



Structure Modulates Similarity-Based Interference in Sluicing: An Eye Tracking study

Jesse A. Harris *

Department of Linguistics, University of California, Los Angeles, Los Angeles, CA, USA

In cue-based content-addressable approaches to memory, a target and its competitors are retrieved in parallel from memory via a fast, associative cue-matching procedure under a severely limited focus of attention. Such a parallel matching procedure could in principle ignore the serial order or hierarchical structure characteristic of linguistic relations. I present an eye tracking while reading experiment that investigates whether the sentential position of a potential antecedent modulates the strength of similarity-based interference, a well-studied effect in which increased similarity in features between a target and its competitors results in slower and less accurate retrieval overall. The manipulation trades on an independently established Locality bias in sluiced structures to associate a *wh*-remnant (*which ones*) in clausal ellipsis with the most local correlate (*some wines*), as in *The tourists enjoyed some wines, but I don't know which ones*. The findings generally support cue-based parsing models of sentence processing that are subject to similarity-based interference in retrieval, and provide additional support to the growing body of evidence that retrieval is sensitive to both the structural position of a target antecedent and its competitors, and the specificity or diagnosticity of retrieval CUES.

Keywords: working memory, similarity-based interference, ellipsis, eye tracking, sentence processing

OPEN ACCESS

Edited by:

Colin Phillips,
University of Maryland, USA

Reviewed by:

Masaya Yoshida,
Northwestern University, USA
Andrea Eyleen Martin-Nieuwland,
University of Edinburgh, UK

***Correspondence:**

Jesse A. Harris
jharris@humnet.ucla.edu

Specialty section:

This article was submitted to
Language Sciences,
a section of the journal
Frontiers in Psychology

Received: 10 June 2015

Accepted: 13 November 2015

Published: 18 December 2015

Citation:

Harris JA (2015) Structure Modulates Similarity-Based Interference in Sluicing: An Eye Tracking study. *Front. Psychol.* 6:1839. doi: 10.3389/fpsyg.2015.01839

INTRODUCTION

The rapid formation of non-adjacent syntactic dependencies during online sentence comprehension offers intuitive evidence for the importance of an efficient retrieval system. Two well-studied cases are argument-verb dependencies, in which an argument must be related to its verb no matter the amount of intervening material (1a), and anaphoric dependencies, in which a pronominal element, like *him* or *one*, is associated with another co-referring expression, possibly from among multiple possibilities (1b). The noun phrase *the barber* must be retrieved from memory in each case: either in the subject-trace (gap) position t_1 (1a) or else as a co-referring expression (1b). Only in the former case is the dependency unambiguously determined by structure; the second case is simply the most plausible given the topicality of *the barber* and real world knowledge, as illustrated by fact that likely co-reference possibilities depend on the predicate (1c).

- (1) a. The barber₁ (that John₂ wanted his son₃ to visit) $t_{1/*2,*3}$ died yesterday.
 - b. He_{1>2>3} was 95 years old.
 - c. He_{2>3,#1} was saddened to hear the news.

Findings from both types of cases support recent cue-based parsing models of sentence processing in which all possible antecedents are activated in parallel through a fast, domain-general associative cue-matching procedure (for review of evidence and models see, e.g., Lewis et al., 2006; Van Dyke and Johns, 2012; Caplan and Waters, 2013). In such models, a retrieval cue—e.g., the verb *died* in (1a) or the pronoun *he* in (1b), initiates direct access to all possible targets (possible antecedents or dependencies) from memory. However, not all possible targets must be equally activated within memory: some targets might receive greater activation by virtue of sharing syntactic or semantic features with the retrieval cue, while other targets might receive less activation as a function of temporal decay (Van Dyke and Johns, 2012, for review, though a reviewer points out that decay as the primary source of forgetting is not strongly supported in the general memory literature, as demonstrated by Keppel and Underwood, 1962 and Waugh and Norman, 1965). Further, the allocation of attentional resources in memory is severely constrained, possibly limited to a single item within the focus of attention (e.g., McElree, 2001). Thus, the memory architecture employed in sentence processing strongly resembles the architecture thought to underlie domain-general tasks (e.g., McElree, 2000; McElree et al., 2003; Van Dyke and Lewis, 2003; Lewis and Vasishth, 2005; Lewis et al., 2006).

A major research question within the cue-based parsing literature addresses the extent to which domain-specific knowledge influences the retrieval process; in particular, whether syntactic constraints affect antecedent retrieval, and, if so, at precisely what stage of retrieval. Much of the research conducted thus far has examined whether syntactically inaccessible targets are considered viable candidates for retrieval. The two major schools of thought addressing this issue in the cue-based parsing literature are (i) structure-based accounts, and (ii) unconstrained-cue accounts. In the former, grammatical constraints filter the set of possible antecedents retrieved from memory by constraining the search set to grammatically permissible positions (e.g., Nicol and Swinney, 1989; Sturt, 2003; Xiang et al., 2009; Chow et al., 2014). Thus, items in structurally inaccessible positions would be effectively ignored in the retrieval process, though grammatically illicit items may feed repair processes triggered by retrieval failure (Chow et al., 2014). In contrast, the latter, unconstrained-cue, approach proposes that grammatical constraints are but one of the many possible factors guiding retrieval (e.g., Badecker and Straub, 2002; Lewis and Vasishth, 2005; Chen et al., 2012; Jäger et al., 2015). Badecker and Straub (2002), for example, propose that the set of possible antecedents are restricted not by tree geometry, but a number of other factors, such as attention or discourse importance. Elements in the focus of attention are often identified with discourse topics or a center, though details regarding how to define topics vary considerably (see, e.g., Chafe, 1976; Gundel, 1985; Gundel et al., 1993; Grosz et al., 1995). Under this type of account, syntactically inaccessible items might well interfere with retrieving a target, particularly if such items are highly salient in the discourse, though, of course, structure may be one of the factors that determines discourse salience. Although it should be clear that both viewpoints agree that structural information is a

factor in retrieval, they permit very different types of mechanisms by which structure is utilized.

Unfortunately, results that would clearly arbitrate between the two camps are somewhat mixed. For example, while filler-gap dependencies are sensitive to syntactic islands (for review see, e.g., Sprouse et al., 2012), retrieving pronouns and reflexives has sometimes been found to be susceptible to interference from structurally inaccessible items (e.g., Badecker and Straub, 2002; Jäger et al., 2015; though see Chow et al., 2014 for a failure to replicate Badecker and Straub's results). The general finding is that syntactic information plays an important role in retrieval for at least certain types of dependencies (Van Dyke and McElree, 2006, 2011; Van Dyke, 2007; Dillon et al., 2014; Cunnings et al., 2015), even if it is currently unclear whether that role is to filter possible antecedents or provide weighted or probabilistic cues regarding their likelihood of matching the probe (Badecker and Straub, 2002; Cunnings et al., 2015). Further, it is unclear at this stage whether syntactic information should be treated as value in a feature bundle (e.g., Lewis and Vasishth, 2005), and if configurational relations like *c-command* inherently preclude such a treatment (see Alcocer and Phillips, 2012; Kush, 2013, for discussion).

The present research seeks to widen the empirical coverage of the issue of structural import by leveraging an independently established bias guiding the resolution of correlate-remnant pairs in clausal ellipsis (sluiced) structures. In (2), for example, the remnant *which (ones)* might be paired with either a subject correlate (*a few linguists*) or an object correlate (*some silly examples*), spelled out as (2a) and (2b), respectively.

- (2) [A few linguists] gave [some silly examples], but I don't remember which (ones).
- a. ...I don't remember which linguists. (Subject correlate)
 - b. ...I don't remember which examples. (Object correlate)

Although sluicing ellipsis permits correlates in both subject and object position above, it shows a strong preference for the latter (Frazier and Clifton, 1998). As mentioned, much research in content-addressable retrieval systems in language processing addresses whether grammatical restrictions govern the resolution of various types of anaphoric dependencies, especially the availability of syntactically illicit antecedents (e.g., Sturt, 2003; Martin et al., 2012, 2014; Cunnings et al., 2015; see also Phillips et al., 2011, for review). In contrast, correlates to sluiced remnants are merely heavily biased, rather than constrained grammatically. Thus, sluiced structures provide a potentially revealing counterpoint to studies investigating syntactic barriers to accessibility: if structurally defined preferences, in addition to grammatical restrictions, influence the retrieval process, the retrieval system may use structural information to privilege some products of memory over others. In which case, the retrieval system might be said to avail itself of domain-specific, linguistically defined structural preferences, in addition to hard-coded grammatical principles. In the remainder of the paper, I review the aspects of cue-based parsing that are most relevant for this study, along with the basic assumptions regarding the ellipsis structures explored here. I

then present two experiments which collectively support the idea that retrieving antecedents in sluiced sentences is subject to interference effects (congruent with Martin and McElree, 2011, discussed in detail below) and that the strength of such interference effects depends on the sentential position of possible targets and strength of the cue provided at retrieval.

Cue-based Parsing and Similarity-based Interference

Many models of how linguistic representations are encoded and retrieved during real-time language comprehension have been proposed (for review, see Van Dyke and Johns, 2012; Caplan and Waters, 2013). In contrast to cue-based retrieval systems, traditional models employ an operation that searches through items in memory, typically in a serial fashion (Doshier and McElree, 1992; McElree and Doshier, 1993), as developed for short-term memory in general (Sternberg, 1966, 1975). In these models, the path of the search varies according to whether the search queue starts with the first item encountered, as in forward search, or the most recent item, as in backward search. The basic prediction of search models is that retrieval time should increase as a function of search space within the path: the more items that must be searched through, the greater the search time required to do so. This prediction, however, has not been supported by Speed Accuracy Tradeoff (SAT) experiments which report that speed of reaching a stable judgment or interpretation is unaffected by the size of the putative search space, even though accuracy is (McElree and Doshier, 1989; McElree, 2000, 2006; McElree et al., 2003; Foraker and McElree, 2007, 2011; Martin and McElree, 2008, 2009, 2011). For example, McElree (2001) manipulated the amount of intervening material (underlined) between a relative clause filler *the book* and its object position gap t_1 , shown in (3). Subjects were trained to make an acceptability judgment in response to tones at various post-sentence intervals. Measuring the rate at which responses achieved a stable interpretation (the asymptote) as a function of time, he found that while accuracy decreased as more material intervened between the filler and the gap, the speed at which the asymptote was reached did not (see also McElree et al., 2003).

- (3) a. This was [the book]₁ that the editor admired t_1 .
 b. This was [the book]₁ that the editor who the receptionist married admired t_1 .
 c. This was [the book]₁ that the editor who the receptionist who quit admired t_1 married.

That accuracy, not the rate of reaching the asymptote, is affected by increasing items in memory not only provides an important argument against search models, it also suggests that the quality of the representations recovered is susceptible to interference from distractors. A content-addressable retrieval system is able to account for this tradeoff by proposing that all items are compared against the target in parallel, not by a search operation, but by an automatic associative cue-matching procedure which compares partial representations of items in memory (possibly encoded as bundles of semantic and grammatical features; Clark

and Gronlund, 1996) against cues from the target (as in ACT-R models, Anderson et al., 2004; Lewis et al., 2006). Thus, linear or structural distance is irrelevant for retrieval times, but additional competitors interfere with the cue-matching process by introducing partial matches, thereby degrading the quality of the retrieved item. In other words, the greater the similarity or overlap between a target and its competitors, the greater the effect of interference (Crowder, 1976; Anderson and Neely, 1996).

Similarity-based interference, i.e., the failure to successfully distinguish a target from similar competitors in retrieval, has now been well documented both in general memory manipulations (Nairne, 2002; Öztekin and McElree, 2007) and in language processing contexts (Lewis, 1996; Gordon et al., 2001, 2002, 2004, 2006; Van Dyke and Lewis, 2003; Van Dyke and McElree, 2006, 2011; Van Dyke, 2011; Autry and Levine, 2014, among others). Dual task paradigms combine the two, so that a subject attempts to retain a list of words in working memory while processing a separate sentence for comprehension. When items in the memory set are similar to critical words in the sentence, performance decreases on both reading speed and comprehension accuracy (Fedorenko et al., 2006), especially if those items overlap with retrieval cues for long distance dependencies (Van Dyke and McElree, 2006).

