect and Quantifier Type could tell us something. While there is a clear difference between *some* and each instance of *every* regardless of modification in terms of the relationship between the quantifier and the determiner, the lack of effect could be driven

by the lack of a significant difference in both modified experiments, but the significant effect from experiment 1 could be the reason why the interaction almost reached significance. If no significant interaction existed at all, and the results of experiment 1 were due to a type I error, perhaps we would not be finding main effects from Quantifier Type, or Illusion as we see above.

Then again, we can only speculate about this. While the meta study can provide some insight, the only way to properly assess these possibilities would be by comparing all conditions within a single experiment. Besides, this leaves some open questions regarding what is happening with the negative quantifier and why it shows greatly decreased ratings compared to other experiments and this could be due to an increase in complexity. As well as to why experiment 6 doesn't show similar results to experiment 5 at least.

5.4 Implications for NPI licensing

5.4.1 Identifying local monotonicity as downward-entailing

From our experiments, we propose that the NPI has a much more important role on its licensing process than was originally assumed. We believe that the upon its inception, it is the NPI that provides hints to the parser that its current context should be downward entailing, and then this prompts the parser to check for an operator in the preceding structure that is able to provide this environment without immediately checking whether the context is in effect DE.

Given this two events, the parser identifies the NPI as licensed and keeps going on with the structure. Thus, the licensed NPI facilitates the identification of the structure in which its located. If the NPI is not located inside the correct DE context, though, the parser is not able to immediately check the contextual information and this gives rise to illusions of grammaticality as evidenced from the results of chapters 2 and 3.

Furthermore, as shown in chapter 4, a licensed NPI can pre-empt the appearance of garden path effects when it is located at the head of the ambiguous reduced relative clause. In order for this to happen, the parser must be able to correctly identify the structural context that contains the subordinate verb, and correctly interpret its tense as past participle instead of as simple past.

Garden path research show that the parser usually prefers² the main clause interpretation of the sentence upon reaching the first verb at the inside of the reduced relative clause (Townsend & Bever, 2001). This interpretation, however, is exclusive to the reduced relative clause structure and not to the full relative clause since the overt head of the relative clause (e.g. *who was*) prevents the ambiguity of the context. Given the NPI's ability to avert the appearance of garden path effects at disambiguation, it is reasonable to assume that the NPI's presence takes on a similar disambiguating function, at least as long as the NPI is licensed.

Crucially, we believe that this provides evidence that the correct identification of the reduced relative clause as well as the assumption by the parser that the NPI is licensed is at least in part due to the NPI itself and not only due to the presence of the universal quantifier as a licenser. While the universal could provide the parser with an expectation of a downward-entailing context, the universal quantifier does not seem to be strong enough to override the structural preferences of the parser by itself, even if it were to provide the parser with a semantically driven structure ahead of time; rather the NPI kicks off series of processes that provides evidence to the parser of the context should be there. Considering that the parser assumes that incoming linguistic elements should be licensed, it could take both sources of information as confirmation that the NPI must be licensed.

² This preferences can be attributed solely to structural preferences, as in the garden path model (Clifton & Ferreira, 1989; Ferreira & Clifton, 1986; Frazier, 1979; Frazier & Rayner, 1982; Rayner et al, 1983), to initial biases for simpler and more frequent structures as in constraint-based models (Tanenhaus & Trueswell, 1995), principles of discourse like in the referential theory (Crain and Steedman, 1985), or immediate probabilistic selection among different alternative options as in the unrestricted race model (Traxler, Pickering, & Clifton, 1998; van Gompel, et al., 2005; van Gompel, Pickering, & Traxler, 2001).

Taken into account the results of all of our experiments, we can establish a time course of the licensing process and illusory licensing. Chapter 3, through the use of self-paced reading presents an early probe of the licensing process; chapter 4 captures intermediate effects that are visible during ambiguity resolution which are also rapidly executed; finally, the use of offline measures in chapter 2 demonstrates the effects of a late probe of illusory licensing.

In this manner, early effects can be seen, however briefly, in self-paced reading experiments in chapter 3, specifically in experiment 5 where we saw that after the appearance of the NPI sentences with the universal quantifier were able to produce increases in reading speeds compared to sentences with the definite article despite both being ungrammatical. Rapid processes initiated at the NPI could let the parser know that its environment should be DE and given the previous appearance of a potential licenser that can provide the desired context the parser can promptly consider the NPI to be licensed. At this point, the process may involve some sort or search for an element that has a desired feature, or may be able to perform the necessary computations to determine the monotonicity of the of the context locally and online.

If search is initiated, intermediate processes involve the prompt discovery of a preceding NPI-licensing element such as the universal quantifier which allow the parser to determine that the NPI is licensed and within its required structural position, which in turn allows it to use this information to identify the form of the following verb by ranking the a determined structure (in this case as a subordinate verb that is part of (reduced) relative clause, facilitating ambiguity resolution (Chapter 4). In the case of illusory licensing, upon reaching the NPI the universal quantifier seems to facilitate the NPIs processing and thus no anomalies are initially detected despite the it being located outside the restrictor clause (chapter 3, experiment 5). Though we did not test for distance effects, a recent, yet unpublished, follow-up study by Alaberkyan, A. (2022-) has found similar effects to those of Parker and Phillips (2016). Results that show that NPI licensing and illusory licensing is sensitive to distance effect in certain configurations reflect

how early and intermediate process are rapidly applied and that the search for an adequate licensor or identification of the correct context is rapid and bound to some local domain such that illusory licensing can happen with a locally available potential licensor but not with nonlocal potential licensor.

