5

GRAMMATICAL ILLUSIONS AND SELECTIVE FALLIBILITY IN REAL-TIME LANGUAGE COMPREHENSION

Colin Phillips, Matthew W. Wagers and Ellen F. Lau

ABSTRACT

Grammatical constraints impose diverse requirements on the relations between words and phrases in a sentence. Research on the online implementation of grammatical constraints reveals a strikingly uneven profile. The parser shows impressive accuracy in the application of some rather complex constraints, but makes many errors in the implementation of some relatively simple constraints. Just as the study of optical illusions has played an important role in the study of visual perception, the parser's highly selective vulnerability to interference and "grammatical illusions" provides a valuable tool for understanding how speakers encode and navigate complex linguistic representations in real time.

1. INTRODUCTION

Grammatical constraints impose many structural and featural requirements on the relations between words and phrases in a sentence, which include constraints on anaphora, agreement, case, and unbounded dependencies, to name but a few. In investigating how these various requirements are implemented during online language processing, we have a powerful tool for understanding how linguistic representations are encoded and navigated in real time.

Research on the online implementation of grammatical constraints goes back at least 25 years, but a number of recent developments have made this line of inquiry particularly interesting. First, in linguistics there is growing interest in the question of how grammatical computations might be understood as real-time mental processes, with proposals emerging from all corners of the linguistic landscape (e.g., Cann, Kempson, & Marten, 2005; O'Grady, 2005; Phillips, 1996, 2003; Phillips & Lewis, in press; Steedman, 2000).

Second, psycholinguistics has seen a resurgence of interest in the question of how structured information is encoded and accessed in memory. In contrast to earlier work that tended to treat working memory as a passive buffer whose most interesting property was its capacity, recent work has paid closer attention to how structural relations are encoded, and to how relevant and irrelevant information is distinguished in memory. This maturing of research on memory for linguistic structure is closely tied to developments in the literature on the cognitive (neuro-)science of memory (for reviews, see Cowan, 2000; Jonides et al., 2008; McElree, 2006; Ricker, AuBuchon, & Cowan, 2010).

Third, findings on the online status of different constraints are yielding a rich profile of grammatical (in)sensitivity. Past research on the psycholinguistics of filler-gap dependencies, anaphora, agreement, thematic binding, and other phenomena has proceeded largely independently in a series of subliteratures. When the findings on different linguistic phenomena are brought together, a number of striking contrasts emerge. Human parsers are quite good at implementing some rather complex grammatical constraints, such as island constraints on filler-gap dependencies, and strikingly bad at respecting some very simple constraints, such as subject-verb agreement. In some cases comprehenders are susceptible to *grammatical illusions*, but in other cases they appear to be immune to illusions.

Here we present a preliminary profile of *selective fallibility* to grammatical illusions in language comprehension. Just as the study of optical illusions has proven to be a valuable tool in understanding visual information processing, we expect that the study of (in)sensitivity to linguistic illusions will be fruitful in research on language.

We begin in Section 2 by summarizing different mechanisms that comprehenders might use to access linguistic material in memory. These mechanisms present a trade-off between speed and structure-sensitivity of processing. In Sections 3, we survey grammatical phenomena where comprehenders appear to

show impressive online sensitivity (island constraints on unbounded dependencies, backwards anaphora and Principle C, and constraints on reflexives), and in Section 4 we describe cases where comprehenders are susceptible to grammatical illusion (subject-verb agreement, case licensing, negative polarity item licensing, and comparatives). The antilocality constraint on pronouns (Principle B) is included in Section 4, although its status remains uncertain. In Section 5, we discuss some possible ways of explaining the selective fallibility profile.

2.   TWO WAYS OF SEARCHING STRUCTURES

Computing grammatical relations requires coordinating information that is not always temporally proximal. The processor must therefore be able to preserve information over time, either by maintaining it concurrently with incoming material or by encoding and storing it for later retrieval. Broadly speaking, retrieval could proceed in either of two ways, one that is strictly structure-sensitive and one that is not. The mechanism of retrieval is thus likely to play an important role in determining how sensitive comprehenders are to grammatical constraints.

The first type of retrieval occurs by means of a search. In a search, pairwise comparisons are made between the information desired and candidate encodings in memory. Crucially these comparisons are made in a controlled order and searches can thereby be grammatically constrained. For example, a syntactic tree can be searched node by node, with the dominance relations between nodes determining the sequence of nodes to be examined (Knuth, 1965). In this manner, a simple algorithm can be formulated for finding a c-commanding antecedent by always preceding up the dominance path, and never examining more than one node down. Similarly, in this approach it is straightforward for an agreeing verb in English to be matched against the features of the appropriate subject, by proceeding up the dominance path to the top of the current clause and then consulting the features of the left-hand daughter of the clausal node. Of course, this kind of structure-guided search is only effective if the speaker has an appropriate structural representation of the sentence.

The second type of retrieval occurs by means of content-addressable access. Content-addressable access occurs by probing the entire memory simultaneously with a set of cues to the desired information (Clark & Gronlund, 1996). Encodings in memory are differentially activated on the basis of their similarity to the cues. Global activation reflects whether or not the desired information is in the memory, and the degree of match for any individual encoding determines the likelihood with which that encoding is sampled and restored into the processing stream. Content-addressable retrieval is driven by the inherent properties of single encodings and it is generally less structure-sensitive than serial structure-guided search, for two reasons. First, some

notions that correlate with structural properties, such as "subject" or "dative" can be used as retrieval cues, but since these merely contribute to the degree of match between the retrieval cues and items in memory, they have no special status, and imperfectly matching items in memory may be retrieved. Second, genuinely relational notions such as c-command cannot easily be recast in terms of the inherent properties of individual memory chunks, and so it is difficult to use relational notions such as c-command as retrieval cues.

The major advantage of retrieval by search is that it can be structurally constrained. It can selectively target information in specific structural locations, and should be able to avoid interference from similar information in structurally irrelevant locations. Consequently, grammatically illicit constituents can be avoided with high accuracy. The downside of searches is that they can be slow and, in particular, it may take some time to discover that the desired information is not present. The major advantage of content-addressable retrieval is that the time to perform the retrieval does not depend upon the size of the representation. It can be determined in constant time whether or not the desired information is present. The downside is that grammatically illicit constituents can be activated, and even partially activated constituents may mislead the comprehender into a false impression of grammaticality. For more extensive discussion of these access mechanisms and how they might apply to language processing, see Lewis and Vasishth (2005), Wagers (2008), and Dillon (2011).

There is evidence that both mechanisms are active in language comprehension (McElree, Foraker, & Dyer, 2003), which is perhaps unsurprising given their complementary strengths and weaknesses. Much attention has been recently devoted to content-addressable retrieval, and the associated phenomenon of similarity-based interference (Gordon, Hendrick, & Johnson, 2001; Lewis, Vasishth, & Van Dyke, 2006; Van Dyke & McElree, 2006). We discuss this phenomenon below in accounting for some kinds of grammatical fallibility. However, it remains an open question how the language comprehension system decides when to use either mechanism or how it arbitrates between conflicting outcomes.

3.   IMMEDIATE GRAMMATICAL SENSITIVITY

We begin by reviewing a series of grammatical constraints that appear to be faithfully respected in online structure building. These successes are a testament to the grammatical sophistication of the human parser, but they become all the more interesting when set against the evidence for fleeting insensitivity to certain other constraints. Also, to the extent that speakers are able to rapidly online implement constraints that make reference to detailed structural notions, this provides good evidence that they are constructing suitably rich representations online. This is important when we try to understand the reasons for failure to

respect other types of structural constraints online. If we have independent evidence for the accuracy of the structural representations, then it becomes more likely that the errors reflect limitations of the search/access mechanisms than limitations of the representations themselves.

3.1.   Island Constraints on Filler-Gap Dependencies

Languages commonly allow words or phrases to be displaced, leaving the canonical position empty. These "filler-gap dependencies" are most familiar in wh-questions, relative clauses, and topicalization, but they are also present in scrambling processes in many languages and in some varieties of comparative constructions. Although a filler may appear arbitrarily far from its associated gap (1), there are many domains that typically block filler-gap dependencies. Following terminology introduced by Ross (1967) these domains are known as "islands." The class of islands includes relative clauses (2a), wh-clauses (2b), factive clauses (2c), subjects (2d), adjuncts (2e), and coordinate structures (2f). Similar domains induce island effects in many languages, but there is also cross-language variation in island effects, and the extent of this variability is a topic of current debate (e.g., Boeckx, 2007; Stepanov, 2007; Yoshida, 2006).

(1)   What does the teacher think that the children expect her to assign __ for homework?

(2)   a.   * What did the agency fire the official that recommended ___?
       b.   * Who do you wonder whether the press secretary spoke with ___?
       c.   * Why did they remember that the corrupt CEO had been acquitted ___?
       d.   * What did the fact that Joan remembered ___ surprise her grandchildren?
       e.   * Who did Susan watch TV while talking to ___ on the phone?
       f.   *What did the Senate approve ___ and the House reject the bill.

An important finding on the processing of filler-gap dependencies is that comprehenders construct dependencies in advance of information that disambiguates the position of the gap. This property is known as *active* dependency formation, and it is likely an instance of the more pervasive phenomenon of anticipatory structure building in language comprehension. For example, much evidence indicates that comprehenders posit gaps at least as soon as they encounter a verb, and before they know the specific location of the gap. One type of evidence comes from reading-time slowdowns in sentences like (3a) when readers encounter an overt object ("us") where a gap was expected, relative to reading times for the same word in a control sentence without a filler-gap dependency (*filled-gap effect*: Crain & Fodor, 1985; Lee 2004; Stowe, 1986), or in processing disruptions at the verb position when the filler is an implausible object of the verb as in (4) (Traxler & Pickering, 1996; see also Garnsey, Tanenhaus, & Chapman, 1989).

(3) a. My brother wanted to know who Ruth will bring <u>us</u> home to __ at
         Christmas.
     b. My brother wanted to know if Ruth will bring <u>us</u> home to Mom at
         Christmas.

(4) That's the {pistol/garage} with which the heartless killer <u>shot</u> the hapless
     man yesterday afternoon.

A number of studies have tested the online effects of island constraints by
investigating whether active filler-gap dependency formation occurs in island
environments. Most studies have concluded that it does not. For example,
Stowe (1986, Experiment 2) found no filled gap effect following a preposition
inside a subject noun phrase (5a), and Traxler and Pickering (1996) found no
plausibility effect inside a relative clause island (5b).