What is less clear is the extent to which structural information modulates the accessibility of a target. On the one hand, grammatically illicit cues appear to license a syntactically dependent element in case of illusory licensing (including negative polarity items, NPI, Vasishth et al., 2008; Xiang et al., 2013, as well as agreement attraction, Wagers et al., 2009; Dillon et al., 2013). To illustrate, comprehenders often accept configurations in which an NPI licenser like *no* simply precedes, rather than *c-commands*, the NPI *ever*, as in *The restaurants that no newspapers have recommended in their reviews have ever gone out of business*, raising the possibility that configurational information may sometimes be ignored. On the other hand, several recent studies show that distractors within the same syntactic position more greatly interfere in the formation of long distance dependencies (Van Dyke and McElree, 2006, 2011; Dillon et al., 2013, 2014). For example, Van Dyke and McElree (2011) observed that syntactic constraints limit the amount of interference exerted by semantically similar distractors, without eliminating interference completely (also Van Dyke, 2007; Van Dyke and Johns, 2012). They attribute these results to greater disturbance from *retroactive interference*, in which retrieval is hampered by a distractor D that separates a cue C from its target T (schematically: T-D-C), than from *proactive interference*, in which retrieval is impeded by distractors processed before the target (D-T-C).

Similarly, Dillon and colleagues find an increased interference effect for structurally distant antecedents of a reflexive *ziji* in Mandarin Chinese—i.e., when a distractor intervened between a target and the reflexive. They propose that structurally local domains restrict the initial area for dependency formation.

- (4) **Local search hypothesis:** The parser uses positional syntactic information during the retrieval of syntactic dependents, and

positional cues serve to restrict retrieval to constituents in some local syntactic domain.

The Local Search Hypothesis contrasts with content-addressable retrieval models that would explain putative effects of Locality in terms of activation decay (see comments in Lewis et al., 2006). Lewis and Vasishth (2005), for example, sharply discriminate items within the focus of attention from those outside of it, which are subject to decay, unless they are reactivated during retrieval processes. A model of this type need not rely on structural information, or even serial order, for retrieval cues. To the extent that syntactic information is utilized, it is established through encoding of morphosyntactic features like [\pm Theme] or [\pm Object] in their feature bundles, which collectively identify appropriate structures for retrieval. Thus, such models are entirely domain-general in the sense that the memory operations active during retrieval of linguistic material are the same as those that are active during other types of retrieval. If correct, this uniformity would be a powerful virtue—why postulate specialized retrieval procedures for linguistic structures when the unique aspects of language parsing could be captured simply in terms of specialized features comprising the representations that form the products of memory? In other words, what would be distinctive about language processing would be not so much the *mechanisms* involved in retrieval as how objects over which such mechanisms operate would be *encoded* in memory.

Nevertheless, the following study lends further support to the general finding that the sentence position for an antecedent is relevant during the retrieval process as linguistic dependencies are resolved, though it cannot by itself resolve whether it is best to conceive of such information as structural in nature over sequential or temporal orderings. The study capitalizes on key properties of sluicing ellipsis, as introduced in the following section.

Sluicing and the Locality Bias

Sluicing describes focus-sensitive clausal ellipsis after a *wh*-question (Ross, 1967, 1969; Chung et al., 1995; Merchant, 2001, among others), such as (5a) below. Following Merchant (2001) and others, I assume an account of sluicing which derives the overt structure through movement of the *wh*-element *who*₁ from its base-generated, clause-internal position to a fronted position followed by optional ellipsis \langle he is meeting *t*₁ \rangle of the remainder of the clause. Thus, the unelided sentence (5b) is the source for the sluice (5a).

- (5) a. John is meeting {someone/a friend} for dinner, but I can't tell you who₁ \langle he is meeting *t*₁ \rangle .
 b. John is meeting {someone/a friend} for dinner, but I can't tell you who₁ he is meeting *t*₁.

Sluicing places restrictions on the types of nouns that can serve as correlate to the remnant, though these restrictions depend on the type of *wh*-element and its restrictor residing in the remnant (Chung et al., 1995; Romero, 1998). For example, proper names and definite nouns are often unacceptable correlates for a *who*-remnant unless it is followed by *else* (6a). In select cases, the

wh-element may co-refer with an adjunct (6b) or argument (6c) correlate that did not appear in the antecedent clause overtly, in an operation that Chung et al. (1995) call “sprouting.”

- (6) a. John is meeting Mary/the president for dinner, but I can't tell you who *(else).
 b. John is meeting Mary/the president for dinner, but I can't tell you where/why/with who.
 c. John ate, but I can't tell you what₁.

Sluicing, along with other forms of ellipsis, has received much attention in recent processing literature (Frazier and Clifton, 1998, 2005, 2011; Carlson et al., 2009; Martin, 2010; Poirier et al., 2010; Dickey and Bungler, 2011; Martin and McElree, 2011, among others). Previous results from processing ellipsis support the expectations of content-addressable retrieval systems, in that retrieval of antecedent material at the ellipsis site appears not to be affected by the size or complexity of the recovered material, an effect explained as either a cost-free copying mechanism (Frazier and Clifton, 2001, 2005; Frazier, 2008) or as a direct pointer in memory (Martin and McElree, 2009). In addition, Martin and McElree (2011) find that increasing the distance to a correlate in sluiced sentences affects retrieval accuracy, not retrieval speed, as predicted by content-addressable systems in which retrieval speed is held constant (as various models of retrieval propose, e.g., McElree, 2000; McElree et al., 2003; Lewis and Vasishth, 2005; Lewis et al., 2006). More detailed comparison to previous studies on sluiced sentences and cue-based parsing is delayed until the General Discussion.

Another important general result is that sluices show a structural preference to associate the remnant with the most local correlate in the antecedent clause, a principle formalized as the Locality bias below (see also Harris and Carlson, 2015, for a similar preference with *let alone* ellipsis).

- (7) **Locality bias:** Associate the remnant of clausal ellipsis with a correlate occupying the structurally most local position.

Initial evidence for the Locality bias came from Frazier and Clifton (1998), who manipulated whether a sluiced sentence contained one or more possible correlates (8). In a self-paced reading study, they found that cases with multiple possible antecedents (8b) were read faster than unambiguous structures (8a). The penalty for (8a) can be attributed, in effect, to a violating the preference for Local correlates¹.

- (8) a. Somebody claimed that the president fired someone but nobody knows who.
 b. Somebody claimed that the president fired Fred but nobody knows who.

¹Fully ambiguous sentences also show a strong Locality bias in silent reading, as explored in the controls for a different study (Harris, 2015). Subjects saw four ambiguous sluiced sentences like “An editor called a journalist, but I can't say which one it was any more” and were asked how they interpreted the sentences. There was an 86% bias toward the most local, object-correlate interpretation. In addition, the vast majority of the subjects ($N = 48$) displayed a complete or majority Locality bias for the items: 3 subjects were at chance, and only 1 subject had a consistent preference for non-local correlates.

Carlson et al. (2009) provide additional support for the Locality bias for sluiced sentences in an auditory questionnaire. They observed that unless the subject was focus-marked by a pitch accent or an *it*-cleft, sentences with object correlates were rated higher than alternatives.

A similar Locality bias has been observed for sluices with sprouted antecedents. Frazier and Clifton (2005) report a naturalness rating and reading time advantage when a verb with an implicit object (*studied*) appeared in second conjunct position (*slept and studied*) as compared to first conjunct position (*studied and slept*), as in *Michael (slept and studied/studied and slept), but he didn't tell me what*². The basic result coheres with the expectations of the Locality bias in that near antecedents confer a processing advantage over far antecedents (see also Martin and McElree, 2011), though it does complicate the prediction that antecedent distance or complexity is not relevant to retrieval. As an aside, the Locality bias is not unique to sluiced sentences. It has been observed in other move-and-delete types of ellipsis, such as replatives (Carlson, 2013) and *let alone* ellipsis (Harris and Carlson, 2015), as well as ambiguous gapping structures (Carlson et al., 2005).

The locus of the Locality bias is open to multiple possibilities. One such possibility derives from the assumption that the processor must ultimately recover elided material for interpretation, presumably by employing default biases and cues from information structure. Another, perhaps not mutually exclusive, possibility suggested by Frazier and Clifton (1998) and Carlson et al. (2005, 2009) is that the most likely correlate is determined by default-focus marking on the most embedded constituent (Selkirk, 1984; Cinque, 1993). In canonical English SVO sentences, the most deeply embedded constituent happens to be the object. For whatever reason, the preferences guiding remnant resolution in sluicing ellipsis appear to diverge from the first-mention bias established for third person pronouns, in which a pronoun is preferentially associated with the subject of a preceding clause (Arnold, 1998, among others).

In any event, there seems to be good evidence that sluiced sentences prefer the most local correlate as the antecedent for the remnant. We now turn to how the expectation for Local correlates might affect content-addressable retrieval systems, as outlined above.

The Current Study

An important advantage of using sluicing ellipsis to address the questions above is that the retrieval cues may be explicitly manipulated by modifying the inner restrictor of the *wh*-element. In example (2), for instance, the correlate-remnant pair can be disambiguated simply by repeating the nominal phrase directly, as in *A few linguists gave some silly examples, but I don't remember which linguists*. Such cases determine which noun functions as the correlate to the remnant by completely specifying the relationship: in such cases, there is *total overlap* between remnant and correlate. The eye tracking experiment below exploits this

²Frazier and Clifton (2005) argued that this asymmetry indicates the presence of syntactic structure within the ellipsis site, as the restriction on accessibility follows from their "conjunction domain hypothesis," an independently motivated syntactic constraint on extraction. See Martin and McElree (2011) for commentary.

possibility by manipulating whether the restrictor is completely specified by a nominal like *which tourists/wines* (cue-rich probe) or partially specified by a pronoun like *which ones* (cue-poor probe), along with whether the indefinite (assumed to be the preferred correlate) appears in the preferred object location (9a) or not (9b). In addition, it manipulated whether a definite noun distractor appeared in the plural, thereby providing partial cue overlap with the indefinite.

- (9) a. The tourists sampled some wine(s), but I don't know which wines/ones.
 b. Some tourist(s) sampled the wines, but I don't know which tourists/ones.

Sluiced sentences like (9) conceivably involve two instances of retrieval: first, the recovery of the elided IP after the remnant, and, second, the pairing between the remnant and the correlate. Regarding the recovery of the elided IP, I adopt an approach in which a syntactic representation is recovered through some sort of cost-free mechanism, such as syntactic copying or recycling (Frazier and Clifton, 2001, 2005) or a pointer in memory (Martin and McElree, 2009), such that the size and complexity of the antecedent clause is essentially irrelevant for retrieval speed (Martin and McElree, 2009, 2011).

Regarding the pairing of the remnant with the correlate, there are several theoretical options to consider, especially with respect to the different types and strengths (diagnosticity) of cues in the remnant. First, we might imagine that the parser forms a dependency between the remnant and correlate selectively, that is, only when the remnant contains a pronoun, as in *ones*, but not when its inner restrictor is fully specified, as in *tourists* or *wines*. In this case, a fully specified restrictor could be interpreted via straightforward composition, without retrieving a correlate. However, this approach is unlikely given results from sprouting in sluicing ellipsis, which show a penalty when there is no overt correlate in antecedent clause (Frazier and Clifton, 1998; Dickey and Bunger, 2011).

The two remaining options would require that a dependency be formed between all types of correlates and remnants, but differ in what type of mechanism establishes it. One option to consider is one in which the nominal in the restrictor obviates the associative cue-matching procedure by forming a direct link to the previous instance of the noun, trivially avoiding cue overload effects altogether. Another option is that establishing a dependency between correlate and remnant evokes an associative cue-matching procedure, as proposed for anaphoric dependencies in general, but mitigates cue overload effects by virtue of the total overlap in cues between the remnant probe and the target correlate. In either case, we would expect the strength of the cue at the remnant to modulate the retrieval process. For concreteness, I adopt the latter approach, acknowledging that the experiments below do not depend on or arbitrate between these two possibilities.

As observed by a reviewer, it may be important that the two types of dependencies are not independent: if a comprehender resolves the remnant to an object-position correlate in (9), she is also committed to a particular syntax for the IP ellipsis, e.g., [*which wines*]₁ *they sampled* t₁/[*which tourists*]₁ t₁ *sampled them*.

While this dependency should be explored in depth, we will not do so here³. I will simply assume that however recovery of the IP ellipsis impacts retrieval of the remnant, the effects will be comparable across conditions.

An important conceptual issue for cue-based parsing models in general is what types of information constitute cues for retrieval. In such models, it is conceivable that any information coded as a feature value in a feature bundle is qualified to serve as a cue for retrieval, though some types of information, especially relational information, might be less amenable to representation by features than others (Kush, 2013). Following recent literature (e.g., Lewis and Vasishth, 2005), I assume that retrieval is driven, at least in part, by features from lexical (gender and number) and morphosyntactic (grammatical roles and case) information derived from context, and that what is retrieved are partial representations of constituents. Further, retrieval occurs whenever an item has to be associated with another item in memory for complete interpretation, including better-studied cases of anaphora resolution, verbal agreement, NPI licensing, and variable binding, although different kinds of dependencies might attend to distinct types of information. I remain largely agnostic about the internal organization of retrieval with respect to other interpretive processes, e.g., whether retrieval is discrete, continuous, or cascaded, as the study was not designed to address these issues, and the results below are consistent with any number of possibilities.

