



# Ellipsis interference revisited: New evidence for feature markedness effects in retrieval

Dan Parker<sup>\*</sup>

Linguistics Program, Department of English, College of William & Mary, United States

## ARTICLE INFO

### Keywords:

Sentence processing  
Verb phrase ellipsis  
Similarity-based interference  
Cue-based retrieval  
Bayesian data analysis  
Computational modeling

## ABSTRACT

An active question in psycholinguistics concerns how we mentally encode and retrieve linguistic information in memory. In particular, it remains unclear what information (“cues”) guide retrieval. Previous work has extensively tested retrieval of noun phrases, but less is known about retrieval of other constituents, including verb phrases (VPs). This study examines retrieval for VP ellipsis to allow for a more comprehensive theory of cues. Four experiments (acceptability, self-paced reading) used an interference paradigm to examine voice information (active, passive) in retrieval. Results revealed a selective profile: passive ellipsis shows interference, but active ellipsis does not. These results are aligned with the markedness asymmetry observed for agreement attraction, where marked features (plural, passive) trigger interference, but unmarked features (singular, active) do not. This analysis motivates a unified account of verbal dependencies where markedness plays a more fundamental role than previously assumed. Lastly, I use ACT-R to demonstrate how markedness effects might arise in a cue-based retrieval architecture and discuss the current findings with respect to the leading theories of interference effects.

## Introduction

Over the past several decades, significant inroads have been made in explicitly characterizing the memory retrieval mechanisms that support real-time sentence processing. A leading proposal claims that when retrieval is engaged for sentence processing, the processor engages a cue-based matching procedure that is susceptible to similarity-based interference (see Jäger et al., 2017, for a review). However, we still know very little about what sorts of information (“cues”) guide retrieval. For instance, most studies to date have focused on retrieval of a noun phrase (NP) for dependency formation (e.g., subject-verb agreement, anaphora), which involves a relatively narrow set of morphosyntactic and semantic features like gender, number, and animacy. By comparison, much less is known about how we recover other types of phrases from memory, like verb phrases (VPs), which are unique in terms of the syntactic, morphological, semantic, and discourse information they encode. One case where retrieval targets a VP is in sentences with VP ellipsis like (1).

- (1) a. Sally finished the assignment, and Samantha did too.  
b. Sally [<sub>VP</sub> finished the assignment], and Samantha did [<sub>VP</sub> finish the assignment] too.

VP ellipsis involves the omission of a redundant VP, such as the VP *finish the assignment* in the second clause of (1a), necessitating retrieval of the previous VP for interpretation. Most studies on the processing of ellipsis have focused either on the conditions under which ellipsis is permitted, the representation of the ellipsis site, or the discourse relations between the antecedent and ellipsis clauses (see Phillips & Parker, 2014, for a review). A subset of these studies (Martin et al., 2012, 2014; Martin & McElree, 2008, 2009, 2011) has pointed to the possibility that ellipsis engages the same cue-based retrieval mechanism proposed for other dependencies, like subject-verb agreement and anaphora.

The claim that ellipsis processing involves cue-based retrieval is based, in part, on findings of similarity-based interference. For instance, Martin (2018) tested sentences like those in (2) using event-related potentials (ERPs) to examine whether retrieval for VP ellipsis is susceptible to interference from voice features (e.g., active, passive) encoded on a verb. The paradigm in (2) manipulated the match between the voice of the ellipsis clause and that of the antecedent (GRAMMATICALITY: grammatical vs. ungrammatical) and the match between the voice of the ellipsis clause and an “attractor” VP that cannot be an antecedent for the ellipsis (ATTRACTOR MATCH: match vs. mismatch). In (2), *Because Jane drank the cocktail that ...* is the antecedent clause, and the attractor VP

<sup>\*</sup> Address: Department of English, College of William & Mary, P.O. Box 8795, Williamsburg, VA 23187, United States.

E-mail address: [dparker@wm.edu](mailto:dparker@wm.edu).

(underlined) is embedded in the relative clause {*the waiter served* | *was served by the waiter*}.

a. **Grammatical, Attractor Match**

Because Jane drank the cocktail that the waiter served, Bill did too...

b. **Grammatical, Attractor Mismatch**

Because Jane drank the cocktail that was served by the waiter, Bill did too...

c. **Ungrammatical, Attractor Match**

Because Jane drank the cocktail that was served by the waiter, Bill was too...

d. **Ungrammatical, Attractor Mismatch**

Because Jane drank the cocktail that the waiter served, Bill was too...

ERP responses showed that the presence of an attractor VP that matched the voice of the ellipsis clause eased processing relative to the attractor mismatch condition, but only in the ungrammatical conditions, resulting in an “illusion of grammaticality” (Phillips et al., 2011). This profile is a behavioral signature of a cue-based retrieval procedure (Jäger et al., 2017).

There are several reasons to revisit previous conclusions about ellipsis interference. First, the antecedent VPs in (2) were embedded in a fronted subordinate clause (e.g., *Because Jane ...*), which may engage predictive processing to locate a position for interpretation in the following ellipsis clause (e.g., *Bill got the cocktail that was served by the waiter because Jane got the cocktail that was served by the waiter*). An active question in psycholinguistics concerns how retrieval interacts with expectation-based parsing (e.g., Futrell et al., 2020; Levy, 2013), and the relative contributions of retrieval vs. predictive processing in fronted clauses like those in (2) remain poorly understood. Second, it is unclear whether the attractor VP in the relative clause is actually barred from serving as an antecedent for the ellipsis.<sup>1</sup> Further research is needed to determine whether such a reading is possible.

Third, it is implicitly assumed in Martin (2018) that active and passive ellipsis should behave similarly at retrieval: they both deploy a voice cue (e.g., +passive, +active) and are susceptible to interference. I believe this assumption might be based, in part, on the formal literature, which claims that active and passive voice are syntactic features encoded on verbs (e.g., Merchant, 2013; see also C. Kim et al., 2011; C. Kim & Runner, 2018), and as such, can be targeted at retrieval using a corresponding voice cue. According to the cue-based theory of retrieval (e.g., R. L. Lewis et al., 2006), which claims that retrieval cues are derived from the context at the retrieval site, the relevant passive/active features are presumably derived from the auxiliary verb of the ellipsis clause (e.g., where the morphology is expressed). However, several studies suggest that active and passive voice features differentially impact the processing and acceptability of ellipsis. For instance, Kim and colleagues (2011) showed that passive ellipsis is less acceptable than active ellipsis (see also Arregui et al., 2006), and Parker (2018) showed using computational modeling that passive ellipsis incurs additional processing costs relative to active ellipsis.

Passive and active ellipsis also differ with respect to their syntax, morphology, and information structure, which might lead to differences at retrieval. In English, active voice is the default unmarked form, whereas passive voice is marked because it involves changes in word order, verbal morphologically, and information structure that marks the subject as topical (e.g., Givón, 1990; Poppels & Kehler, 2019; Rohde & Kehler, 2014; Shibatani, 1985). Crucially, previous studies of interference effects have shown that feature markedness might influence which cues are deployed at retrieval for linguistic dependency formation. For instance, previous work on interference effects in subject-verb agreement processing (“agreement attraction”) has shown a number asymmetry, such that plural verbs show interference from feature-matching

nouns, but singular verbs do not. It has been suggested that this asymmetry might reflect a privative number marking system where the default singular feature has no explicit marking (i.e., represented by the absence of a number feature) and plural is marked (e.g., Bock & Eberhard, 1993; Eberhard, 1997; Harley & Ritter, 2002; Kimball & Aissen, 1971). Retrieval-based accounts of agreement attraction (e.g., Wagers et al., 2009) capture this difference in the cue specification: Plural verbs deploy a number cue (e.g., +plural) at retrieval for agreement processing, increasing the probability that a plural attractor is erroneously retrieved, but singular verbs do not, reducing the possibility of attraction.

These findings raise the question of whether ellipsis processing would show a parallel markedness asymmetry with respect to interference effects for active and passive voice. This question is addressed in the present study.

### The present study

The goal of the current study is to contribute to our understanding of the source and scope of interference effects by systematically investigating the use of voice cues in retrieval for ellipsis processing. To this end, the present study tested for interference effects in VP ellipsis constructions using a modified version of the paradigm in (2), in which passive and active voice constructions were tested independently. Experiments 1 and 2 used untimed acceptability to verify the constraints on the relationship between VP ellipsis and its antecedent. To preview, results show that an attractor VP inside a relative clause cannot serve as an antecedent for active and passive ellipsis alike. Experiments 3 and 4 then tested for interference effects in moment-by-moment processing using self-paced reading. Results revealed a selective profile, such that passive ellipsis shows interference, but active ellipsis does not. This profile is qualitatively similar to the markedness asymmetry observed for agreement attraction, motivating a uniform account of interference effects in the processing of verbal dependencies. Finally, computational modeling using the ACT-R model of sentence processing is used to show how feature markedness effects are predicted in a cue-based retrieval architecture that imposes a privative cue system (e.g., Wagers et al., 2009).

### Data availability

All of the materials for the present study are available on the Open Science Framework (<https://osf.io/ykqtw/>). The repository contains the following materials:

- The stimuli materials from Experiments 1–4
- The data and code from Experiments 1–4
- A detailed description of the ACT-R model presented in the General Discussion
- The code for the ACT-R model presented in the General Discussion

### Experiment 1: Passive ellipsis acceptability

The goal of Experiment 1 was to verify the constraints on VP ellipsis to establish a suitable paradigm to test for retrieval interference effects in moment-by-moment processing. In particular, it is important to first assess whether an attractor VP embedded inside a relative clause can serve as an antecedent for the ellipsis. To this end, Experiment 1 tested speakers’ sensitivity to the attractor VP in both grammatical and ungrammatical configurations using untimed acceptability judgments. Traditionally, untimed “offline” acceptability judgments have been taken to reflect the underlying linguistic constraints (e.g., Gerken & Bever, 1986; Townsend & Bever, 2001), although there are extralinguistic factors that might cause judgments to diverge from grammatical generalizations (e.g., noise, limitations of the general-purpose mechanisms such as memory, cognitive control, and attention that are

<sup>1</sup> This concern was raised by reviewers of the current study for the 2020 AMLaP conference.

used to implement linguistic computations). For instance, two recent reviews of agreement attraction effects found that attractors can cause interference in offline acceptability judgments (Hammerly et al., 2019), but the magnitude of the effect is reduced compared to time-restricted measures (e.g., speeded acceptability judgments), as evidenced in Parker (2019). In short, there appears to be a time-sensitivity to interference effects, such that they seem to weaken over time, possibly due to reanalysis (see Parker, 2019, for discussion).