(5) a. The teacher asked what the silly story about <u>Greg's</u> older brother was
         supposed to mean.
     b. We like the {book/city} that the author who <u>wrote</u> unceasingly and
         with great dedication saw while waiting for a contract.

A possible concern about these findings is that they rely upon a null effect,
possibly caused by general syntactic complexity rather than island constraints
per se. However, this concern is mitigated by a recent demonstration that the
Coordinate Structure Constraint (CSC) has the effect of extending active
dependency formation to both verbs of a coordinate structure. The CSC
requires that extraction from a coordinate structure targets both conjuncts.
Wagers and Phillips (2009) show that plausibility effects are found in the
second conjunct of a coordinate structure (6a), but not in a minimally different
adjunct clause construction (6b).

(6) a. The {wines|cheeses} which the gourmets were energetically
         discussing or slowly sipping during the banquet were rare imports.
     b. The {wines|cheeses} which the gourmets were energetically
         discussing before slowly sipping {the samples|some wine} during
         the banquet were rare imports.

A number of additional studies have followed a related logic to show that island
constraints regulate active dependency formation (Bourdages, 1992; McElree &
Griffith, 1998; Omaki, Lau, Davidson White, & Phillips, submitted; Omaki &
Schulz, in press; Phillips, 2006; Pickering, Barton, & Shillcock, 1994,
Experiment 2). Additionally, a group of ERP studies has demonstrated that
processing disruption occurs as soon as the parser encounters the edge of an
island domain while it is holding an as-yet unsatisfied wh-phrase (Kluender &
Kutas, 1993; McKinnon & Osterhout, 1996; Neville, Nicol, Barss, Forster, &
Garrett, 1991). Taken together, most evidence on the real-time status of island
constraints indicates that the parser successfully avoids constructing illicit filler-
gap dependencies. A small number of studies have presented results that could be

taken as evidence for online violation of island constraints (Clifton & Frazier, 1989; Freedman & Forster, 1985; Pickering et al., 1994, Experiment 1), but these findings are all open to alternative interpretations, and even the studies' authors do not conclude that the results show that the parser ignores island constraints (for discussion, see Phillips, 2006).

Importantly, this generalization involves the parser's actions in environments where it is not forced to violate an island constraint, and should thus be distinguished from the question of what happens in ungrammatical sentences that contain an unambiguous island constraint violation. The observation that comprehenders are able to recover an interpretation from sentences with clear island violations, albeit with varying degrees of difficulty (Hofmeister & Sag, 2010), is certainly interesting. But it does not bear on the question whether the parser actively constructs filler-gap dependencies in island environments. Additionally, the finding that the parser respects island constraints does not, in general, bear on whether those constraints should be understood as explicit grammatical constraints (Sprouse, Wagers, & Phillips, to appear), or should be regarded as epiphenomena, caused by language processing difficulty (Hofmeister & Sag, 2010; Kluender, 2005; Kluender & Kutas, 1993; Pritchett, 1991) or semantic or functional constraints (Deane, 1991; Erteshik-Shir, 1973). Both approaches are consistent with online respect of island constraints. In fact, it may be that specific instances of online violation of island constraints are more relevant to addressing that debate (Phillips, 2006). For a more general review of filler-gap dependency processing, see Phillips and Wagers (2007).

3.2. Structural Constraints on Backwards Anaphora

Pronouns typically follow their antecedents, but there are many circumstances where a pronoun may precede its antecedent (7a). Such pronouns are known as *cataphoric* pronouns. A well-known constraint on cataphoric dependencies is Principle C of the classic binding theory (Chomsky, 1981), which blocks coreference between a pronoun and a referring expression that it c-commands (7b).[1] Principle C is among the better candidates for a universal constraint: its effects have been found in almost every language where it has been tested (Baker, 1991; but cf. Bruening, 2005; Büring, 2005; Davis, 2009),

---

[1]There are at least two well-known classes of exception to Principle C in English, both of which arise in semantically well-defined contexts. The first class involves comparisons of multiple "guises" of the same individual (Heim, 1992; Reinhart, 1983), as in *He$_i$ then did what John$_i$ always did in such situations*. The second class of exceptions arises in cases where the embedded clause describes an event that interrupts the main clause event, as in *He$_i$ was threatening to leave when Billy$_i$ noticed that the computer had died* (Harris & Bates, 2002). Minor modifications to these examples reintroduce the ill-formedness of Principle C violations, for example, **He$_i$ then did what John$_i$ had already done five minutes earlier; *He$_i$ left the house when Billy$_i$ noticed that the computer had died*. Rather than undermining the validity of the Principle C constraint, such cases help to sharpen its formulation. For further discussion, see Kazanina (2005).

and developmental evidence suggests that it constrains children's interpretations of pronouns at a very young age (e.g., Crain & McKee, 1985; Kazanina & Phillips, 2001; Lukyanenko, Conroy, & Lidz, submitted).

(7) a. While he$_i$ was drinking his coffee, John$_i$ perused the classified ads.
     b. *He$_i$ perused the classified ads while John$_i$ was drinking his coffee.

Evidence for the online effects of Principle C closely resembles the evidence for the effects of island constraints. Just as recognition of a wh-phrase initiates active formation of a filler-gap dependency, recognition of a potential cataphoric pronoun appears to initiate an active search for an antecedent. Evidence for this process comes from *gender mismatch effects* in reading-time measures. In sentences like (8) the underlined noun is read more slowly when it mismatches in gender with the preceding pronoun, suggesting that the parser attempts to link the pronoun and the name before it recognizes that the name is an unsuitable antecedent (van Gompel & Liversedge, 2003). Kazanina and colleagues replicated the gender mismatch effect, but showed that it was absent in configurations like (9), where the pronoun-name link would violate Principle C (Kazanina, Lau, Yoshida, Lieberman, & Phillips, 2007). Similar evidence for online effects of constraints on backwards anaphora has been found in Russian (Kazanina & Phillips, 2010) and Japanese (Aoshima, Yoshida, & Phillips, 2009). Similar conclusions about the online effects of Principle C were reached in an earlier study by Cowart & Cairns (1987) based on naming latency measures.

(8) While he was at the party, the {boy|girl} cruelly teased the {girl|boy} during the party games.
(9) Because last semester she was taking classes full-time while {Kathryn|Russell} was working two jobs to pay the bills, Erika felt guilty.

If Condition C robustly constrains online pronoun resolution, and if this constraint is properly understood in terms of structural relations like c-command, then this implies that the parser builds representations that incorporate appropriate c-command notions. This helps to constrain our understanding of cases where the parser is less successful.

3.3. Forwards Anaphora: Locality Constraints on Reflexives

Across languages, reflexive anaphors come in many forms and are subject to a variety of syntactic and discourse constraints (Barss 2003; Cole, Hermon, & Huang, 2000; Huang, 2004; Koster & Reuland, 1992). However, the distribution of postverbal direct object reflexives in English, such as *himself*, *herself*, or *themselves*, is relatively straightforward. The antecedent of the reflexive must be the subject of the same clause (10). This locality restriction on reflexives is captured by Principle A of the classic binding theory or its analogs (Chomsky, 1981; Pollard & Sag, 1992; Reinhart & Reuland, 1993).

(10)   John$_i$ hoped that Bill$_j$ wouldn't blame himself$_{*i/j}$ for the accident.

A number of studies have tested whether the parser considers grammatically inappropriate antecedents upon encountering a reflexive. Most results converge on the conclusion that it does not. Processing a direct object reflexive triggers selective retrieval of the local subject NP, without interference from other NPs. In a cross-modal priming study using sentences like (11), Nicol and Swinney (1989) showed that immediately after participants heard a reflexive they made faster lexical decisions to visually presented associates of the local subject NP ("doctor"), but not to associates of other NPs in the sentence ("boxer," "skier"). In an elegant eye-tracking study using materials like (12), Sturt (2003) showed that initial reading times for reflexive pronouns were affected by (mis)matches with the gender stereotype of the noun phrase in the local subject position, but not by corresponding (mis)matches with other subject positions in the sentence. In one experiment, Sturt found later effects of interference from grammatically inaccessible NPs, but in another experiment he found no such effects. We obtained similar results in an ERP study using materials like (13), where we found that a P600 is elicited by a reflexive that mismatches the stereotypical gender of the local subject, and is not attenuated by the presence of a matching subject NP in a grammatically inaccessible position (Xiang, Dillon, & Phillips, 2009). Similarly, in eye-tracking studies that manipulated the number of the critical NPs rather than their gender (14) we found that re-reading times were affected by the match between the reflexive and the local subject noun, but not by the match with the grammatically inappropriate antecedent embedded inside a relative clause (Dillon, Mishler, Sloggett, & Phillips, submitted). All of these findings suggest that processing of a direct object reflexive triggers immediate retrieval of the local subject NP, as required by Principle A.

(11)   The boxer told the skier that the doctor for the team would blame himself for the recent injury.

(12)   {Jonathan|Jennifer} was pretty worried at the City Hospital. The surgeon who treated {Jonathan|Jennifer} had pricked herself with a used syringe needle. There should be an investigation soon.

(13)   The tough soldier that {Fred|Katie} treated in the military hospital introduced {himself|herself} to all the nurses.

(14)   The diva that accompanied the harpists on stage clearly presented {herself|themselves} with lots of fanfare.

There is at least one exception to the finding that object reflexives trigger retrieval of only grammatically appropriate antecedents. Badecker and Straub (2002) found in a self-paced reading study that reading times two words beyond a reflexive were slowed by the presence of a gender matching NP in a

grammatically inaccessible subject position, as in (15). However, although this disruption appeared in one study (Experiment 3), no corresponding disruption effects were found in three subsequent experiments using reflexives or reciprocal (*each other*) anaphors (Experiments 4–6).

(15)   {Jane|John} thought that Bill owed himself another opportunity to solve the problem.

In a series of visual world eye-tracking studies, Runner and colleagues have explored the antecedents that are considered for reflexives inside "picture NPs," using examples like *Pick up Ken. Have Ken touch Harry's picture of himself* (Runner, Sussman, & Tanenhaus, 2006). Looking patterns following the reflexive suggested consideration of the nonlocal antecedent *Ken*. However, as Runner and colleagues point out, these findings may indicate that reflexives in picture NPs have the status of "binding-theory-exempt logophors," rather than challenging the findings from studies of argument reflexives.