Assuming an associative cue-matching procedure and a preference for local antecedents, the reading experiment below was designed to test the following two basic predictions:

P1. Locality: The most local antecedent, in this case the object noun, is favored for retrieval.

P2. Nominal Advantage: Nominal restrictors (*which tourists/which wines*) include a rich set of cues specifying retrieval, and thus facilitate retrieval over cue-poor probes like *which ones*.

The most important prediction, however, is one in which distractors outside the local (object) domain are subject to varying degrees of interference; a strong effect of interference is

predicted only in case of partial overlap, as fully specifying the inner restrictor with a nominal should eliminate the effects of cue overload, either by delivering the appropriate correlate directly, or via total overlap between the remnant and the target correlate.

P3. Structure-Dependent Interference: A retrieval penalty for violating Locality

- i. arises when a distractor in the preferred (object) position shares features with the remnant, and
- ii. increases if retrieval is initiated by cue-poor probes (*ones*).

Prior research investigating the effect of structural constraints on retrieval has often used the gender feature in a feature mismatch paradigm (e.g., Clifton et al., 1999; Badecker and Straub, 2002; Sturt, 2003; Chow et al., 2014). Manipulating gender agreement between the remnant and the antecedents was not possible here, given that English does not encode gender for impersonal pronouns like *ones*. Therefore, we must first show that plural definite nouns are viable correlate competitors for unambiguous *which* remnants, like *which tourists* or *which wines*, the central task of the following experiment. An affirmative finding will support the assumption that the plurality feature sufficiently induces similarity-based interference effects in the formation of correlate-remnant pairs with pronouns like *ones* in the next experiment. In addition, it will address the assumption regarding whether the indefinite determiner *some* marks the preferred correlate for sluices with *which*-remnants as opposed to the definite determiner.

As a final terminological note, the present use of “interference” diverges somewhat from a common use in the literature, in which the distractor is not a grammatical antecedent, or otherwise inaccessible (e.g., Van Dyke, 2007; Phillips et al., 2011). If both nouns in the matrix are acceptable as antecedents, the manipulation might be best cast in terms of a “fan effect” in which multiple non-referents interfere with dependency resolution (Anderson, 1974; Anderson and Reder, 1999; Autry and Levine, 2014). However, as the effects in either case would ideally be driven by the same underlying types of retrieval mechanisms, I retain the use of interference here, in hopes of expanding the empirical range of strongly biased, though not strictly speaking ungrammatical, structural preferences.

EXPERIMENT 1

A forced-choice completion test was first conducted over the Internet in order to determine the extent to which a plural definite noun competes with a plural indefinite as a correlate. I take such cases to be indicative of similarity-based interference effects, although they may differ in kind from other types of interference. A further question is whether the extent to which plurality makes a definite noun an appealing correlate is affected by its structural position.

Method

Participants

Twenty-nine subjects were recruited using Amazon Mechanical Turk, an Internet-based service where individuals complete

³The dependency between retrieval of a correlate and the syntax of the ellipsis site raises a number of pertinent issues regarding the relationship between the Locality Bias and parallelism between the matrix and the second conjunct. The interpretation of sluicing is known to be sensitive to parallelism (Carlson, 2002; Dickey and Bunker, 2011), as is conjunction generally (Frazier et al., 1984). If the processor takes the remnant phrase to be an object, perhaps on the basis Minimal Attachment (Frazier, 1978), the processor might prefer an object correlate to create structural parallelism between clauses. In an unpublished auditory forced-choice completion study ($N = 48$), Katy Carlson and I manipulated the surface position of the remnant to appear either as a subject (*A waiter talked to a guest, but ... which waiter/which guest isn't clear*) or as an object (*A waiter talked to a guest, but ... it's not clear which waiter / which guest*), along with the location of a pitch accent (subject or object). In addition to a general 66% bias for local object-position correlates, we found that subject position remnants (*which guest / which waiter isn't clear*) failed to elicit more subject responses than object position remnants did, and that pitch accent placement was the primary determinant of continuation choice. Further, accenting the object resulted in 8% more object continuations with a surface subject remnant, suggesting that linguistic focus (manifested here in the form of pitch accent), not syntactic parallelism, is the driving force behind the bias for local correlates in standard sluicing constructions.

short tasks online for payment. One subject self-identified as a non-native English speaker, and was removed from analysis. A pretest evaluated subjects' competency with three difficult to interpret questions. Three subjects were removed for answering one or more of these questions incorrectly. Four catch items were included to identify inattentive subjects, but no subject was removed on this basis. However, an additional subject was removed for counterbalancing purposes, leaving 24 subjects in the final data set. All subjects were compensated \$4 for their participation, regardless of native language or performance. This experiment, along with the following, were carried out with prior Internal Review Board approval from Pomona College. All subjects gave written informed consent before starting the experiment, and were permitted to remove themselves at any time from the procedure without penalty.

Materials

The 2×2 experimental design crossed Indefinite Location (Object indefinite, Subject indefinite) with Definite Number (Plural, Singular). The levels of the Indefinite Location condition were determined by its syntactic position in the matrix clause. In other words, there were two sequences of determiner in the matrix clause: either (i) a definite subject (singular or plural) followed by a plural indefinite object, or (ii) a plural indefinite subject followed by a definite object (singular or plural). The Plural condition was created from the Singular condition simply by adding the plural marker to the definite noun phrase, e.g., *tourist ~ tourists* or *wine ~ wines*. All critical nouns except one (*fireman ~ firemen*) were regular plurals. Twenty-four quartet fragments like (10) were constructed below. Items are reported in Appendix A of Supplementary Material.

(10) Object indefinite

- a. *Plural definite*: The tourists sampled some wines, but I've forgotten...
- b. *Singular definite*: The tourist sampled some wines, but I've forgotten...

Subject indefinite

- c. *Plural definite*: Some tourists sampled the wines, but I've forgotten...
- d. *Singular definite*: Some tourists sampled the wine, but I've forgotten...

Two forced-choice completions (11) were provided under the fragments in (10). The response options always agreed in plurality with the preceding sentence fragment, e.g., *tourists/wines* in (10a,c), *tourist/wines* in (10b), and *tourists/wine* in (10d). Answers were presented in a different random sequence for each subject.

(11) Forced-choice options

- i. Subject correlate response: which tourist(s).
- ii. Object correlate response: which wine(s).

After a short guided practice consisting of three sample sentences, subjects were presented with an additional 52 items from unrelated experiments with various structures,

12 non-experimental fillers, in which both responses were acceptable, and four catch items permitting only a single correct answer, for a total of 92 items.

Procedure

Items were presented in an individually randomized and fully counterbalanced order, so that subjects saw one and only one sentence fragment from each quartet. Subjects were instructed to rely on their intuitions to select whichever response would make "the resulting sentence sound the most natural." Subjects were given an hour to complete the task, but all subjects finished within 40 min, with an average of 25 min per subject. In addition, encrypted versions of IP addresses were recorded to identify subjects who may have taken the experiment more than once. No such cases were observed.

Results

One item (item 1 in Appendix A of Supplementary Material) contained a typo and was removed from analysis. Data analysis was conducted in R version 3.1.2 (R Development Core Team, 2014). Mean percent and standard deviations for *subject completion* responses by condition are presented in **Table 1**.

Conditions were given sum (deviation) coding so that the hypothetically simplest condition, the Object Correlate—Singular condition, was treated as the statistical baseline. The response data was modeled as a logistic linear mixed effects regression model using the lme4 package (Bates and Maechler, 2009) with by-subjects and by-items random slopes and intercepts, shown in **Table 2**.

As expected, the choice between Subject and Object correlate response closely corresponded to the location of the indefinite *some*: subject position indefinites garnered greater overall Subject correlate responses ($M = 71\%$, $SE = 3$) than object position indefinites ($M = 24\%$, $SE = 3$), $z = 6.77$, $p < 0.001$. The result confirms the intuition that language users prefer indefinites as correlates to remnants in sluiced structures, though

TABLE 1 | Experiment 1: percent *subject* response selected.

Indefinite location	Definite number		Correlate mean	Attraction effect
	Plural	Singular		
Object	28% (4)	19% (3)	24% (3)	−9%
Subject	67% (4)	75% (4)	71% (3)	8%

Standard errors are in parentheses.

TABLE 2 | Experiment 1: results of linear mixed effects regression model.

	Estimate	Std. error	Wald Z	p-estimate
(Intercept)	−0.179	0.362	−0.494	0.621
Correlate location	1.661	0.285	5.834	< 0.001
Definite number	−0.038	0.208	−0.183	0.855
Correlate × Definite number	−0.473	0.199	−2.376	< 0.05

Significant effects are printed in bold.

we should note that the preference is not absolute; see also the Discussion Section of Experiment 2, which acknowledges several complications.

While there was no main effect of Definite Number in this model, there was an interaction between Indefinite Location and Definite Number. In the case of a subject indefinite, more subject completions were observed when the object noun was singular (10d) than plural (10c); in the case of an object indefinite, more subject completions were elicited when the subject noun was plural than when it was singular (10b) than plural (10a), $z = 3.67, p < 0.001$. This reversal is to be expected if the indefinite provides the preferred candidate for the correlate, but a plural distractor interferes with the distinctiveness of the indefinite target.

These patterns are consistent with the theoretically-motivated assumption that *which*-remnants in sluicing prefer the antecedent with the most accessible set of individuals. In this case, the indefinite description *some* makes a set of alternatives salient in the discourse, as opposed to a definite description, which arguably introduces a plural sum of individuals that can be interpreted as a single entity (Link, 1983). Accordingly, responses were transformed to reflect the pairing in which the remnant forms a contrast with the indefinite, as depicted in Table 3. The transformed response data was modeled as a logistic linear mixed effects regression model using the lme4 package (Bates and Maechler, 2009) with by-subjects and by-items random slopes and intercepts and deviation coding as before. The result is shown in Table 4.

The model supports a sole effect of Definite Number, in that a plural definite noun (*the wines/the tourists*) resulted in fewer responses that co-referred with the indefinite target ($M = 69\%$, $SE = 3$) than singular definite distractors ($M = 78\%$, $SE = 3$), $t = -2.38, p < 0.05$, although the indefinite is still generally preferred.

TABLE 3 | Experiment 1: responses by the percentage of cases in which the indefinite was selected as the correlate.

Correlate choice	Plural interference		Correlate mean	Interference difference
	Interference	No interference		
Object	72% (4)	81% (3)	76% (3)	9%
Subject	67% (4)	75% (4)	71% (3)	8%
Interference mean	69% (3)	78% (3)		

Standard errors in parentheses.

TABLE 4 | Experiment 1: results of linear mixed effects regression model on the proportion of transformed responses.

	Estimate	Std. error	Wald Z	p-estimate
(Intercept)	1.661	0.285	5.835	< 0.001
Correlate location	-0.177	0.361	-0.49	0.624
Definite number	-0.474	0.199	-2.38	< 0.05
Correlate × Definite number	-0.038	0.208	-0.181	0.857

Significant effects are printed in bold.

Discussion

The results suggest that the structures are not fully ambiguous: there is a strong preference to associate the remnant with an indefinite correlate. The transformed results also show clear support for general similarity-based interference, in that plural definite distractors, which shared the plurality feature with an indefinite correlate, attracted more remnant resolutions than singular definite distractors.

EXPERIMENT 2

A second experiment was conducted to test the central predictions outlined above. First, by *Locality*, subject position indefinite nouns should elicit an online processing penalty over their more local, object position counterparts. Second, by *Nominal Advantage*, *wh*-restrictors with a fully specified nominal should facilitate the retrieval of their correlates compared to pronominal restrictors like *ones* by virtue of providing a richer feature set for cue-matching. Lastly, by *Structure-Dependent Interference*, plural definite nouns should exert a greater interference effect on retrieval when occupying object position, and, further, that such effects should manifest predominantly when the cues for retrieval are poor.

Method

Participants

Fifty-six native English speaking college students with normal or corrected-to-normal vision were recruited for the experiment, and were compensated \$10 for participation. Nine students were excluded due to excessive blinks leading to extreme data loss, as detailed below, resulting in a final dataset of 47 subjects.

