

## CHAPTER 38

---

# NEGATIVE POLARITY ILLUSIONS

---

HANNA MULLER AND COLIN PHILLIPS

### 38.1. INTRODUCTION

---

DECADES of research have illuminated the syntactic and semantic licensing requirements on negative polarity items (NPIs), but relatively little is known about how these licensing requirements are satisfied in real time as sentences are interpreted. Here we pursue the strategy of using the error profile of the online NPI licensing mechanism to provide clues about the inner workings of that mechanism. This is the logic underlying the study of grammatical illusions, cases where native speakers experience a fleeting perception of acceptability or unacceptability that mismatches their more considered judgments. The selective appearance of these illusions has proven useful in several domains of sentence processing research such as the study of agreement attraction, for instance *\*The key to the cabinets are on the table*. In this chapter, we review the findings on negative polarity illusions, their parallels (and, in some cases, non-parallels) with other grammatical illusions, and the implications of this line of research for understanding the incremental processing of negative sentences as well as negative polarity phenomena more broadly.

#### 38.1.1. NPI illusions: Basic profile

Negative polarity items such as *any, ever, yet, or lift a finger* are licensed when they occur in the scope of negation or similar operators such as *no, not, few, rarely, or doubt*, often described as the class of downward-entailing operators (Ladusaw 1979). Hence the NPI *ever* is licensed in (1) because it is in the scope of the main clause subject *no bill*. It is not licensed in (3), because the sentence includes no licenser. Nor is it licensed in (2)—although the potential licenser *no senators* is present in the sentence, it fails to take scope over the NPI because it is embedded inside a relative clause modifier of the subject. NPI illusions, first demonstrated in German in Drenhaus, Saddy, and Frisch (2005), involve the fleeting perception of acceptability of an unlicensed NPI in a sentence that contains a potential, but structurally inappropriate licenser, such as (2).

- (1) No bills that the senators voted for will ever become law.
- (2) \*The bills that no senators voted for will ever become law.
- (3) \*The bills that the senators voted for will ever become law.

When speakers have no time limitations on providing judgments, sentences like (2) and (3) receive comparably low ratings, whether the task involves binary or gradient judgments. In other words, linguists' claim that the potential licenser in (2) is irrelevant to licensing of *ever* is readily confirmed by large-scale judgment studies. The contrast in acceptability between (2) and (3) emerges in measures involving faster responses, including the speeded acceptability task, in which comprehenders read sentences at a fixed presentation rate and then have just a couple of seconds to respond with a binary acceptability judgment. In this task sentences like (2) are typically accepted in a significantly larger proportion of trials than (3), though not as frequently as (1) is accepted. Acceptance rates vary across studies for a variety of reasons, but typical acceptance rates in a speeded acceptability task are 80–90% for (1), 30–40% for (2), and 10–20% for (3). This suggests that, at an early stage of interpretation, the ungrammaticality of (2) is less apparent than the ungrammaticality of (3). This discrepancy is taken to indicate the susceptibility of the NPI-licensing computation to errors under particular circumstances. Much work has now gone into trying to figure out what those circumstances are, with the expectation that this will illuminate the normal mechanisms of online NPI licensing.

Since the NPI illusion was discovered, it has been shown to be robust across methods and languages. The effect has been demonstrated in German (Drenhaus, Saddy, and Frisch 2005), English (Xiang, Dillon, and Phillips 2009, 2013; Parker and Phillips 2016; Ng and Husband 2017; de Dios Flores, Muller, and Phillips 2017), Turkish (Yanilmaz and Drury 2018), and Korean (Yun, Lee, and Drury 2017). In addition to speeded acceptability, the illusion has been observed in self-paced reading, eye-tracking, and event-related potential (ERP) experiments. In each of these cases the illusion appears as a reduction in the disruption otherwise associated with encountering an unlicensed NPI as in (3).

Two points should be highlighted. First, NPI illusions are not defined prescriptively, as a divergence between individual judgments and population norms. They are diagnosed based on a divergence between speakers' considered judgments and those same speakers' speeded responses. Second, although the illusions have been found across multiple languages and experimental measures, the cross-linguistic diversity of NPIs and licensing environments is such that we should be cautious about assuming that similar illusions will be found for all languages, all NPIs, or all configurations.

## 38.2. SOME UNPROMISING EXPLANATIONS

.....

A few initially appealing hypotheses about the source of NPI illusions warrant discussion. Throughout most of this chapter we assume that NPI illusions offer a window into the computations that underlie normal, successful licensing of NPIs. However, there are of course many component processes to sentence processing and any one of these could, in

principle, give rise to interpretive errors. Below we sketch a few proposals along those lines and why they are unlikely to be sufficient.

One hypothesis that we often encounter when discussing NPI illusions is that speakers erroneously judge sentences like (2) as acceptable because of an error in signal detection. Although *ever* in its position in (2) is clearly unacceptable when carefully considered, and clearly ungrammatical under most theories, the orthographically and phonologically similar *never* is grammatical in that position. On some proportion of trials, then, participants might simply mis-hear or mis-read the input. While there is some independent evidence that related mis-perceptions occur in other cases (Levy et al. 2009), this explanation suffers two key drawbacks. First, it is unclear why (2) but not (3) suffers this problem, as substituting *never* would make the sentence acceptable in either case. Second, de Dios Flores (2017) showed that although the substitution of *never* in (2) is indeed grammatical, speakers often find it difficult to process because of the close proximity of two negative words (*no* and *never*), and thus reject these sentences on a large proportion of trials. If NPI illusions reflect substitution of *never* for *ever*, then speakers should encounter at least as much difficulty with (2) as with (3), contrary to fact.

Another idea that we commonly encounter is that NPI illusions reflect the presence of multiple locally coherent substrings. For example, (4), like (2), gives rise to illusions. The suggestion is that (4), unlike its baseline (5), contains several locally coherent strings: *the authors that no critics recommended for the award* is, on its own, a perfectly well-formed subject, and *no critics recommended for the award have ever received acknowledgment for a best-selling novel* is, on its own, a perfectly well-formed sentence with a reduced relative clause (it is now the critics who are being recommended for the award). Perhaps the fact that this substring has a parse with a well-formed NPI is the source of NPI illusions.

- (4) \*The authors that no critics recommended for the award have ever received acknowledgment for a best-selling novel.
- (5) \*The authors that the critics recommended for the award have ever received acknowledgment for a best-selling novel.