Finally, late effects include the completion of the restructuring or identification of the restrictor clause and the parser's final decision as to whether the NPI was actually located within the desired context or not, with commensurate success or failure of licensing. Illusory licensing is at this point a measure of whether the parser was able to successfully locate the NPI outside or inside of its desired context.

5.4.2 Implications for illusory licensing

Considering what we discussed above, it is possible that illusory licensing such as the one we observed in chapters 2 and 3 could be driven by the semantic and pragmatic information as we hypothesised in chapter two but that the source of this information was lexically provided by the NPI and not only by the universal quantifier's property having a downward-entailing context to its restrictor or its ability to license NPIs. In this case, the source of NPI illusory licensing could be explained by the parser reaching the NPI and immediately assuming that the NPI must be inside a downward-entailing context.

Considering that quantifiers tripartite structure (Kamp, 1981; Heim, 1982, Von Stechow, 1994) and that they function as operators that bind two other arguments (their restrictor and the nuclear scope), the presence of the universal quantifier could already have prepared the parser for the apparition of a restrictor clause. Given that the universal quantifier, same as the focus operator *only* presupposes the contrast of two sets, it is possible that upon its appearance the parser already assumes the entailing contexts of these arguments, or that the universal quantifier provides information that a downward entailing context is imminent. In this sense, upon reaching the NPI, the parser already has enough information to locally make calculations to

determine the context that the NPI requires and the structure in which it should be located. At the NPI itself, the parser would receive confirmation that this context should be active, especially if there is no other information that could have already prompted it, such as an overt relative clause. This is evidenced in our experiments in chapter 4 where we saw that the moment that the parser reached the NPI inside a reduced relative clause it was able to both very quickly license the NPI and recognise the structure it was located in. If either of these things had not happened, the parser would not have shown to have swiftly made a correct/more informed choice about the subordinate verb and the garden path effect in experiment 8 would not have been mitigated at disambiguation. In addition to this, the NPI did not show additional facilitation for the full relative clause condition with the universal quantifier, which shows that the NPI itself allows the parser to recognise the structure when there is no previous element that already determines the structure of the relative clause (such as the relative pronoun and the copula), provided of course, that the NPI can be licensed in the current context through a licenser like the universal quantifier, for our experiments, that was already present previously in the sentence, or made available through other processes.

On this note, there is plenty of evidence already on how NPIs are licensed in contexts where not only retrieval of specific features is improbable, but also where there is no apparent licenser that could provide the context in the sentences. For example, Xiang et al. (2009), that we have reviewed in Chapters 1 and 2, show that individual differences in pragmatic skills can also modulate the presence of illusory NPI licensing in relation to how prone they can be in drawing pragmatic inferences. Similarly, Giannakidou (1998, 2006) proposes that an NPI can be “rescued” by non-veridicality, that is to say through a property of certain linguistic expressions related to attitudes, beliefs or assertions that indicate a lack of commitment to the truth of the proposition being expressed. Non-veridical expressions including some types of questions, or negation. Additionally, these also include modal verbs, as well as affective/emotive verbs, where the global context of the sentence makes alternative propositions available that are not in the

sentence itself. As such, emotive verbs such as ‘be surprised’ can give rise to negative inferences such as ‘did not expect’, that in turn are able to licence an NPI and which are not available through memory retrieval. More recently, Mendia et al. (2018), have shown that spurious NPI licensing can occur in the absence an overt licenser directly from contexts that force an exhaustive parse. In their judgment experiment they compared sentences such as (56, taken from Mendia et al., 2018) which contain contexts involving *shortfall* (Moxey, 2006), which refers to expectations about what the set of references for a previous noun phrase when there is a mismatch between what is expected and what is actually the case, which put pressure on the discourse towards an exhaustive parse. In (56a) for example, the first sentence there is an expectation to keep on talking about the plants, but there is none after the second one, which is parsed exhaustively. They found that shortfall cases improve the acceptability rates of the NPI without the presence of a licenser and attributed the increase in ratings to the presence of a covert exhaustivity operator EXH. Considering that there was no licenser to be retrieved from these sentences, their experiment provides evidence against a theory of where NPI licensing is a product of item-to-item dependency.

- (56) a. [+EXH], [±EVER]
 Whenever the summer is really dry, Susy expects **all** of her plants to die.
 However, a small number of plants have {**ever/∅**} died.
- b. [-EXH], [±EVER]
 Whenever the summer is really rainy, Susy expects **none** of her plants to die.
 However, a small number of plants have {**ever/∅**} died.

We believe that evidence from our studies also provide a strong argument against an item-to-item dependency of NPI licensing in which a feature such as [+downward entailing] or similar could be involved as retrieval models assume. Since we have not been able to identify any satisfactory features that license an NPI it would not be entirely farfetched to consider alternatives as to how the NPI is being licensed, and evidence of memory retrieval could be solely due to the parser’s integration of contextual information locally.

It could be interesting to test if we can find similar garden path reducing effects in sentences with NPI illusory licensing with *every*. This could provide us with a potential way to explain why in chapter 3 we failed to produce illusory licensing in experiment 6 with the unmodified condition by manipulating the restrictor clause in self-paced measures since the NPI could be giving clues about its required context and the parser considers this as an extension of the restrictor clause.

