Cognition 157 (2016) 321–339

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Original Articles

Negative polarity illusions and the format of hierarchical encodings in memory

Dan Parker [a,*], Colin Phillips [b,c]

[a] Linguistics Program, Department of English, College of William & Mary, United States
[b] Department of Linguistics, University of Maryland, United States
[c] Language Science Center, University of Maryland, United States

ARTICLE INFO

Article history:
Received 5 November 2013
Revised 29 June 2016
Accepted 30 August 2016
Available online 7 October 2016

Keywords:
Sentence processing
Negative polarity
Agreement
Linguistic illusions
Memory
Binding
Representation

ABSTRACT

Linguistic illusions have provided valuable insights into how we mentally navigate complex representations in memory during language comprehension. Two notable cases involve illusory licensing of agreement and negative polarity items (NPIs), where comprehenders fleetingly accept sentences with unlicensed agreement or an unlicensed NPI, but judge those same sentences as unacceptable after more reflection. Existing accounts have argued that illusions are a consequence of faulty memory access processes, and make the additional assumption that the encoding of the sentence remains fixed over time. This paper challenges the predictions made by these accounts, which assume that illusions should generalize to a broader set of structural environments and a wider range of syntactic and semantic phenomena. We show across seven reading-time and acceptability judgment experiments that NPI illusions can be reliably switched "on" and "off", depending on the amount of time from when the potential licensor is processed until the NPI is encountered. But we also find that the same profile does not extend to agreement illusions. This contrast suggests that the mechanisms responsible for switching the NPI illusion on and off are not shared across all illusions. We argue that the contrast reflects changes over time in the encoding of the semantic/pragmatic representations that can license NPIs. Just as optical illusions have been informative about the visual system, selective linguistic illusions are informative not only about the nature of the access mechanisms, but also about the nature of the encoding mechanisms.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Successful language comprehension routinely requires establishing dependencies between non-adjacent words and phrases. These dependencies are subject to diverse syntactic, semantic, and discourse constraints, and often rely on memory access mechanisms to recover the appropriate information from the encoding of the previous context. For instance, to relate the verb *were* in (1) to its subject for subject-verb number agreement, memory access mechanisms must recover the encoding of the plural subject noun *paintings*, and avoid interference from similar information in structurally irrelevant locations, such as the non-subject plural noun *curators*.

(1) The paintings that impressed the curators were recently sold at auction.

Speakers are typically highly accurate in retrieving the appropriate information from memory for dependency formation, but a growing number of studies have reported that memory retrieval in language comprehension is sometimes susceptible to interference from structurally irrelevant items (Drenhaus, Saddy, & Frisch, 2005; Gordon, Hendrick, & Johnson, 2001; Lewis, Vasishth, & Van Dyke, 2006; Van Dyke & McElree, 2006, 2011; Wagers, Lau, & Phillips, 2009). In this paper, we focus on interference effects that have been argued to trigger 'linguistic illusions'. Linguistic illusions are cases where speakers appear to accept incoherent or ungrammatical sentences during the early stages of comprehension, but judge those same sentences as unacceptable after more reflection. These effects can arise during the comprehension of linguistic dependencies, and have been presented as evidence that comprehenders temporarily consider ungrammatical linguistic dependencies (Pearlmutter, Garnsey, & Bock, 1999; Vasishth, Brüssow, Lewis, & Drenhaus, 2008; Wagers et al., 2009; Xiang, Dillon, & Phillips, 2009).

The current study examines the cause of linguistic illusions in order to diagnose the cognitive mechanisms for encoding and accessing linguistic information in memory. Linguistic illusions

* Corresponding author at: Linguistics Program, College of William & Mary, P.O. Box 8975, Williamsburg, VA 23187, United States.
E-mail address: dparker@wm.edu (D. Parker).

http://dx.doi.org/10.1016/j.cognition.2016.08.016
0010-0277/© 2016 Elsevier B.V. All rights reserved.

could, in principle, reflect either an error in how we mentally encode structured linguistic representations in memory, or an error in how we access information in those representations later. Many recent accounts have argued that linguistic illusions are the product of faulty memory access processes, with the additional assumption that the encoding of the sentence remains fixed over time (e.g., Vasishth et al., 2008; Wagers et al., 2009; see also Lewis & Phillips, 2015, and Phillips, Wagers, & Lau, 2011). Importantly, these accounts predict that linguistic illusions should generalize to a broad range of syntactic and semantic environments.

In contrast, we report the results from seven reading-time and acceptability judgment experiments that challenge our current understanding of how linguistic illusions arise. We show that one type of illusion, which depends on semantic/pragmatic licensing mechanisms, shows a fleeting time profile, such that it is present or absent depending on the amount of time from when the potential licensor is processed until the licensee is encountered. But we find that the same profile does not arise for another type of illusion, which depends on a morphosyntactic licensing mechanism. These results are unexpected under existing accounts, which predict that illusions should be rather general. We take these results to suggest that the encoding of emerging semantic representations is not fixed, as previously assumed, but rather, changes over time. Up to now, linguistic illusions have been taken to be especially informative about the access mechanisms used in language comprehension. Here, we argue that they are also informative about the nature of the encoding mechanisms.

1.1. Linguistic illusions

One type of linguistic dependency that is highly susceptible to illusory licensing involves subject-verb agreement (Clifton, Frazier, & Deevy, 1999; Dillon, Mishler, Slogget, & Phillips, 2013; Pearlmutter et al., 1999; Staub, 2009; Tanner, Nicol, & Brehm, 2014; Wagers et al., 2009). Subject-verb number agreement in English and many other languages is subject to specific structural and morphological requirements: the number feature of the verb or auxiliary verb must agree with the number feature of the subject. Agreement attraction arises when comprehenders are temporarily misled during agreement resolution by a feature-matched item that is not the subject of the verb. For instance, Wagers et al. (2009) examined the comprehension of subject-verb agreement in sentences like (2) using self-paced reading and speeded acceptability judgments.[1] The sentences in (2) are ungrammatical due to the number mismatch between the plural verb and its subject. Comprehenders are typically highly sensitive to such errors. However, Wagers and colleagues found that the presence of a plural distractor (e.g., *cabinets* and *runners*) reduced the reading time disruption associated with the number mismatch and increased acceptability judgments, relative to the singular distractor condition. The eased processing and increased acceptability suggests that comprehenders were misled by the structurally irrelevant plural distractor, giving rise to an illusion of acceptability.

(2) a. ∗The key to the *cabinet(s)* unsurprisingly were rusty after many years of disuse.
b. ∗The *runner(s)* who the driver see each morning always wave.

Agreement illusions are not simply cases of proximity concord (e.g., Quirk, Greenbaum, Leech, & Svartvik, 1985) or local coherence (e.g., Tabor, Galantucci, & Richardson, 2004), as illusions are observed when the distractor does not intervene between the verb and its subject, as in (2b). Nor is the effect simply a consequence of dialectal variation, as speakers agree on the unacceptability of sentences like (2) when given ample time, and all speakers are prone to the illusion.

Furthermore, the illusion cannot reflect misrepresentation or faulty encoding of the subject, as has been previously claimed (e.g., Eberhard, Cutting, & Bock, 2005). If comprehenders simply misrepresented the number feature of the subject, we might expect them to experience "illusions of ungrammaticality", where sentences with grammatical agreement would be misperceived as ungrammatical. However, comprehenders generally do not experience illusions of ungrammaticality (Wagers et al., 2009). Wagers et al. (2009) argued that this grammatical asymmetry is expected if agreement illusions are due to properties of faulty memory access mechanisms, rather than misrepresentation or faulty encoding of the subject phrase. Under this account, encountering a plural-marked verb triggers a retrieval process that probes all items in memory at once, in parallel, for a match to the required structural and morphological cues, e.g., [+subject] and [+plural]. In sentences with ungrammatical agreement like (2), the competition between the true subject and the distractor is relatively even, since both items only partially match the retrieval cues. On some portion of trials, the distractor may be incorrectly retrieved due to a partial-match to [+plural]. Misretrieval of the plural distractor can give the comprehender the false impression that agreement is licensed, resulting in an illusion of acceptability. In sentences with grammatical agreement, by contrast, the distractor is less likely to interfere because the fully matched subject should outcompete partial matches. Crucially, this contrast would be unexpected if comprehenders simply misrepresented or incorrectly encoded the subject phrase.

Another linguistic illusion that has sometimes been argued to reflect misretrieval involves negative polarity items (NPIs). NPIs are expressions like *ever*, *any*, *yet*, *lift a finger*, and *a rat's ass* that are generally only acceptable in sentences that contain a downward entailing operator in a structurally higher position (Ladusaw, 1979). Negation is the canonical example of a downward entailing operator (see Giannakidou, 2011, for a review of the contexts that license NPIs). For instance, the NPI *ever* in (3a) is licensed because it appears in the scope of the negative phrase *no diplomats*. The scope of negation for current purposes corresponds roughly to the 'c-command' domain of negation, i.e., the structural sister of the negation in a syntactic tree and any element contained within the structural sister.[2] When negation is absent, as in (3b), or is not structurally higher than the NPI, as in (3c), the NPI is not licensed.

(3) 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.

[1] In a self-paced reading task, participants use button presses to control the presentation for each word of a sentence. In a speeded-acceptability judgment task, sentences are presented one word at a time at a fixed rate. At the end of the sentence, participants have 2–3 s to make a yes/no response about the perceived acceptability of the sentence. Both tasks are widely used in psycholinguistic research.

[2] There are cases that call for an elaboration of the c-command generalization. For example, in the sentence *Nobody's mother has ever served ice cream for dinner*, the NPI *ever* is licensed even though it is not syntactically c-commanded by the negation. In this case, it appears that the entire NP *nobody's mother* counts as the relevant licensor. Nothing in the current study depends on these elaborations.

Sentences like (3b-c) are reliably judged to be unacceptable when participants are given ample time (Xiang, Dillon, & Phillips, 2006). But in time-sensitive measures, sentences like (3c) are often treated as if they were acceptable, giving rise to an illusion of acceptability. This effect has been reported for German and English using eye-tracking (Vasishth et al., 2008), self-paced reading (Xiang, Grove, & Giannakidou, 2013; Xiang et al., 2006), speeded acceptability judgments (Drenhaus et al., 2005) and event-related potentials (Drenhaus et al., 2005; Xiang et al., 2009; Yanilmaz & Drury, 2014).

Vasishth et al. (2008) argued that NPI illusions reflect a partial-match memory retrieval error. Under this account, NPI licensing, much like agreement, is treated as a structural item-to-item dependency that is formed by retrieving an individual negative licensor from the preceding context using structural and semantic cues like [+c-command] and [+negative]. In sentences like (3c), the irrelevant licensor may be incorrectly retrieved due to a partial match to the cue [+negative], giving the comprehender the fleeting impression that the NPI is licensed.

A competing account argues that the illusion reflects over-application of semantic/pragmatic licensing mechanisms, rather than misretrieval (Xiang et al., 2009). This account is motivated by the fact that in addition to direct licensing by negation, NPIs can also be licensed by the semantic/pragmatic properties of entire propositions. For instance, the NPI *ever* is licensed in a sentence like *I am surprised that John ever finished* by the negative implicature *I expected that John would not finish* (Linebarger, 1987). In the linguistics literature, theories of negative polarity phenomena typically do not make reference to item-to-item licensing, and instead assume that the structural constraint on NPI licensing is a by-product of semantic/pragmatic licensing mechanisms (Chierchia, 2006; Giannakidou, 2011; Horn, 2010; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Linebarger, 1987). Under this view, subject-verb agreement and NPIs involve fundamentally different licensing mechanisms: whereas agreement involves an item-to-item dependency between specific morphosyntactic features (e.g., person, number, gender), NPIs rely on compositional interpretive mechanisms. According to Xiang and colleagues, comprehenders may over-apply the semantic/pragmatic licensing mechanisms in ungrammatical sentences like (3c), where the restrictive relative clause and the embedded negation might trigger an unwarranted pragmatic implicature that spuriously licenses the NPI.

The accounts proposed by Vasishth et al. (2008) and Xiang et al. (2009) differ in their views of how NPI licensing obtains, but agree that the illusion reflects limitations in the implementation of the licensing mechanisms, rather than misrepresentation or faulty encoding of the licensing context. Under both accounts, the illusion may be viewed as a kind of partial-match effect, where the licensing mechanisms are able to evaluate the features of a licensor independently from the position of those features in the structured representation of the sentence.

However, the conclusion that the illusion reflects a faulty implementation of the licensing mechanism is based on a narrow range of findings and specific assumptions about the generality of the illusion, which have not been tested. For instance, existing evidence has come entirely from a single NPI, *ever*, in a restricted licensing configuration. But since the illusion is argued to be the product of faulty access mechanisms engaged after an NPI is encountered, rather than the product of faulty encodings, both accounts have assumed that the illusion should generalize across items and configurations. Thus, an important task to better understand the source of these effects is to determine whether the NPI illusion extends to a broader range of NPIs and licensing configurations.

1.2. The current study

We set out to test the claims that illusions are (i) a general property of NPIs, and (ii) the product of faulty licensing mechanisms. We tested the first claim by comparing the NPIs *ever* and *any*. Results from Experiments 1–3 revealed that while *ever* elicits the illusion, *any* does not, which implies that illusions are not a general property of NPIs. We then tested the second claim that illusions should generalize across licensing configurations. Results from Experiments 4–5 disconfirmed this prediction, as they showed that the illusion can be switched off by moving the NPI further away from the potential licensors.