If the relative clause VP cannot be an antecedent for the ellipsis, as previously claimed (Martin, 2018), acceptability judgments should not be strongly modulated by the presence of a matching attractor in either the grammatical or ungrammatical configurations. Conversely, if the relative clause VP can be an antecedent for the ellipsis, then we should expect a modulation of acceptability, such that sentences with a matching attractor (i.e., an attractor that matches the voice of the ellipsis) are rated substantially higher than their mismatching counterparts (e.g., on a par with a matching target antecedent). These predictions are tested with passive ellipsis in Experiment 1, and then with active ellipsis in Experiment 2.

Participants

One hundred self-reported native English speakers were recruited via Amazon’s Mechanical Turk web service (<https://www.mturk.com>). All participants provided informed consent. Participants were compensated \$3.00 each. The experiment lasted approximately 20 min.

Stimuli and experimental design

Experiments 1 and 2 used a modified version of the paradigm in (2). Passive and active ellipsis constructions were separated for independent investigation, and simple coordinate structures without fronting were used to restrict the possibility of predictive ellipsis resolution. Twenty-four sets of four conditions like those in Table 1 were constructed based on the materials used in Martin (2018). The experiment used a 2 × 2 (mis)match paradigm that manipulated the match between the voice of the ellipsis clause, which was passive in Experiment 1, and that of the antecedent (GRAMMATICALITY: grammatical vs. ungrammatical) and the match between the voice of the ellipsis clause and the attractor VP (underlined) in a relative clause (ATTRACTOR MATCH: match vs. mismatch).<sup>2</sup>

The 24 sets of test sentences were distributed across 4 lists in a Latin square design and combined with 48 filler sentences of similar length

**Table 1**  
Sample item set from Experiment 1. The clause with the attractor is underlined.

| Condition                         | Sentence   |
|-----------------------------------|--|
| Grammatical, Attractor Match      | Jane was recruited for the event that <u>was organized by the villagers</u> , and John was too last Saturday in the afternoon. |
| Grammatical, Attractor Mismatch   | Jane was recruited for the event that <u>the villagers organized</u> , and John was too last Saturday in the afternoon.        |
| Ungrammatical, Attractor Match    | Jane recruited for the event that <u>was organized by the villagers</u> , and John was too last Saturday in the afternoon.     |
| Ungrammatical, Attractor Mismatch | Jane recruited for the event that <u>the villagers organized</u> , and John was too last Saturday in the afternoon.            |

<sup>2</sup> An anonymous reviewer pointed out that interference from the lure would lead to an implausible interpretation (e.g., ... *John was organized by the villagers*). As the results of Experiments 3 and 4 show, the implausibility of the lure did not prevent a mismatch effect. However, it remains an open question how interference impacts the global interpretation of the sentence.

and complexity, such that each participant read a total of 72 sentences. The ratio of grammatical-to-ungrammatical sentences was 1:1. The ungrammatical fillers involved subject-verb agreement errors and mismatched ellipsis involving verbal and nominal gerundive antecedents to obscure the voice manipulation in the test sentences.

Procedure

The experiment was conducted using Ibex (<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). Each sentence was displayed in its entirety on the screen along with the rating scale. Participants could click boxes or use the numerical keypad to enter their ratings. The order of presentation was randomized for each participant. Prior to the experiment, there was a brief instructional check that asked several comprehension questions (e.g., *What button do I press to respond “YES”*) to confirm that participants read the instructions.

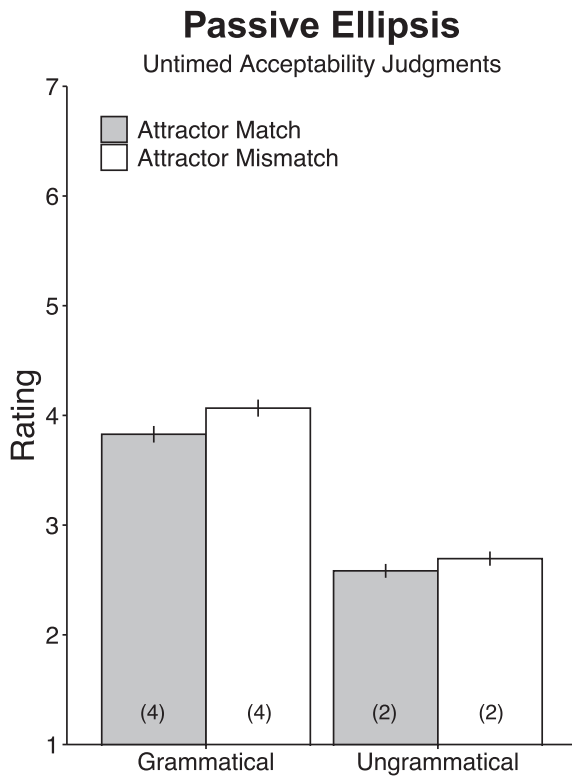
Data analysis

Data were analyzed in a Bayesian framework (e.g., Gelman et al., 2014) to estimate the direction and magnitude of the main effects (GRAMMATICALITY and ATTRACTOR MATCH) and their interaction. Two Bayesian hierarchical ordinal mixed-effects models were constructed using the *brms* package (Bürkner, 2017) in the R statistical computing environment (R Development Core Team, 2018). The first model (“crossed model”) examined the main effects of GRAMMATICALITY (grammatical vs. ungrammatical) and ATTRACTOR MATCH (match vs. mismatch) and their interaction, with rating as the dependent variable. The main effects and their interaction were specified as fixed effects, and the factors were sum-coded ( $\pm 0.5$ ). The second model (“nested model”) examined the effect of attractor match within the grammatical and ungrammatical conditions independently (labeled as “Grammatical Attraction” and “Ungrammatical Attraction” in Fig. 2). The nested model also used sum-coding ( $\pm 0.5$ ). Each model included participants and items as random effects and used a “maximal” random effects structure with a full variance–covariance matrix specification for participants and items (Barr et al., 2013).

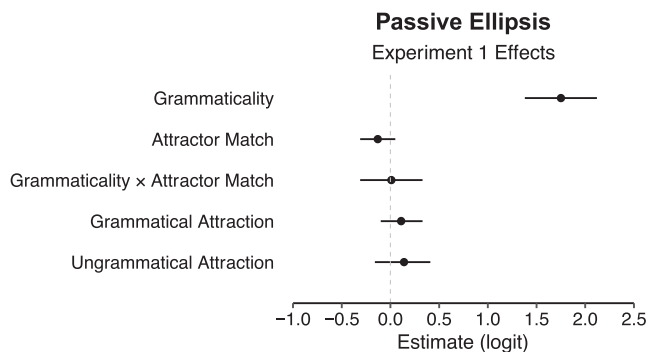
All models used regularizing, mildly informative priors following the methods described in Schad et al. (2021) and Nicenboim and Vasishth (2016). The priors for the fixed effects and their interaction were defined as a standard Normal(0,1). For each effect estimate, the mean of the posterior distribution with the 89% Bayesian credible interval (CrI) is reported. The CrI indicates the range within which we can be 89% certain that the true effect lies (given the data and model). This inferential approach differs from null hypothesis significance testing, as it does not rely on dichotomous decisions about significance (e.g., reject the null hypothesis vs. fail to reject the null hypothesis), but rather allows us to evaluate the evidence for/against an effect being zero in a graded fashion (Nicenboim & Vasishth, 2016). An effect was deemed reliable if the CrI is on one side of zero, and an effect spanning values above and below zero was deemed inconclusive (Kruschke et al., 2012; see also Nicenboim & Vasishth, 2016).

Results

Mean and median ratings by condition are shown in Fig. 1. Results of the Bayesian analyses are reported in Table 2 and a visualization of effect sizes is provided in Fig. 2. The crossed model revealed clear evidence for a main effect of GRAMMATICALITY, as grammatical sentences received higher ratings than ungrammatical sentences. There was no evidence of an effect of ATTRACTOR MATCH or an interaction with GRAMMATICALITY, as the CrIs spanned values above and below zero. There was no evidence of an effect of ATTRACTOR MATCH in either the grammatical conditions or the ungrammatical conditions in the nested comparisons.



**Fig. 1.** Mean ratings for Experiment 1. Error bars indicate standard error of the mean. Median ratings are included in parentheses.



**Fig. 2.** Posterior probability distributions of effect sizes from Experiment 1. The dot corresponds to the mean of the distribution and the lines are the 89% credible interval (CrI). The dashed line indicates an effect of zero.