5.4.3 Implications for ambiguity resolution

When we compared reduced relative clause sentences headed by the universal quantifier that had an NPI with those that didn't have one, we saw that despite the cases without NPIs being simpler and more frequent (according to the parser's preference mentioned above), they presented a very high peak at the verb region. Despite the fact that at this point these sentences should be considered absolutely grammatical, they presented the biggest slowdown than any other condition among universal quantifier conditions specifically at this area. This peak effect at the verb is very short lived, and we see it disappear at the next region.

A similar peak effect can also be observed in determiner sentences but in a slightly different pattern. While in the quantifier cases the distinction happened in reduced relative clauses between sentences that either contained an NPI or not, in determiner cases this effect is dependent on the type of relative clause when an NPI is present. Obviously, in the determiner case the conditions without an NPI are both licensed, but the NPI ones are not. In here we notice that the peak appears inside a reduced relative clause with an NPI and not in one without one. This shows two things; one: the increases in reading times are determined by unlicensed NPI effects, but the peak is determined by the verb in the reduced relative clause which is why we see this difference between both unlicensed NPI cases.

In general, we saw that the verbs inside reduced relative clauses were processed slower than verbs inside their analogous full relative clauses, furthermore, almost all conditions show an increase in reading times at the verb relative to the previous region, except in the universal and

NPI condition. This increase in reading times, especially visible in most of the reduced relative clauses could point an increase in processing difficulty which probably reflect processes related to the selection between alternative interpretations of the verb. Therefore, this could be evidence for parallel models of ambiguity resolution that consider that several structural interpretations are considered simultaneously at the ambiguous verb. The sudden reduction at the next region, though, could provide evidence that the parser makes quick decisions and commits to a preferred structure at that time.

In this sense, the difference we see in reduced relative sentences with a licensed NPI in universal quantifier conditions does not only show that the licensed NPI allows the parser to quickly identify the context of the relative clause, but also it allows the parser in turn to circumvent costly processes related to the selection of the correct or more likely verbal form since the information about the structure in which the verb is contained is made readily available by the NPI. While this could be a consequence of the parser being able to identify that it must be inside a reduced relative clause, the parser could still be required to check for the correct verbal form. Instead in the NPI + RRC condition we see that the process is heavily accelerated. Because of this, we believe that in ambiguous cases like this, the NPI facilitates the identification of the structure and which becomes readily available to the parser and that this is not just an effect of the quantifier's potential prediction of an upcoming restriction..

5.4.4 Implications on memory mechanisms

One of our first questions during this dissertation was whether NPI licensing is a matter of memory retrieval mechanisms, and if so, whether a serial-search or a content-addressable mechanism. Based on evidence from illusory licensing, previous research has shown that NPIs present a selective profile on its sensitivity to illusions. While some studies have shown that, in some configurations, unlicensed NPIs with a structurally irrelevant licensors give rise to illusory licensing (Drenhaus et al., 2005; Vasishth et al., 2008), a result that is consistent with a content-addressable, cue-based parallel memory mechanism, other research has found that in specific

configurations, these illusions can be turned off based on the distance between the NPI and its licenser (Parker & Phillips, 2016).

On this dissertation we have found evidence that the licensing process seems to be primarily driven by semantic and pragmatic processes, nevertheless, it could still be debatable that memory retrieval still has an important role in the licensing on NPIs. Evidence from illusory licensing, shows that the parser can access a licenser that should be structurally irrelevant to the NPI, a property that is predicted by content-addressable, cue based memory mechanisms. Furthermore, based on our experiments in chapter 4, we see that the NPI has to necessarily be licensed in order to have a positive effect in reducing the garden path effect.

The involvement of such a memory mechanism, nevertheless, requires the existence of specific cues and features that the parser must use to identify the correct element required for the licensing process. Our original attempt to test for a potential [Quantifier] feature proved that this was not a possibility since illusions being caused by specific quantifiers and not quantifiers in general.

Furthermore, we discussed that [+Downward Entailing] could not be easily considered a feature considering that both downward-entailing and upward-entailing refer to specific semantic contexts which can span entire propositions. In this sense, it could be too difficult to define how this feature would be assigned and how long would it span. If a word gives an instruction for the context to appear, would this mean that every word after contains this feature? Furthermore, when does it stop or where does it change, and how? If several words contain or are inside this type of context, would the parser retrieve all of them, or only the one that onsets the entailing context? – it seems that proposing such a feature is more problematic than helpful. We believe that rather than looking for specific features that either belong to specific licensers, or to all NPI licensers in general, we are looking at general semantic/pragmatic properties of whole propositions which, at least in the case of quantifiers, could be explained by the tripartite

structure of the quantifiers and which are constructed and calculated through the entire sentence.

A possibility that could be considered is that, in addition to the parser looking for a licenser when it reaches an NPI, the process defining the semantic contexts of the sentence could have already been initiated by the quantifier itself. This way, the parser could interpret the quantifier by also taking into account upcoming structures if the quantifier contains its operational structure (Σ) semantically encoded (57). Upon appearance of the universal quantifier, the parser could be able to have access to this contextual information, which makes the parser aware that a licensing context for the NPI (could be a DE context) is coming. When the parser meets an NPI, it already has the information necessary to license the NPI locally computing the context in which it is found and in turn assuming the structural position of the NPI (i.e. the relative clause).

$$(57) \quad \begin{array}{c} \text{every}_{xy} \\ | \\ \Sigma : \forall xy \quad [\text{Restriction}(x)]_{DE} [\text{Scope}(y)]_{UE} \end{array}$$