Based on these results, we hypothesized that the contrasting profiles could reflect either the structural position of the NPI in the sentence or the amount of time from when the potential licensor is processed until the NPI is encountered. We distinguished between these alternatives in Experiment 6 by holding constant the structural position of the NPI and manipulating the position of a parenthetical clause to vary the time between the potential licensors and the NPI. Once again, we saw the same modulation of the illusion, as the results showed that the illusion could be switched off by increasing the amount of time from when the potential licensors is processed until the NPI is encountered. Importantly, these tests differ from previous tests of the illusion (e.g., Vasishth et al., 2008; Xiang et al., 2009) by focusing on the time before the NPI arrives, rather than the time available to process the NPI. In order to understand the cause of this effect, in Experiment 7, we tested whether the parenthetical phrase manipulation used for NPIs would have a similar impact on agreement illusions. Results showed that the position of the parenthetical phrase did not impact agreement illusions. The contrast between agreement and NPI illusions suggests that the mechanisms responsible for switching the NPI illusion on/off are not shared across all illusions, but rather are specific to the semantic encodings that can license NPIs. Lastly, we modeled our results using an explicit computational model of memory retrieval based on the equations described in Vasishth et al. (2008). The model failed to capture the full pattern of results, which suggests that the on and off behavior of the NPI illusion is unlikely to simply be due to limitations of memory access mechanisms.

These results are surprising, since they are not predicted by existing accounts. We found that the illusion can be switched on or off depending on when the encoding of the licensing context is probed for NPI licensing. These findings suggest that the encoding is not fixed, as previously assumed, but rather, changes over time. Based on these findings, we argue that the semantic encoding changes, such that the internal elements become opaque as candidates for causing illusions with the passage of time. We suggest two possibilities for how the encoding might change, including evolution in the content of the encoding and changes over time in the format of the encoding. Under either account, the selective NPI illusion reflects access to different internal stages of the encoding process, such that the illusion can be reliably switched on or off, depending on when the encoding is probed.

2. Comparison of *ever* and *any*

Existing findings about NPI illusions show that they are robust across experimental measures, labs, and languages (Drenhaus et al., 2005; Vasishth et al., 2008; Xiang et al., 2009, 2013; Yanilmaz & Drury, 2014). But the general conclusions about how speakers access linguistic representations have been drawn from findings based on a single NPI, *ever*, in a narrow range of licensing configurations. Existing accounts predict that all other NPIs should behave similarly with respect to illusory licensing. In Experiments 1–3, we directly compared the NPIs *ever* and *any* to test the

assumptions that illusions are a general property of negative polarity phenomena, and that they reflect limitations of licensing mechanisms that are engaged after an NPI is encountered.

The NPI *any* is similar to the NPI *ever* in that it may be directly licensed by negation in a structurally higher position, as shown in (4).

(4) a. *Nobody* has eaten any of the cake that Mary made.
b. *No student* wants any homework over spring break.

The NPI *any* is the canonical example of an NPI in linguistic analyses. However, it has not featured prominently in previous research on NPI processing because it can receive either a negative polarity interpretation or a "free-choice" interpretation in contexts where a negative polarity interpretation is unavailable. Free choice *any* invites a choice among a set of alternatives, and unlike its NPI counterpart, it is not polarity sensitive, i.e., is acceptable without negation, as shown in (5).

(5) a. Sally will read any book on the shelf.
b. Pick a card, any card.

However, the free-choice reading of *any* can be blocked if the choice among a set of alternatives is unavailable. For instance, the free-choice reading of *any* is blocked in (6a) because the abstract mass noun *satisfaction* as the direct object cannot refer to a set of alternatives, making NPI licensing necessary, as in (6b). If comprehenders reliably interpret *any* as an NPI in contexts like (6b), then these contexts are appropriate for testing whether *any* is susceptible to illusory licensing.

(6) a. ??The criminals felt any satisfaction from the crime.
b. *No criminals* felt any satisfaction from the crime.

2.1. Experiment 1: Untimed acceptability ratings for any and ever

Experiment 1 tested whether *any* and *ever* would show parallel behavior in untimed acceptability ratings, as a way to show that we have successfully identified contexts that force the NPI reading of *any*. Specifically, we compared the effects of the presence and structural location of a potential NPI licensor on ratings for sentences like those in (6b) with *any* and *ever*. We reasoned that if comprehenders show sensitivity to the structural licensing conditions for *any*, just like we see for *ever*, then this suggests that we have successfully identified contexts that force the NPI reading of *any*. Based on previous findings for *ever*, we expected low ratings for sentences that lack an appropriate licensor relative to sentences with a structurally appropriate licensor. On the other hand, if sentences with *any* and *ever* show qualitatively different profiles, then this would imply that we have not blocked the free-choice reading of *any*.

2.1.1. Participants

Participants were 18 native speakers of English who were recruited using Amazon's Mechanical Turk web service (https://aws.amazon.com/mturk). All participants in this and the following experiments provided informed consent, and were screened for native speaker abilities. The screening probed knowledge of the constraints on English tense, modality, morphology, ellipsis, and syntactic islands. Participants in Experiment 1 were compensated $2.50. The experiment lasted approximately 15 min.

2.1.2. Materials

Experimental materials consisted of 36 sets of 6 items, which varied in the presence and location of an NPI licensor (grammatical licensor|irrelevant licensor|no licensor) and the type of NPI (*ever*|*any*). All items consisted of a subject noun phrase (NP) modified by an object relative clause, followed by a main clause predicate that contained the NPI. In the grammatical licensor conditions, the main clause subject NP included the negative determiner *no*, which is a reliable NPI licensor. In the irrelevant licensor conditions, the determiner *no* instead appeared in the relative clause subject. In the no licensor conditions, the licensors were replaced with the definite determiner *the*, which cannot license NPIs. The relative clause was always followed by the auxiliary verb *have*, which highlighted the right edge of the relative clause to ensure that participants would correctly construct a parse in which the irrelevant licensor was not structurally higher than the NPI. The NPI *any* immediately followed the main clause verb and appeared in the determiner position of the direct object. In order to block the free-choice reading of *any*, we used abstract mass nouns as the direct object, which cannot refer to a set of alternatives, and we used simple past tense verbs that favored an episodic interpretation of the sentence. Following previous studies, the NPI *ever* appeared immediately after the auxiliary verb. Care was taken in the construction of the materials to ensure that the sentences within each item set were plausible and identical in all respects aside from the licensing of the NPI. We paid particular attention to ensuring that the subject NPs were similarly plausible across all positions, regardless of the presence or position of the negative determiner. An example set of items is provided in Table 1. The full items list, for all experiments, can be found on the first author's website.

Each participant rated 84 sentences, consisting of 36 NPI sentences and 48 filler sentences. The 36 sets of NPI items were distributed across 6 lists in a Latin Square design. The filler sentences were of similar length and complexity to the NPI sentences, but lacked an NPI. Roughly half of the filler sentences included determiners similar to those in the NPI sentences and in similar positions, to prevent the possibility that participants might develop superficial reading and rating strategies based on the distribution of the determiners in the NPI sentences. Materials were balanced such that, across the experiment, half of the sentences were ungrammatical. The anomalies in the filler sentences comprised a variety of violations, including agreement errors, pronoun-antecedent gender mismatches, and unlicensed verbal morphology.

2.1.3. Procedure

Sentences were presented using Ibex (Alex Drummond, http://spellout.net/ibexfarm/). Participants were instructed to rate the acceptability of each sentence using a 7-point scale ('7' = most acceptable, '1' = least acceptable). Participants were required to complete the experiment in one hour, which gave them adequate time to rate each sentence. All participants completed the task within 30 min. Each sentence was displayed in its entirety on the screen along with the rating scale. Participants could click boxes to enter their rating, or use a numerical keypad. The order of presentation was randomized for each participant.

2.1.4. Data analysis

Data were analyzed using linear mixed-effects models, with fixed factors for experimental manipulations and a fully specified random effects structure, which included random intercepts and slopes for all fixed effects by participants and by items (Baayen, Davidson, & Bates, 2008; Barr, Levy, Scheepers, & Tily, 2013). Fixed factors were coded using Helmert contrasts (Vasishth & Broe, 2011; Venables & Ripley, 1999), which compared the effects of

Table 1
Sample set of items for Experiment 1.

any	GRAMMATICAL LICENSOR	*No authors* [that the critics recommended] have received any acknowledgment for a best-selling novel.
	IRRELEVANT LICENSOR	The authors [that *no critics* recommended] have received any acknowledgment for a best-selling novel.
	NO LICENSOR	The authors [that the critics recommended] have received any acknowledgment for a best-selling novel.
ever	GRAMMATICAL LICENSOR	*No authors* [that the critics recommended] have ever received acknowledgment for a best-selling novel.
	IRRELEVANT LICENSOR	The authors [that *no critics* recommended] have ever received acknowledgment for a best-selling novel.
	NO LICENSOR	The authors [that the critics recommended] have ever received acknowledgment for a best-selling novel.

grammaticality (grammatical vs. ungrammatical) and illusory licensing (irrelevant licensor vs. no licensor) within each NPI type (*ever* vs. *any*). Models were estimated using the *lme4* package (Bates, Maechler, & Bolker, 2011) in the R software environment (R Development Core Team, 2014). If there was a convergence failure or if the model converged but the correlation estimates were high, the random effects structure was simplified following Baayen et al. (2008). An effect was considered significant if the absolute t/z-value was greater than 2 (Gelman & Hill, 2007).

2.1.5. Results

The results of Experiment 1 are presented in Fig. 1. A main effect of grammaticality was observed for both NPIs, as ratings were higher for grammatical sentences relative to ungrammatical sentences (*ever*: $\hat{\beta} = -3.60$, SE = 0.41, $t = -8.631$; *any*: $\hat{\beta} = -3.56$, SE = 0.40, $t = -8.83$). No differences were found within the ungrammatical conditions for either NPI, and no interactions were observed ($ts < 2$).

2.1.6. Discussion

The results of Experiment 1 showed that sentences with *any* and *ever* are rated similarly based on untimed acceptability measures. Importantly, the main effect of grammaticality was highly significant for sentences with *any*, and the magnitude of this effect was similar across NPIs. These results indicate that we were successful in forcing the NPI reading of *any*, and they provide the basis for testing whether *any* is susceptible to illusory licensing using time-sensitive measures in the next set of experiments.

2.2. Experiment 2: Speeded acceptability judgments for any and ever

The goal of Experiment 2 was to test the claims that illusions are a general property of negative polarity phenomena, and that they reflect limitations of licensing mechanisms that are engaged after an NPI is encountered. To this end, we compared *any* and *ever* using speeded acceptability judgments. Speeded acceptability judgments provide a time-sensitive measure that has been shown to reliably elicit linguistic illusions by restricting the amount of time that comprehenders have to reflect on acceptability intuitions (Drenhaus et al., 2005; Wagers et al., 2009; Xiang et al., 2006). As in previous studies, an illusion is predicted to manifest as increased rates of acceptance for sentences involving an irrelevant licensor, relative to sentences with no licensor.

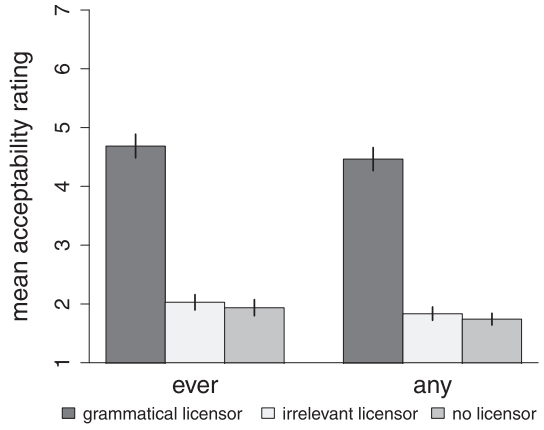

Fig. 1. Mean untimed acceptability ratings and standard error by participants for Experiment 1.

2.2.1. Participants

Participants were 18 native English speakers from the University of Maryland community. Participants were either compensated $10 or received credit in an introductory linguistics course. All participants were naïve to the purpose of the experiment. The speeded acceptability task lasted approximately 20 min, and was administered as part of a one-hour session involving unrelated experiments.

2.2.2. Materials

Experimental materials consisted of the same 36 sets of 6 items as in Experiment 1, with the same filler sentences.

2.2.3. Procedure

Sentences were presented on a desktop PC using Ibex. Sentences were presented one word at a time in the center of the screen in a rapid serial visual presentation (RSVP) paradigm at a rate of 300 ms per word. Participants were instructed to read each sentence carefully, and to judge whether each sentence was an acceptable sentence that a speaker of English might say. A response screen appeared for 3 s at the end of each sentence during which participants made a 'yes/no' response by button press. The forced choice judgments used in this experiment differed from the Likert scale ratings used in Experiment 1. This difference was motivated by the task. Whereas Likert scale ratings are more common with untimed measures, forced choice judgments are more common with speeded presentations. If participants waited longer than 3 s to respond, they were given feedback that their response was too slow. The order of presentation was randomized for each participant.

2.2.4. Data analysis

Data analysis followed the same steps as in Experiment 1 using logistic mixed-effects models, which are suitable for dependent measures that are categorical (i.e., 'yes' or 'no').

2.2.5. Results

Fig. 2 shows the percentage of 'yes' responses for the 6 experimental conditions. A main effect of grammaticality was observed for both NPIs, as grammatical sentences were more likely to be accepted than ungrammatical sentences (*ever*: $\hat{\beta} = 2.69$, SE = 0.62, $z = 4.30$; *any*: $\hat{\beta} = 3.04$, SE = 0.72, $z = 4.18$). But results for the ungrammatical conditions diverged. Contrasting profiles were observed for ungrammatical sentences with *any* and *ever*, reflected in an interaction between illusory licensing and NPI type

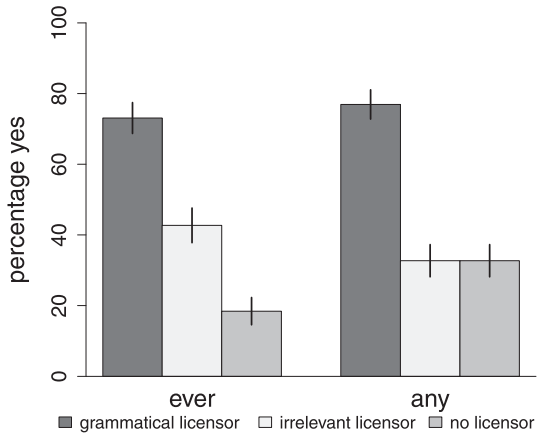

Fig. 2. Speeded acceptability judgments for Experiment 2. Error bars indicate standard error of the mean.