**Table 2**

Bayesian analysis of judgments from Experiment 1. The table reports the posterior means of the fixed effects and the corresponding 89% Bayesian credible intervals [CrI]. Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                           | Posterior mean [89% CrI] |
|----------------------------------|--------------------------|
| <b>Crossed Model</b>             |                          |
| Grammaticality                   | <b>1.75 [1.38, 2.12]</b> |
| Attractor Match                  | -0.13 [-0.31, 0.05]      |
| Grammaticality × Attractor Match | 0.01 [-0.31, 0.33]       |
| <b>Nested Model</b>              |                          |
| Grammatical Attraction           | 0.11 [-0.10, 0.33]       |
| Ungrammatical Attraction         | 0.14 [-0.16, 0.41]       |

## Discussion

The goal of Experiment 1 was to confirm the constraints on passive VP ellipsis (i.e., whether a passive VP inside of the relative clause can serve as the antecedent for passive ellipsis) using untimed acceptability judgments. Results showed a clear effect of grammaticality, such that grammatical sentences with a matching target VP were rated higher than ungrammatical sentences with a mismatching target. However, there was no indication that a voice matching attractor VP inside the relative clause modulated ratings for either the grammatical or ungrammatical conditions. In so far as untimed acceptability judgments reflect the underlying linguistic constraints, as is traditionally assumed (e.g., Gerken & Bever, 1986; S. Lewis & Phillips, 2015; Schütze, 2016; Townsend & Bever, 2001), these results suggest that that a VP inside a relative clause is not a possible antecedent for passive ellipsis, as claimed in Martin (2018). We now turn to the same test for active ellipsis.

## Experiment 2: Active ellipsis acceptability

Experiment 2 used the same (mis)match paradigm from Experiment 1, but instead used active ellipsis. As in Experiment 1, if a voice-matching VP inside the relative clause can serve as an antecedent for the ellipsis, then we predict sentences with a matching active attractor VP to be rated higher than their mismatching counterparts. If the relative clause VP cannot serve as an antecedent, no modulation is expected.

## Participants

One hundred self-reported native English speakers were recruited via Amazon's Mechanical Turk web service. All participants provided informed consent. Participants were compensated \$3.00 each. The experiment lasted approximately 20 min. Two participants failed the instruction check, leaving data from a total of 98 participants for analysis.

## Stimuli and experimental design

Experiment 2 used the same experimental items and fillers from Experiment 1, except that the ellipsis was active, as shown in Table 3.

## Procedure

Experiment 2 used the same procedure that was used in Experiment 1.

## Data analysis

Data analysis for Experiment 2 followed the same steps as in Experiment 1. Experiment 2 also included an additional model to compare the effects of GRAMMATICALITY, ATTRACTION, VOICE (passive vs. active), and their interactions across Experiments 1 and 2 to evaluate

**Table 3**

Sample item set from Experiment 2. The clause with the attractor is underlined.

| Condition                         | Sentence   |
|-----------------------------------|--|
| Grammatical, Attractor Match      | Jane recruited for the event that <u>the villagers organized</u> , and John did too last Saturday in the afternoon.            |
| Grammatical, Attractor Mismatch   | Jane recruited for the event that <u>was organized by the villagers</u> , and John did too last Saturday in the afternoon.     |
| Ungrammatical, Attractor Match    | Jane was recruited for the event that <u>the villagers organized</u> , and John did too last Saturday in the afternoon.        |
| Ungrammatical, Attractor Mismatch | Jane was recruited for the event that <u>was organized by the villagers</u> , and John did too last Saturday in the afternoon. |

any potential differences between active and passive ellipsis in acceptability judgments. This model used the same coding scheme and priors used in the previous models (e.g.,  $\pm 0.5$  for VOICE,  $\text{Normal}(0,1)$ ).

## Results

Mean and median ratings by condition are shown in Fig. 3. Results of the Bayesian analyses are reported in Table 4 and a visualization of effect sizes is provided in Fig. 4. The crossed model revealed clear evidence for a main effect of GRAMMATICALITY, such that grammatical sentences were rated higher than ungrammatical sentences. There was no evidence of an effect of ATTRACTOR MATCH or an interaction with GRAMMATICALITY, as the CrIs spanned values above and below zero. There was no clear evidence of an effect of ATTRACTOR MATCH in either the grammatical conditions or the ungrammatical conditions in the nested comparisons.

Results of the cross-experiment comparison (Experiment 1 vs. Experiment 2) are reported in Table 5. There was a main effect of GRAMMATICALITY such that grammatical sentences received higher ratings than ungrammatical sentences. There was also a main effect of VOICE, such that sentences with active ellipsis received higher ratings than sentences with passive ellipsis. However, there was no evidence that passive and active ellipsis behaved differently with respect to interference in acceptability judgments.

## Discussion

Experiment 2 replicated the results of Experiment 1 for active ellipsis. There was a clear effect of grammaticality, but no evidence that an active attractor VP modulates acceptability. Crucially, for active and passive ellipsis alike, the availability of a feature-matching VP inside a relative clause does not appear to improve judgments relative to their mismatching counterparts. These results suggest that a matching VP in this position is not a possible antecedent for the ellipsis, as claimed in

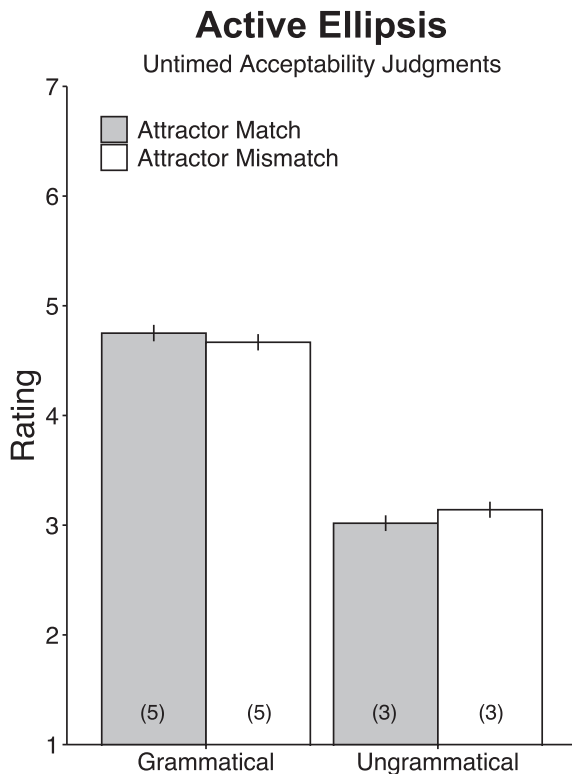


Fig. 3. Mean ratings for Experiment 2. Error bars indicate standard error of the mean. Median ratings are included in parentheses.

Table 4

Bayesian analysis of judgments from Experiment 2. The table reports the posterior means of the fixed effects and the corresponding 89% Bayesian credible intervals [CrI]. Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                           | Posterior mean [89% CrI] |
|----------------------------------|--------------------------|
| <b>Crossed Model</b>             |                          |
| Grammaticality                   | <b>1.72 [1.23, 2.21]</b> |
| Attractor Match                  | 0.08 [−0.15, 0.32]       |
| Grammaticality × Attractor Match | 0.27 [−0.09, 0.62]       |
| <b>Nested Model</b>              |                          |
| Grammatical Attraction           | −0.22 [−0.51, 0.07]      |
| Ungrammatical Attraction         | 0.05 [−0.22, 0.33]       |

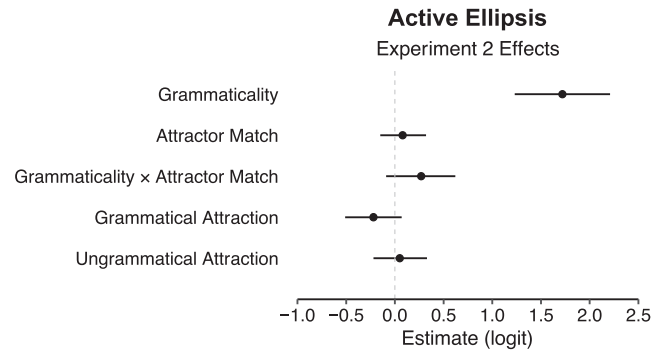


Fig. 4. Posterior probability distributions of effect sizes from Experiment 2. The dot corresponds to the mean of the distribution and the lines are the 89% credible interval (CrI). The dashed line indicates an effect of zero.

Table 5

Results of the cross-experiment comparison (Experiment 1 vs. Experiment 2). Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                              | Posterior mean [89% CrI] |
|-------------------------------------|--------------------------|
| Grammaticality                      | <b>1.78 [1.45, 2.11]</b> |
| Attraction                          | −0.01 [−0.13, 0.12]      |
| Voice                               | <b>0.70 [0.36, 1.03]</b> |
| Grammaticality × Attraction         | 0.15 [−0.09, 0.39]       |
| Grammaticality × Voice              | 0.02 [−0.50, 0.54]       |
| Attraction × Voice                  | 0.24 [−0.09, 0.56]       |
| Grammaticality × Attraction × Voice | 0.38 [−0.20, 0.78]       |

Martin (2018).

Importantly, the results of Experiments 1 and 2 are consistent with previous studies on the acceptability of voice (mis)matched ellipsis. For instance, Arregui et al. (2006) and Kim et al. (2011) tested sentences in which the ellipsis and its antecedent either matched in voice (“voice matched”) or mismatched (“voice mismatched”) using various acceptability measures (e.g., Likert scale, magnitude estimation). These studies revealed three key effects:

- Active voice matched sentences (e.g., Experiment 2 grammatical conditions) are more acceptable than passive voice matched sentences (e.g., Experiment 1 grammatical conditions).
- Voiced mismatched sentences with an active ellipsis and passive antecedent (e.g., Experiment 2 ungrammatical conditions) are more acceptable than voice mismatched sentences with a passive ellipsis and an active antecedent (e.g., Experiment 1 ungrammatical conditions).
- Voice matched sentences are more acceptable than mismatched sentences for active and passive ellipsis alike.

All three of these effects were observed in the current study, providing additional verification that the results of Experiments 1 and 2

are reliable. Taken together, the results of Experiments 1 and 2 confirm that we have a suitable interference paradigm to examine how active and passive voice cues influence retrieval processes in moment-by-moment processing.