Thus, most evidence suggests that the processing of simple argument reflexives in English is insensitive to structurally inappropriate antecedents, indicating that the parser engages a retrieval process that selectively targets the subject of the current clause. This conclusion seems unremarkable, until it is set against the various examples that we turn to next of processes that are less grammatically sensitive. In particular, the grammatical accuracy of reflexive processing is an interesting counterpoint to the widespread interference effects found in agreement processing, although both processes target a local subject NP.

4.   GRAMMATICAL ILLUSIONS

In contrast to the grammatical constraints that act as robust constraints on structure generation, preventing illicit dependencies from even being considered, a number of grammatical constraints appear to have a more delayed impact on language processing, leading to cases of *grammatical illusions*. In this section, we summarize a number of the areas where grammatical illusions have been found to date. Importantly, the illusions are highly selective, occurring reliably only under specific circumstances.

4.1.   Selective Fallibility in Agreement Comprehension

Verb agreement in English and many other languages is subject to a straightforward constraint: the morphological features of the finite verb (or auxiliary) must agree with the corresponding features on the subject NP. Nevertheless, many studies have shown that subject-verb agreement is highly susceptible to interference effects. Importantly, however, the interference effects are quite selective.

Elicited production studies confirm what is clear from casual observation: speakers often produce verbs that agree with an inappropriate noun phrase. For example, in (16) the auxiliary agrees not with the singular subject of the clause (*key*) but instead with an embedded plural noun (*cabinets*). Agreement attraction errors of this kind occur frequently in naturalistic speech and writing, and they have been repeatedly confirmed in language production experiments (Bock & Miller, 1991; Solomon & Pearlmutter, 2004; Vigliocco & Nicol, 1998; for a review see Eberhard, Cutting, & Bock, 2005).

(16)   The key to the cabinets are on the table.

An analog of the production errors is found in comprehension studies. Reading times are normally disrupted by agreement errors, but in sentences like (16) the error is often overlooked, leading to reduced disruption effects in average reading times (Pearlmutter, Garnsey, & Bock, 1999; Wagers, Lau, & Phillips, 2009), and reduced accuracy in acceptability rating studies (Clifton, Frazier, & Deevy, 1999; Häussler & Bader, submitted). When a sentence like (16) is judged acceptable we refer to it as an *illusion of grammaticality*. These illusions are both more general and more restricted than one might conclude from examples like (16) alone.

First, attraction effects are not restricted to cases in which a verb agrees with a linearly proximate noun, and hence the term "proximity concord" (Francis, 1986; Quirk, Greenbaum, Leech, & Svartvik, 1985) is misleading. Powerful attraction effects can be caused by nouns that are nonadjacent to the verb. Production studies on complex N-PP1-PP2 noun phrases have shown that a noun inside PP1 is more likely to induce attraction errors than a noun inside PP2 (Franck, Vigliocco, & Nicol, 2002), and that verb-adjacent nouns induce fewer attraction errors when they are more distant structurally from the true subject NP (Bock & Cutting, 1992; Solomon & Pearlmutter, 2004). This has been interpreted as evidence for percolation of number features from a nonsubject node to the highest node in the subject NP (Eberhard et al., 2005; Franck et al., 2002). Additionally, nonlocal attraction effects are found in comprehension and production in configurations where a relative clause verb agrees with the head of the relative clause, rather than the subject of the relative clause (17) (Bock & Miller, 1991; Clifton et al., 1999; Franck, Lassi, Frauenfelder, & Rizzi, 2006; Franck, Soare, Frauenfelder, & Rizzi, 2010; Wagers et al., 2009). Hence, attraction effects should not be viewed as cases of "local coherence" effects (Tabor, Galantucci, & Richardson, 2004).

(17)   The runners who the driver see each morning always wave.

Second, attraction effects are morphologically selective. In English, attraction effects are caused by plural nouns but not by singular nouns. Cross-language studies suggest that this reflects the broader generalization that morphologically "marked" features on nouns can induce agreement attraction, whereas "default" features do not (Harrison, Hartsuiker, Branigan, & Pickering 2005; Badecker & Kuminiak, 2007).

Third, agreement attraction effects are grammatically selective. Our comprehension studies in English have repeatedly found that attraction leads to illusions of grammaticality (18a), but it does not lead to illusions of ungrammaticality (18b), that is, cases where correct agreement might appear to be bad due to interference from another NP. We interpret this grammatical asymmetry in agreement attraction as evidence that attraction reflects cue-based retrieval processes that are triggered only when a verb fails to bear the expected agreement features (Wagers et al., 2009). In cases where the anticipated correct agreement is encountered on the verb, no further retrieval of subject features is required, and hence no errors occur. In contrast, when comprehenders encounter incorrect agreement on the verb, this triggers a retrospective check, based on cue-based retrieval, and this leads to attraction effects due to the fact that the attractor noun partially matches the retrieval cues. An additional study suggests that these processes are independent of the interpretive processes that thematically bind subjects and verbs (Wagers, Lau, Stroud, & Phillips, 2008).

(18)   a.   The key to the cabinets unsurprisingly were on the table.
       b.   The key to the cabinets unsurprisingly was on the table.

Clear attraction effects for subject-verb agreement have been found in exactly the same configurations that fail to show corresponding effects for reflexive licensing. In a within-subjects comparison, we found that the same materials that failed to yield interference effects for reflexives (14) elicited strong attraction effects for subject-verb agreement (19) (Dillon et al., submitted).

(19)   The diva that accompanied the harpists on stage clearly {was|were} flawless on the high notes.

4.2.   Selective Fallibility in Case Licensing

An interesting example of selective fallibility is found in German case licensing (Bader, Meng, & Bayer, 2000). (20a–b) are both embedded clauses in German where the case-marking on the subject NP is incompatible with the clause-final verb, yet the two errors are detected at rather different rates. Speakers reliably reject sentences like (20a), in which a dative-marked subject is followed by an active verb that does not license dative case on its subject. But speakers show a strikingly high rate of acceptance for the equally ungrammatical (20b), where a nominative-marked subject is followed by a passive verb (indicated by the auxiliary "wurde") that requires a dative subject.[2]

---

[2]The 61% rejection rate is specifically for conditions with the feminine determiner *die*, which is ambiguous between nominative and accusative. Rejection rates were higher for conditions with the masculine determiner *der* (79%), which is case unambiguous. Collapsing across genders, rejection rates were 91% and 70% for (20a–b), respectively.

(20) a. * ... dass der Mutter$_{[DAT]}$ das Buch geschickt hat   90% rejection

         " ... that the mother sent the book"

      b. * ... dass die Mutter$_{[NOM]}$ das Buch geschickt wurde   61% rejection

         " ... that the mother was sent the book"

This pattern of illusory case licensing has similarities and differences with the profile found in agreement attraction. The examples in (20) involve illusions of grammaticality, as in agreement attraction, but unlike agreement attraction there is no structurally inappropriate element that the illusory licensing can be attributed to. However, the contrast between marked dative NPs and unmarked nominative NPs parallels a key property of agreement attraction. When speakers encounter a subject bearing marked dative case they are highly sensitive to the case-licensing requirements of the subject. Similarly, in sentences where the true subject NP is plural, speakers are relatively immune to agreement attraction. Case illusions arise following "unmarked" nominative subjects, just as agreement illusions are most common following unmarked singular subjects.

4.3. Forwards Anaphora: Antilocality Constraints on Pronouns

We have already seen evidence that Binding Principles A and C constrain the online generation of interpretations for reflexives and backward anaphora, respectively. The effects of Binding Principle B on the real-time interpretation of forward anaphora are less straightforward. Existing studies show a mixture of evidence for immediate versus delayed effects of the constraint.

To a first approximation, Principle B blocks interpretations in which a pronoun takes a clause-mate antecedent, capturing the fact that coreference is reliably judged to be unacceptable in sentences like (21a), and that only the nonlocal antecedents for the pronoun are acceptable in (21b).

(21) a. John$_i$ likes him$_i$.
      b. The man$_i$ that John$_j$ had dinner with suspects that Bill$_k$ is stalking him$_{i/j/*k}$.

The full description of the constraint is somewhat more complicated, since there are special discourse circumstances where pronouns do appear to allow local antecedents (Evans 1980; Higginbotham, 1983). In (22a) the pronoun "him" is at least marginally acceptable with "Bill" as its antecedent, and it is certainly better than (22b). Local coreference is also acceptable in cases like (23) that assert a link between two different "guises" (i.e., independent mental representations) of an individual. Based on these and other considerations, it has been argued that Principle B applies only to cases of bound variable anaphora like (24). It has been further suggested that cases of local coreference are grammatically well formed, and that the unacceptability of examples like

(21) reflects independent pragmatic constraints (Heim, 1992; Reinhart, 1983; for recent reviews see Büring, 2005; Elbourne, 2008). This complex state of affairs has been blamed for children's notorious difficulty in excluding local coreference interpretations for pronouns, which is the topic of a sizeable developmental literature (e.g., Chien & Wexler, 1990), although this literature shows substantial variability in children's performance, possibly as a function of task design (for review see Conroy, Takahashi, Lidz, & Phillips, 2009).

(22) a. ?   I know what Mary, Sue, and Bill$_i$ have in common. Mary likes him$_i$, Sue likes him$_i$, and Bill$_i$ likes him$_i$, too.
      b. *   I know Mary, Sue, and Bill$_i$ have in common. Mary likes him$_i$, Sue likes him$_i$, and Bill$_i$ likes himself$_i$, too.

(23)   Bill$_i$ is probably him$_i$. (*Attempting to pick out Bill from an old class photo.*)

(24) a. Every boy$_i$ thinks that Sue likes him$_i$.
      b. *   Every boy$_i$ likes him$_i$.

Studies of the online effects of Principle B show mixed results. The adult literature has focused almost exclusively on coreference rather than on bound variable anaphora. Some studies suggest that Principle B blocks consideration of local antecedents, but other results suggest that local antecedents are fleetingly considered and then filtered out after a brief delay. It is tempting to view this variability as comparable to the variation that one finds in the developmental literature, but whereas findings of interpretations that violate Principle B are rather common in the child literature and much experimental care is needed to remove such effects (Conroy et al., 2009), evidence of (temporary) interpretations that violate Principle B is more elusive in adult online studies.