Materials

Twenty-four sextets were constructed from the items in Experiment 1, modified so that there were three animate subject correlate (12a–c) and three inanimate object correlate (12d–f) conditions. In both cases, there was a condition with a definite plural distractor and a fully-specified nominal in the *wh*-restrictor, e.g., *tourists* or *wines* (12a,d), a definite plural distractor and the plural pronoun *ones* (12b,e), and a definite singular distractor and the plural pronoun *ones* (12c,f). Note that in (12a–c) the indefinite noun appears in the object, and so by hypothesis the local and preferred, position. The pipe symbol “|” indicates how materials were later divided into seven regions for analysis. All conditions were identical after the remnant region

- (12) a. |The tourists |sampled |some wines, |but I’ve forgotten |which wines, ...
- b. |The tourists |sampled |some wines, |but I’ve forgotten |which ones, ...
- c. |The tourist |sampled |some wines, |but I’ve forgotten |which ones,
- d. |Some tourists |sampled |the wines, |but I’ve forgotten |which tourists, ...
- e. |Some tourists |sampled |the wines, |but I’ve forgotten |which ones, ...
- f. |Some tourists |sampled |the wine, |but I’ve forgotten |which ones, ...
 |since they all |seem the same to me.

Lexical level characteristics of length and frequency were computed for nouns in subject and object position. Subject ($M = 6.83$; $SE = 0.27$) and object nouns ($M = 6.96$, $SE = 0.41$) did not differ on number of characters, $t_{(23)} = -0.24$, $p = 0.82$. Two measures of frequency were obtained from the English Lexicon Project (Balota et al., 2007). Subject ($M = 8.12$; $SE = 0.32$) and object nouns ($M = 8.86$, $SE = 0.42$) did not differ on log HAL frequency, $t_{(23)} = -1.56$, $p = 0.13$. Further, subject ($M = 2.26$; $SE = 0.15$) and object nouns ($M = 2.50$, $SE = 0.14$) match on frequency calculated from SUBTLEX, $t_{(23)} = -1.22$, $p = 0.23$. In addition, the length of the remnant region was always included as a predictor in models of that region.

A reviewer notes that the spillover regions may not have been informative with respect to the intended interpretation. However, items were intentionally designed so that properties of the inner restrictor of the remnant provided the only disambiguating information. Further, spillover regions were consistent within an item across all conditions, and thus are unlikely to explain any effects. However, as noted above, it may be fruitful to explore the influence of the structure in unelided counterparts, as in *I don't know which ones (they sampled/were sampled)*. Not all interesting contrasts could be presented in a single experiment, for fear of reducing statistical power or saturating readers with too many similar constructions. Another concern was that the example above contains the ambiguous pronoun *they* after the remnant. However, the item above is unique in that respect. As shown in the Appendix of Supplementary Material, no other item contains a pronoun of any sort.

Procedure

The experiment was presented using EyeTrack, the UMass Amherst presentation software (<http://www.psych.umass.edu/eyelab/>). Materials were presented in a sound isolated room on a 32-bit Dell Optiplex tower, running Windows 7, with peripheral programs and the Internet connection turned off. Text was presented as a single line in black 11pt monospaced font against a white background. The monitor was situated such that approximately three characters subtended 1° of visual angle. Eye movements were recorded on an SR Research Eyelink 1000 eye tracker, mounted on the table approximately 50 cm away from a 19" Mitsubishi Diamond Pro 900u flat-screen CRT monitor running at 170 Hz. Sampling rate was set to 1000 Hz. Drift correct was performed between each trial. Subjects were instructed to read naturally and for comprehension, and were encouraged to take breaks as often as needed.

All items were followed by comprehension questions probing the subject's interpretation (13). Questions were presented in CAPS to clearly differentiate comprehension questions from experimental materials.

(13) WHAT DID I FORGET?

- i. Subject response: WHICH TOURIST(S)
- ii. Object response: WHICH WINE(S)

Subjects selected the answer from among two possible choices on a Microsoft USB Sidewinder gamepad. Question responses were not considered in the analysis below. Experiments lasted approximately 40 min on average.

Results

Individual trials were removed if the participant blinked once or more during the first pass on the remnant region. No trials were removed if blinks occurred in another region or during re-reading of the remnant. Individual trials were also removed if excessive blinking led to significant track loss, or if track loss occurred for some other reason during the experiment (<4% of total trials).

Additionally, short (under 80 ms) and long (over 1200 ms) fixation times were removed from the data, as were trials with blinks on the remnant and track losses using the program EyeDoctor (<http://www.psych.umass.edu/eyelab>). Several standard eye tracking measures were used in the analysis (Rayner, 1998), computed with the DOS version of EyeDry analysis software: first *pass durations* (also known as gaze duration), the sum of all fixation durations within a region before leaving that region in any direction, *go past time*, the time spent after first entering a region to first moving past the region to the right, *percentage of regressions out of* and *percentage of regressions into* a region, *second pass time*, the time spent rereading a region once the region has been exited to the right including zero times indicating failure to re-read, and *total time*, the sum of all fixation times in a region during any point in reading (see, e.g., Staub and Rayner, 2007, for a concise review of these measures). Means and standard errors for these measures are presented in **Table 5** below.

Linear mixed effects regression models were used for all statistical analyses. Fixation and reading time measures (first pass, go past, second pass, and total times) were analyzed with linear regression models, whereas proportion data (regressions in and out of a region) were analyzed with logistic regression models. As models with maximal random effects error structures (as recommended by Barr et al., 2013) typically failed to converge, all models reported here were specified with by-subject and by-items random intercepts, but not with random slopes. Fixed effect predictor contrasts were assigned deviation coding that best cohered with the conceptual aims of the study. To assess the presence of a Locality bias, object noun correlates were coded as the baseline for the Correlate position predictor. To evaluate the effect of Interference, cue-poor probes (*which ones*) without plural interference were treated as the baseline for the Interference predictor, so that the model tests for similarity-based interference effects for nominal and pronominal cues over a no interference condition with a pronominal cue.

Instead of reporting the statistical results of each measure individually, the results are discussed in terms of the predictions of interest, noting when other effects were present. All significant effects for the measures of interest are reported.

Locality

Evidence for the Locality bias was observed in multiple measures of the eye movement record. The earliest evidence was found in first pass times immediately at the remnant (Region 5). indefinite subject correlates ($M = 274$ ms, $SE = 5$) elicited longer first pass times than indefinite object correlates ($M = 249$ ms, $SE = 5$), $t = 3.88$. No other effects were observed in first pass times. Additional evidence for the Locality bias appears in later eye movement measures, as well. Indefinite subject correlates elicited

TABLE 5 | Experiment 2: means and standard deviations for all eye tracking measures.

		Subject	Verb	Object	But I don't know	Which X	Spill over	Final region
FIRST PASS								
Interference nominal	Object	263 (15)	319 (13)	384 (16)	368 (14)	281 (9)	325 (13)	391 (18)
	Subject	292 (16)	320 (18)	327 (14)	375 (16)	300 (10)	335 (14)	405 (23)
Interference pronoun	Object	282 (16)	317 (14)	372 (15)	373 (15)	230 (8)	336 (14)	394 (18)
	Subject	298 (16)	292 (14)	318 (15)	352 (14)	257 (7)	338 (15)	393 (20)
No interference pronoun	Object	282 (21)	336 (17)	366 (14)	368 (15)	234 (8)	333 (15)	408 (20)
	Subject	295 (16)	321 (15)	298 (15)	370 (15)	264 (9)	339 (14)	418 (22)
GO PAST								
Interference nominal	Object	263 (15)	414 (18)	502 (23)	385 (15)	338 (15)	387 (29)	818 (66)
	Subject	326 (29)	380 (21)	438 (25)	389 (17)	363 (15)	350 (19)	834 (54)
Interference pronoun	Object	282 (16)	381 (18)	527 (29)	392 (17)	267 (14)	352 (17)	854 (66)
	Subject	298 (16)	402 (29)	472 (30)	374 (17)	303 (18)	352 (16)	1071 (94)
No interference pronoun	Object	282 (21)	464 (25)	559 (34)	405 (21)	279 (16)	377 (24)	908 (64)
	Subject	295 (16)	378 (20)	448 (29)	394 (18)	317 (20)	400 (25)	975 (70)
SECOND PASS								
Interference nominal	Object	90 (12)	101 (16)	39 (9)	99 (18)	55 (11)	129 (17)	NA
	Subject	84 (12)	104 (17)	29 (7)	79 (13)	42 (11)	140 (16)	NA
Interference pronoun	Object	79 (12)	116 (15)	64 (13)	56 (12)	28 (7)	122 (18)	NA
	Subject	149 (21)	173 (21)	84 (16)	88 (16)	61 (16)	151 (19)	NA
No interference pronoun	Object	128 (14)	154 (23)	73 (14)	68 (12)	37 (8)	123 (15)	NA
	Subject	110 (16)	137 (17)	87 (15)	87 (16)	45 (10)	148 (16)	NA
TOTAL TIMES								
Interference nominal	Object	249 (17)	439 (23)	446 (19)	458 (23)	344 (16)	431 (20)	473 (23)
	Subject	323 (27)	408 (25)	365 (17)	449 (20)	353 (16)	436 (22)	504 (26)
Interference pronoun	Object	270 (18)	429 (23)	472 (23)	412 (20)	253 (12)	438 (23)	496 (25)
	Subject	361 (24)	473 (29)	425 (25)	439 (21)	311 (17)	466 (22)	522 (29)
No interference pronoun	Object	277 (20)	521 (32)	496 (27)	441 (20)	263 (13)	424 (20)	505 (24)
	Subject	328 (23)	449 (23)	411 (24)	451 (23)	307 (15)	465 (20)	546 (25)
REGRESSIONS OUT								
Interference nominal	Object	NA	18 (3)	24 (3)	2 (1)	15 (3)	5 (2)	42 (4)
	Subject	NA	12 (3)	20 (3)	2 (1)	12 (3)	1 (1)	44 (4)
Interference pronoun	Object	NA	13 (3)	26 (3)	2 (1)	7 (2)	2 (1)	45 (4)
	Subject	NA	17 (3)	26 (3)	2 (1)	7 (2)	2 (1)	49 (4)
No interference pronoun	Object	NA	23 (3)	24 (3)	3 (1)	10 (2)	5 (2)	47 (4)
	Subject	NA	14 (3)	25 (3)	3 (1)	11 (2)	5 (2)	52 (4)
REGRESSIONS IN								
Interference nominal	Object	68 (6)	32 (4)	5 (2)	22 (3)	8 (2)	35 (4)	NA
	Subject	54 (6)	29 (4)	4 (2)	16 (3)	4 (1)	37 (4)	NA
Interference pronoun	Object	62 (6)	35 (4)	7 (2)	12 (3)	3 (1)	36 (4)	NA
	Subject	66 (5)	38 (4)	11 (2)	13 (3)	4 (2)	36 (4)	NA
No interference pronoun	Object	78 (4)	34 (4)	10 (2)	16 (3)	4 (2)	36 (4)	NA
	Subject	60 (5)	34 (4)	13 (3)	18 (3)	5 (2)	42 (4)	NA

longer go past times than object correlates on the remnant region ($M_{Subject} = 328$ ms, $SE = 10$; $M_{Object} = 295$ ms, $SE = 9$), $t = 2.72$, and on the sentence final region, ($M_{Subject} = 960$ ms, $SE = 43$; $M_{Object} = 860$ ms, $SE = 38$), $t = 2.12$. Fixed effects of models substantiating the above effects are provided in **Table 6**.

Further, violating Locality manifested in a persistent penalty for total times, as indefinite subject correlates were significantly longer in the sentential subject region ($M_{Subject} = 337$ ms, $SE = 40$; $M_{Object} = 265$, $SE = 11$), $t = 4.65$, the remnant ($M_{Subject} = 324$, $SE = 9$; $M_{Object} = 287$, $SE = 8$), $t = 3.40$, and the final region ($M_{Subject} = 524$ ms, $SE = 15$; $M_{Object} = 491$ ms, $SE = 14$), $t = 2.20$. Models computed for total times are provided in **Table 7**. In addition, there was a marginally significant penalty for indefinite subject correlates ($M = 147$ ms, $SE = 10$) compared to indefinite object correlates ($M = 125$, $SE = 10$) on the spill-over region in second pass re-reading times, $t = 1.95$. A summary of the main effects on the remnant is provided in **Figure 1**.