We could consider that this structure is maintained active through the sentence, similar to how a verb's argument structure remains active during sentence processing³, and accessed automatically when the parser meets the NPI, which allows it to license or reject NPIs automatically regardless of whether the NPI is located inside the appropriate licensing context. In this case, the application of this semantic structure could be constrained to specific domains and upon crossing a semantic boundary by reaching another context-providing linguistic element it could be overridden by the new structural context, which could help explain other NPI

³ The Principle of Minimal Attachment suggests that the parser is sensitive to the argument structure of the verb, as it biases attachment decisions towards the more typical argument structure (Frazier and Fodor, 1978). Additionally, evidence from event-related brain potentials (ERPs) has shown that the processing of argument structure information is associated with early ERP components (e.g. N400 and P600), suggesting that argument information is processed quickly and automatically during sentence comprehension (Bornkessel-Schlesewsky et al., 2005).

phenomena such as intervention effects or even certain distance effects, and could explain why we saw, for example, the effect of the NPI disappear in chapter 4 after the disambiguation regions, where only the garden path effect seemed to remain. Additionally, this could explain how properties of the quantificational restriction are applied during both illusory and normal licensing. While a lack of defined boundaries could allow the parser to mistakenly interpret some instances of an unlicensed universal quantifier (as we saw in experiments 1 and 5), the appropriate delimitation of such boundaries could facilitate acceptability or licensing of the NPI. As we saw, the appearance of an overt relative clause, allowed the parser to quickly identify the NPI as licensed when it was preceded by the universal quantifier, but not by the determiner. Additionally, the parser was able to very quickly identify the licensed NPI inside reduced relative clause when it was preceded by the universal quantifier and also use this information to select the correct verbal form and prevent garden path effects. Moreover, this information was only useful in the reduced relative clause where there were no other indications of the context and the presence of the NPI didn't show any indication of facilitation for the FRC where the context was already defined by the overt elements of the relative clause. While in chapter 4 we only used grammatical sentences with the universal quantifier, we could imagine that similar mechanisms drive spurious licensing of the NPI. Illusions may arise when the NPI does not have previous information about the contextual delimitations of the restrictor which could explain why we saw an increase in acceptability and facilitation in SPR in experiments 1 and 5, respectively when the universal quantifier was not modified. If the parser was able to use the appearance of the NPI to determine the structure, in some cases it might license it incorrectly. While evidence that an overt relative clause receives no facilitation from the NPI could explain that the boundaries of the restrictor being correctly delimited can reduce this possibility by delimiting a domain in which these effects are no longer prevalent.

Because of these reasons, we believe that our theory can help explain both licensing of NPIs with the universal quantifier, as well as provide a model that can explain why and when illusory

licensing can arise, while at the same time providing strong evidence that NPI licensing cannot be a matter of item-to-item dependency.

5.5 Future Research:

Follow up studies to the research presented in this dissertation should continue to try to determine a fine-grained time course of the licensing processes of NPIs. An online method that could prove to be useful for this purpose is eye-tracking. Measuring eye movements could allow us to observe patterns such as first pass reading times and regressions, as well as where do participants look for information when the context proves to be difficult or ambiguous. By applying this method to sentences similar to those in chapters 2 and 3, where we added prenominal and postnominal modifiers to illusory licensing with the universal quantifier could help us better understand where the effect of the manipulation of the restrictor clause actually takes effect.

Follow up studies to those in chapter 4 could include testing for similar sentences to those in experiment 8, but with an added adjective to the restrictor clause in order to observe whether the early determination of a contrasting set cancels out the facilitation effects of the NPI and allows the reduced relative clause to be ambiguous again thus allowing garden path effects to appear again at the disambiguation regions. Additionally, by looking at whether illusory licensing in similar configurations to experiment 8 can also subvert the appearance of garden path sentences we could be able to better observe whether illusory licensing is resolved before disambiguation of the garden path occurs or if this happens at an even later stage. So far, it seems that effects of an unlicensed NPI seem to stop at the disambiguation (where they were prevalently displaced by the garden path effect entirely). It could be possible that this boundary could also be applied for resolution of illusory licensing. The addition of questions about the meaning of the sentence could also be added to the experiment to see whether participants have considered other structures despite facilitation from the NPI and if this remain in

comprehenders memory, or whether the presence of the NPI categorically makes the garden path effect disappear.

5.6 Conclusion:

The processing of Negative Polarity Items is a central topic in psycholinguistic research and still subject to much debate. In this dissertation we have attempted to tackle the issues of NPI licensing and illusory licensing through the use of the universal quantifier *every* as an unexplored vehicle to enhance our understanding how these long-distance dependency is processed. The current research provides implications to the licensing process and the time course of this process, illusory licensing, the role of the NPI in ambiguity resolution, and implications into how memory mechanisms may be involved in NPI licensing. In our view, the NPI has a much more active role than has previously been assumed in determining the monotonicity of the semantic context and how semantic information is may be taken into account by the parser. We also present a possible theory that can explain how both NPI licensing with the universal quantifier works that can also account for appearance of illusory licensing with the universal quantifier *every*. We also believe that our theory provides evidence that item-to-item dependencies are not possible with NPI licensing, and that we must rather look for alternative solutions that do not involve a retrieval-based account.

While this is just a start, we believe that the current dissertation offers an interesting view and contributes to the field by providing insights into the negative polarity research, as well as new questions for future research.