Again, the parallels with other findings in the sentence processing literature make this hypothesis initially appealing (e.g. Tabor, Galantucci, and Richardson 2004). However, the compatibility of (4) with a reduced relative clause parse is a quirk of this and a few other examples, not a general property of the sentences that yield NPI illusions. For example, in (2) the equivalent substring would be *no senators voted for will ever become law*. This may technically be grammatical, but it reads with the difficulty of a garden path sentence such as *The horse raced past the barn fell*. Comprehenders are not generally known to resort to reduced relative clause parses as a strategy for getting out of difficulty. The sentences used to demonstrate NPI illusions in German are even less amenable to this explanation. The substring in (7) is wholly ungrammatical.

- (6) \*Ein Mann, der keinen Bart hatte, war jemals glücklich.  
a man who no beard had was ever happy  
'A man who had no beard was ever happy.'

- (7) \*Keinen Bart hatte, war jemals glücklich.  
no beard had was ever happy

Since the local coherence hypothesis predicts NPI illusions in only a subset of the sentence types where they have been observed, we do not consider it a viable explanation.

A third hypothesis that we often encounter is that the error in detecting the ungrammaticality of sentences like (2) is related to the scope of the negative quantifier. If the quantified subject of the relative clause takes scope in the main clause via quantifier raising, then a c-command relation between licenser and NPI may obtain (at some level of representation). This hypothesis comes in many flavors, including suggestions that (a) quantifier scope is inherently hard to compute, so participants might sometimes just guess when judging the acceptability of sentences like (2); (b) sentence (2) is in fact grammatical because the quantifier takes wide scope, but it is difficult to process; (c) although a wide-scope interpretation of the quantifier in (2) is ungrammatical, quantifiers do sometimes take exceptional scope, and so the parser entertains this hypothesis for a brief period of time. Of these, we consider (c) the only potentially viable version, and we return to it in section 38.3.2. Options (a) and (b) are unlikely explanations for the following reasons.

If the presence of a quantifier adds processing difficulty, leading speakers to guess randomly when judging the acceptability of sentences like (2), we should expect that all sentences containing quantifiers should receive an acceptance rate closer to 50%, relative to non-quantificational controls, including fully well-formed sentences. However, acceptance rates for the grammatical control condition in (1) are typically near ceiling. The mere presence of a quantifier does not cause speakers to guess more frequently than they otherwise would. One might accommodate this objection by narrowing the hypothesis to quantifiers in embedded clauses. However, the filler items in NPI illusion experiments often include embedded negative quantifiers, to prevent participants from adopting unnatural parsing strategies, and we do not observe reduced accuracy for these control sentences. Furthermore, this hypothesis only straightforwardly predicts illusions in speeded acceptability studies; it is not clear why or how a guessing strategy that is prompted by the presence of a quantifier should lead to an effect on reading times or ERPs at the NPI.

If NPI-illusion type sentences are, in fact, grammatical due to the possibility of scoping the quantifier out of the relative clause, but result in degraded acceptability due to parsing difficulty, we might expect slower, more careful judgments to show an even greater likelihood of acceptance. This is not the case. Typically the longer one thinks about (2), the worse it sounds. Of course, there are other such cases in the literature—multiple center embeddings are a classic case of a grammatical but unacceptable structure, and such sentences never seem to become acceptable, regardless of time constraints. One can, however, figure out what center embeddings would mean if they were acceptable. The same cannot be said of NPI illusions. We know of no formal investigation into what interpretation participants arrive at when they accept these sentences, but a quantifier-raising explanation makes some predictions: the meaning of (2) should be, approximately, ‘it is not the case that there exists some *x* such that *x* is a senator and the bill that *x* voted for will ever become law.’ This does not seem to align with what (2) means, if indeed (2) means anything at all. A final concern is that if raising a negative quantifier out of a relative clause is grammatical, we should expect to see other reflexes of these interpretations, such as the

ability to bind a pronoun in the main clause, contrary to fact (see Kush, Lidz, and Phillips 2015, Experiment 2a). For these reasons we think it unlikely that these sentences are in fact grammatical due to quantifier scope.

We thus find no strong evidence to support claims that NPI illusions arise due to problems in representing the perceptual input, problems in ruling out a locally coherent parse, or (some) problems of quantifier scope. Explanations that place blame on the online NPI licensing mechanism, which will be discussed in section 38.3, seem more plausible.

### 38.3. TWO PROMISING APPROACHES

---

One additional clarification is necessary before these hypotheses can be considered. The following explanations take for granted that NPI illusions reflect failure to detect a grammatical error, but it is in principle possible that the process is one of successful error detection followed by repair. Under the first scenario the licensing mechanism itself is error-prone and may allow ungrammatical structures, while under the second scenario it is not, and instead the repair strategy allows ungrammatical structures. Either scenario could lead to the observed elevated end-of-sentence acceptance rates.

Xiang, Dillon, and Phillips (2009) investigated this issue using ERPs, which are well-suited to this question because of the extremely fine-grained time-course information they provide. If the illusion profile observed in behavioral data reflects a repair process that is initiated only after the violation is successfully detected, we should expect to see a stage in the ERP response when the intrusion condition and ungrammatical baseline pattern together, to the exclusion of the grammatical baseline. This is not the case. Rather, the earliest point in time where the ungrammaticality of (3) is recognized (i.e. the earliest point where the ERP response to this condition diverges from the ERP response to sentences like (1)) is the same as the earliest point where the illusion arises (i.e. when (2) and (3) diverge). This finding suggests that for NPI illusion sentences like (2), there is no early stage at which the ungrammaticality is detected.

We now turn to the two classes of hypotheses that have received the most attention in the literature on NPI illusions: accounts that place the blame on the properties of the memory architecture engaged in resolving all long-distance dependencies, and accounts that place the blame elsewhere, typically in interpretive mechanisms.

#### 38.3.1. Memory mechanisms

Vasishth et al. (2008) assume that NPI licensing involves retrieving a licenser in memory upon encountering an NPI, and they propose that NPI illusions reflect noisy memory retrieval processes, motivating this claim using modeling evidence in the ACT-R framework. We now take a quick detour to describe the memory architecture that is assumed. Elements in memory, including terminal nodes of a tree (roughly, words) and non-terminals, are stored as ‘chunks’ or bundles of feature-value pairings. The relationships between nodes are encoded as features on those chunks, so that the representation of a sentence is a collection of chunks, each of which encodes its links to other chunks. Those

chunks can then be retrieved from memory based on retrieval cues that match the features on the chunks. Retrieval success is a function of both feature match and the chunk's level of activation. A memory access process involving more than one retrieval cue can lead to a *partial match* when a chunk in memory matches some but not all of the feature specifications of the retrieval cues. Multiple partial matches can potentially occur for any given attempt at retrieval. Partial matches play a key role in this account of NPI illusions because a partial match can lead to successful retrieval as long as the activation of the chunk is sufficiently high. ACT-R is a computational model of human cognition that implements a version of this theory of memory access (Lewis and Vasishth 2005).