However, a few measures showed a cost for indefinite object correlates. There were more regressions into the sentential subject region for items with indefinite object correlates ($M = 70\%$, $SE = 3$) than indefinite subject correlates ($M = 60\%$, $SE = 3$), $t = 2.64$, $p < 0.01$. The increased rate of regressions into the subject region may correspond to increased global re-reading for indefinite object correlates, as opposed to regressing back into a specific region. In addition, indefinite object correlates elicited longer total times than indefinite subject correlates did in the region containing the object noun ($M_{Object} = 471$ ms, $SE = 13$; $M_{Subject} = 401$ ms, $SE = 13$), $t = -4.73$. This effect could be explained if total times corresponded to additional re-reading of the correlate. However, it is unclear whether such an explanation can be strongly maintained without supporting evidence from regressions in and second-pass reading measures, of which there is little evidence.

Nominal Advantage

For the second prediction, we expect that nominal restrictors (*which wines/tourists*) should receive a processing benefit over pronominal restrictors (*which ones*), due to greater specificity of the retrieval cue (also known as “cue diagnosticity”; see Martin and McElree, 2009, 2011; Van Dyke, 2011). Indeed, we find the

expected advantage for nominal probes in a variety of measures. In go past times, there was a 117 ms advantage for nominal restrictors over pronominal ones on the final region, $t = -2.42$; see the *Interference nominal* row in **Table 6**. There were fewer regressions into the object region when a nominal cue in the remnant followed ($M = 5\%$, $SE = 1$) as compared to a pronominal cue in the restrictor ($M = 11\%$, $SE = 2$), $t = -2.93$, $p < 0.01$.

Further, the advantage for nominal cue conditions was considerable in second pass re-reading times of every region of the matrix clause. Nominal restrictors elicited shorter second pass times in the subject region ($M_{Nominal} = 87$, $SE = 9$; $M_{Pronominal} = 119$, $SE = 11$), $t = -2.61$, the verb region ($M_{Nominal} = 103$, $SE = 12$; $M_{Pronominal} = 145$, $SE = 14$), $t = -3.21$, and the object region ($M_{Nominal} = 34$, $SE = 6$; $M_{Pronominal} = 80$, $SE = 10$), $t = -4.31$; see **Figure 2**. There

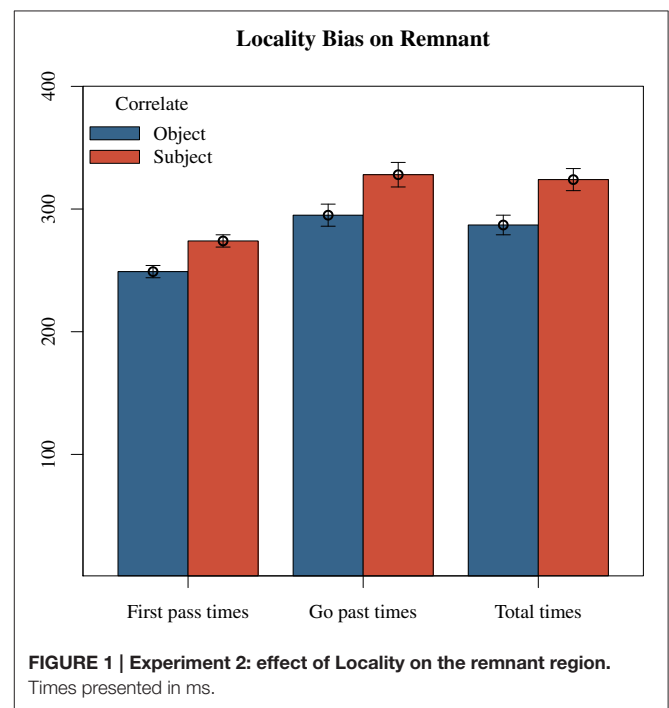


FIGURE 1 | Experiment 2: effect of Locality on the remnant region. Times presented in ms.

TABLE 6 | Experiment 2: linear mixed effects regression models for first fixation and go past times on Remnant and go past times on the sentence final region.

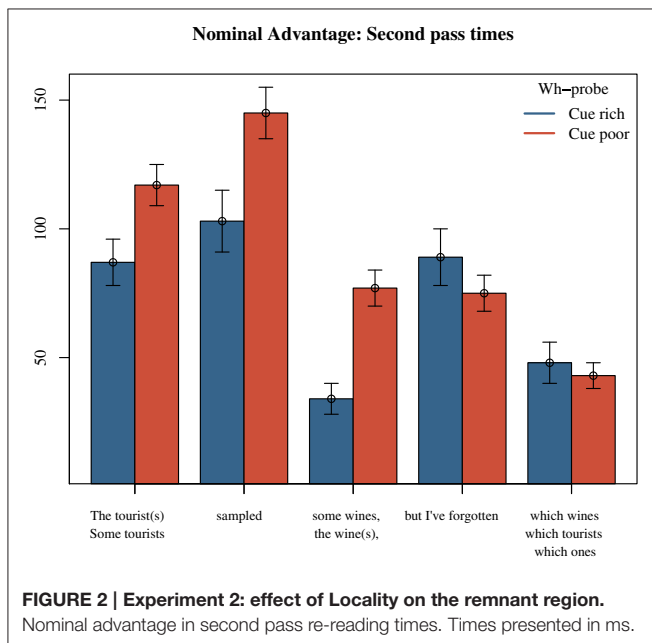
	First pass times on remnant			Go past times on remnant			Go past times on final region		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value	Estimate	Std. error	t-value
(Intercept)	231.863	18.748	12.368	266.994	35.923	7.432	899.026	80.558	11.16
Interference nominal	-7.037	5.745	-1.225	22.199	15.588	1.424	-83.699	34.589	-2.42
Interference pronoun	18.517	8.139	2.275	-18.053	10.95	-1.649	54.613	41.688	1.31
Locality	12.827	3.303	3.883	17.02	6.254	2.721	53.623	25.251	2.124
Length	5.502	3.453	1.593	8.485	6.644	1.277	NA	NA	NA
Interference nominal × Locality	1.357	4.674	0.29	-2.685	8.812	-0.305	-48.508	40.662	-1.193
Interference pronoun × Locality	-2.378	4.653	-0.511	1.398	8.865	0.158	58.372	45.79	1.275

Effects with t-values above |2| are shown in bold.

TABLE 7 | Experiment 2: linear mixed effects regression models for total times on the sentence initial, remnant, and sentence final regions.

	Total times on subject region			Total times on remnant			Total times final region		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value	Estimate	Std. error	t-value
(Intercept)	301.190	24.644	12.222	257.591	32.441	7.940	503.846	36.847	13.674
Interference nominal	-15.269	11.066	-1.380	23.878	14.036	1.701	-20.769	11.479	-1.809
Interference pronominal	13.851	11.069	1.251	-12.632	9.842	-1.284	5.633	11.483	0.491
Locality	36.310	7.816	4.646	19.117	5.620	3.402	17.834	8.107	2.200
Length	NA	NA	NA	9.412	5.996	1.570	NA	NA	NA
Interference nominal × Locality	-2.205	11.079	-0.199	-14.451	7.968	-1.814	0.348	11.494	0.030
Interference pronominal × Locality	12.028	11.071	1.086	10.805	7.958	1.358	-3.258	11.485	-0.284

Effects with *t*-values above |2| are shown in bold.



were also shorter total times for nominal probes in total times for the subject region ($M_{Nominal} = 423$ ms, $SE = 17$; $M_{Pronominal} = 485$ ms, $SE = 20$), $t = -2.60$.

One region witnessed effects other than were expected: there were more regressions into Region 4 for Nominal than Pronominal conditions ($M_{Nominal} = 19\%$, $SE = 2$; $M_{Pronominal} = 15\%$, $SE = 1$), $t = -2.28$, $p < 0.05$. There is no ready explanation of this small effect, as the content of the region was identical in all conditions.

Structure-Dependent Interference

The central prediction of an interaction between similarity-based interference and structural position of the correlate was supported by later eye movement measures. Importantly, a penalty was predicted only for pronominal probes in the remnant like *ones*, but not when the cue was fully specified, as in the case of nominal probes. As expected, there was a greater penalty for second pass times for definite plural object nouns (*the wines*) and a subject correlate (*some tourists*) when the *wh*-restrictor was a

pronominal. In the subject region, the penalty for conditions with pronominal restrictors was significantly greater for indefinite subject correlates ($d = 70$ ms) than for indefinite object correlates ($d = -18$ ms), $t = 3.54$. Similar effects obtained in the verb region, with a 57 ms penalty for subject correlates over object correlates ($d = -17$ ms), $t = 2.26$. Both of these effects are shown in the final row in **Table 8**.

On the remnant region, there was again a greater penalty for indefinite subject correlates ($d = 33$) than for indefinite object correlates ($d = 8$), $t = 2.11$; see **Table 9** and **Figure 3**. What's more, nominal restrictors showed a small 13 ms second pass time advantage for indefinite subject correlates and plural distractor objects, $t = -2.08$ in second pass times; **Table 9**. Finally, in the verb region, there was a greater total times penalty for indefinite subject correlates and plural distractors with pronominal retrieval cues ($d = 72$ ms) than the no interference baseline ($d = 44$ ms), $t = 2.62$.

Discussion

Results from the reading experiment support all three predictions of interest. There was early and sustained support for violating the Locality bias, and an advantage for Nominal cues that manifested in go past and second pass measures. These two effects interacted with respect to interference and cue-specificity: there were greater interference effects when the distractor occupied the preferred object position, such that the effect was enhanced when cues for retrieval at the remnant were partial. The results are thus compatible with previous findings of retroactive interference (e.g., Van Dyke and McElree, 2006; Van Dyke and Johns, 2012), but also adds support to the growing body of evidence that retrieval is modulated by the position of the antecedents (e.g., Van Dyke and McElree, 2011; Chow et al., 2014; Dillon et al., 2014; Kush and Phillips, 2014, among others). While the results are clearly consistent with the central predictions of a cue-based parsing system in which the location of targets in the sentence is not ignored during retrieval, the issue of how precisely to utilize such information within cue-based parsing models is far from settled. I return to this question the General Discussion.

As noted by a reviewer, evidence for the central predictions manifested at somewhat different time courses, although we should exercise caution when assigning linking assumptions to eye movement measures (Clifton et al., 2007). Whereas,

TABLE 8 | Experiment 2: nominal advantage in second pass re-reading times.

	Second pass on subject region			Second pass on verb region			Second pass on object region		
	Estimate	Std. error	t-value	Estimate	Std. error	t-value	Estimate	Std. error	t-value
(Intercept)	107.007	13.958	7.666	131.059	21.721	6.034	63.048	11.175	5.642
Interference nominal	-20.473	7.849	-2.608	-29.972	9.345	-3.207	-29.309	6.794	-4.314
Interference pronominal	7.929	7.849	1.01	15.963	9.348	1.708	11.965	6.794	1.761
Locality	8.137	5.543	1.468	7.064	6.6	1.07	4.291	4.797	0.894
Interference nominal × Locality	-12.287	7.856	-1.564	-7.381	9.357	-0.789	-10.67	6.8	-1.569
Interference pronominal × Locality	27.778	7.851	3.538	21.11	9.349	2.258	7.044	6.795	1.037

Effects with *t*-values above |2| are shown in bold.

evidence for Locality and Nominal Advantage appeared in various measures, support for Structure-Dependent Interference was only observed in the “late” measure of second-pass times, as subjects re-read portions of the sentence. This delayed effect is compatible with multiple interpretations, including a multiple stage model of anaphoric processing in which measures occurring later in the eye movement record could reflect processing at a secondary stage of discourse integration, perhaps along the lines of Garrod and Sanford’s (1990, 1994) bonding and resolution model. If this were the case, Structure-Dependent Interference might reflect difficulty interpreting the link between a poorly specified remnant and a correlate, rather than retrieval difficulty. Alternatively, that the effect appears relatively late in the eye movement record could be attributed to a lag resulting from poor quality matches. In this case, the integration difficulty would directly reflect increased interference from distractors in structurally preferred positions. The results do not arbitrate between these, or any number of, additional possibilities, which must instead be resolved through careful experimental design.

A reviewer proposed that several of the sentences in (12) are ambiguous, in that *which ones* may also co-refer with definite plural nouns. The above design depended on the assumption that definite nouns fail to provide an appropriate antecedent, as discussed in connection with examples (5–6). Yet, there may be a few systematic exceptions to the generalization that the remnant cannot correspond to a definite correlate. Discussion in the literature centers around contrasts like (13) below, which shows that a d-linked *which* remnant can take a definite noun as a correlate, but a simple *wh*-phrase like *what* cannot (Chung et al., 1995; Dayal and Schwarzschild, 2010).

- (13) a. John announced he had eaten the asparagus. We didn’t know which asparagus.
 b. * John announced he had eaten the asparagus. We didn’t know what.