   A number of studies have found that adults do not appear to consider local antecedents for pronouns that violate Principle B. Nicol and Swinney (1989) reported in an early cross-modal priming study that presentation of the pronoun in (25) led to priming of associates of the two acceptable nonlocal antecedents ("boxer," "skier"), but not the illicit local antecedent ("doctor"). At least two other studies using a self-paced reading method have found immediate online effects of Principle B. Clifton, Kennison, and Albrecht (1997, Experiment 2) compared reading times in sentences pronouns that either matched or mismatched the number of a local subject NP (26). They found that the number mismatch impacted reading times in conditions with a possessive pronoun, for which the subject is a grammatically accessible antecedent, but not in conditions with a direct object pronoun. Lee and Williams (2008) tested effects of gender match between pronouns and accessible or inaccessible antecedents, using a similar design to Sturt's studies on reflexives described above (Sturt, 2003), using sentences like (27). In two studies they found that there was a consistent reading time disruption at the

word immediately following the pronoun when it mismatched the gender of the grammatically accessible subject, but they also obtained less consistent later effects of mismatch with the inaccessible subject. A recent series of four studies using self-paced reading and eye-tracking measures, and a variety of different types of antecedent (names, simple vs. complex referring expressions, quantificational NPs) and context found consistent effects of Principle B: reading times were disrupted by gender mismatch with an accessible subject, but not by the gender features of an inaccessible subject (Chow, Lewis, & Phillips, 2011).

In contrast, Badecker and Straub (2002) report three experiments where they found effects of illicit local antecedents for pronouns. In sentences like (28), which always contained a grammatically accessible antecedent for the pronoun, they found that reading times following the pronoun were longer when the illicit local antecedent matched the gender of the pronoun than when it mismatched the pronoun, that is, a "multiple match" effect. They reasoned that the interference effect could only have arisen if the local subject NP was fleetingly considered as an antecedent for the pronoun. In contrast to many of the other studies on online effects of Principle B, these studies primarily used a post-sentence probe recognition task. Comprehension questions unrelated to the pronoun interpretation appeared on 25% of trials. Since the probe recognition task only requires participants to judge whether the probe word appeared in the preceding sentence, the task does not encourage participants to properly comprehend the sentence, and provides little indication of whether they were successful. A subsequent study by Clifton and colleagues failed to replicate the multiple match effect involving inaccessible pronoun antecedents (Clifton et al., 1999),[3] and recent studies using a similar logic found no multiple match effect (Chow et al. 2011). Most of these studies did not use a probe recognition task, but one of the studies by Chow and colleagues used Badecker and Straub's task and identical target sentences and still failed to find a multiple match effect. A similar conclusion to Badecker and Straub has been drawn based on a reading-time slowdown in (29a) relative to (29b) (Kennison, 2003, Experiment 2). One might object that this effect reflects the lack of a suitable antecedent for the pronoun in the sentence, which might force comprehenders to take extreme measures, such as considering a grammatically illicit antecedent, paralleling objections that have been raised about evidence for the Delay of Principle B Effect in children (Conroy et al., 2009; Elbourne, 2005). However, even if this is the source of the effect, it is notable that the parser is able to consider this illegal interpretation, since corresponding violations do not seem to occur to speakers when they process reflexives. In a series of studies based on visual world eye-tracking, Runner and colleagues

---

[3]Although the studies by Badecker and Straub appeared a few years after Clifton and colleagues' findings, they were mostly carried out in the early 1990s.

concluded that illicit antecedents for pronouns in picture NPs are briefly considered and then ruled out (Runner et al., 2006). However, this conclusion relies on the assumption that looks to a picture of a recently introduced referent reflect consideration of that referent as an antecedent for the pronoun, a link that should be treated with caution, as Runner and colleagues point out (p. 219).

(25)   The boxer told the skier that the doctor for the team would blame him for the recent injury.

(26)   The {supervisor | supervisors} paid {him | his assistant} yesterday to finish typing the manuscript.

(27)   The {surgeon | midwife} said that Richard introduced {him | her} during the opening reception of the medical conference.

(28)   John thought that {Bill | Beth} owed him another chance to solve the problem.

(29) a. Carl watched him yesterday during the open rehearsals of the school play.
      b. They watched him yesterday during the open rehearsals of the school play.

In light of the accumulated evidence it is unclear whether Principle B belongs in the list of grammatical constraints that are temporarily violated online. The clearest evidence for online Principle B violations comes from the studies by Badecker and Straub, but most evidence points to a different conclusion. If it turns out that comprehenders do not consider pronoun antecedents that violate Principle B online, this is particularly interesting, as there are a number of reasons why one might expect this constraint to have only delayed effects. First, the constraint is generally analyzed as reflecting a complex interaction of syntactic and pragmatic constraints (e.g., Grodzinsky & Reinhart, 1993; Reinhart, 2006). Second, even a simplified version of Principle B is a negative constraint, which allows a pronoun to take any semantically compatible antecedent *except* a clausemate NP. Third, since pronouns can take such a wide range of discourse antecedents, a plausible resolution strategy would involve first accessing any feature-matching referents in the discourse representation, and only later filtering out structurally illicit matches. Such an approach would plausibly give rise to temporary Principle B violations. The online status of Principle B remains unresolved.

4.4. Negative Polarity Items (NPIs)

Negative polarity items (NPIs) such as *any*, *ever*, *lift a finger*, or *a damn thing* have attracted much attention in the semantics and pragmatics literatures, where interest has focused on the challenge of defining the conditions under which they are licensed. As a first approximation, NPIs are licensed in

configurations where they are c-commanded by negative-like elements (30a). When the negative-like element is absent (30b) or does not c-command the NPI (30c), the NPI is not licensed.

(30) a. No professor will ever say that.
  b. *   A professor will ever say that.
  c. *   A professor that no student likes will ever say that.

However, the class of potential licensors is notoriously difficult to characterize. In addition to transparently "negative-like" elements such as *no*, *few*, *rarely*, and *deny*, NPIs are licensed in less obviously negative-like contexts, such as conditionals, questions, and the first argument of quantifiers (31).

(31) a. If John ever shows up, he will learn that he is fired.
  b. Who has ever been able to answer a question like that?
  c. Everybody who John has ever met ends up finding him fascinating.

In a classic study, Ladusaw (1979) proposed that the common factor underlying NPI-licensing environments is that they are *downward entailing* contexts. A proposition is downward entailing if it entails a more specific proposition. For example, *John didn't buy a car* entails that *John didn't buy a German car*, whereas *John bought a car* does not entail that *John bought a German car*. NPI licensors canonically create downward entailing contexts. More recently, there is a growing consensus among many semanticists that NPI licensing reflects an interaction between the lexical properties of NPIs and the semantics and pragmatics of entire propositions (e.g., Chierchia, 2006; Fauconnier, 1975; Israel, 2004; Kadmon & Landman, 1993; Ladusaw, 1992), rather than a direct structural dependency between an NPI and a specific licensing element. The c-command requirement on NPI licensing is thus understood not as an explicit grammatical constraint, but rather a consequence of how semantic interpretation licenses inferences.

  If NPI licensing is a direct consequence of normal compositional interpretive mechanisms, as argued by current theories, then we might expect incremental interpretation to engender successful and accurate online NPI licensing. However, recent experimental findings indicate that this is not the case. Drenhaus and colleagues found in a speeded grammaticality judgment study that German speakers reliably judge sentences that lack an NPI licensor to be unacceptable, but that acceptance rates are higher for sentences that contain a negatively quantified NP that fails to c-command the NPI (Drenhaus, Saddy, & Frisch, 2005). These findings have been corroborated in German and English using eye-tracking (Vasishth, Brüssow, Drenhaus, & Lewis, 2008), ERPs (Xiang et al., 2009), self-paced reading, and a variety of visual and auditory judgment paradigms (Xiang, Dillon, & Phillips, 2006). In a typical speeded judgment study in English, we find that the presence of a non c-commanding licensor (32c) increases acceptance rates by 15%–30%, relative to sentences with no licensor at all (32b), although acceptance rates are still much lower than for appropriately licensed NPIs (32a).

(32) a.   No bills that the senators voted for will ever become law.
    b.   *   The bills that the senators voted for will ever become law.
    c.   *   The bills that no senators voted for will ever become law.

Vasishth and colleagues offer illusory NPI licensing as evidence for dependency formation mechanisms that involve cue-based retrieval of elements in content-addressable memory (Vasishth et al., 2008). They propose that recognition of an NPI triggers a retrieval process that seeks an element that has the properties [+negative] and [+c-commander], and that illusory licensing occurs in cases where a partially matching [+negative] element is detected in memory, and is incorrectly treated as a full match to the retrieval cues, due to noisy representations. However, the retrieval cues proposed in this account are problematic. The use of the [+negative] cue presumes that NPI licensing involves specific item-to-item dependencies, just as in agreement, anaphora, and filler-gap dependencies. This assumption is at odds with the widespread view in semantics of how NPIs are licensed by entire contexts. Additionally, the fact that licensing environments for NPIs are so diverse undermines the feasibility of identifying a unique feature that defines an NPI-licensing environment and that could be used as a retrieval feature. Likewise, the use of a [+c-commander] feature is also problematic. C-command describes the relation between pairs of nodes in a hierarchical structure, and is not an inherent feature of an individual word or phrase when it is encoded in memory. Hence, it is an unlikely candidate as a retrieval cue for an NPI licensor. As an alternative to the account based on mis-retrieval, Xiang and colleagues have offered an account of illusory NPI licensing that attempts to link illusory NPI licensing to the same semantic and pragmatic processes that are proposed to be responsible for successful NPI licensing (Xiang et al., 2009).

An important task for future work on illusory NPI licensing is to better define the scope of the illusions. All previous demonstrations of illusory NPI licensing involve the NPI *ever* and a configuration in which the spurious licensor is an argument NP in a subject-modifying relative clause. It will be important to determine whether illusory licensing extends to a broader range of licensors and contexts. Retrieval-based accounts of illusory licensing predict that the phenomenon should be rather general, since the illusions are attributed to mechanisms that should apply whenever a comprehender attempts to license an NPI. In contrast, evidence that the illusions are more restricted could motivate an alternative account that relies on specific semantic and pragmatic mechanisms.