References

- Aaronson, D., & Scarborough, H. S. (1976). Performance theories for sentence coding: Some quantitative evidence. *Journal of Experimental Psychology: Human Perception and Performance*, 2(1), 56–70. <https://doi.org/10.1037/0096-1523.2.1.56>
- Alaberkyan, A. (2022-). *Untitled Master's Thesis*. [Unpublished Master's Thesis]. University of Oxford, Oxford.
- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6, 451-474.
- Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, 128, 186-197.
- Bayeen, H., Davidson, D. J., & Bates, D. M. (2008). lme4: Linear mixed-effects models using S4 classes. R package version 0.999375-28. URL: <http://CRAN.R-project.org/package=lme4>.
- Baayen, R. H., Davidson, D., & Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390-412.
- Badecker & Lewis, 2007. Talk presented at CUNY, San Diego
- Baron-Cohen, S., Wheelwright, S., Skinner, R. *et al.* The Autism-Spectrum Quotient (AQ): Evidence from Asperger Syndrome/High-Functioning Autism, Males and Females, Scientists and Mathematicians. *J Autism Dev Disord* 31, 5–17 (2001).
<https://doi.org/10.1023/A:1005653411471>
- Barr, D.J., Levy, R. Scheepers, C., & Tily, H. J. (2013). Random effects structure for testing interactions in linear mixed-effects models. *Frontiers in Psychology*, 4, 328.
- Bates, D., Maechler, M., & Bolker, B. (2011). lme4: Linear mixed-effects models using S4 classes Retrieved from: <http://CRAN.R-project.org/package=lme4>.
- Bergen, L., Levy, R., & Goodman, N. D. (2014). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*.
- Berwick, R., & Weinberg, A. (1984). *The grammatical basis of linguistic performance: Language use and acquisition*. Cambridge, MA: MIT Press.

- Bever, T. G. (1970). The cognitive basis for linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language* (pp. 279-362). John Wiley & Sons.
- Bock, J. K., & Cutting, J. C. (1992). Regulating mental energy: Performance units in language production. *Journal of Memory and Language*, 31(1), 99-127.
- Bock, K., & Miller, C. A. (1991). Broken Agreement. *Cognitive Psychology*, 23, 45-93.
- Boland, J. E., Tanenhaus, M. K., & Garnsey, S. M. (1990). Evidence for the immediate use of verb control information in sentence processing. *Journal of Memory and Language*, 29(3), 330-345.
- Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. *Structures and beyond* 3, 39-103.
- Chierchia, G. (2006). Broaden your views. Implicatures of domain widening and the 'logicality' of language. *Linguistic Inquiry*, 37, 535-590.
- Chierchia, G. (2013). *Logic in Grammar: Polarity, Free Choice, and Intervention*. Oxford University Press.
- Chierchia, G., Fox, D., & Spector, B. (2008). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. *Semantics: An International Handbook of Natural Language Meaning*.
- Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42(4), 368-407.
- Chomsky, N. (1981). *Lectures on government and binding*. Berlin: Mouton de Gruyter.
- Clifton, C., & Ferreira, F. (1986). Semantic and phonological information in sentence processing. *Journal of Verbal Learning and Verbal Behavior*, 25(1), 141-159.
- Clifton, C., & Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes*, 4, 77-103.
- Clifton C., Frazier, L., & Deevy, P. (1999). Feature manipulation in sentence comprehension. *Rivista di Linguistica*, 11, 11-39.
- Clifton C., Frazier, L. (2010): When are Downward Entailing contexts identified? The case of the domain-widener ever. *Linguist Inquiry (Online)*, 41(4), 681.

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

Dillon, B. (2011). *Structured access in sentence comprehension*. (Unpublished Ph.D. Dissertation). University of Maryland, College Park.

Dillon, B., Chow, W. Y., Wagers, M., Guo, T., Liu, F., & Phillips, C. (2014). The structure-sensitive of access: Evidence from Mandarin Chinese. *Frontiers in Psychology*, 5.

Dillion, B., Mishler, A., Slogget, S., & Phillips, C. (2013). Contrasting intrusion profiles for agreement and anaphora: Experimental and modelling evidence. *Journal of Memory and Language*, 69, 85-103.

Drenhaus, H., Saddy, D., & Frisch, S. (2005). Processing negative polarity items: When negation comes through the backdoor. In S. Kepser & M. Reis (Eds.), *Linguistic evidence – empirical, theoretical, and computational perspectives* (pp. 145-165). Berlin/New York: Mouton de Gruyter.

Eberhard, K. M., Cutting, J. C., and Bock, J. K. (2005). Making syntax of sense: number agreement in sentence production. *Psychol. Rev.* 112, 531–559. <http://doi.org/10.1037/0033-295X.112.3.531>

Graesser A. C., Hoffman N. L., Clark L. F. (1980). Structural components of reading time. *Journal of Verbal Learning and Verbal Behavior*, 19(2), 135–151. [https://doi.org/10.1016/S0022-5371\(80\)90132-2](https://doi.org/10.1016/S0022-5371(80)90132-2)

Fauconnier, G. (1975). Pragmatic scales and logical structure. *Linguistic Inquiry*, 6.3, 353-375.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, 25(3), 348-368.

Franck, J., Vigliocco, G., Antoniou, M., & Kosslyn, S. M. (2002). Representational demands of object decision compared to spatial location decision: Evidence from early blind and sighted adults. *Brain and Cognition*, 48(2-3), 609-613.

Franke, M. (2011). Quantity implicatures, exhaustive interpretations, and rational conversation. *Semantics and Pragmatics* 4, 1-82.