These cases have been given various analyses. Chung et al. (1995) suggest that definite nouns are available as correlates whenever they are compatible with the pragmatic contribution of the remnant. They attribute the contrast in (13) to an intuitive conflict in familiarity between the definite *the asparagus* and the novelty imposed by *what* in (13b), a conflict that (13a) avoids.

In contrast, Dayal and Schwarzschild (2010) identify several cases in which presumed speaker knowledge, rather than

TABLE 9 | Experiment 2: Structure-Dependent Interference effects in second pass re-reading times on the remnant region.

	Second pass times on remnant region		
	Estimate	Std. Error	t-value
(Intercept)	31.428	23.059	1.363
Interference nominal	-1.576	10.212	-0.154
Interference pronominal	2.453	7.186	0.341
Locality	4.722	4.121	1.146
Length	2.58	4.347	0.594
Interference nominal × Locality	-12.126	5.843	-2.075
Interference pronominal × Locality	12.313	5.836	2.110

Effects with *t*-values above |2| are shown in bold.

familiarity, is the distinguishing factor. In their account, (14a) is infelicitous because the speaker has contradicted herself: the knowledge state that permits the speaker to assert that John talked to the detective places the speaker in a sufficient epistemic position to answer the question embedded under the sluice, i.e., *which detective did he talk to?* The corresponding assertion with an indefinite (14b) does not place the speaker in such a specific knowledge state as to warrant a self-contradiction (though see Barker, 2013; Barros, 2013, for recent commentary).

- (14) a. * John talked to the detective. I don’t know which detective (he talked to).
 b. John talked to a detective. I don’t know which detective (he talked to).

They also observe that a definite correlate is sometimes available when it does not carry a uniqueness presupposition, as illustrated by the examples in (15) where there is no requirement that there is a singular, identifiable train or particular hospital in the context (akin to so-called weak definites, e.g., Carlson and Sussman, 2005; Aguilar Guevara, 2014).

- (15) a. John is going to take the train, but he doesn’t know yet which train (he is going to take).
 b. They took him to the hospital. She wouldn’t tell us which hospital (they took him to).

As would be predicted, a definite that corresponds to a unique individual within a given context, as in *the Chief of Police* in (16), cannot serve a correlate to the remnant.

- (16) *Ed reported the matter to the Chief of Police but Joe couldn't figure out which chief of police (he reported the matter to).

Although it is unclear whether these examples are as acceptable when fully elided, e.g., *John is going to take the train, but he doesn't know which*, the experimental items were reviewed to determine whether the definite noun could sensibly be interpreted as a correlate to the remnant⁴. Two possible types of cases were observed. The first case involved collective nouns whose members could perform an action on behalf of others in the group (the ? mark reflects my judgment that these sentences are somewhat degraded):

- (17) a. ? The trustees donated 10 million to the university, but I don't know which ones.
 b. ? The professors wrote a letter to the dean, but it doesn't matter which ones.

For example, (17a) could be interpreted so that some of the trustees but not others were responsible for the donation. Here, the definite *the trustees* is interpreted as a collective entity, in which the donation was performed on behalf of the entire group. Although not many items in the experiment permit such a reading, one contender is *The nurses threatened to strike over some contracts, but I'm not sure which ones*.

The second case was one in which the remnant *which ones* does not refer directly to the definite noun phrase. Instead, co-reference appears to be coerced through a partitive interpretation taking the plural definite as the maximal set in order to derive a salient subset (a refset) from it (see Moxey and Sanford, 1993, for terminology). For example, such a reading might paraphrase (12d) as *Some tourists sampled the wines, but I don't know which (ones) of the wines they sampled*. The coercion process could posit a silent or elided partitive phrase, as proposed for bare determiner phrases like *Many (of them) sat down* (Gagnon, 2013). Again, very few plausible cases were found in the experimental items. Two possible cases include the example used as illustration throughout the paper (12), and *Some workers loaded the trucks, but I'm not certain which ones*. Such cases are perhaps strengthened by a distributive semantics of the verb, e.g., a sampling wine involves trying some, but not all, of it.

To assess empirically whether ambiguity could explain the effects observed above, I conducted a *post-hoc* by-items analysis of the results from Experiment 1. Averaging across conditions, no item was biased toward the definite noun completion. However, splitting the data by position of the indefinite revealed that eight items were either biased toward (definite subject: 13, 18, 21; definite object: 12, 24) or on par with (object definite: 2, 10, 18) the indefinite as a correlate. For most measures, there were no differences in the overall statistical effects when these items were removed⁵. However, removing potentially ambiguous items

did weaken the interaction between Locality and Interference in second pass times: although the penalty for non-local correlates was still significantly greater for remnants with pronominal restrictors in the sentence-initial region, the interaction did not persist in following region, even though the interaction was still apparent in other measures, including total times. Thus, even though a plausible definite distractor could have engendered a longer lasting interference penalty from the indefinite, it is unlikely to be the primary source of the effects reported here.

As a final note, ambiguity only becomes a genuine confound if it could otherwise explain the effects attributed to another variable. The only sentences that could have been ambiguous are those with cue-poor (*which ones*) remnants and plural definite distractors, i.e., (12b, d). Several other studies of pronominal ambiguity suggest that competing interpretations do not always result in processing penalties (e.g., attachment ambiguity explored in van Gompel et al., 2001, 2005), especially cases involving pronouns (e.g., Greene et al., 1992). Indeed, in the present case of sluicing, Frazier and Clifton (1998) report that ambiguity between subject and object position correlates did not slow readers down, provided that there was an indefinite correlate in the preferred, object position, like *someone* in (8a). Therefore, it is not yet clear how ambiguity would explain the effects I hope to attribute to interference.

Additionally, although the possibility of ambiguity might challenge whether we can truly interpret the effect of a plural definite as interference *per se*, it cannot fully account for the interaction between the plurality of the definite and its structural location. That is, irrespective of whether or not a definite noun is a possible correlate to *which* remnants, ambiguity does not explain why a plural definite in object position would elicit greater reading penalties than in subject position. Nevertheless, potential ambiguities could be more tightly controlled or even exploited (as in Harris, 2013, 2015) in future studies.

GENERAL DISCUSSION

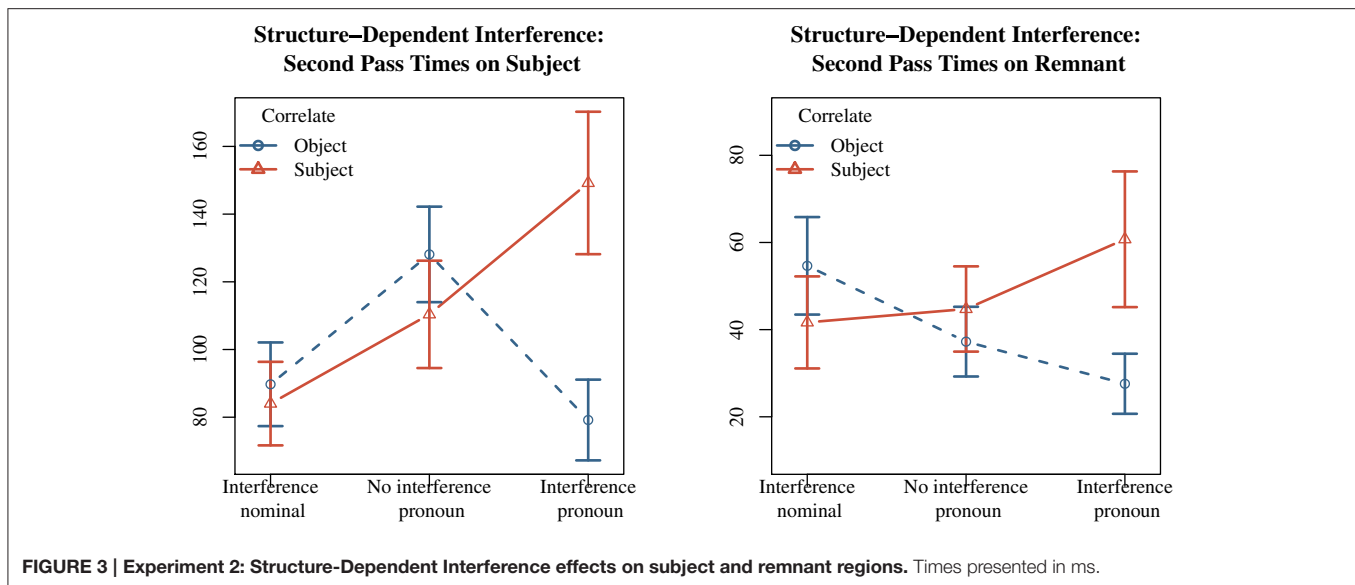
The central question explored in the experiment above was whether positional information modulates similarity-based interference effects in sluicing structures. There was clear evidence that it does. The central manipulation capitalized on the unique syntactic properties of sluices in two ways. First, the Locality bias was employed to impose a preference for structural position of the correlate to a remnant in the elided clause. Second, the lexical content of the inner restrictor of the remnant was manipulated to examine the role of cue-strength in retrieval.

As previously mentioned, this study is not the first to exploit sluiced sentences in an argument in favor of content-addressable retrieval systems. Martin and McElree (2011) utilized two main properties of an object position correlate in sluiced sentences like (14) in SAT and eye tracking. The correlate appeared on its own

potentially ambiguous items were removed, including an interaction supporting Structure-Dependent Interference in the verb region on second pass times, a slowdown for nominal restrictors in go past times at the remnant, and a previously marginal advantage for nominal restrictors reached significance for total times in the sentence-final region.

⁴Thanks to Colin Phillips for discussion of this issue and for providing some of the examples that appear in this section.

⁵Other changes were that the slowdown on the Interference pronoun condition in first pass times was marginal, and the penalty for violating the Locality Bias disappeared for go past and total times in the final region, as well as for second pass times in the spill over region. Several other effects were significant once



or within a conjunct in object position, and when the correlate was contained within a conjunct, which position of the conjunct it occupied (first or second conjunct position). The verbal syntax of only one member of the conjunct, *typed (something)*, provides a correlate, which varied according to whether the object was overt or not, to associate with the remnant *what*.

- (14) a. Michael (slept and) studied (something), but he didn't tell me what₁ <he typed t₁>.
 b. Michael studied (something) (and slept), but didn't tell me what₁ <he typed t₁>.

The design varied the distance between the correlate and remnant, along with the size of the elided material that was to be recovered. In keeping with the findings above, they found that readers spent longer re-reading distant antecedents (14b) than local ones (14a), and suggested, as I have, that interfering antecedents degrades the quality of a match with potential antecedents in memory. However, the materials of their study are quite different from the ones above in three respects. First, as only one conjunct provided a proper correlate (which sometimes had to be sprouted) to the remnant, the experiment lacks the conditions for fully investigating similarity-based interference from other noun phrase distractors. Second, the correlate was always in the object position, thereby satisfying the Locality bias, at least in a broad sense. Third, while the remnant varied according to *wh*-element type (*what*, *which*, and *where*), they did not manipulate the properties of the inner restrictor of the remnant to provide explicit cues to guide the dependency formation. The study above therefore contributes very different, yet congruent, evidence in favor of interference effects in retrieving correlates for sluiced sentences.

It is worth comparing Martin and McElree's study to the present one for another reason, as well. They found that the presence of a conjunct over a single noun in object position did not affect retrieval in either reading time or a SAT task,

and concluded that retrieval processes access the material for ellipsis directly on the basis of its content via a cost-free pointer mechanism, in line with studies on verb phrase ellipsis (Frazier and Clifton, 2001, 2005; Martin and McElree, 2008, 2009). However, it is possible that the mechanisms responsible for retrieving a correlate for the remnant are distinct from those responsible for recovering the elided IP material. Given the previously discussed dependency between resolving the remnant and determining the appropriate syntax of the ellipsis, it stands to reason that the former might be prioritized over the later, rather than attempting to solve two retrieval problems at once. Although it is theoretically possible that the ellipsis site lacks an explicit syntactic representation (Chung et al., 1995), there is good evidence for syntactic structure in sluicing ellipsis from both theoretical (e.g., Merchant, 2001; van Craenenbroeck, 2010) and experimental (e.g., Frazier and Clifton, 2001, 2005; Poirier et al., 2010) literature, in which case retrieving the ellipsis site is unlikely to reduce to simply pairing a correlate to the remnant of ellipsis.

Finally, one might be concerned that increased temporal distance, and thus decay, between the subject and the remnant might sufficiently explain the Locality bias, thereby eliminating structural information *per se* as a factor in the retrieval process. However, this explanation is unlikely given the results of Poirier et al.'s (2010) cross-modal priming study, in which printed targets related to the subject (*the handyman*) and dative object (*the programmer*) distractors were presented at two probe points in an auditory sentence: immediately after the offset of the remnant *₁ or 500 ms downstream *₂.