### Experiment 3: Passive voice in moment-by-moment processing

The goal of Experiments 3 and 4 was to investigate how voice cues are used in retrieval for ellipsis processing. A recent study by [Martin \(2018\)](#) found that voice cues can trigger interference from feature-matching attractor VP. [Martin \(2018\)](#) tested both active and passive cues at retrieval and assumed that both would behave similarly with respect to interference effects. However, several studies have shown an asymmetry between active and passive ellipsis with respect to acceptability and processing dynamics (e.g., [Arregui et al., 2006](#); [Parker, 2018](#)), which raises the question of whether passive and active cues are used in the same way in retrieval for ellipsis processing. Furthermore, previous work on agreement attraction suggests that the markedness of the retrieval cues might modulate susceptibility to interference. The next experiments provide a comparison of interference effects for passive ellipsis (Experiment 3) and active ellipsis (Experiment 4), using the (mis)match paradigm developed in Experiments 1 and 2.

#### Participants

A power analysis indicated that 102 participants would be needed for at least 80% power, assuming an attraction effect with a mean of  $-28$  ms and a standard deviation of 100 ms (values obtained from the meta-analysis of attraction effects conducted by [Jäger et al., 2017](#)). For the current experiment, 120 self-reported native English speakers were recruited via Amazon's Mechanical Turk web service. All participants provided informed consent. Participants were compensated \$3.00 each. The experiment lasted approximately 25 min. Three participants were removed prior to analysis for failing the instructional check, leaving data from a total of 117 participants for analysis.

#### Stimuli and experimental design

Experiment 3 used the same passive ellipsis items that were used in Experiment 1. The 24 sets of test sentences were distributed across 4 lists in a Latin square design and combined with 48 grammatical filler sentences of similar length and complexity, such that each participant read a total of 72 sentences. Each sentence was followed by a comprehension question. Comprehension questions addressed various parts of the sentence to ensure that participants read and interpreted the entire sentence. Comprehension questions never targeted the interpretation of the ellipsis, so as not to draw attention to the critical manipulation. Rather, questions for these sentences targeted the subject-verb relation in the main clause (e.g., *Was it Jane who recruited for the event?*), the subject-verb relation in the relative clause (e.g., *Was it the villagers who organized the event?*), or the temporal and adjectival properties of the sentence.

#### Procedure

Experiment 3 used self-paced reading presented via Ibex. 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 dashes when the next word appeared. On-screen feedback was provided for incorrect answers to the comprehension questions. The order of presentation was randomized for each participant.

#### Data analysis

Experiment 3 followed the same data analysis steps as in

Experiments 1 and 2, using log-transformed reading times as the dependent variable (to aid interpretation, model estimates were back-transformed from the log scale to the millisecond scale). All data were included in the analysis. A Bayes factor analysis was conducted to quantify evidence for the presence of an attraction effect. The null hypothesis ( $H_0$ ) was that there was no attraction effect, and the alternative hypothesis ( $H_1$ ) was that there was an attraction effect. Since the Bayes factor is highly sensitive to the prior distribution ([Nicenboim & Vasishth, 2016](#)), the current analysis used priors obtained from the meta-analysis of attraction effects conducted by [Jäger and colleagues \(2017\)](#). The prior for the effect of attraction was defined as a normal distribution with a mean of  $-0.03$  and a standard deviation of 0.009 (see [Schad et al., 2021](#), for a detailed description of how these values are derived). Reading times were modeled at two regions of interest: the critical region and spillover region. The critical region was the final ellipsis marker *too* and the spillover region was the immediately following word.

#### Results

Mean word-by-word reading times by condition are shown in [Fig. 5](#). Mean reading times by condition and standard errors for the regions of interest are reported in [Table 6](#). Results of the Bayesian analyses are reported in [Table 7](#) and a visualization of effect sizes is provided in [Fig. 6](#). At the critical region, there was an effect of ATTRACTOR MATCH, such that sentences with a feature matching attractor were read faster than sentences with a mismatching attractor. There was also an effect of attraction in the ungrammatical conditions at the critical region. However, this effect should not be considered in the absence of the critical interaction of GRAMMATICALITY with ATTRACTOR MATCH. No other effects were observed in this region. At the spillover region, there was a main effect of GRAMMATICALITY and an interaction of GRAMMATICALITY with ATTRACTOR MATCH. The nested comparisons revealed that this interaction was driven by attraction in the ungrammatical conditions. There was no evidence of attraction in the grammatical conditions. The Bayes factor model comparison revealed an odds ratio of 11:1 in favor of the alternative model ( $H_1$ : an attraction effect) over the null model ( $H_0$ : no effect).

#### Discussion

Experiment 3 isolated the role of the passive voice cue in retrieval for ellipsis processing. Results revealed strong evidence that the passive voice cue triggers interference. Specifically, we observed facilitated reading times for ungrammatical sentences with an attractor VP that matched the passive voice cue, relative to ungrammatical conditions without a feature-matching attractor. These results provide a cross-methodological replication of the interference profile that [Martin \(2018\)](#) elicited using ERPs.

Together with the corresponding offline acceptability judgments from Experiment 1, these results are consistent with the claim that interference effects are most strongly observed in “online” moment-by-moment measures ([S. Lewis & Phillips, 2015](#); [Parker, 2019](#)). But what drives the contrast between online and offline measures with respect to attraction? One possibility suggested by an anonymous reviewer is that for ellipsis, parallelism might be a contributing factor. For instance, the degree of syntactic and semantic similarity between arguments has been shown to influence the resolution of ambiguous gapping structures, which involve a form of ellipsis ([Carlson, 2001](#)). In the stimuli from Experiments 1 and 3, the subject of the ellipsis clause is similar to the subject of the antecedent clause (e.g., both are proper names) and dissimilar to the subject of the relative clause, which is a definite description (e.g., *the villagers*). As such, judgments may have been biased to the main clause (i.e., target) VP on the basis of the similarity between subjects.

There are two reasons why parallelism is an unlikely explanation for the difference between Experiments 1 and 3 with respect to attraction.

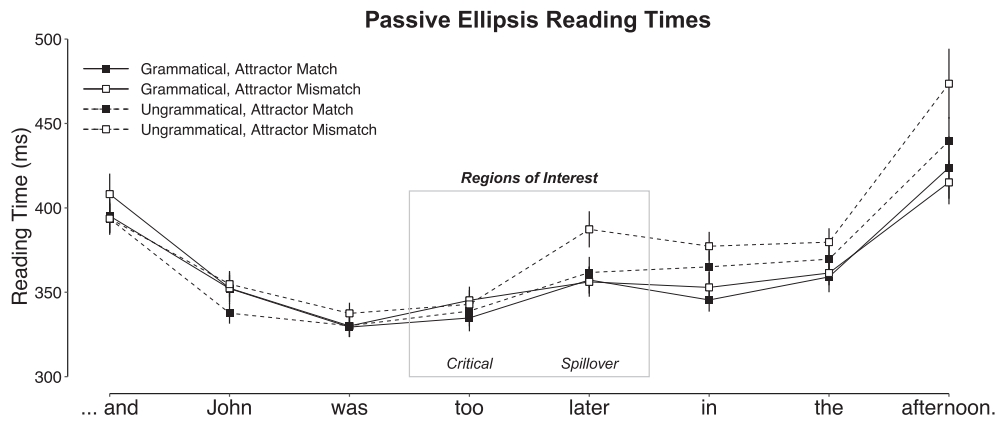


Fig. 5. Mean word-by-word reading times by condition from Experiment 3. Error bars indicate standard error of the mean.

Table 6

Mean reading times in ms by condition and standard error of the mean in parentheses for the regions of interest from Experiment 3.

| Condition                         | Critical Region | Spillover Region |
|-----------------------------------|-----------------|------------------|
| Grammatical, Attractor Match      | 334 (8)         | 357 (8)          |
| Grammatical, Attractor Mismatch   | 345 (8)         | 356 (9)          |
| Ungrammatical, Attractor Match    | 338 (9)         | 361 (9)          |
| Ungrammatical, Attractor Mismatch | 342 (8)         | 387 (11)         |

Table 7

Bayesian analysis of reading times at the critical and spillover regions from Experiment 3. The table reports the posterior means of the fixed effects in ms and the corresponding 89% Bayesian credible intervals [CrI]. Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                           | Posterior mean [89% CrI] |                  |
|----------------------------------|--------------------------|------------------|
|                                  | Critical Region          | Spillover Region |
| <b>Crossed Model</b>             |                          |                  |
| Grammaticality                   | -4 [-13, 5]              | -13 [-24, -3]    |
| Attractor Match                  | -7 [-14, 0]              | -7 [-16, 1]      |
| Grammaticality × Attractor Match | 12 [-3, 27]              | 26 [9, 43]       |
| <b>Nested Model</b>              |                          |                  |
| Grammatical Attraction           | -1 [-11, 9]              | 6 [-6, 17]       |
| Ungrammatical Attraction         | -13 [-23, -3]            | -20 [-32, -8]    |

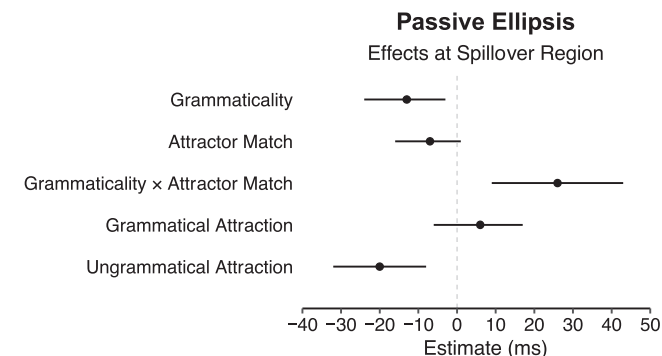


Fig. 6. Posterior probability distributions of effect sizes from Experiment 3. The dot corresponds to the mean of the distribution and the lines are the 89% credible interval (CrI). The dashed line indicates an effect of zero.

First, the fact that we do see sensitivity to the attractor VP in the reading times from Experiment 3 suggests that subject parallelism does not provide a strong constraint on ellipsis resolution in the configurations examined in the current study. Second, it is unclear why parallelism

would impact end-of-sentence judgments but not moment-by-moment processing, especially when parallelism shows an immediate effect in moment-by-moment processing (see Parker, 2017).