($\hat{\beta} = -1.83$, SE = 0.62, $z = -2.91$). This interaction was driven by an illusory licensing effect for *ever*: sentences with an irrelevant licensor were more likely to be accepted than sentences with no licensor ($\hat{\beta} = 1.36$, SE = 0.41, $z = 3.30$). We did not find evidence of a corresponding illusion for *any* ($z < 2$).

2.2.6. Discussion

Experiment 2 tested the claims that illusions are a general property of NPIs by comparing speeded acceptability judgments for *any* and *ever*. Results revealed a striking contrast between these two NPIs: the NPI *ever* showed a robust illusion, replicating previous findings using the same NPI or its German counterpart (e.g., Drenhaus et al., 2005; Xiang et al., 2006), but *any* did not. These results are unexpected under accounts that assume that all NPIs should behave similarly with respect to illusory licensing effects.

A potential concern with the results of Experiment 2 is the high rate of acceptance of ungrammatical sentences with *any*. It is possible, for example, that an illusion for *any* was masked by the increased rate of acceptance of sentences with no licensor. One factor that may have impacted judgments is the free-choice interpretation of *any*. Although the untimed ratings from Experiment 1 suggest that comprehenders ultimately favor a negative polarity reading of *any*, rapid comprehension mechanisms may have temporarily accessed the free-choice interpretation before the restricting abstract mass noun was encountered in the rapid serial presentation. It is possible, then, that in some portion of trials, a lingering free-choice interpretation could have influenced rapid judgments. We believe this is an unlikely possibility since the significant main effect of grammaticality suggests that we were successful in blocking the availability of the free-choice reading of *any*. In other words, if comprehenders systematically accessed a free choice interpretation of *any*, then they should have been untroubled by the absence of negation, contrary to fact. However, we investigate these possibilities further in the next experiment by examining the time course of NPI processing.

2.3. Experiment 3: Self-paced reading times for any and ever

The goal of Experiment 3 was to examine how comprehenders interpret *any* relative to *ever* during moment-by-moment processing. To this end, we tested the same items from Experiments 1–2 using self-paced reading. In self-paced reading, an illusion is predicted to appear as a reduced disruption in reading times for an unlicensed NPI in ungrammatical sentences with an irrelevant licensor, relative to ungrammatical sentences with no potential licensor. A non-illusion, by contrast, is predicted to appear as disrupted reading times for an unlicensed NPI in ungrammatical sentences, with no divergence in reading times between the two ungrammatical sentences.

2.3.1. Participants

Participants were 24 native speakers of English from the University of Maryland community. Participants were either compensated $10 or received credit in an introductory linguistics course. All participants were naïve to the purpose of the experiment. The task lasted approximately 35 min, and was administered as part of a one-hour session involving unrelated experiments.

2.3.2. Materials

Experimental materials consisted of the same 36 sets of 6 items as in Experiments 1 and 2. 72 grammatical fillers were also included, such that each participant read a total of 108 sentences. The filler sentences were of similar length and complexity to the NPI sentences, but lacked an NPI. Roughly half of the filler sentences included determiners similar to those in the NPI conditions and in similar positions to prevent the possibility that participants might develop superficial reading strategies based on the distribution of the determiners in the NPI sentences. Each sentence was followed by a comprehension question. Comprehension questions addressed various parts of the sentence to prevent the possibility that participants might develop superficial reading strategies that would allow them to extract only the information necessary to answer the comprehension question without fully interpreting the sentence. For ungrammatical sentences, comprehension questions addressed only the content prior to the NPI.

2.3.3. Procedure

Sentences were presented on a desktop PC in a moving-window display using Linger (Doug Rohde). Sentences were initially masked by dashes, with white spaces and punctuation intact. Participants pushed the space bar to reveal each word. Presentation was non-cumulative, such that the previous word was replaced with a dash when the next word appeared. Each sentence was followed by a 'yes/no' comprehension question. On-screen feedback was provided for incorrect answers. The order of presentation was randomized for each participant.

2.3.4. Data analysis

Only data from participants with at least 80% accuracy on the comprehension questions were used in the analysis. One participant was excluded due to accuracy below 80%. Average reading times were compared across conditions in the following regions of interest: the word immediately before the NPI (NPI − 1), the NPI (NPI), and two words following the NPI (NPI + 1, NPI + 2). We used the Box-Cox procedure (Box & Cox, 1964) to determine that a natural log would be the appropriate transformation to obtain normally distributed residuals (see Vasishth, Chen, Li, & Guo, 2013, for discussion about the importance of appropriately transforming reading time data). Statistical analyses were carried out over the untrimmed, log-transformed reading time data using linear mixed-effects models. We used Helmert contrast coding for experimental fixed effects, as in Experiments 1–2, with a fully-specified random effects structure, which included random intercepts and slopes for all fixed effects by participants and by items. If there was a convergence failure or if the model converged but the correlation estimates were high, the random effects structure was simplified following Baayen et al. (2008).

2.3.5. Results: ever

Fig. 3 (top) shows the average word-by-word reading times for sentences with *ever*. No effects were observed in the NPI − 1 region

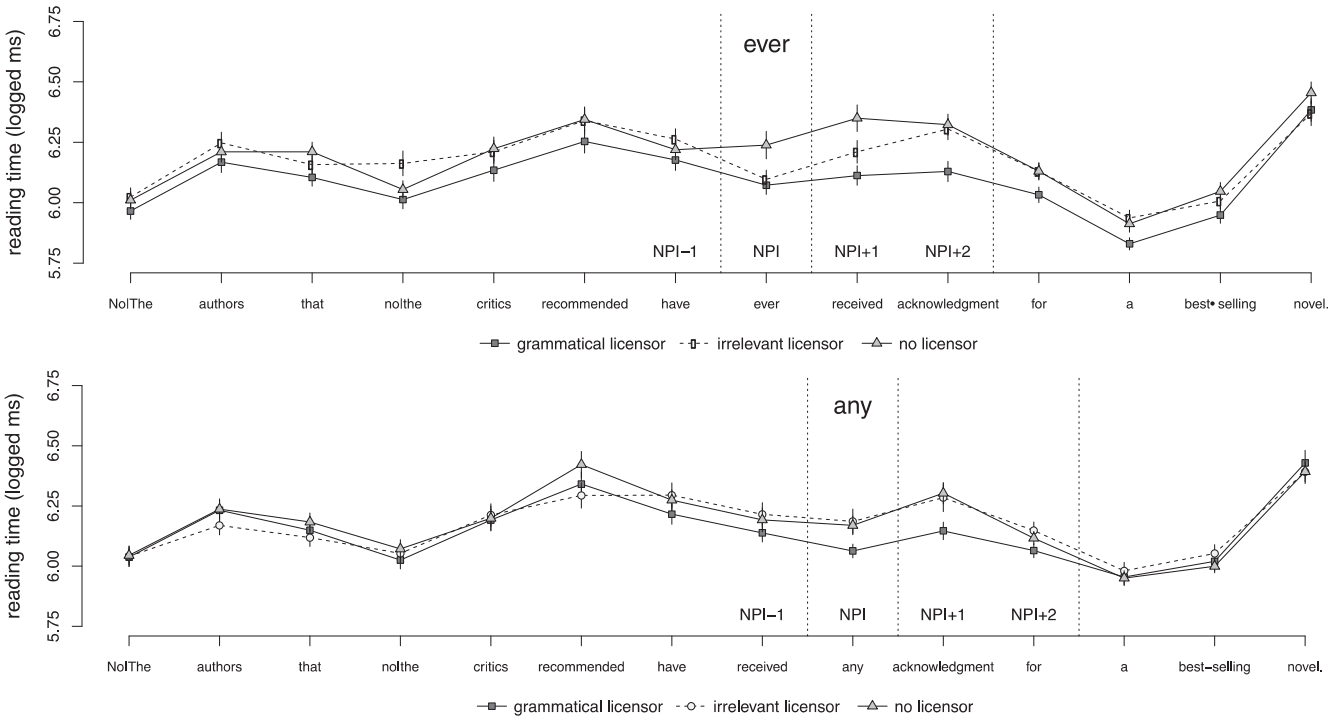

Fig. 3. Word-by-word reading times for sentences with the NPIs *ever* (top) and *any* (bottom) for Experiment 3. Error bars indicate standard error of the mean.

(*t*s < 2). A main effect of grammaticality was observed at the NPI, NPI + 1, and NPI + 2 regions, reflecting slower reading times for ungrammatical conditions relative to the grammatical condition (NPI: $\hat{\beta} = 0.12$, SE = 0.06, $t = 2.07$; NPI + 1: $\hat{\beta} = 0.22$, SE = 0.06, $t = 3.40$; NPI + 2: $\hat{\beta} = 0.22$, SE = 0.06, $t = 3.57$). An illusory licensing effect was observed in both the NPI and NPI + 1 regions, which manifested as a reduced disruption for sentences with an irrelevant licensor relative to sentences with no licensor (NPI: $\hat{\beta} = -0.14$, SE = 0.05, $t = -2.72$; NPI + 1: $\hat{\beta} = -0.13$, SE = 0.06, $t = -2.00$).

2.3.6. Results: any

Fig. 3 (bottom) shows the average word-by-word reading times for sentences with *any*. No effects were observed in the NPI − 1 region (*t*s < 2). A main effect of grammaticality was observed in the NPI, NPI + 1, and NPI + 2 regions, reflecting a slowdown for ungrammatical conditions relative to the grammatical condition (NPI: $\hat{\beta} = 0.15$, SE = 0.05, $t = 2.81$; NPI + 1: $\hat{\beta} = 0.19$, SE = 0.06, $t = 3.10$; NPI + 2: $\hat{\beta} = 0.09$, SE = 0.04, $t = 2.04$). There was no reliable evidence of a reduced reading time disruption due to illusory licensing in any region (*t*s < 2).

2.3.7. Results: Direct comparison of any and ever

A direct comparison of ungrammatical sentences with *ever* and *any* at the NPI regions revealed an interaction between NPI type and illusory licensing ($\hat{\beta} = -0.15$, SE = 0.07, $t = -2.04$). This interaction was driven by the contrasting reading times in the ungrammatical conditions for sentences with *ever*.

2.3.8. Discussion

Experiment 3 used self-paced reading to examine how comprehenders interpret *any* relative to *ever* during incremental processing. This test was important to determine if comprehenders interpret *any* as a free-choice item in our items during rapid comprehension. Results replicated the contrasting profiles for *any* and *ever* seen in speeded acceptability judgments in Experiment 2. As in Experiment 2, the NPI *ever* elicited a reliable illusion, but *any* did not. Importantly, reading times showed rapid detection of unlicensed NPIs for both *any* and *ever* immediately at the NPI regions, which is the earliest point at which we could measure sensitivity to the presence and location of a potential licensor using self-paced reading. Immediate sensitivity to the presence and location of a licensor for sentences with *any* is not expected if comprehenders initially interpreted *any* as a free-choice item, and suggests that the structural conditions on NPI licensing impact rapid comprehension, at least selectively.

To the best of our knowledge, these results provide the first demonstration that the NPI illusion can be selectively switched "on" and "off". The observed contrast is surprising, as it is unexpected under existing accounts that assume that illusions should generalize across items, and that illusions reflect faulty licensing mechanisms. These results indicate that illusions do not apply across-the-board, which further suggests that they cannot simply be due to faulty licensing mechanisms, as these mechanisms are expected to be engaged after any type of NPI is encountered. In Section 6, we provide an explicit demonstration of how the current results diverge from the predictions of existing models using computational simulations. But first, we explore the cause of this behavior in Experiments 4–6.

The contrasting profiles observed for *any* and *ever* could reflect a range of possibilities, including misrepresentation of the input, inherent lexical differences between *any* and *ever*, or differences based on the position or timing of the NPI in the sentence. For instance, assuming a noisy-channel model of sentence comprehension (e.g., Gibson et al., 2013; Levy, 2008), comprehenders may have maintained some uncertainty about whether the input contained the NPI *ever* or its phonological and orthographic near-neighbor *never*. It is possible that in some portion of trials, comprehenders mis-represented *ever* as *never* in contexts that could not support a grammatical interpretation of the NPI. The NPI *any*, by contrast, would be less likely to be misrepresented as one of its near neighbors, e.g., *many*, as this would result in semantic incompatibility with the abstract mass term that followed,

e.g. ??*many satisfaction*. However, this hypothesis does not capture the contrast between the irrelevant licensor and no licensor conditions for sentences with *ever*. Specifically, it does not predict that misrepresentation should apply more in sentences with an irrelevant licensor, which showed greater susceptibility to illusory licensing than sentences with no licensor.

Another possibility is that the contrast could reflect inherent lexical differences between *ever* and *any*. Current theories of negative polarity phenomena (e.g., Chierchia, 2006; Kadmon & Landman, 1993; see Giannakidou, 2011, for a review) treat these items similarly with respect to their semantic/pragmatic properties, but differences in syntactic category could impact how the parser integrates these two items into the existing structure during real-time comprehension. For example, attachment of a determiner like *any* may be relatively straightforward, but there may be uncertainty about the attachment position of the adverb *ever*. This uncertainty may force comprehenders to rely on alternative, "rough-and-ready" licensing heuristics such as the simple presence or absence of a licensor, leading to an illusion. However, we believe that the contrast between *ever* and *any* is unlikely to be a consequence of lexical considerations, given that *ever* and *any* showed identical profiles in untimed judgments in Experiment 1.