- (15) The handyman threw a book to the programmer but I don't know which book *₁ and no one *₂ else seems to know.

There was no difference between decision times for targets related to subject and dative object nouns until position *₂ (which showed a priming effect for the object), suggesting that subject

and the dative object nouns were equally active at the remnant of the ellipsis. Crucially, these effects do not contradict the results of the reading study, since probes related to the indefinite target *a book* could not be tested, given that they were repeated in the inner restrictor of the *wh*-phrase *which book*. If the restrictor were replaced with a cue-poor probe like *which ones*, we would expect an advantage for more local antecedents at, or soon after, the remnant.

Several models of sentence processing could in principle accommodate the findings reported above, models which diverge on how to account for the differences observed between subject and object position correlates. Naturally, the results of a single study cannot determine whether the effect of position reflects temporal precedence, linear distance, or, as I have suggested, structural information. Although various interpretations are possible, structural information has been shown independently to impact the earliest stages of retrieval in several related domains. It stands to reason that retrieval might privilege items located in preferred structural positions, even when the preference is not grammatically controlled. Of course, the nature of the mechanisms that underlie this putative advantage will remain unsettled until an effect of structural privilege is replicated in a design that dissociates structure from other factors, like linear order. Fortunately, sluicing ellipsis offers just the right sort of flexibility to tease such issues apart in the future.

Moreover, uncovering how the processor resolves the multiple dependencies required for interpreting sluiced sentences has only just begun. The configurational possibilities of sluicing ellipsis provide a rich testing ground for disentangling the retrieval processes that are charged with recovering linguistic antecedents and integrating them into a representation as it unfolds during real-time comprehension. While numerous questions remain,

one major challenge is the stage at which semantic and discourse information informs dependency formation in ellipsis, and whether information structural cues or strongly biased contexts can favor potential antecedents in the same way that structural information can. At the minimum, the present study provides additional support for converging evidence for cue-based parsing, and that the mechanisms underlying such retrieval are not wholly blind to the structural location of products in memory.

FUNDING

The author gratefully acknowledges financial support from Pomona College.

ACKNOWLEDGMENTS

This research benefited from conversations with Lyn Frazier, Carson Schütze, Sarah VanWagenen, and the audiences at the 28th Annual CUNY Conference on Human Sentence Processing held at the University of Southern California, as well as at a psycholinguistics lab meeting at the University of Maryland. Many thanks to Brian Dillon, Colin Phillips, and the two reviewers for their generous comments on previous drafts. Any mistakes or errors should be attributed to me alone. I thank Karin Denton for her assistance in running subjects for Experiment 2.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fpsyg.2015.01839>

REFERENCES

- Aguilar Guevara, A. (2014). *Weak Definites. Semantics, Lexicon and Pragmatics*. Ph.D. thesis, Utrecht University, The Netherlands.
- Alcocer, P., and Phillips, C. (2012). *Using Relational Syntactic Constraints in Content-addressable Memory Architectures for Sentence Parsing*. College Park, MD: Ms. University of Maryland.
- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cogn. Psychol.* 6, 451–474. doi: 10.1016/0010-0285(74)90021-8
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of the mind. *Psychol. Rev.* 111, 1036–1060. doi: 10.1037/0033-295X.111.4.1036
- Anderson, J. R., and Reder, L. M. (1999). The fan effect: new results and new theories. *J. Exp. Psychol. Gen.* 128, 186–197. doi: 10.1037/0096-3445.128.2.186
- Anderson, M. C., and Neely, J. H. (1996). “Interference and inhibition in memory retrieval,” in *Handbook of Perception and Memory, Vol. 10, Memory*, eds E. L. Bjork and R. A. Bjork (San Diego, CA: Academic Press), 237–313.
- Arnold, J. E. (1998). *Reference form and Discourse Patterns*. Ph.D. thesis, Stanford University.
- Autry, K. S., and Levine, W. H. (2014). A fan effect in anaphor processing: effects of multiple distractors. *Front. Psychol.* 5:818. doi: 10.3389/fpsyg.2014.00818
- Badecker, W., and Straub, K. (2002). The processing role of structural constraints on interpretation of pronouns and anaphors. *J. Exp. Psychol. Learn. Mem. Cogn.* 28, 748–769. doi: 10.1037/0278-7393.28.4.748
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., et al. (2007). The English lexicon project. *Behav. Res. Methods* 39, 445–459. doi: 10.3758/BF03193014
- Barker, C. (2013). Scopability and sluicing. *Linguist. Philos.* 36, 187–223. doi: 10.1007/s10988-013-9137-1
- Barr, D., Levy, R., Scheepers, C., and Tily, H. J. (2013). Random effects structure in mixed-effects models: keep it maximal. *J. Mem. Lang.* 68, 255–278. doi: 10.1016/j.jml.2012.11.001
- Barros, M. V. (2013). “Harmonic sluicing: Which remnant/correlate pairs work and why,” in *Proceedings of the 23rd Semantics and Linguistic Theory*, ed T. Snider (Santa Cruz, CA: University of California), 295–315.
- Bates, D., and Maechler, M. (2009). *lme4: Linear Mixed-Effects Models using Eigen and R Package Version 0.999375–999331*. Available online at: <https://github.com/lme4/lme4/>
- Caplan, D., and Waters, G. (2013). Memory mechanisms supporting syntactic comprehension. *Psychon. Bull. Rev.* 20, 243–268. doi: 10.3758/s13423-012-0369-9
- Carlson, G., and Sussman, R. (2005). “Seemingly indefinite definites,” in *Linguistic Evidence: Empirical, Theoretical, and Computational Perspectives*, eds S. Kepser and M. Reis (Tübingen: De Gruyter Mouton), 71–85.
- Carlson, K. (2002). *Parallelism and Prosody in the Processing of Ellipsis Sentences*. Outstanding Dissertations in Linguistics Series. Routledge, New York.
- Carlson, K. (2013). The role of only in contrasts in and out of context. *Discourse Process.* 50, 249–275. doi: 10.1080/0163853X.2013.778167
- Carlson, K., Dickey, M. W., Frazier, L., and Clifton, C. Jr. (2009). Information structure expectations in sentence comprehension. *Q. J. Exp. Psychol.* 62, 114–139. doi: 10.1080/17470210701880171
- Carlson, K., Dickey, M. W., and Kennedy, C. (2005). Structural economy in the processing and representation of gapping sentences. *Syntax* 8, 208–228. doi: 10.1111/j.1467-9612.2005.00079.x

- Chafe, W. K. (1976). "Givenness, contrastiveness, definiteness, subjects, topics, and point of view," in *Subject and Topic*, ed C. N. Li (New York, NY: Academic Press), 25–55.
- Chen, Z., Jäger, L., and Vasishth, S. (2012). "How structure-sensitive is the parser? Evidence from Mandarin Chinese," in *Empirical Approaches to Linguistics Theory. Studies in Meaning and Structure*, eds B. Stolerfoht and S. Featherston (Berlin: De Gruyter Mouton), 43–62.
- Chow, W. Y., Lewis, S., and Phillips, C. (2014). Immediate sensitivity to structural constraints in pronoun resolution. *Front. Psychol.* 5:630. doi: 10.3389/fpsyg.2014.00630
- Chung, S., Ladusaw, W. A., and McCloskey, J. (1995). Sluicing and logical form. *Nat. Lang. Semantics* 3, 239–282. doi: 10.1007/BF01248819
- Cinque, G. (1993). A null theory of phrase and compound stress. *Linguist. Inq.* 24, 239–297.
- Clark, S. E., and Gronlund, S. D. (1996). Global matching models of recognition memory: how the models match the data? *Psychon. Bull. Rev.* 3, 37–60. doi: 10.3758/BF03210740
- Clifton, C., Frazier, L., and Deevy, P. (1999). Feature manipulation in sentence comprehension. *Revisita Linguistica* 11, 11–39.
- Clifton, C. Jr., Staub, A., and Rayner, K. (2007). "Eye movements in reading words and sentences," in *Eye Movement Research: A Window on Mind and Brain*, eds R. Van Gompel, M. Fisher, W. Murray, and R. L. Hill (Oxford: Elsevier Ltd.), 341–372.
- Crowder, R. G. (1976). *Principles of Learning and Memory*. Hillsdale, NJ: Erlbaum.
- Cunnings, I., Patterson, C., and Felser, C. (2015). Structural constraints on pronoun binding and coreference: evidence from eye movements during reading. *Front. Psychol.* 6:840. doi: 10.3389/fpsyg.2015.00840
- Dayal, V., and Schwarzschild, R. (2010). "Definite inner antecedents and *wh*-correlates in sluices," in *Rutgers Working Papers in Linguistics*, Vol. 3, eds P. Staroverov, D. Altshuler, A. Braver, C.A. Fasola, and S. Murray (New Brunswick, NJ), 92–114.
- Dickey, M. W., and Bunger, A. C. (2011). Comprehension of elided structure: evidence from sluicing. *Lang. Cogn. Process.* 26, 63–78. doi: 10.1080/01690961003691074
- Dillon, B., Chow, W. Y., Wagers, M., Guo, T., Liu, F., and Phillips, C. (2014). The structure-sensitivity of memory access: evidence from Mandarin Chinese. *Front. Psychol.* 5:1025. doi: 10.3389/fpsyg.2014.01025
- Dillon, B., Mishler, A., Sloggett, S., and Phillips, C. (2013). Contrasting intrusion profiles for agreement and anaphora: experimental and modeling evidence. *J. Mem. Lang.* 69, 85–103. doi: 10.1016/j.jml.2013.04.003
- Dosher, B. A., and McElree, B. (1992). "Memory search: Retrieval processes in short-term and long-term recognition," in *Encyclopedia of Learning and Memory*, ed L. R. Squire (New York, NY: MacMillan), 398–406.
- Fedorenko, E., Gibson, E., and Rohde, D. (2006). The nature of working memory capacity in sentence comprehension: evidence against domain-specific working memory resources. *J. Mem. Lang.* 54, 541–553. doi: 10.1016/j.jml.2005.12.006
- Foraker, S., and McElree, B. (2007). The role of prominence in pronoun resolution: Active versus passive representations. *J. Mem. Lang.* 56, 357–383. doi: 10.1016/j.jml.2006.07.004
- Foraker, S., and McElree, B. (2011). Comprehension of linguistic dependencies: speed–accuracy tradeoff evidence for direct–access retrieval from memory. *Lang. Linguist. Compass* 5, 764–783. doi: 10.1111/j.1749-818X.2011.00313.x
- Frazier, L. (1978). *On Comprehending Sentences: Syntactic Parsing Strategies*. Ph.D. thesis, University of Connecticut, Storrs, CT.
- Frazier, L. (2008). "Processing ellipsis: a processing solution to the undergeneration problem," in *The Proceedings of the 26th West Coast Conference on Formal Linguistics*, eds C. B. Chang and H. J. Haynie (Somerville, MA: Cascadia Proceedings Project), 21–32.
- Frazier, L., and Clifton, C. Jr. (1998). Comprehension of sluiced sentences. *Lang. Cogn. Process.* 13, 499–520. doi: 10.1080/016909698386474
- Frazier, L., and Clifton, C. Jr. (2001). Parsing coordinates and ellipsis: Copy α . *Syntax* 4, 1–22. doi: 10.1111/1467-9612.00034
- Frazier, L., and Clifton, C. Jr. (2005). The syntax-discourse divide: processing ellipsis. *Syntax* 8, 121–174. doi: 10.1111/j.1467-9612.2005.00077.x
- Frazier, L., and Clifton, C. Jr. (2011). "D-linking and memory retrieval: the annoying case of sluicing," in *University of Massachusetts Occasional Papers in Linguistics: Processing Linguistic Structure*. Vol. 38, eds J. A. Harris and M. Grant (Amherst, MA: GLSA Publications), 37–52.
- Frazier, L., Taft, L., Roeper, T., Clifton, C. Jr., and Ehrlich, K. (1984). Parallel structure: a source of facilitation in sentence comprehension. *Mem. Cogn.* 12, 421–430. doi: 10.3758/BF03198303
- Gagnon, M. (2013). "Part and parcel of eliding partitives," in *Proceedings of the 23rd Semantics and Linguistic Theory Conference*, ed T. Snider (Santa Cruz, CA: University of Santa Cruz), 316–335.
- Garrod, S., and Sanford, A. J. (1994). "Resolving sentences in a discourse context: how discourse representation affects language understanding," in *Handbook of Psycholinguistics*, ed M. Gernsbacher (New York, NY: Academic Press), 675–698.
- Garrod, S. C., and Sanford, A. J. (1990). "Referential processing in reading: focusing on roles and individuals," in *Comprehension Processes in Reading*, eds D. A. Balota, G. B. Flores d'Arcias, and K. Rayner (Hillsdale, NJ: Erlbaum), 465–486.
- Gordon, P. C., Hendrick, R., and Johnson, M. (2001). Memory interference during language processing. *J. Exp. Psychol.* 27, 1411–1423. doi: 10.1037/0278-7393.27.6.1411
- Gordon, P. C., Hendrick, R., and Johnson, M. (2004). Effects of noun phrase type on sentence complexity. *J. Mem. Lang.* 51, 97–114. doi: 10.1016/j.jml.2004.02.003
- Gordon, P. C., Hendrick, R., Johnson, M., and Lee, Y. (2006). Similarity-based interference during language comprehension: evidence from eye tracking during reading. *J. Exp. Psychol.* 32, 1304–1321. doi: 10.1037/0278-7393.32.6.1304
- Gordon, P. C., Hendrick, R., and Levine, W. H. (2002). Memory-load interference in syntactic processing. *Psychol. Sci.* 13, 425–430. doi: 10.1111/1467-9280.00475
- Greene, S. B., McKoon, G., and Ratcliff, R. (1992). Pronoun resolution and discourse models. *J. Exp. Psychol.* 18, 266–283. doi: 10.1037/0278-7393.18.2.266
- Grosz, B., Joshi, A., and Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Comput. Ling.* 21, 203–225.
- Gundel, J. (1985). 'Shared knowledge' and topicality. *J. Pragmat.* 9, 83–107. doi: 10.1016/0378-2166(85)90049-9
- Gundel, J., Hedberg, N., and Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language* 69, 274–307. doi: 10.2307/416535
- Harris, J. A. (2013). "Discourse accessibility and structural bias: Processing d-linked phrases in sluices," in *Poster at the 26th Annual CUNY Human Sentence Processing Conference* (Columbia, SC: University of South Carolina).
- Harris, J. A. (2015). *Alternatives on Demand and the Locality Bias: Resolving discourse-linked wh-phrases in sluiced structures*. Los Angeles, CA: University of California.
- Harris, J. A., and Carlson, K. (2015). Keep it local (and final): remnant preferences for let alone ellipsis. *Q. J. Exp. Psychol.* doi: 10.1080/17470218.2015.1062526. [Epub ahead of print].
- Jäger, L. A., Benz, L., Roeser, J., Dillon, B. W., and Vasishth, S. (2015). Teasing apart retrieval and encoding interference in the processing of anaphors. *Front. Psychol.* 6:506. doi: 10.3389/fpsyg.2015.00506
- Keppel, G., and Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *J. Verbal Learn. Verbal Behav.* 1, 153–161. doi: 10.1016/S0022-5371(62)80023-1
- Kush, D. (2013). *Respecting relations: Memory Access and Antecedent Retrieval in Incremental Sentence Processing*. Ph.D. thesis. University of Maryland, College Park, MD.
- Kush, D., and Phillips, C. (2014). Local anaphor licensing in an SOV language: implications for retrieval strategies. *Front. Psychol.* 5:1252. doi: 10.3389/fpsyg.2014.01252
- Lewis, R. L. (1996). Interference in short-term memory: the magical number two (or three) in sentence processing. *J. Psycholinguist. Res.* 25, 93–115. doi: 10.1007/BF01708421
- Lewis, R. L., and Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cogn. Sci.* 29, 375–419. doi: 10.1207/s15516709cog0000_25
- Lewis, R. L., Vasishth, S., and Van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends Cogn. Sci.* 10, 447–454. doi: 10.1016/j.tics.2006.08.007
- Link, G. (1983). "The logical analysis of plurals and mass terms: a lattice-theoretical approach," in *Meaning, Use, and the Interpretation of Language*, eds R. Bäuerle, C. Schwarze, and A. von Stechow (Berlin: Walter de Gruyter and Co), 303–323.