4.5. Comparative Illusions

The best known, yet perhaps least well studied case of a grammatical illusion involves comparative constructions. Sentences like (33) initially strike most listeners as quite acceptable, yet upon further reflection they turn out to be

meaningless.[4] This illusion is so powerful that it hardly needs experimental confirmation, but controlled judgment studies have confirmed that the error is robust (Fults & Phillips, 2004; Wellwood, Pancheva, Hacquard, Fults, & Phillips, 2009). Sentences like (33) are ill-formed because the main clause subject NP *more people* calls for a comparison of two sets of individuals, yet no such comparison set is available in the comparative clause. In more formal terms, the comparative clause lacks a sortal that can host a degree variable. This contrasts with sentences like (34a–b) where a degree variable may hosted either by the bare plural subject (34a) or by the gap in direct object position (34b). In light of the fact that the sentence lacks this basic property, it is remarkable that speakers so commonly fail to notice the error.

(33)   More people have been to Russia than I have.

(34)   a.   More Americans have been to Russia than Canadians have.
       b.   More Americans have been to Russia than I would have thought.

Townsend and Bever (2001) consider these illusions as a consequence of blending of two types of surface forms that are each syntactically and semantically well-formed, such as the two examples in (35), which have substantial lexical overlap, and both express comparisons.[5]

(35)   a.   More people have been to Russia than I.
       b.   People have been to Russia more than I have.

Wellwood and colleagues (2009) propose instead that the illusions are more narrowly related to the semantics of comparison, arising when speakers erroneously interpret the comparative clause as conveying a comparison with a number of events. Under this account, (33) is interpreted as a comparison of the number of events of (other) people going to Russia, and the number of events of the speaker going to Russia. This may be related to the fact that English allows statements about quantities of individuals to be used to convey assertions about quantities of events, for example, *106 million cars crossed the George Washington Bridge in 2007* (Barker, 1999; Krifka, 1990). In support of this account, Wellwood and colleagues show that acceptability ratings are substantially higher in comparatives with repeatable predicates (36a) than in comparatives with nonrepeatable predicates (36b), that is, predicates that an individual is unlikely to satisfy more than once (mean ratings of 5.3 vs. 3.8 on a 7-point scale; fully grammatical counterparts received mean ratings of

---

[4]The first report of these illusions appears to be due to Mario Montalbetti, who in the prologue to his PhD dissertation (Montalbetti 1984) thanks Herman Schultze "for uttering the most amazing */? sentence I've ever heard" (p. 6). Thanks to Jim Higginbotham and Laura Whitton for leading us to this source.

[5]This exposition goes beyond the description in Townsend & Bever's book. Thanks to Tom Bever for further discussion of these cases.

5.5–6.0). Speakers' sensitivity to the semantics of the predicate suggests that the illusion is not merely a consequence of the fact that the sentence consists of a blend of two acceptable sentences. However, it remains unclear whether speakers construct a full interpretation that includes the erroneous interpretation, or whether the acceptability judgments merely reflect a superficial mechanism that detects elements that lend themselves to interpretation in terms of numbers of events, without actually constructing the full interpretation.

(36) a. More undergrads call their families during the
          week than I do. *Repeatable*
      b. More New Yorkers began law school this
          semester than I did. *Nonrepeatable*

The comparative illusion is a little different than other cases of illusory licensing reviewed here. The illusion is so robust that most speakers fall prey to it on most trials, until they attempt to construct a fuller interpretation or until their attention is drawn to the illusion. This contrasts with agreement and NPI-licensing illusions, which arise probabilistically (typically on only 15%–30% of trials), and often are clearest in speeded acceptability tasks or in early online measures. Cases of access to illicit antecedents in pronoun resolution appear to be fleeting at best, such that the errors do not even appear in acceptability judgment tasks, as far as we know. This contrast in the persistence of different types of illicit dependencies could reflect differences in the processes that form the illicit dependencies in the first place. If so, this would motivate a nonuniform account of illusory dependency formation. Alternatively, similar mechanisms may be responsible for the formation of the various dependencies, and the differing time profiles may reflect independent properties of the constructions that cause speakers to notice and retract errors. For example, anaphoric dependencies are interpreted immediately, agreement dependencies appear to have no interpretive consequences, and comparative constructions may allow for "provisional" interpretations. It may be that interpretive processes play a role in identifying and filtering out illicit dependencies.

4.6. Other Possible Illusions

There are a number of other phenomena that might belong in a more comprehensive review of grammatical illusions. In each instance, however, it remains unclear whether the examples are comparable to the other phenomena discussed here. First, there are striking reports of doubly center-embedded sentences with missing verbs, for example, *The patient the nurse the clinic had hired met Jack*, that are rated as more acceptable than fully grammatical counterparts (Frazier 1985; Gibson & Thomas 1999). These examples differ from the others discussed here in the respect that they plausibly reflect complexity-induced overload, and it is not clear what parse is assigned to such dramatically ill-formed sentences. Second, some studies show high acceptability of instances of voice

mismatches in ellipsis, for example, *The dessert was praised by the customer after the critic did already* (Arregui, Clifton, Frazier, & Moulton 2006). In these cases, it remains unclear to what extent the mismatches reflect parsing versus liberal grammatical constraints. Third, recent findings in ERP studies of P600 responses to semantically odd but syntactically well-formed sequences in a number of languages have been interpreted as evidence of mis-binding of arguments to thematic roles that are incompatible with the surface syntax (e.g., *The hearty meal was devouring* . . . , Kim & Osterhout 2005; *For breakfast the eggs would only eat toast and jam*, Kuperberg, Sitnikova, Caplan, & Holcomb, 2003; see also Bornkessel-Schlesewsky et al., 2011; Hoeks, Stowe, & Doedens, 2004; Stroud & Phillips, 2011; van Herten, Chwilla, & Kolk, 2006; Ye & Zhou, 2008; for recent reviews see Bornkessel-Schlesewsky & Schlesewsky, 2008; Kuperberg, 2007). However, it remains unclear whether these effects genuinely reflect thematic mis-binding, or whether they can be explained in terms of more familiar mechanisms that simply detect errors in the surface form of sentences. These ERP effects have sometimes been associated specifically with instances of thematic role reversals, and there is some interesting evidence that a different ERP profile obtains in cases of thematic anomalies that cannot be construed as role reversal (Kim & Osterhout, 2005). However, most evidence suggests that the effects are elicited by a broader range of anomalies involving thematic roles and animacy, and are not selectively elicited by reversals (Stroud, 2008).

5.   SYNTHESIS: SELECTIVE FALLIBILITY

Our review of the online effects of different grammatical constraints yields a profile of "selective fallibility." One set of constraints, including island constraints on filler-gap dependencies, Binding Principle C, and the locality constraint on English reflexives, appears to be implemented accurately, with very limited evidence that the parser constructs illicit dependencies. Another set of constraints has weaker effects on online dependency formation. Illicit local antecedents for pronouns may be fleetingly considered in violation of Principle B, although the evidence for this is equivocal. Illusory licensing of negative polarity items (NPIs) appears to be a robust phenomenon, at least for the specific combinations of NPIs and licensors that have been tested to date. And comparative illusions are so pervasive that most speakers initially find them acceptable. Meanwhile, a third class of grammatical constraints, which includes subject-verb agreement and case licensing, shows an uneven fallibility profile, with some configurations relatively immune to errors and others showing high rates of illusory licensing.

The profile of selective fallibility poses a challenge for most existing accounts of real-time grammatical computation. The cases of rapid and accurate implementation of grammatical constraints are unexpected under accounts that emphasize the parser's use of superficial heuristics (Townsend & Bever,

2001), "good enough" representations (Ferreira & Patson, 2007), its sensitivity to local coherence relations (Tabor et al., 2004), or the pervasiveness of structure-independent interference effects (e.g., Lewis et al., 2006; Van Dyke & McElree, 2006). Meanwhile, the cases of grammatical illusions do not sit comfortably with the view that real-time language processes transparently implement grammatical constraints (e.g., Cann et al., 2005; Kempen & Harbusch, 2002; Phillips, 2004).

It is possible that selective fallibility reflects hard-coded differences between the various grammatical constraints. It could be that some constraints are implemented in the human parsing mechanism while others are not, and that this directly accounts for the observed online differences. But this would not be a very satisfying account, as it would predict that the profile of successes and failures is arbitrary. It would also struggle to explain why we find selective fallibility even within the domain of individual constraints, such as subject-verb agreement. A more attractive approach is to seek a more systematic account of selective fallibility that can predict the online effects of an individual constraint based on its structural description.

A tempting approach is to seek a unique property that distinguishes the fallible from the infallible grammatical constraints. A number of interesting candidate properties suggest themselves, but each accounts for only some of the observed variation.

Directionality. Some grammatical dependencies can be identified based on their left-hand element; others are not identifiable until the right-hand element is reached. This contrast may be important. Filler-gap dependencies and backwards anaphora are among the dependencies that show greatest online sensitivity to structural constraints, and they also share the property of being identifiable based on the left-hand element in the dependency. This property allows the parser to search actively for the right-hand element of the dependency (a gap or an antecedent, respectively), and it may also allow the parser to exclude certain structural domains from its search in advance of bottom-up input. This could be particularly valuable, as a way of avoiding the errors that result from partial cue matches in retrievals from content-addressable memory. Such *prospective* search processes may be very different from the *retrospective* retrieval processes that are required in dependencies that are identifiable only when their right-hand element is reached (e.g., forwards anaphora, NPI licensing). This contrast between prospective and retrospective processes and their contrasting sensitivity to interference has interesting parallels in the domain of cognitive control (e.g., Braver, Paxton, Locke, & Barch, 2009). However, directionality cannot be the sole predictor of grammatical accuracy in parsing, as there are backwards-looking dependencies that show impressive online accuracy (e.g., reflexive binding) and forwards-looking dependencies that exhibit clear illusions (e.g., comparatives).

Locality. The grammatical accuracy of dependencies involving reflexives is one of the clearest cases of structure-sensitivity in retrieval. It is possible that

the local nature of the structural condition on English argument reflexives may be responsible for this accurate retrieval. From the perspective of a cue-based retrieval mechanism, it is difficult to implement a c-command requirement on retrieval, as c-command is a relational property of pairs of positions, rather than an inherent property of individual positions. Hence, the fallibility of grammatical constraints involving c-command requirements has been offered as a virtue of cue-based retrieval models (Vasishth et al., 2008). However, more local relations such as "subject of the current clause" could be straightfor- wardly implemented in a cue-based model, providing a possible account of why a local grammatical constraint might be more accurately implemented (cf. Xiang et al., 2009). However, locality cannot be the sole reason why retrieval of reflexive antecedents is so accurate, since subject-verb agreement involves a similarly local relation, yet it is highly susceptible to illusory licensing.