Ultimately, something must change over time to yield the observed contrast between online and offline measures. One possibility suggested in Parker (2019) is that the difference between online and offline measures reflects error-driven processing, in which the initial attraction error observed in time-sensitive measures triggers a reanalysis procedure that eventually leads to the correct judgment observed in untimed tasks. The results of the current study on ellipsis are consistent with this proposal. For instance, the initial interference effect observed in the reading times from Experiment 3 was quickly followed by increased sensitivity to the mismatch, reflected by a slow-down in the ungrammatical attractor-match condition relative to the grammatical conditions in the regions immediately following the spillover region. This divergence between conditions over time suggests that comprehenders began to reanalyze the sentence, leading to a reduction in interference effects, as observed in the untimed acceptability judgments from Experiment 1. In this respect, ellipsis and agreement pattern similarly with respect to online vs. offline profiles.

There are several other ways in which the effects observed in Experiment 3 resemble agreement attraction effects. First, the respective attractors (passive voice for ellipsis and plural number for agreement) affect the direction and magnitude of the interference effect similarly. The results of the Bayesian analysis suggest that the attraction effect had a mean of -20 ms, putting it on a par with agreement attraction, which shows an average facilitation of -21 ms in reading time measures (Jäger et al., 2017). Second, the interference effect observed for passive ellipsis shows the same grammatical asymmetry that is observed for agreement attraction, with interference arising only in ungrammatical configurations. Third, just as the marked plural cue gives rise to interference for agreement processing, we found that the marked passive cue gives rise to interference for ellipsis processing. The theoretical implications of these similarities are described in detail in the General Discussion. However, to fully establish that agreement and ellipsis pattern similarly with respect to feature markedness, we would need to also find that the unmarked counterpart involving active voice does not trigger attraction, just as the unmarked singular feature does not trigger attraction for agreement. This possibility is tested in Experiment 4.

#### Experiment 4: Active voice in moment-by-moment processing

Experiment 4 tested whether active voice triggers interference during retrieval for ellipsis processing using the paradigm verified in Experiment 2. The results were then compared with those from Experiment 3.

## Participants

One hundred twenty self-reported native English speakers were recruited via Amazon's Mechanical Turk web service. All participants provided informed consent. Participants were compensated \$3.00 each. The experiment lasted approximately 25 min.

## Stimuli and experimental design

Experiment 4 used the same active ellipsis items that were used in Experiment 2 and the same grammatical fillers from Experiment 4.

## Procedure

Experiment 4 followed the same steps as in Experiment 3.

## Data analysis

Experiment 4 followed the same data analysis steps as in Experiment 3. As in Experiments 1–2, Experiment 4 included an additional model to compare the effects of GRAMMATICALITY, ATTRACTION, VOICE (passive vs. active), and their interactions to evaluate potential differences between active and passive ellipsis in reading times. A Bayes factor analysis was also included to quantify evidence for differences between active and passive ellipsis with respect to attraction effects in reading times. These models were fit using the same coding scheme and priors used in the previous models (e.g.,  $\pm 0.5$  for VOICE, Normal(0,1)).

## Results

Mean word-by-word reading times by condition are shown in Fig. 7. Mean reading times by condition and standard errors for the regions of interest are reported in Table 8. Results of the Bayesian analyses are reported in Table 9 and a visualization of effect sizes is provided in Fig. 8. No effects were observed at the critical region as the credible intervals spanned values below and above zero. At the spillover region, there was a clear effect of GRAMMATICALITY, such that grammatical sentences were read faster than ungrammatical sentences. However, there was no evidence of an effect of ATTRACTOR MATCH or an interaction between GRAMMATICALITY and ATTRACTOR MATCH. There was also no evidence of attraction in either the grammatical or ungrammatical conditions in the nested comparisons. The Bayes factor analysis revealed an odds ratio of 9:1 in favor of the null model (H0) over the alternative model (H1).

Results of the cross-experiment comparison (Experiment 3 vs. Experiment 4) are reported in Table 10. There was a main effect of GRAMMATICALITY, such that overall, grammatical sentences were read faster than ungrammatical sentences. Crucially, there was an interaction of VOICE  $\times$  ATTRACTION and a three-way interaction of GRAMMATICALITY  $\times$  VOICE

**Table 8**

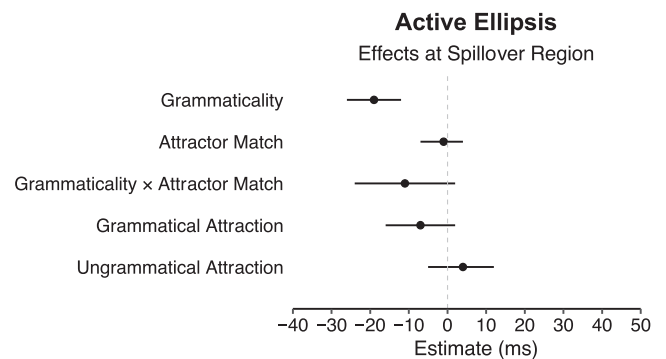
Mean reading times in ms by condition and standard error of the mean in parentheses for the regions of interest from Experiment 4.

| Condition                         | Critical Region | Spillover Region |
|-----------------------------------|-----------------|------------------|
| Grammatical, Attractor Match      | 329 (5)         | 331 (5)          |
| Grammatical, Attractor Mismatch   | 331 (6)         | 344 (6)          |
| Ungrammatical, Attractor Match    | 354 (10)        | 368 (7)          |
| Ungrammatical, Attractor Mismatch | 343 (7)         | 372 (10)         |

**Table 9**

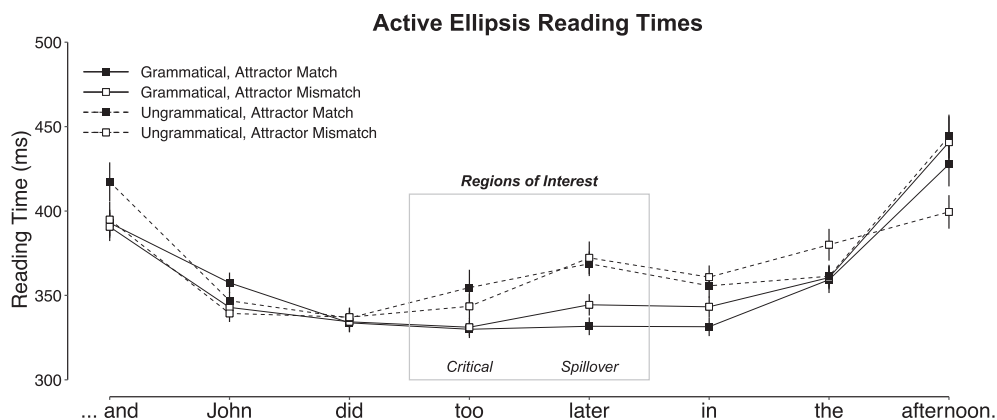
Bayesian analysis of reading times at the critical and spillover regions from Experiment 4. The table reports the posterior means of the fixed effects in ms and the corresponding 89% Bayesian credible intervals [CrI]. Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                                  | Posterior mean [89% CrI] |                  |
|---|--------------------------|------------------|
|   | Critical Region          | Spillover Region |
| <b>Crossed Model</b>                    |                          |                  |
| Grammaticality                          | −6 [−12, 0]              | −19 [−26, −12]   |
| Attractor Match                         | 3 [−3, 9]                | −1 [−7, 4]       |
| Grammaticality $\times$ Attractor Match | −3 [−17, 12]             | −11 [−24, 2]     |
| <b>Nested Model</b>                     |                          |                  |
| Grammatical Attraction                  | 1 [−8, 10]               | −7 [−16, 2]      |
| Ungrammatical Attraction                | 4 [−5, 13]               | 4 [−5, 12]       |



**Fig. 8.** Posterior probability distributions of effect sizes from Experiment 4. The dot corresponds to the mean of the distribution and the lines are the 89% credible interval (CrI). The dashed line indicates an effect of zero.

$\times$  ATTRACTION carried by the attraction effect in the ungrammatical conditions of the passive ellipsis sentences. This contrast is supported by the Bayes factor analysis, which revealed an odds ratio of 4:1 in favor of a difference between passive and active ellipsis with respect to



**Fig. 7.** Mean word-by-word reading times by condition from Experiment 4. Error bars indicate standard error of the mean.

interference effects in the ungrammatical conditions (i.e., the alternative model H1).

## Discussion

Experiment 4 isolated the role of active voice in ellipsis processing. Results revealed the surprising finding that unlike the passive voice cue tested in Experiment 3, active voice does not trigger attraction. Specifically, Experiment 4 revealed a clear effect of grammaticality, which suggests that participants were sensitive to the feature match between the ellipsis and target VP, but there was no evidence for an effect of attraction in any of the statistical comparisons. The contrast between passive and active voice with respect to attraction was supported by an interaction between voice (passive vs. active) and attraction in the ungrammatical conditions across Experiments 3 and 4.

The results from Experiment 4 provide evidence against the reasonable starting assumption in Martin (2018) that active and passive cues should behave similarly in retrieval for ellipsis processing (i.e., both should be deployed, and both should trigger interference). These results are surprising because it seems that a cue for active voice would be just as useful in recovering an antecedent as a cue for passive voice. I suggest that the selective profile with respect to voice interference reflects a markedness effect, like that observed for subject-verb agreement, where only marked forms show interference. This proposal is developed and made explicit with computational modeling in the General Discussion.

## General discussion

### Summary of findings

The goal of the current study was to better understand the source and scope of interference effects in sentence processing. Previous work on retrieval in sentence processing has focused on a relatively narrow set of configurations involving retrieval of NPs, such as the retrieval of a subject NP for subject-verb agreement and anaphora (see Jäger et al., 2017). But retrieval involves much more than NPs. The current study extended research on cue-based retrieval to VPs, focusing on the retrieval for VP-ellipsis, which involves a unique set of cues, such as voice.