- Martin, A. E. (2010). *Memory Operations and Structures in Sentence Comprehension: Evidence from Ellipsis*. Ph.D. thesis, New York University, New York, NY.
- Martin, A. E., and McElree, B. (2008). A content-addressable pointer mechanism underlies comprehension of verb-phrase ellipsis. *J. Mem. Lang.* 58, 879–906. doi: 10.1016/j.jml.2007.06.010
- Martin, A. E., and McElree, B. (2009). Memory operations that support language comprehension: evidence from verb-phrase ellipsis. *J. Exp. Psychol.* 35, 1231–1239. doi: 10.1037/a0016271
- Martin, A. E., and McElree, B. (2011). Direct-access retrieval during sentence comprehension: evidence from sluicing. *J. Mem. Lang.* 64, 327–343. doi: 10.1016/j.jml.2010.12.006
- Martin, A. E., Nieuwland, M. S., and Carreiras, M. (2012). Event-related brain potentials index cue-based retrieval interference during sentence comprehension. *Neuroimage* 59, 1859–1869. doi: 10.1016/j.neuroimage.2011.08.057
- Martin, A. E., Nieuwland, M. S., and Carreiras, M. (2014). Agreement attraction during comprehension of grammatical sentences: ERP evidence from ellipsis. *Brain Lang.* 135, 42–51. doi: 10.1016/j.bandl.2014.05.001
- McElree, B. (2000). Sentence comprehension is mediated by content-addressable memory. *J. Psycholinguist. Res.* 29, 111–123. doi: 10.1023/A:1005184709695
- McElree, B. (2001). Working memory and focal attention. *J. Exp. Psychol.* 27, 817–835. doi: 10.1037/0278-7393.27.3.817
- McElree, B. (2006). “Accessing recent events,” in *The Psychology of Learning and Motivation*, Vol. 46, ed B. H. Ross (San Diego, CA: Academic Press), 155–200. doi: 10.1016/S0079-7421(06)46005-9
- McElree, B., and Doshier, B. A. (1989). Serial position and set size in short-term memory: time course of recognition. *J. Exp. Psychol.* 118, 346–373. doi: 10.1037/0096-3445.118.4.346
- McElree, B., and Doshier, B. A. (1993). Serial retrieval processes in the recovery of order information. *J. Exp. Psychol.* 122, 291–315. doi: 10.1037/0096-3445.122.3.291
- McElree, B., Foraker, S., and Dyer, L. (2003). Memory structures that subserve sentence comprehension. *J. Mem. Lang.* 48, 67–91. doi: 10.1016/S0749-596X(02)00515-6
- Merchant, J. (2001). *The Syntax of Silence: Sluicing, Islands, and the Theory of Ellipsis*. Oxford: Oxford University Press.
- Moxey, L. M., and Sanford, A. J. (1993). *Communicating Quantities: A Psychological Perspective*. Hove, UK: Lawrence Erlbaum Associates, Inc.
- Nairne, J. S. (2002). Remembering over the short-term: the case against the standard model. *Annu. Rev. Psychol.* 53, 53–81. doi: 10.1146/annurev.psych.53.100901.135131
- Nicol, J., and Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *J. Psycholinguist. Res.* 18, 5–19. doi: 10.1007/BF01069043
- Öztek, I., and McElree, B. (2007). Retrieval dynamics of proactive interference: PI slows retrieval by eliminating fast assessments of familiarity. *J. Mem. Lang.* 57, 126–149. doi: 10.1016/j.jml.2006.08.011
- Phillips, C., Wagers, M. W., and Lau, E. F. (2011). “Grammatical illusions and selective fallibility in real-time language comprehension,” in *Experiments at the Interfaces. Syntax and Semantics*, Vol. 37, ed J. Runner (Bingley: Emerald Group Publishing Limited), 147–180.
- Poirier, J., Wolfinger, K., Spellman, L., and Shapiro, L. P. (2010). The real-time processing of sluiced sentences. *J. Psycholinguist. Res.* 39, 411–427. doi: 10.1007/s10936-010-9148-9
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychol. Bull.* 124, 372–422. doi: 10.1037/0033-2909.124.3.372
- R Development Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.
- Romero, M. (1998). *Focus and Reconstruction Effects in WH-Phrases*. Ph.D. thesis, University of Massachusetts Amherst.
- Ross, J. R. (1967). *Constraints on Variables in Syntax*. Ph.D. thesis, MIT, Cambridge, MA.
- Ross, J. R. (1969). “Guess who,” in *Proceedings of the 5th Annual Chicago Linguistic Society*, eds R. I. Binnick, A. Davison, G. M. Green, and J. L. Morgan (Chicago, IL), 252–286.
- Selkirk, E. O. (1984). *Phonology and Syntax: The Relation between Sound and Structure*. Cambridge: MIT Press.
- Sprouse, J., Wagers, M., and Phillips, C. (2012). A test of the relation between working memory capacity and syntactic island effects. *Language* 88, 82–123. doi: 10.1353/lan.2012.0004
- Staub, A., and Rayner, K. (2007). “Eye movements and on-line comprehension processes,” in *The Oxford Handbook of Psycholinguistics*, ed G. Gaskell (Oxford: Oxford University Press), 327–342.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science* 153, 652–654. doi: 10.1126/science.153.3736.652
- Sternberg, S. (1975). Memory-scanning: new findings and current controversies. *Q. J. Exp. Psychol.* 27, 1–32. doi: 10.1080/14640747508400459
- Sturt, P. (2003). The time-course of the application of binding constraints in reference resolution. *J. Mem. Lang.* 48, 542–562. doi: 10.1016/S0749-596X(02)00536-3
- van Craenenbroeck, J. (2010). *The Syntax of Ellipsis: Evidence from Dutch Dialects*. New York, NY: Oxford University Press.
- Van Dyke, J. A. (2007). Interference effects from grammatically unavailable constituents during sentence processing. *J. Exp. Psychol.* 33, 407–430. doi: 10.1037/0278-7393.33.2.407
- Van Dyke, J. A. (2011). Cue-dependent interference in comprehension. *J. Mem. Lang.* 65, 247–263. doi: 10.1016/j.jml.2011.05.002
- Van Dyke, J. A., and Johns, C. L. (2012). Memory interference as a determinant of language comprehension. *Lang. Linguist. Compass* 6, 193–211. doi: 10.1002/lnc3.330
- Van Dyke, J. A., and Lewis, R. L. (2003). Distinguishing effects of structure and decay on attachment and repair: a retrieval interference theory of recovery from misanalyzed ambiguities. *J. Mem. Lang.* 49, 285–413. doi: 10.1016/S0749-596X(03)00081-0
- Van Dyke, J. A., and McElree, B. (2006). Retrieval interference in sentence comprehension. *J. Mem. Lang.* 55, 157–166. doi: 10.1016/j.jml.2006.03.007
- Van Dyke, J. A., and McElree, B. (2011). Cue-dependent interference in comprehension. *J. Mem. Lang.* 65, 247–263. doi: 10.1016/j.jml.2011.05.002
- van Gompel, R. P. G., Pickering, M. J., Pearson, J., and Liversedge, S. P. (2005). Evidence against competition during syntactic ambiguity resolution. *J. Mem. Lang.* 52, 284–307. doi: 10.1016/j.jml.2004.11.003
- van Gompel, R. P. G., Pickering, M. J., and Traxler, M. J. (2001). Reanalysis in sentence processing: evidence against current constraint-based and two-stage models. *J. Mem. Lang.* 45, 225–258. doi: 10.1006/jmla.2001.2773
- Vasishth, S., Brüssow, S., Lewis, R. L., and Drenhaus, H. (2008). Processing polarity: how the ungrammatical intrudes on the grammatical. *Cogn. Sci.* 32, 685–712. doi: 10.1080/03640210802066865
- Wagers, M. W., Lau, E. F., and Phillips, C. (2009). Agreement attraction in comprehension: representations and processes. *J. Mem. Lang.* 61, 206–237. doi: 10.1016/j.jml.2009.04.002
- Waugh, N. C., and Norman, D. A. (1965). Primary memory. *Psychol. Rev.* 72, 89–104. doi: 10.1037/h0021797
- Xiang, M., Dillon, B., and Phillips, C. (2009). Illusory licensing effects across dependency types: ERP evidence. *Brain Lang.* 108, 40–55. doi: 10.1016/j.bandl.2008.10.002
- Xiang, M., Grove, J., and Giannakidou, A. (2013). Dependency dependent interference: NPI interference, agreement attraction, and global pragmatic inferences. *Front. Psychol.* 4:708. doi: 10.3389/fpsyg.2013.00708

Conflict of Interest Statement: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Harris. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.