Higher order representations. When evaluating evidence for (in)accurate online implementation of grammatical constraints, it is important to remember that when a speaker fleetingly considers an illicit interpretation we cannot be certain that this reflects an illicit *syntactic* dependency. For example, to the extent that studies on the online effects of Principle B show that illicit local antecedents are briefly considered, this could reflect retrieval of a referent from a discourse model rather than construction of an illicit syntactic relation. Similarly, one could argue that the illusory licensing effects observed with comparatives and NPIs arise at the level of semantics or pragmatics, rather than in specifically syntactic representations. However, although differences between levels of representation may be responsible for part of the profile described in Sections 3 and 4, they cannot explain the entire profile. If fleeting violations of Principle B are due to retrieval of antecedents from a discourse model, then we must still explain why the same retrieval processes do not lead to temporary violations of Principles A and C. Also, semantic or discourse representations cannot be responsible for all cases of illicit dependency formation, since agreement attraction effects appear to have no interpretive consequences.

Structural priority. We suggest that a more promising account of the selective fallibility profile is one that incorporates aspects of these different partial accounts. We propose that *structural information constrains dependency formation when it has temporal priority over other information.* This is not a claim that structural information has special architectural priority in the language processing system. Rather, it reflects the observation that syntactic information is sometimes available before other types of relevant linguistic information (cf. Frazier & Fodor, 1978). For example, when English speakers encounter a determiner at the left edge of a phrase, they can reliably infer the syntactic category of the phrase (i.e., noun), although they may not know anything more about the semantic or phonological properties of the phrase head. We have argued elsewhere that this may account for the earliness of the ERP responses to certain structural anomalies (Lau, Stroud, Plesch, & Phillips, 2006). In contrast, in cases where structural information relevant to dependency

formation becomes available simultaneously with nonstructural information, then the risk of illicit dependency formation arises, due to the possibility that nonstructural retrieval cues will match the features of structurally irrelevant items in memory.

The first situation where structural information has priority arises in dependencies whose left-hand element provides reliable information that can guide prospective search for the right-hand element of the dependency. This may account for the reliable online implementation of island constraints and Principle C. Once the parser has identified the left-hand element of a filler-gap dependency, it can initiate an active search for a gap site, and it can mark island domains as irrelevant to its search before it encounters most of the material inside the island. Similarly, processing of a cataphoric pronoun can initiate an active search for an antecedent, and the parser can designate the c-command domain of the pronoun as irrelevant to its search, as required by Principle C, in advance of processing the specific content of that domain.

In the case of subject-verb agreement, we have proposed that the observed grammatical asymmetry may reflect the trade-off between prospective and retrospective search (Wagers et al., 2009). When the parser encounters a subject NP, it can reliably project the agreement features of the upcoming verb or auxiliary. When it encounters an inflected word that matches those features it can confirm its prediction, with no need for additional retrieval of the subject NP, and hence no danger that features from an inappropriate noun will interfere. This accounts for the lack of agreement attraction in grammatical sentences (no "illusions of ungrammaticality"). In contrast, if the parser encounters a verb with agreement features that mismatch its prediction, then this may trigger cue-based retrieval processes that search for an appropriate agreement controller. We proposed that those cue-based retrieval processes are responsible for agreement attraction in ungrammatical sentences ("illusions of grammaticality"). Although we have focused here on comprehension, related accounts of the widespread agreement attraction errors in production have been proposed (Badecker & Kuminiak, 2007; Solomon & Pearlmutter, 2004). Following ideas laid out by Solomon and Pearlmutter, we suggest that the key element that speech planning and retrieval in comprehension share is that they are sequence management tasks in which elements may be simultaneously active.

Comparative illusions present an interesting case, since the determiner *more* provides a clear early signal that the grammar of comparative constructions is relevant to the sentence. We might therefore expect prospective search processes to make comparatives immune to illusions in the same way that filler-gap dependencies and backwards anaphora constructions are. This is clearly not the case. However, comparatives may be different due to the diversity of comparative constructions in English, which means that the parser cannot form a definitive prediction about the right-hand element of the comparison. The comparatives in (37) all have the same main clause, but are completed either as a phrasal comparative (37a), or as a clausal comparative

with a bare plural subject (37b), or as a clausal comparative with an elided object that is dependent on the entire main clause for its interpretation (37c).

(37) a. More people have been to Russia than me.
   b. More people have been to Russia than sheep have.
   c. More people have been to Russia than I could imagine.

Hence, the diversity of comparative constructions may prevent the parser from initiating a constrained prospective search process. However, this observation is insufficient to explain the existence of comparative illusions. Additional assumptions are needed to account for the parser's apparent willingness to entertain comparisons that relate a number of individuals to a number of events. As mentioned above, this could be related to well-formed instances in English where object quantification is used to convey event quantification (Barker, 1999; Krifka, 1990). But this still leaves open the question of why comprehenders fail to notice the overextension of this device.

   Our generalization about the relative timing of structural and nonstructural information could account for the online fallibility of Principle B, if indeed these effects are real. The parser's first indication that Principle B is relevant to its actions is the recognition of a pronoun, which must then initiate a retrospective search for an antecedent. Since the structural relation between a pronoun and its antecedent is almost entirely free, it is natural to assume that a pronoun initiates a cue-based search for an antecedent that shares its person, number, and gender features, and hence it would not be surprising for this search to detect nouns that match those cues, even when they violate Principle B. We should note, however, that this account does not explain why the effects of illicit antecedents for pronouns are found inconsistently across studies, in contrast to the robustness of agreement attraction effects.

   Illusory licensing of NPIs could be understood in the same way, as a consequence of cue-based retrieval mechanisms (Vasishth et al., 2008). However, as pointed out above, this approach predicts that it should be a rather general phenomenon, extending across a wide range of NPIs and illicit licensing configurations. Furthermore, it remains unclear how a cue-based search could identify the full range of NPI-licensing environments. An important goal for future work is to establish whether illusory NPI licensing is better characterized in terms of general retrieval processes, or in terms of narrower pragmatic inference mechanisms (Xiang et al., 2009). The existing findings on illusory NPI licensing are confined to comprehension, and it is unclear whether similar errors might arise in production.

   The grammatical accuracy observed in the processing of English argument reflexives contrasts sharply with the highly fallible nature of subject-verb agreement processing, despite the fact that both processes target the same structural position, that is, the subject of the local clause. The accuracy of reflexive processing clearly cannot be attributed to predictive processes. More

work is needed in order to understand the contrast between agreement and reflexive processing, but we tentatively suggest that argument reflexives are immune to interference from structurally inaccessible antecedents because antecedents are retrieved using only structural cues. In effect, we are suggesting that the person, gender, and number features of reflexives like *himself*, *herself*, and *themselves* play no role in the search for antecedents, and are thus equivalent to the feature-neutral reflexives in closely related Germanic languages, like Dutch *zich* or German *sich*. This does not, of course, mean that speakers fail to notice inappropriate antecedents for reflexives, only that they do not use the features of the reflexives as a retrieval cue. This assumption raises further questions, but it is consistent with evidence from child English, where it has been found that children are sensitive to the number and gender features of pronouns but not to the corresponding features of reflexives (Chien & Wexler, 1990; Zukowski, McKeown, & Larsen, 2008).

The one remaining instance of selective fallibility discussed here is the contrast in German speakers' sensitivity to illicit case licensing configurations. Speakers reliably detect errors when a dative-marked subject is followed by a verb that fails to license dative case, but often fail to notice when a verb that requires a dative-marked subject is preceded by a nominative subject (Bader et al., 2000). The contrast between marked dative and unmarked nominative case is reminiscent of the contrast between marked plural and unmarked singular forms in agreement processing. Additionally, it is consistent with the rest of our account that processing of a dative-marked subject NP should initiate a search for a verb that licenses a dative subject. However, the illusion of grammaticality reported for passive verbs that appear with a nominative subject instead of a dative subject (20b) is not predicted by our account. In contrast to the other examples of illusory licensing reviewed here, there is no interfering element that could be held responsible for the illusion. One possibility is that the nominative subject in (20b) is initially analyzed as appropriately licensed when the parser reaches the participle *geschickt* "sent," which is compatible with active and passive continuations alike, and that this case licensing is not consistently revised once the subsequent auxiliary *wurde* disambiguates the clause as a passive.

6.   INTERIM CONCLUSION

We have attempted a preliminary synthesis of findings on the effects of grammatical constraints on real-time language processing. Many previous studies have examined the online effects of individual constraints, but the results of this research are not well known in linguistics, and in psycholinguistics they are distributed across a number of subliteratures that have had only limited influence on one another. We hope to have shown here that it is useful to bring together the findings from different domains, as they allow us to build

a clearer picture of how speakers mentally encode and navigate linguistic representations in real time.

We find many situations where online language processing is highly sensitive to detailed grammatical constraints and immune to interference from inappropriate material. These findings make it difficult to maintain that speakers respond to the pressures of online sentence comprehension by resorting to superficial or "good enough" representations (Ferreira & Patson, 2007; Townsend & Bever, 2001), and they present a challenge for models that emphasize the pervasiveness of interference in sentence processing (e.g., Lewis et al., 2006). Yet in many other situations, we find that online processes are susceptible to interference and to grammatical illusions. In order to explain this "selective fallibility" profile, we have argued that speakers build richly structured representations as they process a sentence, but that they have different ways of navigating these representations to form linguistic dependencies. The representations can be navigated using either structural information or using structure-insensitive retrieval cues. In order to explain why structural constraints dominate in some situations but are at least temporarily overridden in others, one does not need to assume architectural priority for structural information. Rather, structural constraints may impact linguistic dependency formation most strongly in situations where relevant structural information is available in advance of potentially interfering material in the bottom-up input.

It should also be clear from this survey, however, that work in this area has barely begun, and that the rich profile that is emerging raises many questions that are ripe for future investigation.

ACKNOWLEDGMENTS

REFERENCES

Aoshima, S., Yoshida, M., & Phillips, C. (2009). Incremental processing of coreference and binding in Japanese. *Syntax*, *12*, 93–134.