The empirical starting point for the current study was the recent finding reported in Martin (2018) that non-target VPs that match the voice of the ellipsis clause trigger attraction, a form of interference. A critical assumption of this study was that passive and active cues should behave similarly in retrieval (i.e., both are deployed and both trigger interference). However, there were several reasons to think that active and passive ellipsis might behave differently based on evidence from acceptability judgments (Arregui et al., 2006) and computational modeling (Parker, 2018). Previous findings of markedness effects in retrieval, such as the singular-plural asymmetry for agreement attraction (e.g., Bock & Eberhard, 1993; Eberhard, 1997; Harley & Ritter, 2002; Wagers et al., 2009), further suggest that passive and active ellipsis might behave differently with respect to attraction due to

differences in the markedness of passive vs. active ellipsis features. In particular, studies on agreement attraction have shown that marked plural agreement give rises to attraction, but unmarked singular agreement typically does not (but cf. Hammerly et al., 2019). Based on these findings, it was hypothesized in the current study that the marked passive forms would give rise to attraction, but the unmarked active form would not.

The current study addressed these issues in a series of four experiments. Experiments 1 and 2 used untimed acceptability judgments to verify the constraints on VP-ellipsis. Experiments 3 and 4 then used the paradigms vetted in Experiments 1 and 2 to test for attraction effects in moment-by-moment processing, with independent tests for active and passive ellipsis. Self-paced reading measures revealed a clear contrast: passive ellipsis is susceptible to attraction, but active ellipsis is not. Below, I discuss the implications of these findings for current theories of attraction effects and show how the current findings can be captured in a cue-based retrieval architecture using computational modeling.

### Theoretical contribution of current findings

The current findings have several implications for our understanding of the scope and source of interference effects in sentence processing. First, from a methodological standpoint, the finding of attraction for passive ellipsis in self-paced reading measures in the current study provides a cross-methodological replication of the findings reported in Martin (2018). Martin used a match-(mis)match paradigm that involved a mix of passive and active VP ellipsis and observed a clear attraction effect in ERP measures, reflected as a modulation of the P600 amplitude for ungrammatical sentences with a matching attractor VP. The current study narrows the conclusions of this work by showing that attraction from voice features is limited to passive ellipsis.

Second, the current study shows that attraction effects across verbal dependencies (ellipsis and subject-verb agreement) are closely aligned in particular ways not previously attested in the literature:

1. **Direction.** Both agreement and ellipsis attraction are characterized by eased processing at the retrieval site (i.e., a speed-up) in the presence of a grammatically irrelevant but feature matching attractor.
2. **Magnitude.** A recent Bayesian meta-analysis of attraction effects conducted by Jäger and colleagues (2017) found that agreement attraction triggers on average a facilitation of  $-21$  ms in reading time measures (95% credible interval: [95% CrI:  $-36.4$ ,  $-9$ ]). By comparison, ellipsis attraction, as observed in the current study, triggered a facilitatory effect with nearly the same magnitude:  $-20$  ms [95% CrI:  $-35$ ,  $-6$ ].
3. **Grammatical asymmetry.** Many studies on agreement attraction show a “grammatical asymmetry”, with attraction observed in ungrammatical, but not grammatical configurations. The current study on ellipsis also showed a grammatical asymmetry in the same direction as that observed for subject-verb agreement, with attraction found only in ungrammatical configurations.
4. **Feature markedness asymmetry.** Many studies on agreement attraction show a feature markedness effect, such that agreement attraction arises for marked plural verbs, but not unmarked singular verbs. In the current study, ellipsis showed a parallel feature markedness asymmetry, with attraction arising for marked passive sentences, but not unmarked active sentences.

The tempting conclusion to draw from these descriptive similarities is that there is a homogenous underlying cause for attraction effects in subject-verb agreement and ellipsis processing. However, there is considerable debate over the source of attraction effects, in particular, regarding the question of whether attraction reflects an error in the encoding or retrieval process. I discuss below how the current findings relating to the grammatical asymmetry (point 3 above) and the feature

**Table 10**

Results of the cross-experiment comparison (Experiment 3 vs. Experiment 4). Effects with 89% CrIs that do not cross 0 are in bold.

| Effect                              | Posterior mean [89% CrI] |                       |
|-------------------------------------|--------------------------|-----------------------|
|                                     | Critical Region          | Spillover Region      |
| Grammaticality                      | <b>-5 [-10, 0]</b>       | <b>-16 [-23, -10]</b> |
| Attraction                          | -2 [-6, 2]               | -5 [-10, 1]           |
| Voice                               | -8 [-22, 5]              | -7 [-22, 8]           |
| Grammaticality × Attraction         | 5 [-5, 14]               | 8 [-3, 18]            |
| Grammaticality × Voice              | 1 [-5, 7]                | 3 [-4, 9]             |
| Attraction × Voice                  | <b>-5 [-9, 0]</b>        | -3 [-8, 2]            |
| Grammaticality × Attraction × Voice | 8 [-3, 18]               | <b>18 [8, 29]</b>     |

markedness asymmetry (point 4) contribute to the debate about the theoretical framework that best captures attraction effects.

There are two leading accounts of attraction in comprehension. One account locates the source of attraction effects in the *retrieval* processes that are used to form linguistic dependencies (e.g., Wagers et al., 2009). According to this account, attraction reflects misretrieval of a non-target item in a cue-based memory architecture (e.g., R. L. Lewis & Vasishth, 2005; Van Dyke & McElree, 2006). The other account locates the source of attraction effects in the *encoding* of items in memory (e.g., Bock & Eberhard, 1993; Eberhard, 1997; Eberhard et al., 2005; Franck et al., 2002; Hammerly et al., 2019; Patson & Husband, 2015; Staub, 2009, 2010). On this view, attraction arises due to spreading activation or movement (“percolation”) of number information in a sentence (e.g., Bock & Eberhard, 1993; Eberhard, 1997; Eberhard et al., 2005; Franck et al., 2002), resulting in an equivocal representation of the subject’s number marking that disrupts agreement computation.

Previously, the grammatical asymmetry observed for subject-verb agreement has been presented as decisive evidence favoring a retrieval-based account. According to retrieval-based accounts (e.g., Wagers et al., 2009), attraction does not arise in grammatical configurations with singular agreement because (1) the privatively specified agreement system does not deploy a singular number cue, eliminating the possibility of attraction from singular attractors, and (2) the full matching target subject will outcompete items that match only a subset of the cues (i.e., partial-matches). By contrast, encoding-based accounts predict symmetrical effects because the equivocal representation of number on the subject should impact grammatical and ungrammatical sentences similarly (Phillips et al., 2011; Wagers et al., 2009).

However, Hammerly and colleagues (2019) recently used forced choice acceptability judgments (‘Yes’/‘No’) to show that symmetrical attraction effects arise in grammatical and ungrammatical configurations when the bias towards acceptable responses (i.e., a ‘Yes’ judgment) is eliminated. These results favor an encoding-based account that assumes a continuous valuation of number on the subject NP. Further evidence against a retrieval-based account comes from a recent study by Avetisyan et al. (2020) showing that distinctive case marking on NPs does not modulate attraction effects in comprehension. This finding is expected under encoding-based accounts, which lack a mechanism for case information to impact number representation on the subject phrase. Similarly, Schlueter et al. (2018) observed attraction from coordinated singular attractors, which lack a plural marking. This finding is unexpected under retrieval-based accounts, which require a match on the plural retrieval cue for attraction to occur.

At present, the finding that ellipsis shows a grammatical asymmetry like that previously observed for agreement does not arbitrate between the competing perspectives on the source of attraction effects. The asymmetry can be argued to reflect a retrieval-based effect in the same way that has been argued in the past for agreement attraction. However, further tests of response bias, for instance, might reveal a symmetrical effect like that reported in Hammerly et al. (2019). Such a finding would favor an encoding-based account. I leave the task of investigating the role of response bias in ellipsis processing to future work.

One finding from the current study that does favor one of the accounts is the markedness asymmetry between passive and active ellipsis. First, attraction for passive ellipsis is not easily accommodated under an encoding-based framework, as it seems unlikely that the encoding of voice on the target antecedent would be disrupted by spreading activation of passive voice on the attractor in such a way that would mislead ellipsis processing. For instance, revaluing an active VP as passive as the result of spreading activation/feature movement would require extensive modification to the encoding of the VP beyond simply altering the voice valuation of the head (as proposed for number), including reanalysis of thematic assignments (subject → object) and the introduction of additional passive morphology (e.g., *be* + past participle).

Second, and more importantly, the passive/active asymmetry observed for ellipsis is directly predicted by retrieval-based accounts

that impose a privative feature system (e.g., Wagers et al., 2009), as I show in the next section.

#### *How to capture feature marked effects in a cue-based retrieval architecture*

For both ellipsis and subject-verb agreement dependencies, marked features (e.g., passive voice, plural number) trigger attraction effects, but unmarked features (e.g., active voice, singular number) do not. For agreement, number is traditionally represented in grammatical theories as privative and categorical (e.g., see den Dikken, 2011, for a review). On this view, pluralitive languages like English make a categorical, two-way distinction such that an item is either singular or plural, and the marked plural number is represented by the presence of a [+PL] feature, whereas unmarked singular is represented by the absence of a number feature (rather than as [+SG], for instance). Voice receives a parallel treatment: it is categorical (active or passive) and the different forms are distinguished by the presence/absence of particular features (e.g. passive constructions introduce syntactic and morphological markings like an auxiliary verb, *-en*, *by*-phrase that are absent in active constructions) (e.g., Givón, 1990; Poppels & Kehler, 2019; Rohde & Kehler, 2014; Shibatani, 1985).<sup>3</sup>

I suggest that the feature markedness effects observed for ellipsis and agreement reflects a privative cue specification at retrieval. For agreement, the marked plural verb deploys a number cue [+PL], but the unmarked singular verb does not (Wagers et al., 2009). Likewise, for ellipsis, the marked passive ellipsis would deploy a voice cue [+passive], but the unmarked active counterpart would not. That is, for both dependencies, there is a cue deployed for the marked form (passive, plural), but not for the unmarked form (active, singular). This cue specification leads to attraction effects from matching non-target items for the marked forms, but no such effects are expected for the unmarked forms, since the relevant cue is not deployed at retrieval.