Alternatively, the contrast could reflect differences based on the position of the NPI. For example, reactivation of the subject NP upon encountering the main clause verb could have reduced interference from the irrelevant licensor for *any*, which appeared immediately after the verb. No such activation bias would be available for *ever* since it appeared before the main verb. Additionally, differences in the position of the NPI in the sentence impact the amount of time that comprehenders have to interpret the licensing context, and it is possible that the on/off behavior of the illusion depends on the timing of when the encoding of the licensing context is probed for NPI licensing. Specifically, the later position of *any* relative to *ever* gave comprehenders more time to interpret the licensing context, compute the scope relations for the licensors, and encode that information in memory, facilitating detection of an unlicensed NPI. We attempt to distinguish these alternatives in Experiments 4–6.

3. The impact of position on the NPI illusion

Experiments 4–5 were designed to distinguish two sets of possibilities. On the one hand, the contrasting profiles observed for *ever* and *any* in Experiments 2–3 could reflect lexical misrepresentation of the NPIs or inherent lexical differences between the NPIs. On the other hand, the contrasting profiles could reflect differences based on the position or timing of the NPI. To distinguish these two sets of alternatives, we held the NPI constant, testing only *ever*, and we manipulated whether the NPI appeared immediately before the verb or in a later position in a sentential complement clause. If the contrasting profiles observed for *ever* and *any* in Experiments 2–3 reflect misrepresentation or lexical differences, then we should find illusory licensing effects across both positions of *ever*. If, however, the contrasting profiles reflect differences based on the position or timing of the NPI in the sentence, then we should see the same modulation of the illusion observed in Experiments 2–3.

3.1. Experiment 4: Speeded acceptability judgments for main clause and embedded clause ever

3.1.1. Participants

Participants were 18 native speakers of English from the University of Maryland. Participants were either compensated $10 or received credit in an introductory linguistics course. All participants were naïve to the purpose of the experiment. The task lasted approximately 20 min, and was administered as part of a one-hour session involving unrelated experiments.

3.1.2. Materials

Experimental materials consisted of 36 sets of 6 items, which varied in the presence and location of an NPI licensor (grammatical licensor|irrelevant licensor|no licensor) and the position of the NPI *ever* relative to the potential licensors (main clause|embedded clause). The structure of the licensing context was identical to that of Experiments 1–3. The NPI *ever* appeared either immediately before the main clause verb (main clause) or in a sentential complement clause (embedded clause). Unlike in Experiments 1–3, the relative clause was not followed by an auxiliary verb to ensure that the sentences would sound natural in the long-distance NPI licensing configurations. This decision was based on the intuition that the propositional attitude verbs that govern the embedded clause sound unnatural with an auxiliary verb, e.g., . . . *ever anticipated that* . . . vs. (?) . . . *have ever anticipated that* . . .. The sentences within each item set were identical in all other respects. Extra care was taken to ensure that the NPIs were plausible in both short- and long-distance licensing configurations, and that the subject NPs were also plausible in all positions, regardless of the presence or location of the negative determiner. An example set of items is provided in Table 2.

Each participant read 108 sentences, consisting of 36 NPI sentences and 72 filler sentences. The 36 sets of NPI items were distributed across 6 lists in a Latin Square design. The filler sentences were of similar length and complexity to the NPI sentences, but lacked an NPI. Roughly half of the filler sentences included determiners similar to those in the NPI sentences to prevent the possibility that participants might develop superficial reading strategies. Materials were balanced such that across the experiment, half of the sentences were ungrammatical.

3.1.3. Procedure and data analysis

Experiment 4 used speeded acceptability judgments, following the same procedure used in Experiment 2. Data analysis followed the same steps as in Experiment 2.

Table 2
Sample set of items for Experiments 4–5.

Main clause *ever*	GRAMMATICAL LICENSOR	*No journalists* [that the editors recommended for the assignment] ever thought that the readers would understand the complicated situation.
	IRRELEVANT LICENSOR	The journalists [that *no editors* recommended for the assignment] ever thought that the readers would understand the complicated situation.
	NO LICENSOR	The journalists [that the editors recommended for the assignment] ever thought that the readers would understand the complicated situation.
Embedded clause *ever*	GRAMMATICAL LICENSOR	*No journalists* [that the editors recommended for the assignment] thought that the readers would ever understand the complicated situation.
	IRRELEVANT LICENSOR	The journalists [that *no editors* recommended for the assignment] thought that the readers would ever understand the complicated situation.
	NO LICENSOR	The journalists [that the editors recommended for the assignment] thought that the readers would ever understand the complicated situation.

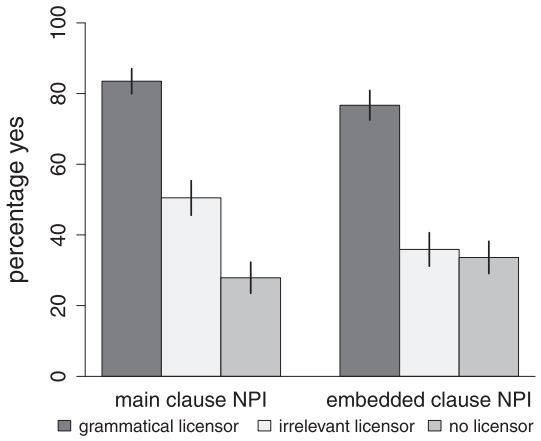

Fig. 4. Speeded acceptability judgments for Experiment 4. Error bars indicate standard error of the mean.

3.1.4. Results

Fig. 4 shows the percentage of 'yes' responses for the 6 experimental conditions. Results showed a main effect of grammaticality for both NPI positions, as grammatical sentences were more likely to be accepted than ungrammatical sentences (main clause: $\hat{\beta} = -3.64$, SE = 0.75, $z = -4.83$; embedded clause: $\hat{\beta} = -3.00$, SE = 0.62, $z = -4.82$). Results for the ungrammatical conditions diverged. Contrasting profiles were observed for ungrammatical sentences with a main clause NPI and embedded clause NPI, as there was an interaction between illusory licensing and NPI position ($\hat{\beta} = -1.05$, SE = 0.45, $z = -2.30$). This interaction was driven by an illusory licensing effect for main clause NPIs ($\hat{\beta} = 1.24$, SE = 0.41, $z = 3.01$). We did not find evidence of a corresponding illusion for embedded clause NPIs ($z < 2$).

3.1.5. Discussion

Experiment 4 used speeded acceptability judgments to test whether the contrasting profiles observed in the previous experiments reflect lexical considerations (i.e., misrepresentation of the NPIs or inherent differences between NPIs) or the position/timing of the NPI. We achieved this by holding the NPI constant and by manipulating whether the NPI appeared near the potential licensors or further away in an embedded clause. Speeded acceptability judgments revealed the same modulation of illusions observed in the *any*/*ever* comparison: *ever* showed a robust illusion when it appeared in the main clause, but not when it appeared in an embedded clause. These results suggest that the contrasting profiles are a consequence of either the position or timing of the NPI in the sentence, relative to the potential licensors.

3.2. Experiment 5: Self-paced reading times for main clause and embedded clause *ever*

Experiment 5 used self-paced reading to examine the time course of the contrasting profiles for main clause and embedded clause *ever* observed in Experiment 4.

3.2.1. Participants

Participants were 30 native speakers of English from the University of Maryland. Participants were either compensated $10 or received credit in an introductory linguistics course. All participants were naïve to the purpose of the experiment. The task lasted approximately 35 min, and was administered as a part of one-hour session involving unrelated experiments.

3.2.2. Materials

Experimental materials consisted of the same 36 sets of 6 items as in Experiment 4. 72 grammatical fillers were also included, such that each participant read a total of 108 sentences. Each sentence was followed by a 'yes/no' comprehension question.

3.2.3. Procedure and data analysis

Experiment 5 used self-paced reading, following the same procedure used in Experiment 3. Data analysis followed the same steps as in Experiment 3. The Box-Cox procedure determined that a natural log was the appropriate transformation.

3.2.4. Results: main clause *ever*

Fig. 5 (top) shows the average word-by-word reading times for sentences with *ever* in the main clause position. No effects were observed in the NPI − 1 or NPI regions (*t*s < 2). A main effect of grammaticality was observed in the NPI + 1 region, carried by a

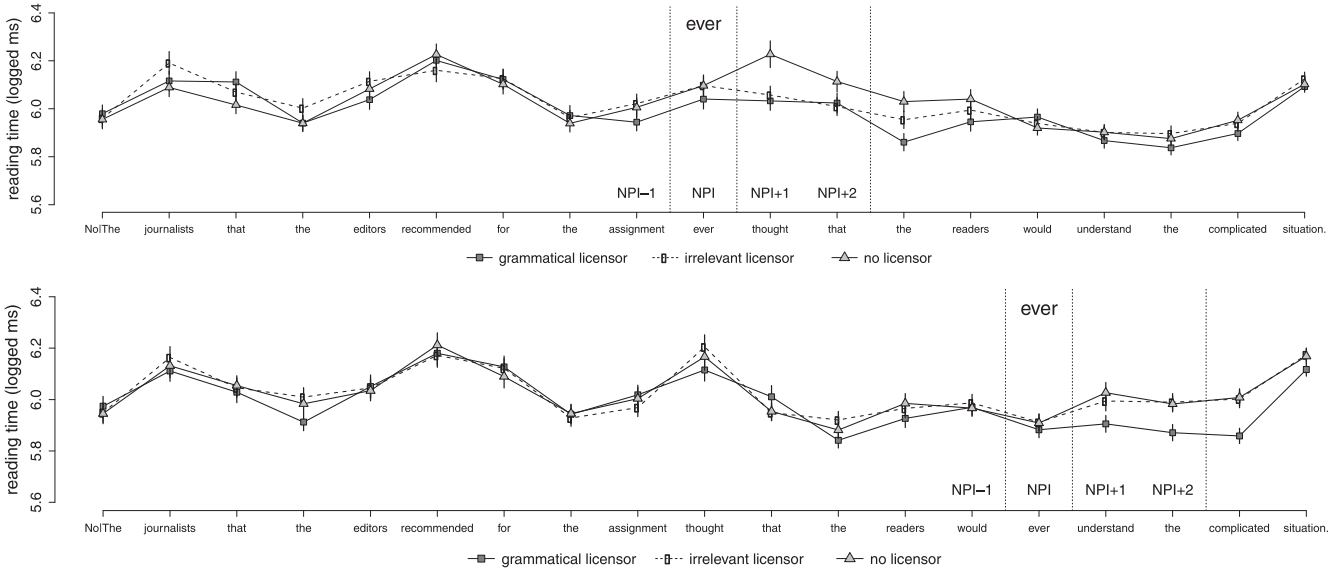

Fig. 5. Word-by-word reading times for sentences with *ever* in the main clause (top) and in the embedded clause (bottom) for Experiment 5. Error bars indicate standard error of the mean.

slowdown for the ungrammatical no licensor condition relative to the grammatical condition (NPI: $\hat{\beta} = 0.14$, SE = 0.05, $t = 2.50$). An illusory licensing effect was observed in NPI + 1 and NPI + 2 regions, and was reflected in a reduced disruption for sentences with an irrelevant licensor relative to sentences with no licensor (NPI + 1: $\hat{\beta} = -0.16$, SE = 0.05, $t = -3.39$; NPI + 2: $\hat{\beta} = -0.10$, SE = 0.04, $t = -2.23$).

3.2.5. Results: embedded clause ever

Fig. 5 (bottom) shows the average word-by-word reading times for sentences with *ever* in the embedded clause position. No effects were observed in the NPI − 1 or NPI regions ($ts < 2$). A main effect of grammaticality was observed in the NPI + 1, and NPI + 2 regions, carried by a slowdown for the ungrammatical no licensor condition relative to the grammatical condition (NPI + 1: $\hat{\beta} = -0.14$, SE = 0.04, $t = 3.01$; NPI + 2: $\hat{\beta} = 0.15$, SE = 0.04, $t = 3.40$). There was no reliable evidence of a reduced disruption due to illusory licensing in any region ($ts < 2$).

3.2.6. Results: direct comparison of main clause and embedded clause ever

A direct comparison of ungrammatical sentences at the NPI spill-over regions revealed an interaction between NPI position and illusory licensing ($\hat{\beta} = 0.13$, SE = 0.06, $t = 2.11$). This interaction was driven by the contrasting reading times in the two ungrammatical conditions for sentences with *ever* in the main clause position.

3.2.7. Discussion

Experiment 5 used self-paced reading to further investigate whether the contrasting profiles observed in the previous experiments reflect lexical considerations (i.e., misrepresentation of the NPIs or inherent differences between NPIs) versus differences based on the position/timing of the NPI in the sentence. Self-paced reading times replicated the contrasting profiles observed in Experiment 4: *ever* showed a robust illusion when it appeared in the main clause, but not when it was further away in an embedded clause. Taken together, the results of Experiments 4–5 do not support the hypotheses that the selectivity of the NPI illusion depends on the lexical status of the NPI or that comprehenders selectively misrepresent *ever*. The finding that a single NPI can exhibit contrasting profiles with respect to illusory licensing suggests that additional factors beyond lexical differences or lexical misrepresentation are likely responsible for the observed effects.

The results of the position manipulation are surprising, as they are not predicted by existing accounts. Existing accounts predict illusions to generalize across licensing configurations, since the illusion is attributed to faulty licensing mechanisms that are engaged after an NPI is encountered. However, we found that the illusion can be reliably switched on and off depending on the position of the NPI relative to the potential licensors. These results, taken together with the contrasting profiles observed in Experiments 2–3, suggest that NPI illusions cannot simply be due to faulty licensing mechanisms that are engaged after the NPI is encountered.

