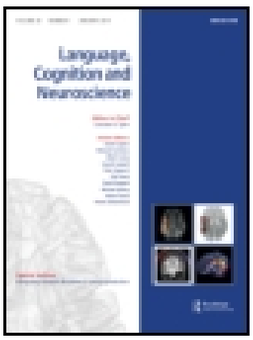

Routledge
Taylor & Francis Group

Language, Cognition and Neuroscience

ISSN: 2327-3798 (Print) 2327-3801 (Online) Journal homepage: http://www.tandfonline.com/loi/plcp21

Grammatical predictions reveal influences of semantic attraction in online sentence comprehension: evidence from speeded forced-choice sentence continuations

Les Sikos, Cecily Jill Duffield & Albert E. Kim

To cite this article: Les Sikos, Cecily Jill Duffield & Albert E. Kim (2016): Grammatical predictions reveal influences of semantic attraction in online sentence comprehension: evidence from speeded forced-choice sentence continuations, Language, Cognition and Neuroscience, DOI: 10.1080/23273798.2016.1186808

To link to this article: http://dx.doi.org/10.1080/23273798.2016.1186808

LANGUAGE, COGNITION AND NEUROSCIENCE, 2016
http://dx.doi.org/10.1080/23273798.2016.1186808

Routledge
Taylor & Francis Group

Grammatical predictions reveal influences of semantic attraction in online sentence comprehension: evidence from speeded forced-choice sentence continuations

Les Sikos[a,b], Cecily Jill Duffield[a,b] and Albert E. Kim[a,c]

[a]Institute of Cognitive Science, University of Colorado, Boulder, CO, USA; [b]Department of Linguistics, University of Colorado, Boulder, CO, USA; [c]Department of Psychology and Neuroscience, University of Colorado, Boulder, CO, USA

ABSTRACT

Research in language processing has established that semantic information can influence (a) online sentence interpretation when syntactic cues are indeterminate, and (b) offline judgments of syntactically unambiguous but semantically anomalous sentences. The question of whether semantic information can influence *online* sentence interpretation when semantic cues oppose *unambiguous* syntax remains unanswered. We investigate this question using two speeded forced-choice sentence continuation studies. Under time constraints, participants continued sentence fragments such as "The hearty meal was devouring … " (Attraction Violation) and "The sealed envelope was devouring … " (No-Attraction Violation) with either BY or THE. Although participants had an overall preference to continue all fragments with THE, they were more likely to choose BY in Attraction Violation than No-Attraction Violation sentences. When choosing THE, participants took longer in the Attraction Violation condition. These results suggest that semantically attractive interpretations can be pursued online even in the face of unambiguous but contradictory syntactic cues.

ARTICLE HISTORY
Received 21 September 2014
Accepted 29 April 2016

KEYWORDS
Sentence processing; syntax; semantics; semantic attraction; semantic P600

A central question in the psychology of language concerns the relative contributions of syntactic and semantic knowledge during online comprehension. Language is characterised by rule-like regularities, and it is clear that such surface-level syntax provides crucial information for sentence interpretation. This is illustrated by the ability of grammatical cues to fundamentally shape interpretations in sentences like "The boy ate the apple" and "The apple ate the boy". In these examples, the syntactic structure drives a crucial semantic decision about who is doing what to whom in the final interpretation of the sentence (i.e. the transitive construction maps the subject of the sentence to the agent of the event). It is also clear that semantic knowledge can render some interpretations more plausible than others, independently of the syntax (e.g. boys eat apples, but apples do not typically eat boys). Over the years there has been substantial debate over the role of syntactic and semantic knowledge in the control of real-time language processing. Here we investigate the degree to which online compositional interpretations are guided by semantic cues.

Previous work

Influential older models of real-time language processing argued that syntax controls and/or precedes the integration of semantic information, and have produced evidence that syntactic knowledge guides online processing commitments (e.g. Ferreira & Clifton, 1986; Frazier & Fodor, 1978) and sentence interpretations (Christianson, Hollingworth, Halliwell, & Ferreira, 2001), even when semantic cues might seem to point to an alternate interpretation. Such models were grounded on the assumption that syntactic processing is fast, automatic, and capable of establishing a quick initial structuring of the input, and that this speed and automaticity was enabled in part by modular encapsulation that prevented semantic influences on the initial commitments (Fodor, 1983).