To offer proof of concept, I simulated the proposed privative cue specification using the ACT-R model of sentence comprehension (R. L. Lewis & Vasishth, 2005), which is a prominent model used to study retrieval interference effects (e.g., Nicenboim & Vasishth, 2018). ACT-R (Adaptive Control of Thought-Rational; Anderson et al., 2004) is a general cognitive architecture based on independently motivated principles of memory and cognition. It has been applied to investigate a wide range of cognitive behavior involving memory access, attention, executive control, and learning. The ACT-R model of sentence processing applies the cognitive principles embodied in the general ACT-R framework to the task of sentence processing.

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 mechanism inspired by the attribute-value matrices from Head-driven Phrase Structure Grammar (Pollard & Sag, 1994). Chunks are encoded as bundles of feature-value pairs. Features are specified for lexical content (e.g., morpho-syntactic and semantic features), syntactic information (e.g., category, case), and local hierarchical relations (e.g., parent, daughter, sister). Values for features include symbols (e.g.,  $\pm$ singular,  $\pm$ passive) or pointers to other chunks (e.g., NP1, VP2).

Linguistic dependencies, such as subject-verb agreement and ellipsis,

<sup>3</sup> One issue that arises for this treatment of voice is the existence of a third type of voice in English known as “middle voice,” as in *The book reads easily*. Middle voice is a non-active voice that is distinct from passive. If middle voice is treated as a syntactic feature distinct from active and passive, it is unclear how to situate it within the proposed feature calculus. However, if middle voice is encoded as “non-active,” rather than as a separate feature, as proposed in the formal literature (see Alexiadou, 2014 for discussion), then it may be possible to extend the current proposal for a privative feature system to “non-active” voice (e.g., passive and middle voice) as the overt feature and active voice as the absence of a voice feature.

are formed using a general retrieval mechanism that evaluates all items in memory, in parallel, using a set of retrieval cues that target specific features of individual memory chunks. Retrieval cues are derived from the current word, the linguistic context, and grammatical constraints, and correspond to a subset of the features of the target (R. L. Lewis et al., 2006). 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. Attraction effects are explained in this model as misretrieval of an attractor phrase that partially matches the retrieval cues (Dillon et al., 2013; Vasishth et al., 2008; Wagers et al., 2009).

The simulations for the current study were conducted using a variant of the ACT-R model of sentence processing based on the equations described in Lewis and Vasishth (2005) and Vasishth, Brissow, Lewis, and Drenhaus (2008). The code for the model was originally developed by Badecker and Lewis (2007). This is the same implementation used in previous studies on interference effects (e.g., Dillon et al., 2013; Kush & Phillips, 2014; Parker, 2018; Parker & Phillips, 2017). Following previous studies, I adopted the standard assumption that longer retrieval latencies entail longer reading times.

Three models were constructed to simulate the processing of passive and active ellipsis: (i) a model for passive ellipsis that included a passive voice cue (+passive), (ii) a model for active ellipsis that did not include a voice cue (set to NULL) to simulate a privative feature system, and (iii) a non-privative model for active ellipsis that included a voice cue (+active) for comparison with the privative model. Following previous studies that have modeled retrieval for dependency formation (e.g., Dillon et al., 2013; Engelmann et al., 2019; Jäger et al., 2020; Kush & Phillips, 2014; Nicenboim & Vasishth, 2018; Parker, 2018; Parker & Phillips, 2017; Vasishth et al., 2008), each model included syntactic cues for category, clause, and depth of embedding. All models used the same set of syntactic cues and the default parameters. The only difference was the specification of the voice cue in the manner described above. The syntactic cues provided a perfect match to the target VP but mismatched the attractor VP along all dimensions. The possibility of misretrieval (i.e., attraction) arises when the voice cue (if there is one) matches the attractor VP and mismatches the target VP. The use of both structural and non-structural cues constitutes an “unconstrained” retrieval procedure, as proposed for other dependencies like subject-verb agreement and reflexives (Jäger et al., 2015). This assumption, along with the possibility of partial-matching (i.e., retrieval based on a match to a subset of the retrieval cues), widens the positions that the retrieval procedure can access, permitting attraction effects (Dillon, 2014). The code for the current simulations is available on the Open Science Framework [see link on title page].

The privative models qualitatively predicted the contrast observed in the empirical data: passive ellipsis showed attraction, but active ellipsis did not. Quantitatively, these models provide a good fit to the observed data, as the 89% CrIs of the model's predictions fall within the 89% CrIs for the observed effects (Fig. 9). By contrast, the profile predicted by the non-privative model that used a voice cue for active ellipsis was not aligned with the observed data (predicted 89% CrI:  $-25[-31, -18]$  vs. observed 89% CrI  $4[-5, 12]$ ). That is, only the models with a privative cue specification for voice were able to capture the profiles observed in the empirical data.

In sum, the proposed analysis treats feature markedness in the same way across dependencies. Specifically, marked features will be deployed as retrieval cues, but unmarked features will not. On this view, there is a close alignment between the way in which (un)marked forms are represented in the grammar (i.e., privative, categorical) and the way in which they are implemented in real-time processing.

## Privative Feature Model vs. Observed Data

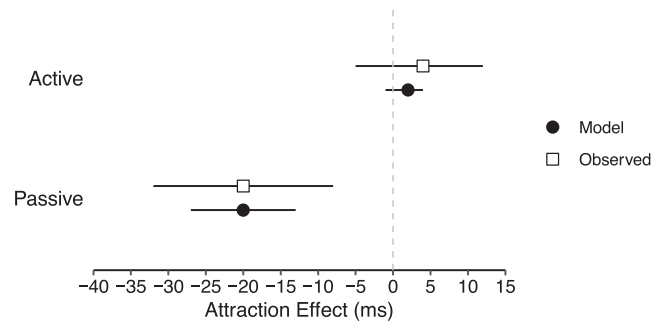


Fig. 9. Comparison of the model predicted attraction effects for passive and active ellipsis and the observed attraction effects from Experiments 3 and 4. The dot/square corresponds to the posterior mean and the lines are the 89% credible interval (CrI). The dashed line indicates an effect of zero.

Lastly, it is important to discuss how the current retrieval-based account of ellipsis fits with existing formal theories of ellipsis. Formal theories of ellipsis typically fall into one of two categories. Syntactic theories assume that the content of the ellipsis site involves detailed structure (e.g., Fiengo & May 1994; Lasnik, 2001; Merchant, 2001, 2008; Ross, 1969; Williams, 1977), whereas referential theories assume that the ellipsis site involves a null proform/pointer, akin to other types of referential expressions, such as pronouns, which lacks internal syntactic structure (e.g., Culicover & Jackendoff, 2005; Ginzburg & Sag, 2000; Hardt, 1999; Martin & McElree, 2008; Tanenhaus & Carlson, 1990). The proposed retrieval-based account fits most naturally with a pointer-style analysis of ellipsis resolution, which is consistent with recent experimental work showing that real-time ellipsis resolution is mediated by a pointer mechanism (e.g., Martin & McElree, 2008). However, the current data does not rule out the possibility that there is detailed structure at the ellipsis site, which could be reconstructed as a post-retrieval operation. I leave the task of distinguishing these possibilities to future work.

## Conclusion

The current study investigated interference effects in the processing of ellipsis constructions. Results revealed a surprising contrast, showing that passive ellipsis is susceptible to attraction effects, but active ellipsis is not. It was suggested that the selective profile with respect to voice interference reflects a markedness effect, like that observed for agreement attraction, where only marked forms show interference. According to this proposal, the marked passive ellipsis construction deploys a voice cue, whereas the unmarked active counterpart does not, leading to the observed profiles. These results point to a uniform account of markedness effects across dependencies and demonstrate that interference effects across dependencies are more similar than previously assumed. Lastly, it was shown how such effects arise in a cue-based retrieval architecture.

## CRedit authorship contribution statement

**Dan Parker:** Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Funding acquisition.

## Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

I would like to thank Adrian Staub, the three anonymous reviewers, and the members of the Computational & Experimental Linguistics Lab at William & Mary for their feedback on this work. I would also like to thank Irene Williams and Hayley Shankle for their help in preparing the stimuli for this study. This work was funded by NSF award BCS-1843309 to Dan Parker.