Applying this to NPI illusions, Vasishth et al. characterize the computation that leads to the misperception of acceptability as retrieval from memory of the structurally irrelevant licenser, which occurs because of that chunk's partial match with the retrieval cues on the NPI. They use +c-commanding and +negative as the retrieval cues. *C-command* is a structural notion that refers to the relation between a node and its sister, and any node contained within the sister. It is closely related to the notion of logical scope. Vasishth et al. acknowledge that a more developed theory should involve a better understanding of the structural cue than simply labelling it as c-command, but they leave that issue for future development. With these retrieval cues, on some proportion of trials retrieval yields success via a partial match with the non-c-commanding negative DP, and because the search was successful, no error signal indicates to the comprehender that the sentence is ungrammatical. Note that this hypothesis relies on very general principles of memory and retrieval processes, and so it predicts that illusions should occur whenever a dependency resolution mechanism gives rise to partial matches. In fact, similar explanations have been proposed for agreement attraction, an illusion of subject-verb agreement (Wagers, Lau, and Phillips 2009). This predicted generality is not, in fact, borne out; we will see in sections 38.4 and 38.5 that NPI illusions are surprisingly constrained in their profile. Amending the memory-based account to accommodate the specificity of the illusion is not straightforward.

### 38.3.2. Interpretive mechanisms

The other prominent hypothesis, or family of hypotheses, that has been proposed is that NPI illusions arise because of errors in the processes by which NPI-containing sentences are interpreted. The details of what computations are involved in interpreting NPIs and how those computations go wrong vary between accounts.

#### 38.3.2.1. *Quantifier scope*

As we discussed above, one such hypothesis places blame for the illusion in the interpretation of quantifiers. In an NPI illusion sentence, the licenser is typically a negative quantifier and that quantifier is unable to properly license the NPI because of its structural position inside of the relative clause. However, it is well known that the interpretation of a quantifier does not always match its surface syntactic position. This is demonstrated by the ambiguity of sentences like (8), and the possibility of (9), in which the pronoun *it* seems to be bound by the quantifier *every* although *every* is inside a nominal modifier.

- (8) Every girl likes some boy.  
a. For every girl  $x$  there is some boy  $y$  such that  $x$  likes  $y$  (each girl may like a different boy)  
b. For some boy  $x$ , every girl  $y$  is such that  $y$  likes  $x$  (all the girls like the same boy)
- (9) [Someone from every <sub>$i$</sub>  city] hates it <sub>$i$</sub> . (May 1985)

Since a quantifier's scope is not determined by its surface position, the sentence comprehender may, upon encountering a quantifier, initially consider all in-principle scopal possibilities. The brief consideration of a wide-scope interpretation of the quantifier in an NPI illusion sentence leads to the initial perception that the NPI is in fact within the scope of the quantifier. As we will see in section 38.4.2, this hypothesis makes some accurate predictions regarding intrusive non-quantificational licensors. However, pinning NPI illusions on quantifier scope may be problematic since (a) negative quantifiers are not able to take exceptionally wide scope in the ways shown by (8) and (9), and (b) wide scope out of a relative clause is not generally possible, and NPI illusions have typically been shown using relative clauses. If the parser has access to these two grammatical facts, it should never consider a wide-scope interpretation of the quantifier in NPI illusion sentences. There would have to be a strong enough bias for wide-scope interpretations, in general, that the online comprehender disregards these grammatical facts. We know of no evidence for such a bias.

### 38.3.2.2. "Rescuing" by contrastive implicatures

One concrete proposal is due to Xiang, Dillon, and Phillips (2009), who argue that (a) successful NPI licensing can be driven by negative pragmatic inferences, and (b) in the sentences that typically yield NPI illusions similar negative inferences can sometimes arise. Under some theories of NPI licensing, emotive factives like *surprised* license NPIs because they license negative inferences, although they do not explicitly encode negativity (Giannakidou 2006b). For example, (10) licenses the inference in (11). (11) clearly places the NPI within the scope of negation. The claim is that the close relationship between (10) and (11) allows the NPI in (10) to be "rescued."

- (10) I'm surprised we have any sugar.  
(11) I thought we didn't have any sugar.

The motivation for a separate pragmatic rescuing operation for NPIs comes from both the theoretical literature on NPI licensing, particularly comparison of English and Greek (see Giannakidou 2006b for an overview) and ERP findings suggesting a qualitative difference between NPIs licensed by emotive factives and those licensed by more overtly negative licensors (Xiang, Grove, and Giannakidou 2016; see Giannakidou and Etxeberria 2018 for further discussion of these findings). NPIs licensed by *no*, *only*, or *very few* yield a difference in both N400 and P600 amplitude when compared to NPIs with no licensor, whereas NPIs licensed by emotive-factives only differ from unlicensed NPIs in the N400 time window.

This could reflect an additional reanalysis process, possibly pragmatic in nature, which occurs for NPIs not preceded by a canonically negative licenser.

Xiang et al. argue that a similar comparison to pragmatic inferences is responsible for NPI illusions. Sentence (12) is thought to give rise to an implicature. Since the speaker has specified a particular set of bills—those that no senators voted for—it can be assumed that the speaker does not have warrant to make the more general claim in (13). Thus there is a suggestion that there are some other bills—those that the senators *did* vote for—which would make (13) false, and so we get the inference from (12) to (14). Since (14) places the NPI within the scope of negation, the same mechanism that is responsible for NPI rescuing in (10) can apply.

(12) The bills that no senators voted for will (ever) become law.

(13) The bills will (ever) become law.

(14) The bills that the senators voted for will not (ever) become law.

The specifics of when a negative inference leads to true grammaticality and when it leads to an illusion are not spelled out. A difficulty for this hypothesis lies in defining the circumstances that give rise to these inferences such that they are not predicted to occur for ungrammatical baseline conditions, which also contain relative clauses that *could* give rise to contrastive negative inferences. A further complication is that Xiang, Dillon, and Phillips (2009) found that unlicensed NPIs and NPIs in illusion configurations differ in the P600 time window, which is unexpected under the interpretation that the P600 reflects pragmatic reanalysis of unlicensed NPIs. We return to some of these issues in section 38.4.1.

### 38.3.2.3. *Covert exhaustification*

The hypothesis that pragmatic inferences lie at the heart of the interpretive mistake is only one of a family of hypotheses that highlight the role of interpretation in NPI illusions. Another version of this idea is the hypothesis that NPI illusions arise because comprehenders infer a covert exhaustification operator which dominates the NPI (Mendia, Poole, and Dillon 2018). The overt exhaustification operator *only* is in fact an NPI licenser, so the thinking is that the silent version might be considered as a possible licenser. Some evidence in favor of this hypothesis comes from the relative increase in acceptability of *ever* in sentences with no clear licenser, but a contextually driven bias toward exhaustive interpretations as in (15), compared to those without a bias toward exhaustive interpretations as in (16).