Arregui, A., Clifton, C., Frazier, L., & Moulton, K. (2006). Processing elided VPs with flawed antecedents. *Journal of Memory and Language*, *55*, 232–246.

Badecker, W., & Kuminiak, F. (2007). Morphology, agreement, and working memory retrieval in sentence production: Evidence from gender and case in Slovak. *Journal of Memory and Language*, *56*, 65–85.

Badecker, W., & Straub, K. (2002). The processing role of structural constraints on the interpretation of pronouns and anaphora. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 748–769.

Bader, M., Meng, M., & Bayer, J. (2000). Case and reanalysis. *Journal of Psycholinguistic Research*, *29*, 37–52.

Baker, M. C. (1991). On some subject-object non-asymmetries in Mohawk. *Natural Language and Linguistic Theory*, *9*, 537–576.

Barker, C. (1999). Individuation and quantification. *Linguistic Inquiry*, *30*, 683–691.

Barss, A. (2003). *Anaphora: A reference guide*. Malden, MA: Blackwell.

Bock, J. K., & Cutting, J. C. (1992). Regulating mental energy: Performance units in language production. *Journal of Memory and Language*, *31*, 99–127.

Bock, J. K., & Miller, C. A. (1991). Broken agreement. *Cognitive Psychology*, *23*, 45–93.

Boeckx, C. (2007). Islands. *Language and Linguistics Compass*, *2*, 151–167.

Bornkessel-Schlesewsky, I., Kretschmar, F., Tune, S., Wang, L., Genç, S., Philipp, M., . . . Schlesewsky, M. (2011). Think globally: Cross-linguistic variation in electro-physiological activity during sentence comprehension. *Brain and Language*, *117*, 133–152.

Bornkessel-Schlesewsky, I., & Schlesewsky, M. (2008). An alternative perspective on "semantic P600" effects in sentence comprehension. *Brain Research Reviews*, *59*, 55–73.

Bourdages, J. S. (1992). Parsing complex NPs in French. In: H. Goodluck & M. S. Rochemont (Eds.), *Island constraints: Theory, acquisition, and processing* (pp. 61–87). Dordrecht: Kluwer.

Braver, T. S., Paxton, J. L., Locke, H. S., & Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proceedings of the National Academy of Sciences*, *106*, 7351–7356.

Bruening, B. (2005). *The Algonquian inverse is syntactic: Binding in Passamaquoddy*. Manuscript, University of Delaware.

Büring, D. (2005). *Binding theory*. Cambridge: Cambridge University Press.

Cann, R., Kempson, R., & Marten, L. (2005). *The dynamics of language*. Syntax and Semantics, Volume 35. San Diego, CA: Academic Press.

Chien, Y.-C., & Wexler, K. (1990). Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language Acquisition*, *1*, 225–295.

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

Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.

Chow, W.-Y., Lewis, S., & Phillips, C. (2011*). Immediate structural constraints on antecedent retrieval in pronoun resolution*. Unpublished manuscript, University of Maryland.

Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin and Review*, *3*, 37–60.

Clifton, C., Jr., & Frazier, L. (1989). Comprehending sentences with long-distance dependencies. In: M. K. Tanenhaus & G. N. Carlson (Eds.), *Linguistic structure in language processing* (pp. 273–317). Dordrecht: Kluwer.

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

Clifton, C., Kennison, S., & Albrecht, J. (1997). Reading the words *her*, *him*, and *his*: Implications for parsing principles based on frequency and on structure. *Journal of Memory and Language*, *36*, 276–292.

Cole, P., Hermon, G., & Huang, C.-T. J. (Eds.). (2000). *Long-distance reflexives*. New York, NY: Academic Press.

Conroy, A., Takahashi, E., Lidz, J., & Phillips, C. (2009). Equal treatment for all antecedents: How children succeed with Principle B. *Linguistic Inquiry*, *40*, 446–486.

Cowan, N. (2000). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.

Cowart, W., & Cairns, H. S. (1987). Evidence for an anaphoric mechanism within syntactic processing: Some reference relations defy semantic and pragmatic constraints. *Memory and Cognition*, *15*, 318–331.

Crain, S., & Fodor, J. D. (1985). How can grammars help parsers? In: D. Dowty, L. Karttunen & A. M. Zwicky (Eds.), *Natural language parsing: Psycholinguistic, computational, and theoretical perspectives* (pp. 94–128). Cambridge: Cambridge University Press.

Crain, S., & McKee, C. (1985). The acquisition of structural restrictions on anaphora. In: S. Berman, J-W. Choe & J. McDonough (Eds.), *Proceedings of the 15th Annual North Eastern Linguistics Society* (pp. 94–110). Amherst, MA: GLSA Publications.

Davis, H. (2009). Cross-linguistic variation in anaphoric dependencies: Evidence from the Pacific Northwest. *Natural Language and Linguistic Theory*, *27*, 1–43.

Deane, P. (1991). Limits to attention: A cognitive theory of island phenomena. *Cognitive Linguistics*, *2*, 1–63.

Dillon, B. (2011). *Structured access in sentence comprehension*. Doctoral dissertation. University of Maryland, College Park, MD.

Dillon, B., Mishler, A., Sloggett, S., & Phillips, C. (submitted). Contrasting interference profiles for agreement and anaphora: Experimental and modeling evidence. Ms., University of Massachusetts, Amherst, and University of Maryland, College Park.

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

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

Elbourne, P. (2005). On the acquisition of Principle B. *Linguistic Inquiry*, *36*, 333–365.

Elbourne, P. (2008). The interpretation of pronouns. *Language and Linguistics Compass*, *2*, 119–150.

Erteshik-Shir, N. (1973). *On the nature of island constraints*. Unpublished doctoral dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Evans, G. (1980). Pronouns. *Linguistic Inquiry*, *11*, 337–362.

Fauconnier, G. (1975). Polarity and the scale principle. *Proceedings of the 11th meeting of the Chicago Linguistics Society*. pp. 188–199.

Ferreira, F., & Patson, N. D. (2007). The 'good enough' approach to language comprehension. *Language and Linguistics Compass*, *1*, 71–83.

Francis, W. N. (1986). Proximity concord in English. *Journal of English Linguistics*, *19*, 309–317.

Franck, J., Lassi, G., Frauenfelder, U., & Rizzi, L. (2006). Agreement and movement: A syntactic analysis of attraction. *Cognition*, *101*, 173–216.

Franck, J., Soare, G., Frauenfelder, U., & Rizzi, L. (2010). Object interference in subject-verb agreement: The role of intermediate traces of movement. *Journal of Memory and Language*, *62*, 166–182.

Franck, J., Vigliocco, G., & Nicol, J. (2002). Attraction in sentence production: The role of syntactic structure. *Language and Cognitive Processes*, *17*, 371–404.

Frazier, L. (1985). Syntactic complexity. In: D. Dowty, L. Karttunen & A. Zwicky (Eds.), *Natural language processing: Psychological, computational, and theoretical perspectives* (pp. 129–189). Cambridge, UK: Cambridge University Press.

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

Freedman, S. E., & Forster, K. I. (1985). The psychological status of overgenerated sentences. *Cognition*, *19*, 101–131.

Fults, S., & Phillips, C. (2004, March). The source of syntactic illusions. Poster presented at the 17th Annual CUNY Conference on Human Sentence Processing. College Park, MD.

Garnsey, S. M., Tanenhaus, M. K., & Chapman, R. M. (1989). Evoked potentials and the study of sentence comprehension. *Journal of Psycholinguistic Research*, *18*, 51–60.

Gibson, E., & Thomas, J. (1999). Memory limitations and structural forgetting: The perception of complex ungrammatical sentences as grammatical. *Language and Cognitive Processes*, *14*, 225–248.

Gordon, P. C., Hendrick, R., & Johnson, M. (2001). Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1411–1423.

Grodzinsky, Y., & Reinhart, T. (1993). The innateness of binding and coreference. *Linguistic Inquiry*, *24*, 69–101.

Harris, C. L., & Bates, E. A. (2002). Clausal backgrounding and pronominal reference: A functionalist approach to c-command. *Language and Cognitive Processes*, *17*, 237–269.

Harrison, A., Hartsuiker, R., Branigan, H., & Pickering, M. (2004, March). Agreement processing in a complex number system. Poster presented at the 17th Annual CUNY Conference on Human Sentence Processing. College Park, MD.

Häussler, J., & Bader, M. (submitted). Agreement checking and number attraction in sentence comprehension: Insights from German relative clauses.

Heim, I. (1992). Anaphora and semantic interpretation: A reinterpretation of Reinhart's approach. In: U. Sauerland & O. Percus (Eds.), *The interpretive tract* (pp. 205–246). Cambridge, MA: MIT Working Papers in Linguistics.

Higginbotham, J. (1983). Logical form, binding, and nominals. *Linguistic Inquiry*, *14*, 547–593.

Hoeks, J. C. J., Stowe, L. A., & Doedens, G. (2004). Seeing words in context: The interaction of lexical and sentence level information during reading. *Cognitive Brain Research*, *19*, 59–73.

Hofmeister, P., & Sag, I. (2010). Cognitive constraints and island effects. *Language*, *86*, 366–415.

Huang, Y. (2004). Anaphora and the pragmatics-syntax interface. In: L. R. Horn & G. S. Ward (Eds.), *The handbook of pragmatics* (pp. 288–314). Oxford: Blackwell.

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

Jonides, J., Lewis, R. L., Nee, D. E., Lustig, C. A., Berman, M. G., & Moore, K. S. (2008). The mind and brain of short-term memory. *Annual Review of Psychology*, *15*, 1–32.

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

Kazanina, N. (2005). *The acquisition and processing of backwards anaphora*. PhD dissertation, University of Maryland.

Kazanina, N., Lau, E. F., Lieberman, M., Yoshida, M., & Phillips, C. (2007). The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language*, *56*, 384–409.

Kazanina, N., & Phillips, C. (2001). Coreference in child Russian: Distinguishing syntactic and discourse constraints. In: A. Do, L. Dominguez & A. Johansen (Eds.), *Proceedings of the 25th Annual Boston University Conference on Language Development* (pp. 413–424). Somerville, MA: Cascadilla Press.

Kazanina, N., & Phillips, C. (2010). Differential effects of constraints in the processing of Russian cataphora. *Quarterly Journal of Experimental Psychology*, *63*, 371–400.