The key finding from Experiments 4–5 is that the on/off behavior of the NPI illusion depends on the position of the NPI relative to the potential licensors. It remains unclear why the position of the NPI would determine the presence or absence of the illusion. One possibility is that the on/off behavior could be a consequence of the material between the licensors and the NPIs that the differences in position introduce. As suggested earlier, reactivation of the subject NP by the main verb could reduce interference from the irrelevant licensor for NPIs that occupy a position after the main verb. Alternatively, the cause of the effect could be extended time between the potential licensors and the NPI. For example, increased distance between the licensors and the NPI gives comprehenders more time to appropriately interpret and encode the licensing context, which could facilitate detection of a subsequent unlicensed NPI. Under this view, the presence or absence of an illusion would depend on the timing of when the encoding of the licensing context is probed for NPI licensing. We attempt to distinguish these possibilities in the next experiment.

4. The impact of time on the NPI illusion

Experiment 6 used speeded acceptability judgments to determine whether the on/off behavior of the NPI illusion is a consequence of the structural position of the NPI or the timing of the NPI relative to the potential licensors.[3] To distinguish these hypotheses, we held constant the structural and linear position of the NPI *ever*, and manipulated the position of a parenthetical phrase to vary the time between the potential licensors and the NPI, as illustrated in Table 3. If the structural or linear position of the NPI is the key to switching the illusion on/off, then the placement of the parenthetical phrase should not impact the illusion, since the position of the NPI was held constant. However, if the key is the timing of the NPI, then we might observe the same modulation of the illusion seen earlier.

4.1. Experiment 6: Speeded acceptability judgments for the timing manipulation

4.1.1. Participants

Participants were 18 native speakers of English who were recruited using Amazon Mechanical Turk. Participants were compensated $2.50. The experiment lasted approximately 15 min.

4.1.2. Materials

Experimental materials consisted of 36 sets of 6 items, which varied in the presence and location of an NPI licensor (grammatical licensor|irrelevant licensor|no licensor) and the position of a parenthetical phrase (intervening|non-intervening). The structure of the licensing context was identical to that of Experiments 1–5. Across all conditions the NPI *ever* appeared in the same structural and linear position immediately before the main verb. A 4–7 word parenthetical phrase appeared either at the beginning of the sentence (non-intervening) or between the main auxiliary and the NPI (intervening). The intervening position was chosen to signal the intended main clause predicate attachment of the parenthetical clause. Care was taken in the construction of the materials to ensure that the sentences were plausible and identical in all respects aside from the licensing of the NPI. We paid particular attention to ensuring that the parenthetical phrase was similarly plausible across all positions. To ensure that the parenthetical phrase did not specifically highlight the licensors, the parenthetical phrase never directly engaged or referred to any component of the complex subject NP. Following recent work by Dillon and colleagues, we assume that the content of parenthetical structures is processed independently of the embedding structure, and hence does not interfere with the material in the embedding structure (Dillon, Clifton, & Frazier, 2014).

Each participant read 108 sentences, consisting of 36 NPI sentences and 72 filler sentences. The 36 sets of NPI items were dis-

[3] Experiment 6 is not accompanied by a self-paced reading task. Based on the parallel findings between speeded acceptability judgments and self-paced reading times found in Experiments 2–5 and in previous studies, we did not anticipate a difference across methodologies, and a follow-up experiment using self-paced reading was not conducted.

Table 3
Sample set of items for Experiment 6.

Non-intervening parenthetical	GRAMMATICAL LICENSOR	As the editors mentioned, *no authors* [that the critics recommended for the assignment] have ever received a pay raise.
	IRRELEVANT LICENSOR	As the editors mentioned, the authors [that *no critics* recommended for the assignment] have ever received a pay raise.
	NO LICENSOR	As the editors mentioned, the authors [that the critics recommended for the assignment] have ever received a pay raise.
Intervening parenthetical	GRAMMATICAL LICENSOR	*No authors* [that the critics recommended for the assignment] have, as the editor mentioned, ever received a pay raise.
	IRRELEVANT LICENSOR	The authors [that *no critics* recommended for the assignment] have, as the editor mentioned, ever received a pay raise.
	NO LICENSOR	The authors [that the critics recommended for the assignment] have, as the editor mentioned, ever received a pay raise.

tributed across 6 lists in a Latin Square design. The filler sentences were of similar length and complexity to the NPI sentences, but lacked an NPI. Roughly half of the filler sentences included determiners similar to those in the NPI sentences and in similar positions to prevent superficial reading strategies. Materials were balanced such that across the experiment half of the sentences were ungrammatical.

4.1.3. Procedure and data analysis

Experiment 6 used speeded acceptability judgments, following the same procedure used in Experiments 2 and 4. Data analysis followed the same steps as in Experiments 2 and 4.

4.1.4. Results

Fig. 6 shows the percentage of 'yes' responses for the 6 experimental conditions. Results showed a main effect of grammaticality for sentences with an intervening and non-intervening parenthetical phrase, as grammatical sentences were more likely to be accepted than ungrammatical sentences (intervening parenthetical phrase: $\hat{\beta} = -3.95$, SE = 0.58, $z = -6.71$; non-intervening parenthetical phrase: $\hat{\beta} = -5.29$, SE = 1.75, $z = -3.02$). As in the previous experiments, results for the ungrammatical conditions diverged. Contrasting profiles were observed for ungrammatical sentences with intervening and non-intervening parenthetical phrases, as there was an interaction between illusory licensing and parenthetical phrase position ($\hat{\beta} = -0.95$, SE = 0.45, $z = -2.07$). This interaction was driven by an illusory licensing effect for sentences with a non-intervening parenthetical phrase ($\hat{\beta} = 0.74$, SE = 0.29, $z = 2.55$). We did not find evidence of a corresponding illusion for sentences with an intervening parenthetical phrase ($z < 2$).

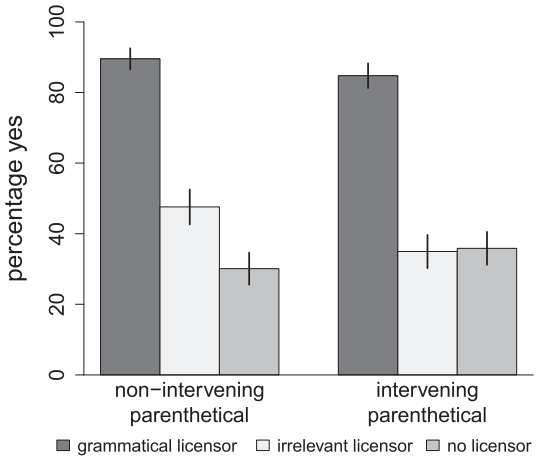

Fig. 6. Speeded acceptability judgments for Experiment 6. Error bars indicate standard error of the mean.

4.1.5. Discussion

The goal of Experiment 6 was to determine whether the on/off behavior of the NPI illusion is a consequence of the structural position of the NPI or the timing of the NPI relative to the potential licensors. To this end, we held constant the linear and structural position of the NPI, and manipulated the position of a parenthetical phrase, thus varying the time between the potential licensors and the NPI. Speeded acceptability judgments revealed the same modulation of the illusion reported in Experiments 2–5: sentences with a non-intervening parenthetical phrase showed the illusion, but sentences with an intervening parenthetical phrase did not. These results suggest that the presence of the NPI illusion depends on the amount of time between the licensors and the NPI, rather than the structural or linear position of the NPI in the sentence.

Thus far, we have found consistent findings across experiments. We have discovered that the NPI illusion can be reliably switched on/off, which is unexpected under existing accounts. Existing accounts predict that the illusion should generalize across lexical items and licensing configurations, since the illusion has been attributed to faulty licensing mechanisms that are engaged after an NPI is encountered. However, Experiment 6 showed that the illusion can be reliably switched on/off, depending on the relative position of the potential licensors and the NPI. In Section 6, we provide an explicit demonstration of how the current results diverge from the predictions of existing models using computational simulations. But first, we present the results of a follow-up study that was designed to test whether other linguistic illusions can be switched on/off in a similar fashion, to better understand the cause of this behavior.

5. The impact of time on agreement illusions

Experiments 2–6 showed that the NPI illusion can be switched off by simply moving the NPI further away from the potential licensors. In order to understand the cause of this effect, it is important to know whether other linguistic illusions can be modulated in a similar fashion. To this end, Experiment 7 tested whether the manipulation of parenthetical phrase placement used with NPIs in Experiment 6 would have a similar impact on agreement illusions. If the mechanisms responsible for switching the NPI illusion on/off are shared across illusions, then we might expect the passage of time to switch off agreement illusions just as it did for NPI illusions.

5.1. Experiment 7: Speeded acceptability judgments for the agreement timing manipulation

5.1.1. Participants

Participants were 18 native speakers of English who were recruited using Amazon Mechanical Turk. Participants were compensated $2.50. The experiment lasted approximately 15 min.

5.1.2. Materials

Experimental materials consisted of 48 sets of 8 items, modified from Wagers et al. (2009) and Pearlmutter et al. (1999). The experimental conditions consisted of a 2 × 2 × 2 factorial design, which crossed the factors grammaticality (grammatical|ungrammatical), distractor number (singular|plural), and parenthetical phrase position (non-intervening|intervening). In all conditions, the subject head noun was modified by a prepositional phrase (PP) that contained the irrelevant licensor, and the agreeing verb was a present tense form of *be*. Grammaticality was manipulated by varying the number of the verb such that it matched or mismatched the number of the subject (grammatical = *was*, ungrammatical = *were*). Distractor number was manipulated such that the number of the distractor either matched or mismatched the number of the agreeing verb (plural|singular). A 4–7 word parenthetical phrase was placed either at the beginning of the sentence (non-intervening) or between a post-PP adverb and the agreeing verb (intervening). The parenthetical phrase never directly engaged or referred to any component of the complex subject NP. An example set of items is provided in Table 4.

One difference between the agreement items and the NPI items in Experiments 1–6 is that the NPI items included one grammatical condition, whereas the agreement items included two grammatical conditions. The additional grammatical conditions were included to test for the grammatical asymmetry observed in previous studies. As discussed in the Introduction, agreement illusions are typically observed only in ungrammatical sentences (Dillon et al., 2013; Wagers et al., 2009), a finding that provides the primary evidence that illusions reflect limitations of memory access mechanisms, rather than misrepresentation of the subject NP.

Each participant read 144 sentences, consisting of 48 agreement sentences and 96 filler sentences. The 48 sets of agreement items were distributed across 8 lists in a Latin Square design. The filler sentences were of similar complexity to the agreement sentences. Materials were balanced such that across the experiment half of the sentences were ungrammatical.

5.1.3. Procedure and data analysis

Experiment 7 used speeded acceptability judgments, following the same procedure used in Experiments 2, 4, and 6. Data analysis followed the same steps as in Experiments 2, 4, and 6.

5.1.4. Results

Fig. 7 shows the percentage of 'yes' responses for the 8 experimental conditions. A main effect of grammaticality was observed for sentences with non-intervening and intervening parenthetical phrases, as grammatical sentences were more likely to be accepted than ungrammatical sentences (non-intervening: $\hat{\beta} = -4.99$, SE = 0.60, $z = -8.30$; intervening: $\hat{\beta} = -3.71$, SE = 0.46, $z = -7.98$). No effects were observed within the grammatical conditions for sentences with a non-intervening parenthetical phrase ($z < 2$). An illusion of ungrammaticality was observed for grammatical sentences with an intervening parenthetical phrase, as grammatical sentences were more likely to be rejected when a plural distractor was present ($\hat{\beta} = 1.25$, SE = 0.56, $z = 2.20$). In the ungrammatical conditions, illusions of grammaticality were observed for sentences with intervening and non-intervening parenthetical phrases (non-intervening: $\hat{\beta} = 1.64$, SE = 0.47, $z = 3.44$; intervening: $\hat{\beta} = 1.77$, SE = 0.48, $z = 3.69$).

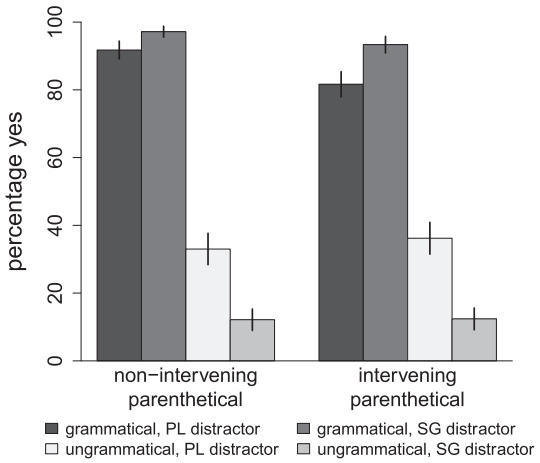

Fig. 7. Speeded acceptability judgments for Experiment 7. Error bars indicate standard error of the mean.

5.1.5. Discussion

The motivation behind Experiment 7 was to better understand the cause of the on/off behavior of the NPI illusion by testing whether the mechanisms responsible for modulating the NPI illusion are shared across different linguistic illusions. Specifically, we tested whether the manipulation of parenthetical phrase placement used in Experiment 6 would have a similar impact on agreement illusions as it did for NPI illusions. The results from Experiment 7 revealed that the same manipulation had no impact on agreement illusions.

These results indicate that the mechanisms responsible for switching the NPI illusion on/off are not shared across illusions, which further suggests that NPI illusions and agreement illusions have different causes, as argued in previous research (Xiang et al., 2009, 2013; see Introduction for discussion). Earlier research found that agreement and NPI illusions are qualitatively similar in

Table 4
Sample set of items for Experiment 7.