(15) Whenever the summer is really dry, Susy expects all of her plants to die. However, a small number of the plants have ever died.

(16) Whenever the summer is really rainy, Susy expects none of her plants to die. However, a small number of the plants have ever died.

As with the contrastive inference hypothesis, a difficulty for this theory lies in identifying the factors that drive the comprehender to infer the exhaustive operator, such that they occur for illusion sentences but not for ungrammatical baseline sentences.



#### 38.3.2.4. *Licensing as integration into context*

One final proposal is that illusions arise because of problems in rapidly transitioning out of the relative clause, which has a downward-entailing and NPI-compatible meaning. Under this hypothesis, online NPI licensing is not implemented by means of retrieval of an individual licensor but rather by means of integration of the semantic content of the NPI into a phrasal meaning, which results in successful integration only if the phrasal meaning has the appropriate characteristics, such as downward-entailment. It has been suggested that the contribution of an NPI is to strengthen a negative claim (Kadmon and Landman 1993). When an NPI is encountered soon after strong negative claims are suggested in the relative clause, the comprehender may mistakenly integrate the NPI into this contextual meaning, despite syntactic information that signals that this is unallowed.

This hypothesis correctly predicts the two most startling restrictions on NPI illusions—their sensitivity to intervening material and their disappearance with licensors that merely allow but do not encourage strong, exceptionless inferences at the clause level. We now turn to these surprising patterns and their implications for the hypotheses discussed so far.

### 38.4. RESTRICTIONS ON NPI ILLUSIONS

---

#### 38.4.1. Licensor distance

Although the NPI illusion effect appears to be robust across measures and languages, instances where the illusion fails to arise in specific contexts are particularly informative. Increasing numbers of studies have found that the conditions under which NPI illusions arise are more specific than simply the presence of a non-c-commanding licensor. Some of the accounts discussed above predict a more even profile for the illusion than has been observed, while others predict that the phenomenon should be sensitive to factors that do not appear to strongly impact the illusion.

One striking limit on the illusion is that it disappears when the lure is suitably far away from the NPI (Parker and Phillips 2016). Sentences with either an embedding clause or a parenthetical intervening between the intrusive licensor and the NPI were found to not yield illusions (17a and 18a), whereas closely matched controls with less material between the intrusive licensor and the NPI (17b and 18b) elicited clear illusions in both speeded acceptability and self-paced reading measures. Parker and Phillips argue that it is simply the passage of time that makes the illusion go away, based on the assumption that the parenthetical phrase in (18a) has no impact on the relative structural positions of the lure and the NPI.

- (17a) \*The journalists that no editors recommended for the assignment thought that the readers would ever understand the complicated situation. <- *no illusion*
- (17b) \*The journalists that no editors recommended for the assignment ever thought that the readers would understand the complicated situation. <- *illusion*

- (18a) \*The authors that no critics recommended for the assignment have, as the editors mentioned, ever received a pay raise. <- *no illusion*
- (18b) \*As the editors mentioned, the authors that no critics recommended for the assignment have ever received a pay raise. <- *illusion*

These contrasts show that material in between the lure and the NPI makes a difference, but it is unclear from this evidence whether the relevant distance is the distance from the intrusive licenser itself to the NPI, or the distance from the licensing environment to the NPI, that is, from the right edge of the relative clause, the last position at which an NPI could have been fully acceptable. Adding material outside the relative clause increases both of these distances. However, it is clear that previous descriptions of the environments that give rise to NPI illusions—namely, contexts where an unlicensed NPI is preceded by a non-c-commanding licenser—are insufficient. The fallibility of the licensing mechanism is narrower.

It is also not clear from Parker and Phillips's findings how much distance is enough distance to "turn the illusion off." In both of the contrasts in (17–18) the difference between the condition in which illusions are observed and the condition in which illusions are not observed is several words. From these comparisons alone, we cannot determine whether the impact of the intrusive licenser declines gradually with the passage of time or whether there is a well-defined point when the illusion "turns off." However, Parker and Phillips also find a contrast between (19a), which does not give rise to illusions, and (19b), which does. Here the difference in position is only one word.

- (19a) \*The authors that no critics recommended have received any acknowledgment for a best-selling novel. <- *no illusion*
- (19b) \*The authors that no critics recommended have ever received acknowledgment for a best-selling novel. <- *illusion*

However, the position of the NPI is not the only difference between (19a) and (19b), since the identity and the syntactic category of the NPI also differs, that is, *any* vs. *ever*. Prior demonstrations of NPI illusions in both English and German have typically used adverbial NPIs (*ever* and *jemals*), rather than the stereotypical NPI *any*, due to the fact that *any* also allows a free choice reading, for which no licenser is needed. Parker and Phillips avoid this ambiguity by combining *any* with abstract mass nouns such as *acknowledgment* that should block the free choice reading. It is unfortunately impossible to manipulate NPI position and NPI identity independently, and so we cannot be certain that the contrast in (19) is another instance of the distance effects seen in (17) and (18), perhaps showing that the main verb is a particularly important landmark for preventing NPI illusions. Alternatively, there could be item-specific differences between *ever* and *any* that are responsible for the disappearance of the illusion in (19a). However, we regard this as less likely, as the licensing conditions on *ever* and *any* are generally identical.

The effect of distance on NPI illusions is not straightforwardly predicted by either the memory retrieval account or interpretation-based accounts of the illusion. Either account could potentially accommodate the finding, but with different adjustments.

Under the memory retrieval account of NPI illusions we might expect that a non-c-commanding lure should be just as disruptive at long distances as at short distances. Simulations in ACT-R show that the model specifications used by Vasishth and colleagues (2008) do not predict the observed effect of distance (Parker and Phillips 2016). One could invoke a decay mechanism that makes the lure less accessible in memory at greater distances. But a decay mechanism would affect true licensors similarly, predicting that even c-commanding licensors should become less effective as the distance from the NPI increases, contrary to fact. It might be possible to precisely calibrate the decay rate so that NPI illusions would disappear at longer distances but correct NPI licensing would persist at the same distances. Under the memory retrieval account NPI illusions reflect a partial match to retrieval cues whereas correct NPI licensing reflects a full match to the same retrieval cues. This means that at certain combinations of decay rates and distance a non-c-commanding licensor could have sufficiently low activation to prevent retrieval while a c-commanding licensor still could be retrieved. But this account would predict that at even greater distances the c-commanding licensor should fail to be retrieved. We know of no evidence that this happens.