Frazier, L. (1979). *On comprehending sentences: Syntactic parsing strategies*. PhD thesis, University of Connecticut.

Frazier, L. (1994). Sentences (re)analysis. Paper presented at the 7th CUNY Conference on Human Sentence Processing, CUNY, New York, March.

Frazier, L. (1999). *On Sentence Interpretation*. Netherlands: Springer

Frazier, L., & Clifton Jr, C. (1996). *Construal*. MIT Press.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14(2), 178-210.

Gazdar, G. (1979). Pragmatics, implicature, presupposition and logical form. *Academic Press*.

Gelman, A., Hill, J. (2007). *Data Analysis using regression and multilevel/hierarchical models*. Cambridge University Press.

Giannakidou, A. (1998). *Polarity sensitivity as (non)veridical dependency*. Amsterdam, Philadelphia: John Benjamins.

Giannakidou, A. (2006). Only, emotive factives, and the dual nature of polarity dependency. *Language*, 82, 575-603.

Giannakidou, A. (2011). Negative and Positive Polarity Items. In von Stechow, Klaus, Claudia Maienborn and Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning* de Gruyter, 1660-1712.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1-67.

Gordon, P. C., Hendrick, R., & Johnson, M. (2001). Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1411-1423.

Gordon, P. C., Hendrick, R., & Johnson, M. (2004). Effects of noun phrase type on sentence complexity. *Journal of Memory and Language*, 51, 97-104.

Gordon, P.C., Hendrick, R., Johnson, M., & Lee, Y. (2006). Similarity-based interference during language comprehension: Evidence from eye tracking during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1304-1321.

- Gordon, P. C., Hendrick, R., & Levine, W. H. (2002). Memory-load interference in syntactic processing. *Psychological Science*, 13, 425-430.
- Grice, H. P. (1975). Logic and conversation. *Syntax and Semantics*, 41-58.
- Grodner, D., & Gibson, E. (2005). Consequences of the serial nature of linguistic input. *Cognitive Science*, 29, 261-290.
- Guerzoni, E., & Sharvit, Y. (2007). A question of strength: On NPIs in interrogative clauses. *Linguistics and Philosophy*, 30, 361-391.
- Heim, I. (1982). The semantics of definite and indefinite noun phrases. Ph.D. dissertation, University of Massachusetts at Amherst. Distributed by the GLSA. Published in 1989 by Garland (New York).
- Horn, L. R. (1972). *On the semantic properties of logical operators in English*. Ph.D. Thesis, University of California, Los Angeles.
- Horn, L. R. (1976). *On the semantic properties of logical operators in English*. Indiana University Linguistics Club.
- Huang, Y., & Ferreira, F. (2021). Shared neural processing of hierarchical structures in language and music. *Neuropsychologia*, 156, 107821.
- Inoue, A., & Fodor, J. D. (1995). The role of case in parsing: Evidence from Japanese. *Journal of Psycholinguistic Research*, 24, 447-463.
- Israel, M. (2004). The pragmatics of polarity. In L. Horn, & G., Wards (Eds.), *The handbook of pragmatics* (pp. 701-723). Oxford: Blackwell.
- Jacobson, P. (2014). *Compositional Semantics: An introduction to the Syntax/Semantics Interface*. Oxford: Oxford University Press.
- Juliano, C., & Tanenhaus, M. K. (1993). Contingent frequency effects in syntactic ambiguity resolution. In *Proceedings of the 15th Annual Conference of the Cognitive Society*, pp. 593-598. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Juliano, C., & Tanenhaus, M. K. (1994). A constraint-based lexical account of the subject/object attachment preference. *Journal of Psycholinguistic Research*, 23, 459-471.

- Juola, J. F., Ward, N. J., & McNamara, T. (1982). Visual search and reading of rapid serial presentations of letter strings, words, and text. *Journal of Experimental Psychology: General*, *111*(2), 208–227. <https://doi.org/10.1037/0096-3445.111.2.208>
- Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*(4), 329–354. <https://doi.org/10.1037/0033-295X.87.4.329>
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, *111*(2), 228–238. <https://doi.org/10.1037/0096-3445.111.2.228>
- Kadmon, N., & Landman, F. (1993). Any. *Linguistics and Philosophy*, *16*, 353-422.
- Kamp, H. (1981). A theory of truth and semantic representation. In: Paul Portner and Barbara H. Partee (eds.), *Formal Semantics: The Essential Readings*. Oxford: Blackwell, pp. 189-223, 2002.
- Katzir, R. (2007), Structurally-defined alternatives. *Linguistic and Philosophy* *30* (6), 669-690.
- Kazanina, N., Lau, E., Lieberman, M., Yoshida, M., & Philips, C. (2006). The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language*, *56*, 384-409.
- Kieras, D. E. (1981). The role of major referents and sentence topics in the construction of passage macrostructure, *Discourse Processes*, *4*:1, 1-15, DOI: [10.1080/01638538109544503](https://doi.org/10.1080/01638538109544503)
- Klima, E. S. (1964). Negation in English. In J. A. Fodor & J. J. Katz (Eds.). *The structure of language*. Englewood Cliffs, NJ: Prentice-Hall.
- Krifka, M. (1995). The semantics and pragmatics of polarity items in assertion. *Linguistic Analysis*, *15*, 209-257.
- Kush, D. (2013). *Respecting relations: Memory access and antecedent retrieval in incremental sentence processing*. (Unpublished Ph.D. Dissertation). University of Maryland, College Park.
- Ladusaw, W. A. (1979). *Negative polarity items as inherent scope relations*. Ph.D. Dissertation, University of Texas at Austin.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*(3), 1126–1177.