Non-intervening parenthetical	GRAMMATICAL PL DISTRACTOR	According to the janitor, the key to the cabinets probably was destroyed by the fire.
	GRAMMATICAL SG DISTRACTOR	According to the janitor, the key to the cabinet probably was destroyed by the fire.
	UNGRAMMATICAL PL DISTRACTOR	According to the janitor, the key to the cabinets probably were destroyed by the fire.
	UNGRAMMATICAL SG DISTRACTOR	According to the janitor, the key to the cabinet probably were destroyed by the fire.
Intervening parenthetical	GRAMMATICAL PL DISTRACTOR	The key to the cabinets, according to the janitor, probably was destroyed by the fire.
	GRAMMATICAL SG DISTRACTOR	The key to the cabinet, according to the janitor, probably was destroyed by the fire.
	UNGRAMMATICAL PL DISTRACTOR	The key to the cabinets, according to the janitor, probably were destroyed by the fire.
	UNGRAMMATICAL SG DISTRACTOR	The key to the cabinet, according to the janitor, probably were destroyed by the fire.

certain respects, e.g., both arise in contexts that contain a partially-matched distractor and show similar behavioral profiles. These findings led to the attractive generalization in the psycholinguistics literature that these phenomena may have a single underlying cause, such as the product of general memory access processes. However, this generalization is not supported by the linguistics literature, where it is uncontroversial that agreement and NPI licensing involve fundamentally different mechanisms. Specifically, whereas agreement involves an item-to-item structural dependency between specific morphosyntactic features (e.g., person, number, gender), a growing number of studies on NPI licensing have claimed that NPIs rely on semantic and pragmatic mechanisms, rather than a direct structural relation between an NPI and a specific licensor (Chierchia, 2006; Giannakidou, 2011; Horn, 2010; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Linebarger, 1987). Our finding that agreement and NPIs show contrasting profiles aligns with these representational differences, which invites the conclusion that the contrasting profiles reflect a fundamental difference in how we mentally encode and navigate syntactic versus semantic/pragmatic representations.

A concern with the results of Experiment 7 is the illusion of ungrammaticality for grammatical sentences with an intervening parenthetical phrase. Illusions of ungrammaticality are rare, but have been reported for agreement processing (Lago, Shalom, Sigman, Lau, & Phillips, 2015; Wagers, 2008). Illusions of ungrammaticality might be expected if the plural distractor caused faulty encoding of the subject's number feature, as proposed in the percolation model of agreement illusions (Eberhard et al., 2005). However, recent work suggests that illusions of ungrammaticality for agreement likely reflect either processing difficulty that arises in the late stages of agreement processing, or a Type I error (Lago et al., 2015). We believe that this effect reflects a Type 1 error because the effect was absent in sentences with a non-intervening parenthetical phrase and illusions of ungrammaticality are rarely observed.

Another potential concern is that the parenthetical phrases used in Experiment 7 included an NP with features that overlap with the retrieval cues from the auxiliary verb. In the example item set (Table 4), this NP is *the janitor*, which yields a partial match in the grammatical *was* conditions. As such, the grammatical conditions from our test of agreement had two distractors, while the corresponding conditions from the NPI experiments had only one. We think that it is unlikely that this difference affected the contrast between the NPI and agreement experiments for several reasons. First, our focus here is on how the position of the parenthetical phrase impacts susceptibility to illusory licensing in the ungrammatical conditions. Illusory licensing effects are still observed in the ungrammatical conditions, which do not contain the additional distractor. Second, accounts of similarity-based interference (e.g., Gordon et al., 2001; Lewis & Vasishth, 2005; Van Dyke & McElree, 2006) predict that the presence of additional NPs that match the retrieval cues should lead to increased processing difficultly. In contrast, we found that sentences with the additional feature-matching distractor received the highest ratings. Third, recent research suggests that parenthetical content and its embedding utterance may be represented in distinct memory stores, which might limit the degree to which parenthetical material interferes at retrieval when processing main clause content (Dillon et al., 2014).

The key finding from Experiment 7 is that the same on/off profile observed for NPI illusions does not extend to agreement illusions, which suggests that the mechanisms responsible for the modulation of the NPI illusion are not shared across illusions. These findings narrow down the cause of this surprising behavior. In particular, they suggest that the modulation of the NPI illusion is likely tied to the semantic/pragmatic representations and interpretive mechanisms responsible for normal NPI licensing. We develop this proposal further in the General Discussion. But first, we turn to our final study, which used computational simulations to examine how the current results for NPI illusions diverge from the predictions of existing models.

6. Testing model assumptions

We conducted computational simulations to test the assumptions of the Vasishth et al. (2008) model of the NPI illusion, since it is the most explicit model of the phenomenon. Specifically, our goal was to determine whether our NPI findings would fall out from the model. To derive the predictions, we used a variant of the ACT-R model of sentence processing (Adaptive Character of Thought—Rational; Anderson, 1990) based on the equations described in Lewis and Vasishth (2005) and Vasishth et al. (2008).[4] This model is capable of generating precise quantitative fits to observed data, and it has been used to accurately capture NPI illusions in previous research (Vasishth et al., 2008).

6.1. The model

In the model, linguistic items are encoded as "chunks" in a content-addressable memory (Kohonen, 1980), and hierarchical structure arises as a consequence of a pointer mechanisms inspired by the attribute-value matrices from Head-driven Phrase Structure Grammar (HPSG; Pollard & Sag, 1994). Linguistic dependencies are formed using a general retrieval mechanism that probes all items in memory, in parallel, using a set of retrieval cues that target specific features of individual memory chunks. Memory chunks are differentially activated based on their match to the retrieval cues, and the success of retrieving a chunk is proportional to the chunk's overall activation at the time of retrieval. Linguistic illusions are explained in this model as misretrieval of irrelevant chunks that partially match the retrieval cues (Dillon et al., 2013; Vasishth et al., 2008; Wagers et al., 2009).

Previous implementations of the ACT-R model have included a wide range of processing modules for lexical access, memory retrieval, and syntactic parsing (e.g., Lewis & Vasishth, 2005; Vasishth et al., 2008). However, the simulations here focus solely on retrieval latencies, and abstract away from the contribution of ancillary modules by stipulating the chunks in memory and retrievals required to parse the sentence. There are certainly additional processes associated with sentence parsing that contribute to behavioral measures, but for current purposes, we adopt the standard assumption that longer retrieval latencies entail longer reading times.

We simulated the hypothesized retrievals involved in the key manipulations from Experiments 2–6. For each condition, a schedule of constituent creation times and retrievals was estimated from the reading times in Experiments 3 and 5. Retrievals associated with the processing of a given constituent, such as retrieval of a subject by a verb, or retrieval of a negative licensor by an NPI, occurred 200 ms after the creation of the retrieval trigger. An important assumption in this model is that NPI licensing involves retrieval of overt negation in an item-to-item dependency. This assumption differs from leading linguistic accounts, which assume that NPIs are licensed by the semantics and pragmatics of entire propositions, rather than by a direct relation with overt negation (Chierchia, 2006; Giannakidou, 2011; Horn, 2010; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Linebarger, 1987). Nevertheless, we implemented NPI licensing as described in Vasishth

[4] The code for the simulations was originally developed by Badecker and Lewis (2007).

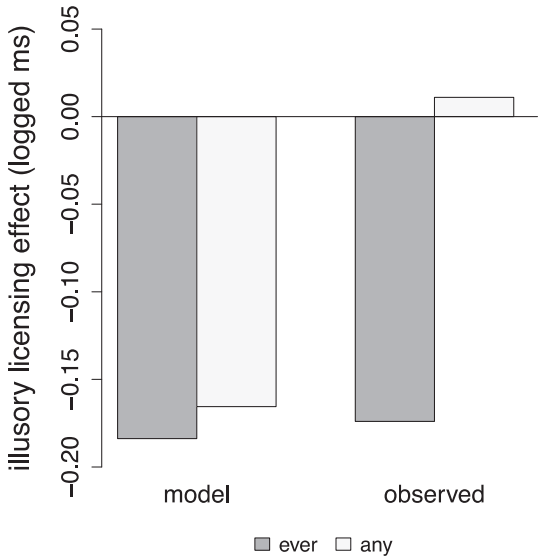

Fig. 8. Comparison of predicted and observed illusions for *ever* and *any* (Experiment 3). Model values correspond to the average across the range of parameter settings.

et al. (2008) to better understand how our findings diverge from the predictions of existing models. The schedules are provided in the Supplementary material. Differences between the manipulations were modeled only as differences in the feature composition of the potential licensors (e.g., ±negative, ±c-command) and the position/timing of the NPI.

To ensure that the modeling results would be robust against variation in the model's parameters, we ran a series of simulations that systematically combined parameter values used in previous research. Values of the *total source activation*, *activation noise*, *maximum associative strength* ("fan"), *decay*, and *match penalty* were manipulated.[5] The scaling factor was held constant at 1.0 across simulations. This method resulted in 3000 different models with unique parameter value combinations. 1000 Monte Carlo simulations were run for each model, providing for each simulation a prediction of the most probable retrieval target and its retrieval latency.

We focused on two measures from the simulations: predicted error rate and retrieval latency. Error rate corresponds to the percentage of runs for which the irrelevant licensor, rather than the target, was retrieved. This measure maps monotonically to speeded acceptability judgments, with higher rates of retrieval error corresponding to increased rates of judgment errors. Retrieval latencies provide a measure of how long on average it took to retrieve the item with the highest probability of retrieval on a given model run, and map monotonically to self-paced reading times, with higher latencies corresponding to higher reading times. These measures were used to compute illusory licensing as the difference in predicted error rate or retrieval latencies between the ungrammatical conditions with an irrelevant licensor and no licensor. Thus, for predicted error rate, a larger positive value corresponds to a stronger illusion, reflecting increased rates of acceptance for sentences with an irrelevant licensor, relative to sentences with no licensor. For predicted retrieval latencies, a smaller negative value corresponds to a stronger illusion, reflecting faster processing for sentences with an irrelevant licensor relative to sentences with no licensor.

[5] The total source activation parameter took one of five values: 0.5, 0.75, 1.0, 1.25, 1.5. Four values were used for the activation noise parameter: 0, 0.1, 0.2, 0.3. Five values were used for the fan parameter: 1.0, 1.25, 1.5, 1.75, 2. Five values were used for the decay parameter: 0, 0.25, 0.5, 0.75, 1.0. Six values were used for the match penalty parameter: −1.0, −0.8, −0.6, −0.4, −0.2, 0.0.

6.2. Modeling results

Simulations of the *ever*/*any* comparison and the position manipulation showed that the illusion was attenuated for a post-verbal NPI (e.g., *any* and *ever* in an embedded clause) relative to a pre-verbal NPI (e.g., *ever* in a main clause). This effect was driven by reactivation of the subject NP at the main verb, which attenuated interference from the irrelevant licensor. However, simulations failed to predict the on/off behavior observed in our experiments, as shown in Figs. 8 and 9. Simulations for the timing manipulation also showed that the illusion was attenuated for sentences with an intervening parenthetical phrase relative to sentences with a non-intervening parenthetical phrase, but again failed to predict the on/off behavior that we observed in our experiments, as shown in Fig. 10.

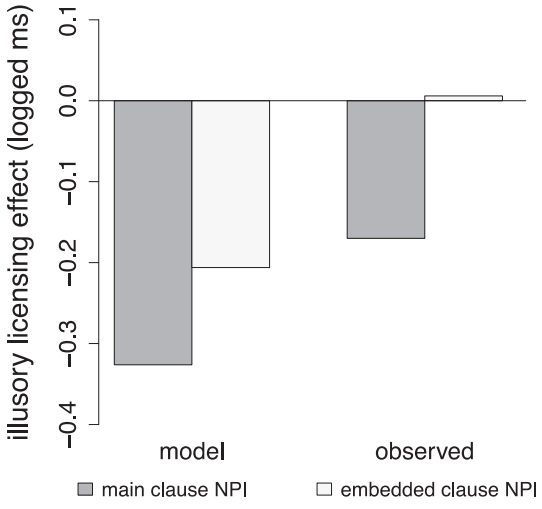

Fig. 9. Comparison of predicted and observed illusions for the clausal manipulation (Experiment 5). Model values correspond to the average across the range of parameter settings.

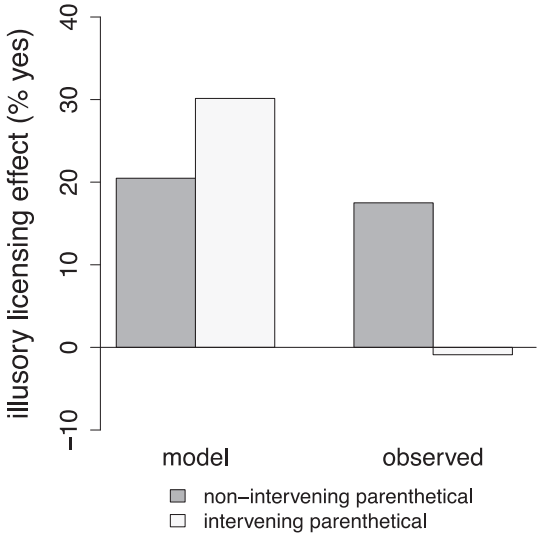

Fig. 10. Comparison of predicted and observed illusions for the NPI timing manipulation (Experiment 6). Model values correspond to the average across the range of parameter settings.

6.3. Discussion

The goal of our simulations was to verify the assumptions of the Vasishth et al. (2008) model of the NPI illusion, and to assess whether our NPI findings would be predicted by the model. Importantly, the model replicated the basic illusory licensing effect reported in previous research (e.g., Vasishth et al., 2008), which suggests that the current implementation was an appropriate choice to test the retrieval-based account of NPI licensing more broadly. We then tested the licensing configurations that we tested in the current study using the same model and the same parameter settings, and found that the model failed to capture our new findings. Specifically, simulations showed that differences in memory dynamics can modulate retrieval outcomes to some degree in the direction of a reduction, but that such effects are not sufficient to capture the on/off behavior that we observed for the NPI illusion. These results confirm that our new findings are not predicted by the current retrieval-based theory of NPI licensing. More generally, these results suggest that our effects are unlikely to be due to limitations of general memory retrieval processes that are engaged whenever an NPI is encountered.