A further challenge to memory-based explanations of NPI illusions comes from an additional contrast found by Parker and Phillips. Although added distance blocks NPI illusions, it seems to have no impact on agreement illusions. If both types of illusions depend on the same memory retrieval processes, then they should be affected similarly by added material. We return to this issue in section 38.5.1.

Interpretation-based accounts of NPI illusions face uneven difficulty in capturing the effect of distance. If NPI illusions reflect illicit scoping of the negative quantifier outside the relative clause then added material should have no effect on the illusion. Alternatively, if NPI illusions reflect erroneous pragmatic inferences, then it is not clear how those erroneous inferences would change with added distance. Similarly, the silent exhaustification account does not straightforwardly predict that added material should prevent NPI illusions. Under any of these accounts additional stipulations might capture the distance effect, but they do not follow from the theory.

The one version of the interpretation-based account of NPI illusions that could easily capture the distance effect is the hypothesis that the illusions arise from the mistaken integration of the NPI into the relative clause, or mistaken effects of the licensing environment created by the relative clause. Under that account, if the NPI is further away from the relative clause boundary then there is no danger that it would be impacted by the RC's licensing environment. This account further predicts that the relevant distance is not the distance between the lure and the NPI, but rather the distance between the relative clause and the NPI. Further work is needed to determine whether this is, in fact, the relevant distance.

### 38.4.2. Licensor type

A surprising additional contrast involves the nature of the intrusive licensor. In self-paced reading and speeded acceptability studies de Dios Flores, Muller, and Phillips (2017) found that illusions occur when the intrusive licensor is a negative quantifier, for example *no*, *very few*, but that simple sentential negation, for example *haven't*, *did not*, does not induce NPI

illusions. This is surprising because sentential negation is the most stereotypical kind of NPI licenser of all, and also likely the most frequent NPI licenser.

- (20) \*The authors that the critics haven't recommended for the award have ever received acknowledgment for a best-selling novel. <- *no illusion*
- (21) \*The authors that no critics have recommended for the award have ever received acknowledgment for a best-selling novel. <- *illusion*

The contrast in (20–21) does not reduce to the distance effect described above, as the sentential negation is closer to the NPI, yet it is less likely to induce an illusion.

The fact that negative quantifiers act as lures for NPI illusions while sentential negation does not is unlikely to reduce to an effect of linear position, since negative quantifiers in object position in the relative clause also induce illusory NPI licensing. In follow-up studies we have found that subject relative clauses as in (22) and (23) elicit qualitatively similar patterns of findings as (20) and (21)—that is, illusions for intrusive quantificational licensers (*very few*) but not simple negation.

- (22) The critics that haven't recommended the authors of alternative genres have ever objected to mainstream literary trends. <- *no illusion*
- (23) The critics that have recommended very few authors of alternative genres have ever objected to mainstream literary trends. <- *illusion*

The contrast between quantificational lures and sentential negation is especially challenging for an account that attributes NPI illusions to noisy memory retrieval mechanisms. Sentential negation should be at least as good a match to retrieval cues for NPI licensing as a negative quantifier, so it should be just as effective a lure.

The hypothesis that NPI illusions reflect an exceptional (and illicit) scope interpretation of the quantifier could capture differences between quantificational lures and sentential negation, as quantifiers might scope out of the relative clause while simple negation does not.

It is unclear how this contrast might be captured under the silent exhaustification account. It has so far not been spelled out how the silent exhaustifier interacts with the overt, non-c-commanding licenser to produce illusions only when that overt licenser is present. This must be clarified in order to assess whether different overt licensers should impact NPI illusions.

Similarly, the hypothesis that NPI illusions reflect licensing by contrastive implicatures does not make explicit claims about how contrastive implicatures relate to the overt non-c-commanding licenser, and so the theory makes uncertain predictions about the impact of different licensers on inducing illusions. Xiang, Dillon, and Phillips (2009) suggest “that speakers may be more likely to generate [negative] inferences if the contrasting referents are made very salient in the discourse. Negative quantifiers can do exactly this.” This suggestion is superficially consistent with Moxey, Sanford, and Dawydiak (2001) and Sanford, Moxey, and Paterson (1996), both of which investigate quantifiers ranging from near-0% (e.g. *few*) to near-100% (e.g. *not quite all*) and find that the complement set is made salient, and can be referred to with a pronoun.

- (24) Few of the football fans went to the game. They watched it on TV instead.

But it is unclear how the salience of the complement set to the embedded quantifier makes salient the complement set to the (non-quantificational) DP being modified. That is, while the DP in (25) should make some set of non-mentioned senators salient, it's not clear that it makes some set of non-mentioned bills salient. With a more explicit theory of this link, we may be able to determine whether (26) should also make non-mentioned bills salient, and by extension whether the contrastive implicature explanation for NPI illusions can accommodate the observed contrast between *no* and *not*.

- (25) The bills that no senators voted for . . .

- (26) The bills that the senators didn't vote for . . .

Lastly, the hypothesis that NPI illusions are because of erroneous integration of the NPI into the relative clause meaning can account for these facts without much further elaboration. Recall that this hypothesis suggests that the degree of fit between the relative clause meaning and the NPI gives rise to illusions. We suspect that clauses with a negative quantifier may be more likely to be used to make strong, negative claims (i.e. the types of claims that NPIs are used in), because they suggest a lack of exceptions to the claim being made. Simple sentential negation, on the other hand, can be used to make strong exceptionless claims but can also be used to make more limited negative claims. Because of this, negative quantifiers may encourage inferences toward strengthening, whereas sentential negation merely allows such inferences. Tentative evidence in support of this generalization comes from the relatively higher probability that *no* will be followed by *ever* compared to the probability that *n't* or *not* will be followed by *ever*, in the COCA corpus (de Dios Flores, Muller, and Phillips 2017).

### 38.4.3. Clause type

Parker and Phillips (2011) report an additional contrast that could help to distinguish accounts of NPI illusions. Lure NPs appear to induce stronger illusions when they are the subject of a restrictive relative clause (28) than when they are the subject of a complement clause (27). However, this contrast is currently less firmly established than some others discussed here, so it should be taken with caution.