- Levy, R. (2013). Memory and surprisal in human sentence comprehension. In G. Jäger, J. Rogers, & J. E. Green (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 2479–2484). Cognitive Science Society.
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29, 375-419.
- Lewis, R. L., Vasishth, S., & Van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Science*, 10, 447-454.
- Lewis, S., & Phillips, C. (2014). Aligning grammatical theories and language processing models. *Journal of Psycholinguistic Research*.
- Linebarger, M. (1980). *The grammar of negative polarity*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Linebarger, M. (1987). Negative polarity and grammatical representation. *Linguistics and Philosophy*, 10, 325-387.
- MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes*, 32, 692-715.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological review*, 109, 35.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Syntactic ambiguity resolution as lexical ambiguity resolution. In C. Clifton, L. Frazier, & K. Rayner (Eds), *Perspectives on sentence processing*, pp. 123-153. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- McElree, B. (2000). Sentence comprehension is mediated by content-addressable memory structure. *Journal of Experimental Psychology: General*, 122, 291-315.
- McElree, B. (2001). Working Memory and focal attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 817-835.
- McElree, B., (2006). Accessing recent events. In B. H. Ross (Ed.). *The psychology of learning and motivation*. San Diego: Academic Press.
- McElree, B., Foraker, S., & Dyer, L., (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, 48, 67-91.

- Mendia, J. A., Poole, E., & Dillon, B. (2018). Spurious npi licensing and exhaustification. *Semantics and linguistic theory*, 28, 233–250.
- Merlo, P. (1994). A corpus-based analysis of verb continuation frequencies for syntactic processing. *Journal of Psycholinguistic Research*, 23, 435–457.
- Mitchell, D. C., Corley, M., & Garnham, A. (1992). Effects of context in human sentence parsing: Evidence against a discourse-based proposal mechanism. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 69-88.
- Mitchell, D. C., & Green, D. W. (1978). The Effects of Context and Content on Immediate Processing in Reading. *Quarterly Journal of Experimental Psychology*, 30(4), 609–636. <https://doi.org/10.1080/14640747808400689>
- Nairne, J. S., (1988). A framework for interpreting recency effects in immediate serial recall. *Memory & Cognition*, 16, 343-352.
- Nairne, J. S., (1990). A feature model of immediate memory. *Memory & Cognition*, 18, 251-269.
- Ni, W., Crain, S., & Shankweiler, D. (1996). Sidestepping garden paths: Assessing the contributions of syntax, semantics, and plausibility in resolving ambiguities. *Language and Cognitive Processes*, 11, 283-334.
- Nicol, J. L., Forster, K. I., & Veres, C. (1997). Subject-verb agreement processes in comprehension. *Journal of Memory and Language*, 36(4), 569-587.
- Nicol, J., & Sweeney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research*, 18, 5-19.
- Nieuwland, M. S., Ditman, T., and Kuperberg, G. R. (2010). On the incrementality of pragmatic processing: an ERP investigation of informativeness and pragmatic abilities. *J. Mem. Lang.* 63, 324–346. <https://doi.org/10.1016/j.jml.2010.06.005>
- Parker, D. J. (2014). *The cognitive basis for encoding and navigating linguistic structure*. Ph.D. Dissertation.
- Parker, D., & Phillips, C. (2016). Negative polarity illusions and the format of hierarchical encodings in memory. *Cognition*, 157, 321-339.
- Pearlmutter, N., Garnsey, S., & Bock, K. (1999). Agreement process in sentence comprehension. *Journal of Memory and Language*, 41, 427-456.

- Phillips, C. (2006). The real-time status of island phenomena. *Language*, 82, 795-823.
- Phillips, C. (2013). *Encoding and navigating structured representations: Three recent surprises*. University of Connecticut: Talk at the 44th Meeting of the North East Linguistics Society (NELS).
- Phillips, C., Wagers, M., & Lau, E. F. (2011). Grammatical Illusions and selective fallibility in the real-time language comprehension. In J. Runner (Ed.), *Experiments at the Interfaces* (pp. 147-180). Bingley, UK: Emerald Publications.
- Pollard, C., & Sag, I. (1992). Anaphors in English and the scope of binding theory. *Linguistic Inquiry*, 23, 261-303.
- Posner, M. I. (1978). *Chronometric Explorations of Mind*. Englewood Heights, NJ: Erlbaum.
- Quian, T., Kukona, A., & Frazier, L. (2018). Early and late sensitivity to plausibility in sentence processing by Mandarin-Chinese speakers. *Journal of Memory and Language*, 99, 48-60.
- R Development Core Team (2014). R: A language and environment for statistical computing. Austria: R Foundation for Statistical Computing. Retrieved from: <http://www.R-project.org>.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85.2, 59-108.
- Rayner, K., Kambe, G., Duffy, S. A., & Albrecht, J. E. (1983). The effect of clause wrap-up on eye movements during reading. *Quarterly Journal of Experimental Psychology*, 35, 407-420.
- Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, 24, 657-720.
- Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and philosophy* 27 (3), 367-391.
- Sikos, L., Tomlinson, S., Traut, H., and Grodner, D. (2013). "Incremental computation of scalar implicatures: an ERP study," in *Poster Presented at the 26th Annual CUNY Sentence Processing Conference* (Columbia, SC)
- Slattery, T. J., Sturt, P., Christianson, K., Yoshida, M., & Ferreira, F. (2013). Lingering misinterpretations of garden-path sentences arise from competing syntactic representations. *Journal of Memory and Language*, 69, 104-120.
- Solomon, E. S., & Perlmutter, N. J. (2004). Semantic integration and syntactic planning in language production. *Cognitive Psychology*, 49, 1-46.
- Spivey-Knowlton, M.J., & Sedivy, J.C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, 55, 227-267.