7. General discussion

7.1. Summary of results

The present study used NPI illusions as a model system for examining how we mentally encode and navigate complex linguistic representations. Linguistic illusions could, in principle, reflect either an error in how we encode structured linguistic representations, or an error in how we access information in those representations. Recent accounts have argued that the NPI illusion reflects limitations of the access mechanisms that are responsible for normal NPI licensing, with the additional assumption that the encoding of the sentence remains fixed over time (e.g., Vasishth et al., 2008; Xiang et al., 2009). However, the conclusion that the illusion reflects a faulty implementation of the licensing mechanism is based on a relatively narrow range of findings and a specific set of assumptions about the generality of the effect, which have not been tested. For instance, existing evidence has come entirely from a single NPI *ever* in a restricted configuration. But since the illusion is argued to be the product of faulty licensing mechanisms that are engaged after an NPI is encountered, all accounts have assumed that the illusion should generalize across items and configurations.

The starting point for our study was to test the claims that illusions are (i) a general property of NPIs and (ii) a product of faulty licensing mechanisms. We tested the first claim by directly comparing the NPIs *ever* and *any* in Experiments 1–3. Converging results from self-paced reading times and speeded acceptability judgments revealed that while *ever* elicits the illusion, *any* does not. These results imply that illusions are not a general property of NPIs. We then tested the second claim in Experiments 4–5. Our results disconfirmed the prediction that illusions should generalize across configurations, as we found that the illusion can be switched off by moving the NPI further away from the potential licensors.

Based on these results, we hypothesized that the on/off behavior of the illusion could reflect either the structural position or the timing of the NPI relative to the potential licensors. We distinguished these alternatives in Experiment 6 by holding constant the structural and linear position of the NPI and manipulating the placement of a parenthetical clause to vary the amount of time between the potential licensors and the NPI. Once again, we saw the same modulation of the illusion: the illusion disappeared when the parenthetical clause intervened between the irrelevant licensor and the NPI. These results suggest that the key to switching the illusion on/off is the amount of time between the potential licensors and the NPI. In order to understand the cause of this effect, in Experiment 7 we tested whether the parenthetical phrase manipulation used for NPIs would have a similar impact on agreement illusions. We found that the placement of the parenthetical phrase did not impact agreement illusions. The contrast between agreement and NPI illusions challenges the attractive generalization that both illusions are the product of a memory retrieval error, and suggests that the mechanisms responsible for switching the NPI illusion on/off are not shared across all illusions. Lastly, we simulated our results using an explicit computational model of memory retrieval based on the model described in Vasishth et al. (2008). The model failed to capture the full pattern of results, which suggests that the on/off behavior of the NPI illusion cannot simply be due to limitations of memory retrieval mechanisms.

The results of the present study inform our understanding of the source and scope of linguistic illusions in several ways. First, the results show that at least one type of illusion involving NPIs is much more selective than previously assumed, occurring reliably under specific conditions. The finding that NPI illusions are highly selective is surprising because it is not predicted by existing accounts. Second, the selectivity of the illusion indicates that the effect cannot simply be a consequence of faulty licensing mechanisms. If the illusion was purely due to faulty licensing mechanisms that are engaged whenever an NPI is encountered, we would expect the illusion to generalize across items and configurations. However, our findings disconfirmed this prediction.

Lastly, we discovered that agreement illusions cannot be switched on and off under the same conditions. Psycholinguistic studies have often treated agreement and NPIs as similar types of structural dependencies. However, the consensus in the linguistics literature is that agreement and NPIs involve fundamentally different licensing mechanisms. For instance, whereas subject-verb agreement depends on a morphosyntactic licensing mechanism, NPIs are licensed by the compositional-semantic properties of entire propositions, rather than by a direct relation with a specific lexical item (Chierchia, 2006; Giannakidou, 2011; Horn, 2010; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Linebarger, 1987). The fact that this representational difference aligns with different processing profiles suggests that there isn't a homogenous cause for agreement and NPI illusions. Rather, we suggest that the contrasting processing profiles likely reflect fundamental differences in how we mentally encode and navigate syntactic and semantic representations. We develop this account below and discuss how our proposal relates to existing models of sentence comprehension.

7.2. Changing encodings

The key finding from the present study is that NPI illusions can be reliably switched on and off depending on when the encoding of the licensing context is probed for NPI licensing. This finding points to the status of the encoding as the source of our effects, rather than faulty mechanisms that are engaged whenever an NPI is encountered. Specifically, we argue that the encoding is not fixed, as previously assumed, but rather, changes over time. At one moment, irrelevant items inside the licensing context are transparently accessible as candidates for causing illusions. Then, at a later point in time, those same irrelevant items become opaque as candidates for causing illusions. If the encoding changes with the passage of time, we might expect different behaviors at different points in time depending on when the encoding is probed, as observed in our experiments.

We suggest two ways in which the encoding could change over time, neither of which has been considered in previous research. These possibilities include evolution in the content of the semantic

encoding that is probed for NPI licensing, and changes in the format of the encoding. Under both accounts, the on/off behavior of the NPI illusion would reflect access to different internal stages of the encoding process, such that the illusion can be reliably switched on/off depending on when the encoding is probed. We discuss each of these possibilities in turn below.

The first possibility is that the content of the encoding that is probed for NPI licensing may evolve over time. In the linguistics literature, it is widely assumed that NPIs are licensed by the semantics and pragmatics of entire propositions (Chierchia, 2006; Giannakidou, 2011; Horn, 2010; Israel, 2004; Kadmon & Landman, 1993; Krifka, 1995; Linebarger, 1987). Real-time computation of the compositional-semantic relations takes some amount of time, and it is reasonable to assume that the computation might involve multiple steps. It is, therefore, possible that the licensing mechanisms tap into the results of the intermediate stages of that interpretation process, yielding different profiles depending on whether the NPI is encountered earlier or later in the computation of the meaning. It is possible then that the different stages of the corresponding encoding processes can be revealed by varying when the interpretation is probed for NPI licensing. When the interpretation is probed immediately after the potential licensing context is first encountered, comprehenders might not have had enough time to finish computing and encoding the semantic/pragmatic relations that support NPI licensing. In this scenario, comprehenders may be forced to rely on incomplete inferences based on simple heuristics, such as the simple presence or absence of a licensor, which may give rise to an illusion. This component of the proposal builds on the earlier suggestion that NPI illusions reflect noisy pragmatic inferences (Xiang et al., 2009, 2013). By contrast, no illusion is expected after the appropriate semantic/pragmatic interpretation is fully encoded, preventing subsequent illusions.

Another possibility is that the interpretation is rapidly and accurately computed, but that the format of its encoding changes over time. An assumption of accounts that rely on partial-match interference (e.g., Vasishth et al., 2008) is that the individual features of a linguistic representation can always be independently evaluated, giving rise to partial-match interference effects at retrieval. This is a property of some encoding schemes, such as the one assumed in the Lewis and Vasishth model (Lewis & Vasishth, 2005; Vasishth et al., 2008), but not others. Many encoding schemes, including tensor-product variable bindings (e.g., Smolensky, 1990) and other vector-based models (e.g., Kanerva, 1994, 1996, 1997; Plate, 1991, 1994, 2003; Rachkovskij & Kussel, 2001), assume that the format of a representation changes over time. Vector-based models generally distinguish between "atomic" localist representations, where individual feature values are explicitly represented and transparently accessible, and "complex" distributed representations, which are constructed from atomic representations by a binding operation, e.g., convolution, addition, permutation, etc. This binding operation creates a new encoding that is completely dissimilar to its bound features. In this state, the atomic features are no longer accessible, and the bound representation must exhibit an all-or-none match in order to be recovered, preventing the possibility for partial-matching.[6]

The distinction between localist and distributed representations could be extended to explain the on/off behavior of the NPI illusion. Specifically, we suggest that the encoding of a sentence interpretation is built in two stages. The parser begins by constructing a localist representation of the sentence in which component feature values are made explicit and transparently accessible, as assumed in the Lewis and Vasishth model (Lewis & Vasishth, 2005). In this stage, individual features, such as +negation, can be independently evaluated, creating the opportunity for partial match interference. In the second stage, those same features are bound together to form a distributed representation that interfaces with the interpretive system and pragmatic inferencing. In this stage, the individual features are no longer independently evaluable and are opaque for causing illusions, since bound representations can only be recovered holistically, i.e., without partial matching. For example, when processing sentences like *The journalist that no editors recommended for the assignment (ever) thought that the readers would (ever) understand the complicated situation* (from Experiments 4–5), the parser may bind the semantic features, such as the embedded negation, to their position in the structure at various points during the sentence, creating a holistic representation of the subject that must be recovered without partial-matching. If the NPI appears at a time when the component feature values of a representation are still accessible, such as in the main clause position, then features such as +negation may be directly used for licensing, leading to an illusion. However, if the NPI appears after some delay, such as in the embedded clause position, then the licensing process refers to the bound representation, which more faithfully represents the licensing conditions, and then no illusion is found.

Of course, this is a preliminary account, and its full consequences need further examination. For example, the two-stage encoding proposal implies that NPIs may be licensed in two different ways, depending on the encoding that is consulted at the point of licensing.

The present results do not decide between the possibilities that the on/off behavior of the NPI illusion reflects changes in the content of the interpretation or changes in the format of its encoding. However, there are some recent findings that narrow down the space of possibilities and rule out some simple alternative accounts of our findings, such as the failure to appropriately interpret negation. For instance, recent work by Nieuwland and Kuperberg (2008) and Urbach, Delong, and Kutas (2015) rule out the simple alternative that negation is not incorporated into speakers' interpretations rapidly enough during incremental processing. Nieuwland and Kuperberg compared ERP responses for sentences in which negation was pragmatically licensed with sentences in which negation was pragmatically unlicensed (e.g., *With proper equipment, scuba-diving isn't very safe/dangerous* ... vs. *Bulletproof vests aren't very safe/dangerous* ...) and found N400 effects on the critical words (underlined). Importantly, they also found that in conditions with pragmatically licensed negation, the N400 effect was reduced in the congruous condition that was compatible with world knowledge compared to the incongruous condition (e.g., *dangerous* vs. *safe*), which suggests that negation can be incrementally interpreted under some conditions. Urbach and colleagues tested sentences with the negative quantifier *few* and found a reduced N400 effect for sentences that were compatible with world knowledge compared to sentences that were not (e.g., *Few kids prefer vegetables/sweets* ...). Importantly, these studies provide some evidence that negation can be integrated quickly enough and to a sufficient degree to impact the processing of upcoming words, which is an important first step for licensing an NPI by negation.

Recent research also rules out the possibility that comprehenders fail to encode scope relations for negatively quantified expressions. For instance, Kush and colleagues examined retrieval associated with the resolution of bound variable pronouns using eye-tracking while reading (Kush, Lidz, & Phillips, 2015). Bound-variable pronouns are superficially similar to NPIs in that they may be licensed by a negatively quantified NP in a structurally higher position, as in *No cyclist suspected that the spectators loathed him*. But unlike NPI licensing, the bound-variable relation is

[6] Crucially, the component items are not forever inaccessible. The sub-components of a representation may be recovered later via decoding processes that unpack the contents of the representation.

classified as a true item-to-item structural dependency. Kush and colleagues tested comprehenders' sensitivity to the presence and location of a negative licensor using an interference paradigm similar to the one used in the current study. Reading times revealed that comprehenders can rapidly distinguish structurally irrelevant licensors from structurally relevant licensors during dependency formation. Importantly, these findings suggest that comprehenders can reliably encode and rapidly access the structural relations for negatively quantified expressions to prevent interference from structurally irrelevant material during dependency formation, which may also be necessary for NPI licensing.

The findings from these studies argue against a simplistic account of our NPI findings in which illusory licensing reflects a failure to incorporate negation or failure to encode the necessary structural relations for negatively quantified expressions. These findings suggest that certain aspects of the interpretation are computed rapidly and remain stable over time, which is less obviously compatible with the hypothesis that the selective NPI illusion reflects changes over time in the content of the interpretation. We suggest instead that the selective illusion is likely a consequence of changes in the format of the semantic encoding.

In order to give a detailed account of how the encoding changes, we need to determine the cues for consolidation of the semantic representation. There are a number of factors that could affect the encoding of the interpretation, including key pieces of linguistic material and functional time and memory pressures. For instance, given the stringent limitations on working memory, comprehenders may be forced to periodically consolidate the encoding of an expanding representation to reduce processing load and conserve memory resources. However, the *ever*/*any* contrast from Experiments 2–3 makes it less likely that the observed effects are due only to time and memory pressures. Another possibility is that the introduction of key pieces of linguistic material, such as a main verb, affect the encoding of the sentence. For instance, encountering the main verb in a sentence like *The authors that no critics recommended have ever received* ... could force comprehenders to 'wrap-up' their interpretation of the subject phrase and consolidate its encoding before subsequent material is interpreted. The results of Experiments 2–5 (i.e., *ever*/*any* comparison and NPI position manipulation) are consistent with a central role for the verb. However, the results of the timing manipulation in Experiment 6 do not favor a verb-centric account, since a break in a sentence created by a parenthetical phrase can trigger the same modulation of the illusion seen in Experiments 2–5. Ultimately, we think that the introduction of key pieces of linguistic material is a likely trigger for important changes in the semantic encoding of the sentence, but this must be a special case of a more general process that can be triggered by other material, or perhaps by the mere passage of time.

7.3. Relation to existing characterizations of the parser

Our study contributes to an extensive literature on memory for linguistic structure, including earlier work on how linguistic material is integrated and stored in working memory. For instance, psycholinguistic studies by Sachs (1967) and Bransford and Franks (1971) were broadly interested in comprehenders' ability to notice different types of syntactic and semantic changes in a previously interpreted sentence, which amounts to testing susceptibility to different types of whole-sentence 'lures'. Building on this earlier work, we focused on a shorter time scale and asked the more specific question of whether the licensing of specific grammatical items would be susceptible to the presence of word-sized lures that are present, but in the wrong configuration.