- (27) \*The analyst's prediction that no stock would fall overnight was ever taken seriously by the financial executives. <- *reduced or absent illusion*

- (28) \*The bicycles that no experienced cyclists bought for their daily training have ever used aluminum gears. <- *illusion*

Parker and Phillips found that clause type had no effect on (un)acceptability of NPI licensing in a standard untimed acceptability task, but it appeared to make a difference in a self-paced reading task. In the standard relative clause configuration in (28) they

replicated the common finding of reduced disruption in reading times following unlicensed *ever*. In the complement clause configuration in (27) they found no clear evidence of a corresponding reduction. However, they did not provide corroborating evidence from speeded acceptability judgment, which typically provides clearer evidence on NPI illusions, and the evidence that the two conditions differ from one another is limited.

Parker and Phillips offer the contrast as evidence in favor of the pragmatic licensing account of NPI illusions, based on the idea that a complement clause is less likely to generate a contrastive implicature.

A contrast between relative clauses and complement clauses would be challenging for accounts based on noisy retrieval mechanisms or on exceptional scope mechanisms. A lure in a complement clause should be just as likely to match retrieval cues and it should be at least as capable of scoping out of the embedded clause. The silent exhaustification account should predict that complement clauses are at least as likely to induce illusions as relative clauses. The hypothesis that erroneous NPI licensing involves integrating the NPI into the nearby context does not clearly predict a difference as a function of clause type, as the boundary of the licensing environment should be equally clear in the complement clause and the relative clause.

### 38.5. COMPARISON WITH OTHER PHENOMENA

The parser’s selective fallibility has been useful for understanding other aspects of sentence processing. The high-level similarity between the conditions that give rise to agreement attraction, for example *The key to the cabinets are rusty*, and those that give rise to NPI illusions have led to the suspicion of a common underlying cause. Both involve apparent attempts to resolve a dependency that is sensitive to *c-command* using non-*c-commanding* material, presumably because of the lure of morphological feature matches. Of course, subject–verb agreement involves a stricter structural condition than just *c-command*, but the parallels are tempting nonetheless. Explanations that capitalize on these parallels, such as the cue-based retrieval explanation discussed above, have the potential advantage of covering a wide range of phenomena with relatively few ad hoc assumptions. Positing a different explanation for every illusion or processing error is unappealing for obvious reasons.

However, NPI illusions require more than just a non-*c-commanding* negative word, making the parallels more superficial than highly general explanations predict. Factors like the type of licenser (*didn’t* vs. *no*), the distance from the NPI to the licenser or licensing domain, and to some extent the type of clause containing the licenser (complement clause vs. relative clause) have all been shown to be important for eliciting NPI illusions. Explicit comparisons of NPI illusions and agreement attraction have revealed that agreement attraction is less constrained in its profile. Furthermore, work on the online processing of other dependencies that depend on *c-command*, such as bound pronoun interpretation, have shown that a morphological feature match with a non-*c-commanding* lexical item sometimes has no impact on the computation of a long-distance dependency. Furthermore, there are several well-known cases of comprehenders’ failure to detect anomalies where

memory retrieval using morpho-syntactic cues play little if any role, including Moses illusions and Escher sentences. While there are clearly common factors across the contexts and dependencies that give rise to linguistic illusions, further work is needed to identify precisely the factors that determine when the sentence comprehender is vulnerable to errors and when it is not.

### 38.5.1. Subject–verb agreement

Three factors that may impact NPI illusions and agreement attraction differently have been investigated: individual differences in pragmatic reasoning ability, the distance effect discussed above, and *at-issue* status of the clause containing the licensor. In light of these findings and the more general fact that subject–verb agreement and NPI licensing are not thought to be captured by the same grammatical mechanism, we think it unlikely that a single explanation can explain both phenomena.

First, Xiang and colleagues (2013) showed that an individual’s susceptibility to NPI illusions, but not agreement attraction, correlates with that individual’s pragmatic reasoning ability as measured by the Autism-Spectrum Quotient (AQ). Specifically, participants with lower scores on the communication skills subsection of the AQ questionnaire, which indicates *better* communication skills, were more likely to judge an NPI illusion sentence as acceptable in a speeded acceptability task, compared to those with higher AQ scores (i.e. worse communication skills). No corresponding relationship was found for agreement attraction. However, it is worth noting that pragmatic reasoning ability accounts for a relatively small proportion of the illusion effect, as NPI illusions still occur for the high-AQ participants, just to a lesser extent. Also, no relationship with AQ score was shown for the self-paced reading data, which could reflect the greater variability in reading-time data.

Parker and Phillips (2016) similarly found that a factor that matters for NPI illusions does not appear to be relevant for agreement attraction. They showed that the distance effect discussed above, whereby NPI illusions disappear when the NPI is far from the lure, does not arise for agreement attraction. Number agreement effects are comparably large regardless of the time between the intervening plural noun and the erroneously plural verb.

Ng and Husband (2017) identified a factor that impact agreement attraction rates but does not appear to matter for NPI licensing. Specifically, they manipulated whether the intrusive element was contained in *at-issue* (relative clause, as in (29)) or *not-at-issue* (appositive, as in (30)) content, using self-paced reading.

(29) The authors [that no critics recommended] have ever received acknowledgment for a best selling novel.

(30) The authors [the ones that no critics recommended] have ever received acknowledgment for a best selling novel.

For agreement attraction, illusions only occurred when the intrusive plural was contained in a relative (*at-issue*) clause. For NPIs, the *at-issue* status of the clause and the grammaticality of the NPI appeared to have an additive effect. This suggests that agreement

computations consider only at-issue content whereas an NPI in need of a licenser may consider at-issue and not-at-issue content equally plausible candidates. However, due to the lack of an ungrammatical baseline condition in the NPI conditions, it's difficult to say for certain whether the magnitude of the illusion is being manipulated or merely the grammaticality effect.

The important takeaway from these three comparisons is that the factors that modulate the size of the NPI illusion appear to be different than the factors that modulate the size of the agreement attraction effect. This is a problem for any hypothesis that treats these phenomena as consequences of the same processing problem, for example erroneous retrieval of a partially feature-matching lexical item in memory. If this retrieval process were, in fact, responsible for both types of illusion, then we should see that manipulations of the same characteristics affect both, contrary to fact. We may take this to indicate that the memory retrieval account is plausible as an explanation of only agreement attraction, or only NPI illusions, or neither. Note that this does not amount to the claim that there is no memory retrieval process involved in licensing NPIs. It is surely the case that processing any multi-word utterance involves many successive memory access operations. Rather, the claim is that imperfect matches between retrieval cues and items in memory do not drive the illusion.

### 38.5.2. Pronoun resolution

Though the comparison with pronouns has been examined less extensively than the comparison with agreement attraction, there are some interesting findings in this domain. Like NPIs, pronouns are subject to a *c*-command constraint when they receive a bound variable interpretation. Only a *c*-commanding antecedent can bind the pronoun. Also, reflexive anaphors require a *c*-commanding antecedent. (In many languages, including English, the restriction is narrower: a reflexive requires a local *c*-commanding antecedent.) In both of these cases, previous research suggests that the parser is rather good at ignoring non-*c*-commanding lures.