- Sturt, P. (2003). The time-course of the application of binding constraints in reference resolution. *Journal of Memory and Language*, 48, 542-562.
- Sturt, P., Haan, R. de, & Steinhauer, K. (2014). *Eyes on the prize: Late closure but not garden-path effects in gaze-mediated sentence processing*. *Journal of Memory and Language*, 76, 79-94.
- Tabossi, P., Spivey-Knowlton, M.J., McRae, K., & Tanenhaus, M.K. (1994). Semantic effects on syntactic ambiguity resolution: Evidence for a constraint-based resolution process. In C. Umiltà & M. Moscovitch (Eds), *Attention and performance XV*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension. In J. L. Miller & P. D. Eimas (Eds.), *Handbook of Perception and Cognition, Vol. 11: Speech, language and communication*. New York: Academic Press.
- Taraban, R., & McClelland, J.L. (1990). Parsing and comprehension: A multiple-constraint view. In D.A. Balota, G.B. Flores d'Arcais, & K. Rayner (Eds), *Comprehension processes in reading*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Townsend, D.J. and Bever, T.G. (2001) *Sentence Comprehension: The Integration of Rules and Habits*. MIT Press, Cambridge.
- Traxler, M. J. (2019). *Introduction to psycholinguistics: Understanding language science*. John Wiley & Sons.
- Traxler MJ, Pickering MJ, Clifton C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, 39, 558–592.
- Trueswell, J.C., & Tanenhaus, M.K. (1994). Toward a constraint-based lexicalist approach to syntactic ambiguity resolution. In C. Clifton, L. Frazier, & K. Rayner (Eds), *Perspectives on sentence processing*, pp. 155–179. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Trueswell, J.C., Tanenhaus, M.K., & Garnsey, S.M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, 33, 285–318.
- Trueswell, J. C., Tanenhaus, M. K., & Kello, C. (1993). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(3), 528–553.

- van Dyke, J. A. (2011). Interference effects from grammatically unavailable constituents during sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(6), 1394–1412.
- van Dyke, J., A., & McElree, B. (2006). Retrieval interference in comprehension. *Journal of Memory and Language*, 55, 157-166.
- van Gompel RPG, Pickering MJ, Pearson J, Liversedge SP. (2005). Evidence against competition during syntactic ambiguity resolution. *Journal of Memory and Language*, 52, 284–307.
- van Gompel RPG, Pickering MJ, Traxler MJ. (2001). Reanalysis in sentence processing: Evidence against constraint-based and two-stage models. *Journal of Memory and Language*, 45, 225–258.
- van Rooij, R., & Schulz, R. (2004). Exhaustive interpretation of complex sentences. *Journal of logic, language and information* 13 (4), 491-519.
- Vasishth, S., & Broe, M. (2011). *The foundations of statistics: A simulation-based approach*. Berlin: Springer.
- Vasishth, S., Brüßow, S., Lewis, R. L., & Drenhaus, H. (2008). Processing Polarity: how the ungrammatical intrudes on the grammatical. *Cognitive Science*, 32, 685-712.
- Vasishth, S., Drenhaus, H., & Srinivasan, N. (2010). Psycholinguistics meets formal grammar: The syntax of German light verb constructions. *Journal of Memory and Language*, 63, 167-180.
- Venables, W. N., & Ripley, B. D. (1999). *Modern applied statistics with S-plus*. New York: Springer.
- Vigliocco, G., & Nicol, J. (1998). Separating hierarchical relations and word order in language production: Is proximity concord syntactic or linear? *Cognition*, 68, 13-29.
- von Fintel, K. (1994). *Restrictions on Quantifier Domains*. Ph.D. Dissertation, Graduate Student Linguistics Association (GLSA), University of Massachusetts, Amherst.
- Wagers, M., Lau, E.F., & Phillips, C. (2009). Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language*, 61, 206-237.
- Wagers, M. W., & Phillips, C. (2014). Expectation and locality effects in Japanese scrambling: Evidence from real-time sentence processing. *Journal of Memory and Language*, 75, 117-142.

Xiang, M., Dillon, B. W., & Phillips, C. (2006). Testing the strength of the spurious licensing effect for the negative polarity items. *Talk at the 19th CUNY conference on human sentence processing, New York.*

Xiang, L., Dillon, B., & Phillips, C. (2009). Illusory licensing effects across dependency types: ERP evidence. *Brain and Language*, 108, 40-55.

Xiang, M., Grove, J., & Giannakidou, A. (2013). Dependency dependent interference: NPI interference, agreement attraction, and global pragmatic inferences. *Frontiers in Psychology.*

Xiang, M., Grove, J., Giannakidou, A. (2015). *Semantic and pragmatic processes in the comprehension of negation: an event related potential study of negative polarity sensitivity.*

Yanilmaz, A., & Drury, J. (2014). Intrusion effects on NPI licensing in Turkish: Does the parser ignore the grammar? *Poster at the 27th CUNY conference on human sentence processing.* The Ohio State University.