Kempen, G. & Harbusch, K. (2002). Performance grammar: A declarative definition. In: A. Nijholt, M. Theune, & H. Hondorp (Eds.), *Computational linguistics in the Netherlands 2001* (pp. 148–162). Amsterdam: Rodopi.

Kennison, S. (2003). Comprehending the pronouns *her*, *him*, and *his*: Implications for theories of referential processing. *Journal of Memory and Language*, *49*, 335–352.

Kim, A., & Osterhout, L. (2005). The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language*, *52*, 205–225.

Kluender, R. (2005). Are subject islands subject to a processing account?. In: V. Chand, A. Kelleher, A. J. Rodriguez & B. Schmeiser (Eds.), *Proceedings of the 23rd West Coast Conference on Formal Linguistics* (pp. 475–499). Somerville, MA: Cascadilla Press.

Kluender, R., & Kutas, M. (1993). Subjacency as a processing phenomenon. *Language and Cognitive Processes*, *8*, 573–633.

Knuth, D. E. (1965). On the translation of languages from left to right. *Information and Control*, *8*, 607–639.

Koster, J., & Reuland, E. (1992). *Long-distance anaphora*. Cambridge: Cambridge University Press.

Krifka, M. (1990). Four thousand ships passed through the lock: Object-induced measure functions on events. *Linguistics and Philosophy*, *13*, 487–520.

Kuperberg, G. K. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research*, *1146*, 23–49.

Kuperberg, G. R., Sitnikova, T., Caplan, D., & Holcomb, P. J. (2003). Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research*, *17*, 117–129.

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

Ladusaw, W. A. (1992). Expressing negation. In: C. Barker & D. Dowty (Eds.), *Proceedings of the 2nd conference on Semantics and Linguistic Theory* (pp. 237–259.). Ithaca, NY: Cornell Linguistics Circle.

Lau, E. F., Stroud, C., Plesch, S., & Phillips, C. (2006). The role of structural prediction in rapid syntactic analysis. *Brain and Language*, **98**, 74–88.

Lee, M.-W. (2004). Another look at the role of empty categories in sentence processing (and grammar). *Journal of Psycholinguistic Research*, **33**, 51–73.

Lee, M.-W., & Williams, J. (2008). *The role of grammatical constraints in intra-sentential pronoun resolution*. Manuscript, London Metropolitan University and University of Cambridge.

Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, **29**, 375–419.

Lewis, R. L., Vasishth, S., & Van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Sciences*, **10**, 447–454.

Lukyanenko, C., Conroy, A., & Lidz, J. (submitted). Infants' adherence to Principle C: Evidence from 30-month olds.

McElree, B. (2006). Accessing recent events. In: B. H. Ross (Ed.), *The psychology of learning and motivation* (46, pp. 155–200). San Diego, CA: Academic Press.

McElree, B., Foraker, S., & Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, **48**, 67–91.

McElree, B., & Griffith, T. (1998). Structural and lexical constraints on filling gaps during sentence comprehension: A time-course analysis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 432–460.

McKinnon, R., & Osterhout, L. (1996). Constraints on movement phenomena in sentence processing: Evidence from event-related potentials. *Language and Cognitive Processes*, **11**, 495–823.

Montalbetti, M. (1984). *After binding*. Doctoral dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Neville, H., Nicol, J. L., Barss, A., Forster, K. I., & Garrett, M. F. (1991). Syntactically based sentence processing classes: Evidence from event-related brain potentials. *Journal of Cognitive Neuroscience*, **3**, 151–165.

Nicol, J., & Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research*, **18**, 5–19.

O'Grady, W. (2005). *Syntactic carpentry: An emergentist approach to syntax*. New York, NY: Routledge.

Omaki, A., Lau, E. F., Davidson White, I., & Phillips, C. (submitted). Hyper-active gap filling: Pre-verbal object gap creation in English filler-gap dependency processing.

Omaki, A., & Schulz, B. (In press). Filler-gap dependencies and island constraints in second language sentence processing. *Studies in Second Language Acquisition*.

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

Phillips, C. (1996). *Order and structure*. PhD dissertation, Massachusetts Institute of Technology, Cambridge, MA.

Phillips, C. (2003). Linear order and constituency. *Linguistic Inquiry*, **34**.

Phillips, C. (2004). Linguistics and linking problems. In: M. Rice & S. Warren (Eds.), *Developmental language disorders: From phenotypes to etiologies* (pp. 241–287). Mahwah, NJ: Erlbaum.

Phillips, C. (2006). The real-time status of island phenomena. *Language*, **82**, 795–823.

Phillips, C., & Lewis, S. (to appear). Derivational order in syntax: Evidence and architectural consequences. In C. Chesi (Ed.), *Directions in derivations*. Waltham, MA: Academic Press.

Phillips, C., & Wagers, M. (2007). Relating structure and time in linguistics and psycholinguistics. In: M. G. Gaskell (Ed.), *Oxford handbook of psycholinguistics* (pp. 739–756). Oxford: Oxford University Press.

Pickering, M., Barton, J. S., & Shillcock, R. (1994). Unbounded dependencies, island constraints, and processing complexity. In: C. Clifton, , L. Frazier & K. Rayner (Eds.), *Perspectives on sentence processing* (pp. 199–224). London: Lawrence Erlbaum.

Pollard, C., & Sag, I. A. (1992). Anaphors in English and the scope of the binding theory. *Linguistic Inquiry*, *23*, 261–303.

Pritchett, B. (1991). Subjacency in a principle-based parser. In: R. C. Berwick (Ed.), *Principle-based parsing: Computation and psycholinguistics* (pp. 301–345). Dordrecht: Kluwer.

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

Reinhart, T. (1983). *Anaphora and semantic interpretation*. London: Croom Helm.

Reinhart, T. (2006). *Interface strategies*. Cambridge, MA: MIT Press.

Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, *24*, 261–321.

Ricker, T. J., AuBuchon, A. M., & Cowan, N. (2010). Working memory. *Wiley Interdisciplinary Reviews in Cognitive Science*, *1*, 573–585.

Ross, J. R. (1967). *Constraints on variables in syntax*. Doctoral dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Runner, J. T., Sussman, R. S., & Tanenhaus, M. K. (2006). Processing reflexives and pronouns in picture noun phrases. *Cognitive Science*, *30*, 193–241.

Solomon, E., & Pearlmutter, N. (2004). Semantic integration and syntactic planning in language production. *Cognitive Psychology*, *49*.

Sprouse, J., Wagers, M, & Phillips, C. (to appear). A test of the relation between working memory capacity and syntactic island effects. *Language*.

Steedman, M. (2000). *The syntactic process*. Cambridge, MA: MIT Press.

Stepanov, A. (2007). The end of CED? Minimalism and extraction domains. *Syntax*, *10*, 80–126.

Stowe, L. A. (1986). Parsing WH-constructions: Evidence for on-line gap location. *Language and Cognitive Processes*, *3*, 227–245.

Stroud, C. (2008). *Structural and semantic selectivity in the electrophysiology of sentence comprehension*. Unpublished doctoral dissertation. University of Maryland, College Park, MD.

Stroud, C., & Phillips, C. (2011). Examining the evidence for an independent semantic analyzer: An ERP study in Spanish. *Brain and Language* [epub ahead of print]. DOI: 10.1016/j.bandl.2011.02.001.

Sturt, P. (2003). The time-course of the application of binding constraints in reference resolution. *Journal of Memory and Language*, *48*, 542–562.

Tabor, W., Galantucci, B., & Richardson, D. (2004). Effects of merely local coherence on sentence processing. *Journal of Memory and Language*, *50*, 355–370.

Townsend, D. J., & Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. Cambridge, MA: MIT Press.

Traxler, M. J., & Pickering, M. J. (1996). Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language*, *35*, 454–475.

Van Dyke, J. A., & McElree, B. (2006). Retrieval interference in sentence comprehension. *Journal of Memory and Language*, *55*, 157–166.

van Gompel, R. P. G., & Liversedge, S. P. (2003). The influence of morphological information on cataphoric pronoun assignment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 128–139.

van Herten, M., Chwilla, D. J., & Kolk, H. H. J. (2006). When heuristics clash with parsing routines: ERP evidence for conflict monitoring in sentence perception. *Journal of Cognitive Neuroscience*, *18*, 1181–1197.

Vasishth, S., Brüssow, S., Lewis, R., & Drenhaus, H. (2008). Processing polarity: How the ungrammatical intrudes on the grammatical. *Cognitive Science*, *32*, 685–712.

Vigliocco, G., & Nicol, J. (1998). Separating hierarchical relations and word order in language production. Is proximity concord syntactic or linear?. *Cognition*, *68*, 13–29.

Wagers, M. (2008). *The structure of memory meets memory for structure in linguistic cognition.* Doctoral dissertation. University of Maryland, College Park, MD.

Wagers, M. W., Lau, E. F., Stroud, C., & Phillips, C. (2008, March). Agreement and the subject of confusion. Talk presented at the 21st Annual CUNY Conference on Human Sentence Processing. Chapel Hill, NC.

Wagers, M. W., Lau, E. F., & Phillips, C. (2009). Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language*, *61*, 206–237.

Wagers, M. W., & Phillips, C. (2009). Multiple dependencies and the role of the grammar in real-time comprehension. *Journal of Linguistics*, *45*, 395–433.

Wellwood, A., Pancheva, R., Hacquard, V., Fults, S. W., & Phillips, C. (2009, March). The role of event comparison in comparative illusions. Poster presented at the 22nd Annual CUNY Conference on Human Sentence Processing. Davis, CA.

Xiang, M., Dillon, B., and Phillips, C. (2006, March). Testing the strength of the spurious licensing effect for negative polarity items. Talk at the 19th Annual CUNY Conference on Human Sentence Processing. New York, NY.

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

Ye, Z., & Zhou, X. (2008). Involvement from cognitive control in sentence comprehension: Evidence from ERPs. *Brain Research*, *1203*, 103–115.

Yoshida, M. (2006). *Constraints and mechanisms in long-distance dependency formation.* Doctoral dissertation. University of Maryland, College Park, MD.

Zukowski, A., McKeown, R., & Larsen, J. (2008). A tough test of the locality requirement for reflexives. In: H. Chan, H. Jacob & E. Kapia (Eds.), *Proceedings of the 32nd annual Boston University Conference on Language Development* (pp. 586–597). Somerville, MA: Cascadilla Press.