## References

- Alexiadou, A. (2014). Active, middle, and passive: The morpho-syntax of Voice. *Catalan Journal of Linguistics*, 13, 19–40.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychology Review*, 111, 1036–1060.
- Arregui, A., Clifton, C., Jr., Frazier, L., & Moulton, K. (2006). Processing elided VPs with flawed antecedents. *Journal of Memory and Language*, 55, 232–246.
- Avetisyan, S., Lago, S., & Vasishth, S. (2020). Does case marking affect agreement attraction in comprehension? *Journal of Memory and Language*, 112, 1040872.
- 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. In 20th CUNY Conference on human sentence processing. San Diego: University of California.
- Barr, D. J., Levy, R. P., Scheepers, C., & Tily, H. J. (2013). Random effects structure for testing interactions in linear mixed-effects models. *Frontiers in Psychology*, 4, 328.
- Bock, K., & Eberhard, K. M. (1993). Meaning, sound and syntax in English number agreement. *Language and Cognitive Processes*, 8, 57–99.
- Bürkner, P.-C. (2017). Brms: An r package for Bayesian multilevel models using stan. *Journal of Statistical Software*, 80, 1–28.
- Carlson, K. (2001). The effects of parallelism and prosody in the processing of gapping structures. *Language and Speech*, 44, 1–26.
- Culicover, P. W., & Jackendoff, R. (2005). *Simpler syntax*. Oxford University Press.
- den Dikken, M. (2011). Phi-feature inflection and agreement: An introduction. *Natural Language and Linguistic Theory*, 29, 857–874.
- Dillon, B. (2014). Syntactic memory in comprehension of reflexive dependencies: An overview. *Language and Linguistics Compass*, 8, 171–187.
- Dillon, B., Mishler, A., Sloggett, S., & Phillips, C. (2013). Contrasting intrusion profiles for agreement and anaphora: Experimental and modeling evidence. *Journal of Memory and Language*, 69, 85–103.
- Eberhard, K. M. (1997). The marked effect of number on subject–verb agreement. *Journal of Memory and Language*, 36, 147–164.
- Eberhard, K. M., Cutting, J., & Bock, K. (2005). Making syntax of sense: Number agreement in sentence production. *Psychological Review*, 112, 531–559.
- Engelmann, F., Jäger, L. A., & Vasishth, S. (2019). The effect of prominence and cue association on retrieval processes: A computational account. *Cognitive Science*, 12, Article e12800.
- Fiengo, R., & May, R. (1994). *Indicies and identity*. MIT Press.
- Franck, J., Vigliocco, G., & Nicol, J. (2002). Subject-verb agreement errors in French and English: The role of syntactic hierarchy. *Language and Cognitive Processes*, 17, 371–404.
- Futrell, R., Gibson, E., & Levy, R. P. (2020). Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing. *Cognitive Science*, 44.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). Chapman and Hall/CRC.
- Gerken, L., & Bever, T. G. (1986). Linguistic intuitions are the result of interactions between perceptual processes and linguistic universals. *Cognitive Science*, 10, 457–476.
- Ginzburg, J., & Sag, I. (2000). *Interrogative investigations*. CSLI Publications.
- Givón, T. (1990). *Syntax: A functional typological introduction*. John Benjamins.
- Hammerly, C., Staub, A., & Dillon, B. (2019). The grammatical asymmetry in agreement attraction reflects response bias: Experimental and modeling evidence. *Cognitive Psychology*, 70–104.
- Hardt, D. (1999). Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy*, 22, 187–221.
- Harley, H., & Ritter, E. (2002). Person and number in pronouns: A feature-geometric analysis. *Language*, 78, 482–526.
- Jäger, L. A., Benz, L., Roeser, J., Dillon, B., & Vasishth, S. (2015). Teasing apart retrieval and encoding interference in the processing of anaphors. *Frontiers in Psychology*, 6, 1–18.
- Jäger, L. A., Engelmann, F., & Vasishth, S. (2017). Similarity-based interference in sentence comprehension: Literature review and Bayesian meta-analysis. *Journal of Memory and Language*, 94, 305–315.
- Jäger, L. A., Merten, D., Van Dyke, J. A., & Vasishth, S. (2020). Interference patterns in subject-verb agreement and reflexives revisited: A large-sample study. *Journal of Memory and Language*, 111. <https://doi.org/10.1016/j.jml.2019.104063>
- Kim, C., Koble, G., Runner, J., & Hale, J. (2011). The acceptability cline in VP-ellipsis. *Syntax*, 14, 318–354.
- Kim, C., & Runner, J. (2018). The division of labor in explanations of verb phrase ellipsis. *Linguistics and Philosophy*, 41, 41–85.
- Kimball, J., & Aissen, J. (1971). I think, you think, he think. *Linguistic Inquiry*, 2, 241–246.
- Kohonen, T. (1980). *Content-addressable memories*. Springer-Verlag.
- Kruschke, J. K., Aguinis, H., & Joo, H. (2012). The time has come: Bayesian methods for data analysis in the organizational sciences. *Organizational Research Methods*, 15, 722–752.
- Kush, D., & Phillips, C. (2014). Local anaphor licensing in an SOV language: Implications for retrieval strategies. *Frontiers in Psychology*, 5, 1–12.
- Lasnik, H. (2001). When can you save a structure by destroying it? In M. Kim, & U. Strauss (Eds.), *Proceedings of the North East linguistic society* (Vol. 31, pp. 301–320). Amherst, MA: GLSA Publications.
- Levy, R. P. (2013). Memory and surprisal in human sentence comprehension. In R. P. G. van Gompel (Ed.), *Current issues in the psychology of language sentence processing* (pp. 78–114). Psychology Press.
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29, 375–419.
- Lewis, R. L., Vasishth, S., & Van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Science*, 10, 447–454.
- Lewis, S., & Phillips, C. (2015). Aligning grammatical theories and language processing models. *Journal of Psycholinguistic Research*.
- Martin, A. E. (2018). Cue integration during sentence comprehension: Electrophysiological evidence from ellipsis. *PLOS ONE*, 13(11).
- Martin, A. E., & McElree, B. (2008). A content-addressable pointer mechanism underlies comprehension of verb-phrase ellipsis. *Journal of Memory and Language*, 58, 879–906.
- Martin, A. E., & McElree, B. (2009). Memory operations that support language comprehension: Evidence from verb-phrase ellipsis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 1231–1239.
- Martin, A. E., & McElree, B. (2011). Direct-access retrieval during sentence comprehension: Evidence from sluicing. *Journal of Memory and Language*, 64, 327–343.
- Martin, A. E., Nieuwland, M. S., & Carreiras, M. (2012). Event-related brain potentials index cue-based retrieval interference during sentence comprehension. *NeuroImage*, 59(1859–1869).
- Martin, A. E., Nieuwland, M. S., & Carreiras, M. (2014). Agreement attraction during comprehension of grammatical sentences: ERP evidence from ellipsis. *Brain & Language*, 135, 42–51.
- Merchant, J. (2001). *The syntax of silence: Sluicing, islands, and the theory of ellipsis*. Oxford University Press.
- Merchant, J. (2008). Variable island repair under ellipsis. In K. Johnson (Ed.), *Topics in ellipsis*. Cambridge University Press.
- Merchant, J. (2013). Voice and ellipsis. *Linguistic Inquiry*, 44, 77–108.
- Nicenboim, B., & Vasishth, S. (2016). Statistical methods for linguistic research: Foundational ideas—Part II. *Language and Linguistics Compass*, 10, 591–613.
- Nicenboim, B., & Vasishth, S. (2018). Models of retrieval in sentence comprehension: A computational evaluation using Bayesian hierarchical modeling. *Journal of Memory and Language*, 99, 1–34.
- Parker, D. (2017). Processing multiple gap dependencies: Forewarned is forearmed. *Journal of Memory and Language*, 97, 175–186.
- Parker, D. (2018). A memory-based explanation of antecedent-ellipsis mismatches: New insights from computational modeling. *Glossa: A Journal of General Linguistics*, 3, 1–27.
- Parker, D. (2019). Two minds are not always better than one: Modeling evidence for a single sentence analyzer. *Glossa: A Journal of General Linguistics*, 4, 1–31.
- Parker, D., & Phillips, C. (2017). Reflexive attraction in comprehension is selective. *Journal of Memory and Language*, 94, 272–290.
- Patson, N. D., & Husband, M. (2015). Misinterpretations in agreement and agreement attraction. *Quarterly Journal of Experimental Psychology*, 69, 950–971.
- Phillips, C., & Parker, D. (2014). The psycholinguistics of ellipsis. *Lingua*, 151, 78–95.
- Phillips, C., Wagers, M. W., & 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). Emerald Publications.
- Pollard, C., & Sag, I. (1994). *Head-driven phrase structure grammar*. University of Chicago Press.
- Poppels, T., & Kehler, A. (2019). Reconsidering asymmetries in voice-mismatched VP-ellipsis. *Glossa: A Journal of General Linguistics*, 4, 1–22.
- R Development Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Rohde, H., & Kehler, A. (2014). Grammatical and information-structural influences on pronoun production. *Language, Cognition and Neuroscience*, 29, 912–927.

- Ross, J. R. (1969). Guess who? In R. I. Binnick, A. Davidson, G. M. Green, & J. L. Morgan (Eds.), *Proceedings of the fifth annual meeting of the Chicago linguistics society* (pp. 252–286). Chicago: University of Chicago.
- Schad, D. J., Nicenboim, B., Bürkner, P.-C., Betancourt, M., & Vasishth, S. (2021). Workflow techniques for the robust use of Bayes factors. *Psychological Methods*, 26.
- Schlueter, Z., Williams, A., & Lau, E. F. (2018). Exploring the abstractness of number retrieval cues in the computation of subject-verb agreement in comprehension. *Journal of Memory and Language*, 99, 74–89.
- Schütze, C. T. (2016). *The empirical base of linguistics: Grammaticality judgments and linguistic methodology*. Language Science Press.
- Shibatani, M. (1985). Passives and related constructions: A prototype analysis. *Language*, 61, 821–848. <https://doi.org/10.2307/414491>
- Staub, A. (2009). On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language*, 60, 308–327.
- Staub, A. (2010). Response time distributional evidence for distinct varieties of number attraction. *Cognition*, 114, 447–454.
- Tanenhaus, M. K., & Carlson, G. (1990). Comprehension of deep and surface verb phrase anaphora. *Language and Cognitive Processes*, 5, 257–280.
- Townsend, D. J., & Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. MIT Press.
- Van Dyke, J. A., & McElree, B. (2006). Retrieval interference in sentence comprehension. *Journal of Memory and Language*, 55, 157–166.
- Vasishth, S., Brüssow, S., Lewis, R. L., & Drenhaus, H. (2008). Processing polarity: How the ungrammatical intrudes on the grammatical. *Cognitive Science*, 32, 685–712.
- Wagers, M. W., Lau, E. F., & Phillips, C. (2009). Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language*, 61, 206–237.
- Williams, E. (1977). Discourse and logical form. *Linguistic Inquiry*, 8, 101–139.