However, our proposal that the encoding of the sentence changes over time is more closely related to earlier suggestions that the parser structures incoming material in multiple stages (e.g., Frazier & Fodor, 1978; Kimball, 1973; Townsend & Bever, 2001; see also Abney, 1991; Fodor, Bever, & Garrett, 1974; Whitney, 2004; but cf. Marslen-Wilson & Tyler, 1980). Existing models differ in their claims about the nature of the units that are shunted between the different stages and the timing of the transitions, but they all assume that the complete representation of a sentence is built in multiple stages, with one stage temporally prior to the other. For example, in the classic 'Sausage Machine' model proposed by Frazier and Fodor (1978), the parser first constructs a shallow representation of the input before shunting constituents off to a more compact store where they are combined into a complete structure for interpretation. Importantly, this division of labor captures certain "shortsighted" errors, where comprehenders fail to recognize different ways to attach a constituent. According to Frazier and Fodor, such effects arise because certain properties of the existing structure have been shunted to the next stages of processing, restricting the range of attachment sites that are available for immediate processing.

Our proposal for multiple stage encoding shares some key assumptions with these models. For instance, we assume with others that the division of labor into multiple stages is a consequence of functional time and memory pressures that arise when processing a sentence (e.g., Frazier & Fodor, 1978). Dividing the process of parsing a sentence into sub-stages is but one solution to keeping the demands on working memory within their limits as a parse is extended. More specifically, we suggest like others that the division of labor into multiple steps can explain why certain types of linguistic information are not immediately accessible for ongoing parsing operations.

Our proposal moves beyond existing models in the suggestion that the format of the encodings might be fundamentally different at each stage. Models such as the Sausage Machine, assume that the format of the first stage encoding is qualitatively similar to the format of the second stage encoding. In contrast, we suggest that the format of the first versus second stage encodings is fundamentally different, such that it impacts the transparency of certain types of linguistic information encoded in memory. In short, both proposals assume that sentence representations are built in multiple stages to explain why comprehenders cannot recognize certain aspects of the existing structure, at least temporarily, but the phenomena that these accounts seek to explain are very different. For Frazier and Fodor, multiple-stages are invoked to explain why comprehenders initially fail to recognize an attachment site. In contrast, we proposed a multiple-stage encoding process, as one of several possibilities, to explain why comprehenders are initially misled by grammatically irrelevant material.

More broadly, it is also worth highlighting that multiple stage encoding schemes have been invoked to capture superficially similar findings in other cognitive domains. For example, research on visual processing suggests that there is an initial stage of object recognition during which individual object features, such as shape and color, are encoded independently of each other, giving rise to illusory conjunctions. This first stage of encoding is followed by a feature-integration stage, where the separate, independently accessible features are consolidated into a single, unitized encoding of the object in visual memory, preventing illusory conjunctions. This two-stage encoding process is explicitly characterized by the feature integration theory of attention (Treisman & Gelade, 1980; Treisman & Schmidt, 1982; Treisman, Sykes, & Gelade, 1977; see also Chun & Potter, 2000; Holcombe & Clifford, 2012; Treisman, 1996; Wolfe, 2007, 2012). A potentially interesting line of research would be to examine the extent to which multiple stage encoding processes are similar across cognitive domains.

8. Conclusion

In this study, we used linguistic illusions as a tool to examine how we encode and access linguistic information in memory. Existing accounts have assumed that illusions reflect limitations of the memory access mechanisms, and have made the additional assumption that the encoding of the sentence remains fixed over time. Our findings challenged the predictions made by these accounts, which assume that illusions should generalize across items and configurations. Specifically, we found that one type of illusion involving NPIs can be systematically switched on/off depending on when the encoding is probed for NPI licensing. To the best of our knowledge, this is the clearest case of what it takes to switch the illusion on/off. We took these findings to suggest that the encoding of the licensing context is not fixed, but rather, changes over time. Previously, selective illusions have been taken to be informative about the memory access mechanisms. Here, we took them to be informative also about the nature of structured representations in memory and the nature of the encoding mechanisms.

Acknowledgments

We would like to thank Brian Dillon, Norbert Hornstein, Michael Israel, Dave Kush, Sol Lago, Ellen Lau, Jeff Lidz, Luiza Newlin-Łukowicz, Shravan Vasishth, Alexander Williams, and Ming Xiang for helpful discussions, and Glynis MacMillan for assistance with running the experiments. This work was supported in part by NSF-IGERT grant DGE-0801465 to the University of Maryland and by NSF grant BCS-0848554 to CP.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.cognition.2016.08.016.

References

Abney, S. P. (1991). Parsing by chunks. In R. Berwick, S. P. Abney, & C. Tenny (Eds.), *Principle-based parsing: Computation and psycholinguistics* (pp. 257–278). Dordrecht: Kluwer Academic Publishing.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

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, W., & Lewis, R. (2007). A new theory and computational model of working memory in sentence production: Agreement errors as failures of cue-based retrieval. *Talk at the 20th CUNY conference on human sentence processing*. San Diego: University of California.

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.

Box, G. E., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society Series B (Methodological), 26*, 211–252.

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology, 2*, 331–350.

Chierchia, G. (2006). Broaden your views: Implicatures of domain widening and the "logicality" of language. *Linguistic Inquiry, 37*, 535–590.

Chun, M. M., & Potter, M. C. (2000). A two-stage model for multiple target detection in RSVP. *Jounral of Experimental Psychology: Human Perception and Performance, 21*, 109–127.

Clifton, C., Frazier, L., & Deevy, P. (1999). Feature manipulation in sentence comprehension. *Rivista di Linguistica, 11*, 11–39.

Dillon, B., Clifton, C., Jr., & Frazier, L. (2014). Pushed aside: Parentheticals, memory, and processing. *Language, Cognition, and Neuroscience, 29*, 483–498.

Dillon, B., Mishler, A., Slogget, S., & Phillips, C. (2013). Contrasting intrusion profiles for agreement and anaphora: Experimental and modeling 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: de Gruyter.

Eberhard, K., Cutting, J., & Bock, K. (2005). Making syntax of sense: Number agreement in sentence production. *Psychological Review, 112*, 531–559.

Fodor, J. A., Bever, T., & Garrett, J. (1974). *The psychology of language: An introduction to psycholinguistics and generative grammar*. New York: McGraw-Hlll.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition, 6*, 291–325.

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

Giannakidou, A. (2011). Positive polarity items and negative polarity items: Variation, licensing, and compositionality. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (pp. 1660–1712). Berlin: Mouton de Gruyter.

Gibson, E., Paintadosi, S. T., Brink, K., Bergen, L., Lim, E., & Saxe, R. (2013). A noisy-channel account of crosslinguistic word-order variation. *Psychological Science, 24*(7), 1079–1088.

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

Holcombe, A. O., & Clifford, C. W. G. (2012). Failures to bind spatially coincident features: Comment on Di Lollo. *Trends in Cognitive Science, 16*, 402.

Horn, L. R. (2010). Polarity after the Thirty Years War: What's the point? Semantics Seminar, Rutgers University. November 1.

Israel, M. (2004). The pragmatics of polarity. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 701–723). Oxford: Blackwell.

Kadmon, N., & Landman, F. (1993). Any. *Linguistics and Philosophy, 16*, 353–422.

Kanerva, P. (1997). Fully distributed representation. In *Proceedings of the 1997 real world computing symposium* (pp. 358–365). Tsukuba-city, Japan: Real World Computing Partnership.

Kanerva, P. (1994). The spatter code for encoding concepts at many levels. In M. Marinaro & P. G. Morasso (Eds.), *Proceedings of the international conference on artificial neural networks (ICANN 1994)* (pp. 226–229). Springer-Verlag.

Kanerva, P. (1996). Binary spatter-coding of ordered K-tuples. In C. von der Malsburg, W. von Seelen, J. C. Vorbruggen, & B. Sendhoff (Eds.), *Proceedings of the international conference on artificial neural networks (ICANN 1996)* (pp. 869–873). Berlin: Springer.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition, 2*, 15–47.

Kohonen, T. (1980). *Content-addressable memories*. Berlin; New York: Springer-Verlag.

Krifka, M. (1995). The semantics and pragmatics of weak and strong polarity items. *Linguistic Analysis, 25*, 209–257.

Kush, D., Lidz, J., & Phillips, C. (2015). Relation-sensitive retrieval: Evidence from bound variable pronouns. *Journal of Memory and Language, 82*, 18–40.

Ladusaw, W. A. (1979). *Negative polarity items as inherent scope relations* (Dissertation/Thesis). Austin, TX: University of Texas.

Lago, S., Shalom, D., Sigman, M., Lau, E., & Phillips, C. (2015). Agreement processes in Spanish comprehension. *Journal of Memory and Language, 82*, 133–149.

Levy, R. (2008). A noisy-channel model of rational human sentence comprehension under uncertain input. In *Proceedings of the 2008 conference on empirical methods in natural language processing, Honolulu* (pp. 234–243).

Lewis, S., & Phillips, C. (2015). Aligning grammatical theories and language processing models. *Journal of Psycholinguistic Research, 44*, 27–46.

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.

Linebarger, M. (1987). Negative polarity and grammatical representation. *Linguistics and Philosophy, 10*, 325–387.

Marslen-Wilson, W., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. *Cognition, 8*, 1–71.

Nieuwland, M. S., & Kuperberg, G. R. (2008). When the truth is not too hard to handle: An event-related potential study on the pragmatics of negation. *Psychological Science, 19*, 1213–1218.

Pearlmutter, N., Garnsey, S., & Bock, K. (1999). Agreement processes in sentence comprehension. *Journal of Memory and Language, 41*, 427–456.

Phillips, C., Wagers, M., & Lau, E. F. (2011). Grammatical illusions and selective fallibility in real-time language comprehension. In J. Runner (Ed.). *Experiments at the interfaces* (Vol. 37, pp. 147–180). Bingley, UK: Emerald Publications.

Plate, T. (1994). *Distributed representations and nested compositional structure* (Dissertation/Thesis). Toronto, Canada: University of Toronto.

Plate, T. (2003). *Holographic reduced representation: Distributed representation of cognitive structure*. Stanford: CSLI.

Plate, T. (1991). Holographic reduced representations: Convolution algebra for compositional distributed representations. In J. Mylopoulos & R. Reiter (Eds.), *Proceedings of the 12th international joint conference on artificial intelligence (IJCAI)* (pp. 30–35). San Mateo, CA: Kaufmann.

Pollard, C., & Sag, I. A. (1994). *Head-driven phrase structure grammar*. Stanford, CA: CSLI Publications.

Quirk, R., Greenbaum, S., Leech, G., & Svartvik, J. (1985). *A comprehensive grammar of the english language*. New York: Longman.

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.

Rachkovskij, D. A., & Kussel, E. M. (2001). Binding and normalization of binary sparse distributed representations by context-dependent thinning. *Neural Computation, 2*, 411–452.
Sachs, J. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics, 2*, 437–442.
Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist networks. *Artificial Intelligence, 46*, 159–216.
Staub, A. (2009). On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language, 60*, 308–327.
Tabor, W., Galantucci, B., & Richardson, D. (2004). Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language, 50*, 355–370.
Tanner, D., Nicol, J., & Brehm, L. (2014). The time-course of feature interference in agreement comprehension: Multiple mechanisms and asymmetrical attraction. *Journal of Memory and Language, 76*, 195–215.
Townsend, D. J., & Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. Cambridge, Mass: MIT Press.
Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.
Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the percpetion of objects. *Cognitive Psychology, 14*, 107–141.
Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology, 6*, 171–178.
Treisman, A., Sykes, M., & Gelade, G. (1977). Selective attention and stimulus integration. In S. Dornica (Ed.), *Attention and performance VI: Proceedings of the sixth international symposium on attention and performance* (pp. 333–361). Hillsdale, NJ: Erlbaum.
Urbach, T. P., Delong, K. A., & Kutas, M. (2015). Quantifiers are incrementally interpreted in context, more or less. *Journal of Memory and Language, 83*, 79–96.
Van Dyke, J. A., & McElree, B. (2006). Retrieval interference in sentence comprehension. *Journal of Memory and Language, 55*, 157–166.
Van Dyke, J. A., & McElree, B. (2011). Cue-dependent interference in comprehension. *Journal of Memory and Language, 65*, 247–263.
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., Chen, Z., Li, Q., & Guo, G. (2013). Processing Chinese relative clauses: Evidence for the subject-relative advantage. *PLoS ONE*, 8.
Venables, W. N., & Ripley, B. D. (1999). *Modern applied statistics with S-plus*. New York: Springer.
Wagers, M. (2008). *The structure of memory meets memory for structure in linguistic cognition* (Dissertation/Thesis). University of Maryland.
Wagers, M., Lau, E. F., & Phillips, C. (2009). Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language, 61*, 206–237.
Whitney, C. S. (2004). *Investigations into the neural basis of structured representations* (Dissertation/Thesis). University of Maryland.
Wolfe, J. (2007). Guided Search 4.0: Current progress with a model of visual search. In W.Gray (Ed.), Integrated models of cognitive systems, Oxford, New York (pp. 99–119).
Wolfe, J. (2012). The binding problem lives on: Comment on Di Lollo. *Trends in Cognitive Science, 16*, 307–308.
Xiang, M., Dillon, B. W., & Phillips, C. (2006). Testing the strength of the spurious licensing effect for negative polarity items. *Talk at the 19th CUNY conference on human sentence processing, New york*.
Xiang, M., Dillon, B., & Phillips, C. (2009). Illusory licensing effects across dependency types: ERP evidence. *Brain & 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*.
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.