For example, Kush, Lidz, and Phillips (2015) manipulated the structural relation between an antecedent and a pronoun in sentences like (a) and (b). The 'but'-clause in (a) has a high attachment site, so the quantificational phrase *any janitor* fails to *c*-command the pronoun *he*. This led to processing disruption, relative to (b) where the 'when'-clause has a low attachment site, placing the pronoun in the *c*-command domain of the quantifier. No corresponding contrast obtained when the quantifier was replaced with referential *the janitor*, which can co-refer with the pronoun without *c*-command.

- (31a) Kathi didn't think **any janitor** liked performing his custodial duties, but **he** had to clean up messes left after prom anyway.
- (31b) Kathi didn't think **any janitor** liked performing his custodial duties, when **he** had to clean up messes left after prom anyway.

Reflexive anaphors show an even closer parallel to NPI illusion configurations. They require a *c*-commanding antecedent, and that antecedent must match the reflexive in person, number, and gender. In study after study, a feature-matching but non-*c*-commanding



potential antecedent to a reflexive, such as ‘the managers’, yielded little or no facilitated processing or increased acceptability, relative to a baseline with non-matching potential antecedents (Clifton, Frazier, and Deevy 1999; Cunnings and Sturt 2014; Dillon et al. 2013; Nicol and Swinney 1989; Sturt 2003; Xiang, Dillon, and Phillips 2009). However, Parker and Phillips (2017) found that when the structurally appropriate antecedent is an especially poor match to the reflexive, mismatching in multiple features, the resulting reading time profile resembles illusory licensing. Similar findings were obtained by Sloggett (2017). However, the fact that these studies had to work so hard to induce illusions reinforces the conclusion that the parser is less tempted by non-c-commanding antecedents for reflexive anaphors than it is by non-c-commanding NPI licensors.

(32) \*The executive [who oversaw the managers] doubted **themselves** on most decisions.

Summarizing, although the NPI illusion literature may give the impression that c-command constraints are easily violated in resolving linguistic dependencies, findings about bound variable pronouns and reflexives suggest that the parser is capable of ignoring non-c-commanding lures. Any theory of NPI illusions that places blame on a general insensitivity to c-command relations must therefore address these contrasting profiles.

The contrasts may lend some support to interpretation-based accounts of NPI illusions, since those accounts invoke mechanisms that are specific to NPI licensing, and do not predict a broad array of parallel phenomena. However, there is another possible source for the contrasts that remains to be tested. As we saw in section 38.4.1 above, when the determiner NPI *any* appears following a main verb it is not susceptible to an illusion (Parker and Phillips 2016). This is the same position where bound variable pronouns and reflexive anaphors have been found to be resistant to non-c-commanding lures. The ideal comparison of NPIs with pronouns and reflexives would place the anaphor in the same position where the NPI *ever* has been found to elicit illusions. We know of no way to do this in English, so a cross-linguistic investigation will be necessary to address this question.

### 38.5.3. Other illusions

While the comparison with anaphora and bound pronouns suggests that not all c-command computations are error-inducing, we see from other phenomena that comprehension errors often have little to do with c-command. For example, argument role information has been argued to be difficult to access quickly under some circumstances, resulting in misunderstandings of sentences such as “the dog was bitten by the man” (Ferreira 2003), or expectations for upcoming words that overlook argument role constraints (Chow et al. 2018). So-called ‘Moses illusions’ involve an inappropriate but conceptually related lexical item, which often goes undetected by the comprehender: “How many of each animal did Moses bring on the Ark” is often met with the answer “two,” rather than “Moses didn’t have an ark.” And so-called “Escher sentences” or “comparative illusions,” such as “More people have been to Russia than I have,” are often initially, and sometimes persistently, perceived as acceptable, natural sentences, despite being semantically incoherent. It is only after careful consideration that the strangeness of the comparative can be detected. The underlying computations responsible for these phenomena

are currently not well understood (though Wellwood et al. 2018 makes important progress in identifying the profile of Escher sentences).

The specifics of the cue-based retrieval hypothesis are both too constrained and too broad to capture the profile of these illusions, but we believe there is some merit in the generalization that partial matches to some sought-out representation underlie many illusions. For example, we have suggested that NPI illusions occur only with quantificational licensors because the licensing contexts they create are highly compatible with an NPI, and perhaps even predictive of one. The syntactic information indicating that a nearby NPI cannot be evaluated in that context (because it is outside that syntactic domain) may be insufficient to overcome the lure of this highly compatible representation. Similarly, argument role reversals occur only when the interpretation generated by the reversal is a canonical event (i.e. dogs more typically bite men than men bite dogs), and the mispredictions that make comprehenders vulnerable to this misinterpretation have been found to have a similar temporal profile to NPI illusions (Chow et al. 2018). For Escher sentences, the more available an event-counting interpretation is for the first clause, the higher the likelihood that the sentence will be accepted, presumably because the event described by the second clause is semantically (though not syntactically) compatible with this comparison (Wellwood et al. 2018). We are optimistic that further work in these areas will illuminate key characteristics that determine when the sentence comprehender encounters difficulty.

## 38.6. BROADER IMPLICATIONS

---

In sum, we have seen that NPI illusions show a surprisingly restricted profile, and that this profile is not shared by other linguistic phenomena that are sensitive to a *c*-command constraint. These findings favor some of the the accounts we began with in section 38.3 more than others. First, a memory-based account is difficult to defend, given that its strength lies in its generality (i.e. we need a theory of memory anyway) but that generality predicts illusions to be more widespread than is in fact the case. As for interpretation-based accounts, the two clearest contrasts—the distance effect observed by Parker and Phillips (2016) and the licensor effect observed by de Dios Flores, Muller, and Phillips (2017)—require additional stipulations for hypotheses that place blame on contrastive implicatures and covert exhaustification, but are straightforwardly accommodated by a theory that derives NPI illusions from the mistaken integration of an NPI into a nearby negative context. The hypothesis that quantifier scope is at fault has mixed success—it clearly predicts a contrast in licensors but requires further stipulations to account for the distance effect.

Having evaluated the main hypotheses on the table for NPI illusions in English, we now turn to cross-linguistic phenomena, and finally the implications of this work for the processing of negation.

### 38.6.1. Cross-linguistic word order differences

While most work on NPI illusions has focused on English, the phenomenon has been reported in German, Turkish, and Korean as well. As mentioned in section 38.1.1 the first demonstration

of NPI illusions was in German, and the currently known facts in German fairly closely track findings in English. The Turkish and Korean cases are necessarily less similar, as both of these languages allow NPIs to be licensed by a negative word that *follows* the NPI.

Yanilmaz and Drury (2018) found evidence from acceptability judgments and ERPs for NPI illusions in Turkish with the NPI *kimse* (*anyone*). This finding contrasts with the English findings in several ways: first, Parker and Phillips (2016) found no illusions for the equivalent English NPI *any*. Second, Yanilmaz and Drury's stimuli placed the intrusive licenser inside a complement clause, not a relative clause, whereas Parker and Phillips (2011) argued that the strength of the illusion in English is weaker or absent in complement clauses. Third, Turkish sentences with verbal (non-quantificational) negation yielded illusions, unlike the English pattern, where illusions have to-date only been induced by negative quantifiers (de Dios Flores, Muller, and Phillips 2017). It is not clear how the hypotheses outlined above should be evaluated with respect to Turkish, since the word order facts are so different. For example, the cue-based retrieval hypothesis turns on the importance of feature-based specifications of chunks in memory and retrieval of those chunks based on feature matches. In Turkish, however, the licenser does not need to be plucked from memory; instead, when an NPI is encountered, the licenser needs to be found in the subsequent material. It is unclear how the mechanisms invoked in retrieval-based accounts or in interpretation-based accounts of NPI illusions should generalize to the sequencing of information in Turkish.

A similar pattern of illusory NPI licensing has been reported in Korean, where the licenser follows the NPI as in Turkish (Yun, Lee, and Drury 2017). The authors argue that in Korean the illusions are modulated by prosodic marking of phrase boundaries, but the effects are statistically marginal.

In light of the little that is already known about NPI illusions in head-final languages, together with the systematic modulation of NPI illusions found in English, and the rich semantic literature on the diversity of NPI types and licensing configurations cross-linguistically, there is clearly much scope for learning more about NPI licensing from a cross-linguistic investigation of NPI illusions.

### 38.6.2. Processing negation

There is a long-standing debate around the time-course of processing negation. Some work has suggested that negative sentences must be processed in two stages: first, the affirmative proposition expressed by the sentence minus the negative word is mentally represented, and then the truth conditions of this representation are reversed (e.g. Wason 1959; Fischler et al. 1983). This idea is motivated in part by data showing that participants are slower and less accurate when evaluating whether a negative sentence is true, compared to when evaluating whether an affirmative sentence is true (e.g. Clark and Chase 1972; Just and Carpenter 1976), suggesting an additional processing step. Challenges to this hypothesis have typically taken the form of demonstrations that negative sentences can in fact be processed quite quickly, when other factors are controlled for (Nieuwland and Kuperberg 2008; Tian, Breheny, and Ferguson 2010; Burnsky et al. 2017). These findings tell us that generating negative representations isn't always hard but they don't tell us much about what those representations look like. We think NPIs can shed some light on this.

There is a robust effect of grammaticality on reading times for NPIs, even when the NPI occurs sentence-medially. Because the difference between licensed and unlicensed NPIs relies on the difference between negated and not negated sentences, at least some representation of negation must have already been generated sentence medially, contrary to the claims of two-stage accounts. However, the grammaticality effect on its own cannot tell us to what extent the contribution of negation has been processed, only that there has been *some* recognition of the presence of negation. A two-stage theory of negation processing can accommodate immediate NPI licensing effects if what requires two stages is the computation of the *semantic* representation of the sentence, and the licensing process operates purely at the level of strings or surface syntax. This raises the question of whether immediate NPI licensing relies on access to the compositional meaning of the licensing clause.

The hypotheses discussed here make fundamentally different claims about the answer to this question. Memory-based accounts, for example, require no representation of meaning beyond the lexical meaning of the licenser. On the other hand, the claim that NPI illusions arise because of the mistaken integration of the meaning of the NPI with the meaning of the relative clause clearly assumes that relative clause meanings can be generated before the full proposition is processed. If this idea is on the right track, further specification of which properties of the relative clause contribute to the probability of an illusion will be hugely important for understanding which aspects of negation can be computed rapidly and which require access to the complete propositional meaning.

It is worth noting that in evaluating whether online licensing attempts rely on compositional meanings we cannot trivially infer whether the grammatical constraint on NPIs is a constraint on meanings. The question of whether NPI licensing is a syntactic or semantic phenomenon is, of course, a point of active debate. But it is in principle possible that an NPI's grammatical status is a consequence of the truth conditions of the proposition containing it but the online sentence comprehender searches for c-commanding negative words as a first-pass 'heuristic' for determining licensing. The opposite is also possible—the grammatical constraint on NPIs lies in the syntax, but the online sentence comprehender uses an inferred speaker message as a first-pass check. Because of the possibility of this type of misalignment, we cannot make strong claims about the nature of NPI licensing constraints from patterns in online NPI licensing. However, if we assume the simplest possible linking hypothesis—that the parser implements the grammar in the same terms that the grammar specifies the constraint—the hypotheses discussed here do make different claims about the nature of NPI licensing. The memory-based account's difficulty in accounting for the full range of illusion data may suggest that interpretation-based accounts of illusions (and, by extension, interpretation-based accounts of NPI licensing) are more appropriate than syntactic accounts. But further exploration of the assumed linking hypothesis is needed.

## 38.7. CONCLUSION

---

In this chapter we have reviewed leading approaches to explaining NPI illusions. This includes accounts that regard NPI licensing as fundamentally a process of retrieving

specific items from memory, and accounts that regard NPI licensing as fundamentally an interpretive problem. The surprisingly selective profile of illusions provides important clues, and challenges to existing accounts. NPI illusions seem to disappear when additional material precedes the NPI. They occur primarily with negative quantifiers. They might also be sensitive to the type of clause that contains the lure (relative clause vs. complement clause), though this finding is less well documented. The selective profile of NPI illusions also diverges from the fragility profile of other linguistic dependencies, such as subject–verb agreement and bound variable anaphora. This suggests that the source of NPI illusions may be more closely related to specific properties of NPI licensing. We proposed that a promising approach to NPI illusions is one that attributes them to interpretive processes that persist shortly after comprehenders exit a true NPI licensing environment. For reasons that remain to be clarified, negative quantifiers create an interpretive environment that is especially compatible with NPI licensing, and comprehenders do not immediately register that the semantic licensing properties no longer apply once they exit the licensing environment and start working on another part of the sentence structure.

## ACKNOWLEDGMENTS

.....

This work was supported in part by NSF NRT award DGE-1449815 to the University of Maryland and by a University of Maryland Flagship Fellowship to HM. We are especially grateful to Iria de Dios Flores for useful discussions of the issues covered here.