However, a number of studies have shown that comprehenders can use semantic information to guide interpretation when syntactic cues are indeterminate. For instance, sentences like "The witness examined by the judge turned out to be unreliable" contain a temporary syntactic ambiguity between a main clause (e.g. the witness did the examining) and a reduced relative clause (e.g. the witness was examined). This ambiguity is resolved at the word *by*, where the sentence becomes anomalous if the comprehender has committed to the incorrect main clause parse. Online measures show that such sentences are harder to

CONTACT Les Sikos sikos@coli.uni-saarland.de

process when semantic cues support the incorrect parse, as above, than when semantic cues support the correct parse (e.g. Tabossi, Spivey-Knowlton, McRae, & Tanenhaus, 1994; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995; Trueswell, Tanenhaus, & Garnsey, 1994). For example, in "The **evidence** examined by the judge turned out to be unreliable", knowledge that *evidence* cannot examine things facilitates a reduced relative clause interpretation. This suggests that when syntactic cues are ambiguous, and unable to clearly assign thematic roles, semantic information is used to do so. More recently, eye-movement data have indicated that comprehenders rapidly fixate objects in a scene upon hearing a verb that predicts those objects (e.g. The boy will eat …) prior to actually hearing the noun phrase that unambiguously refers to the object ("the cake"; e.g. Kamide, Altmann, & Haywood, 2003). Such evidence suggests that semantic properties of the current word can lead to predictions about semantic and syntactic features of upcoming words, before they have been encountered. Overall, the findings mentioned above show that semantic information can influence real-time language processing when syntactic cues are indeterminate. However, they do not indicate what influence is exerted by semantic cues in the presence of unambiguous syntactic cues. Because the semantic interpretations in these studies are always compatible with a plausible syntactic parse, these studies do not directly pit semantic cues against unambiguous syntactic structures. By and large, unambiguous syntactic cues are assumed to drive interpretations.

There is, however, suggestive evidence from offline judgments that semantic cues can influence interpretation even when opposed by unambiguous syntactic cues. Ferreira (2003) asked participants to listen to syntactically unambiguous sentences containing thematic role reversals like "The dog was bitten by the man" and then to answer questions about who did what to whom in the sentence. Participants tended to make errors such as identifying the dog as the agent of the biting action, suggesting that they had assigned thematic roles in the most semantically plausible manner, rather than following the syntactic cues. Ferreira (2003) concluded that we often rely on heuristics to yield a "good enough" representation of an utterance's meaning, rather than necessarily building a representation that is syntactically faithful to the input (see also Ferreira & Patson, 2007). Similar findings and proposals have, in fact, existed for a long time. Work on patients with aphasia has indicated that patients with syntactic deficits may sometimes interpret such sentences using semantic heuristics. For instance, under some circumstances patients may interpret sentences like "The apple ate the boy" by reversing the thematic roles assigned by the unambiguous syntax in order to achieve a plausible interpretation (Caramazza & Zurif, 1976; Saffran, Schwartz, & Linebarger, 1998). In typical adults and children, observers take longer to verify the truth value of pictures matched to reversible sentences (e.g. "The dog is chasing the cat") than to non-reversible sentences (e.g. "The girl is watering the flowers") (Slobin, 1966). A key limitation of the evidence described above is that it has involved judgments made after sentence completion and/or patients with language deficits, which limits the evidence's ability to illuminate real-time language processing computations in neurologically typical brains.

Some controversial support for the idea that online interpretive commitments can be controlled by semantic cues even in the face of opposition from syntactic cues comes from studies using event-related brain potentials (ERPs). For instance, Kim and Osterhout (2005) found that sentences like "The hearty meal was devouring … " elicited P600 effects rather than N400 effects at the verb, relative to controls. This result differed from previous ERP work in two ways. First, semantic processing difficulty is known to modulate the N400 component (Kutas & Hillyard, 1980, 1984), yet no N400 enhancement was found. Second, P600 effects are reliably elicited by syntactic processing difficulty (Hagoort, Brown, & Groothusen, 1993; Osterhout, Holcomb, & Swinney, 1994), but the syntactic cues in such sentences are well formed. Kim and Osterhout (2005) argued that online processing of the verb was influenced by a "semantic attraction" to a highly plausible but ungrammatical interpretation—that is, the subject noun (*meal*) is assigned to the theme role of the verb (*devour*). On this account, the semantically attractive interpretation directly opposes the surface syntactic cues, causing the unambiguous and well-formed syntax to appear anomalous (i.e. *devour**ing*** should be *devour**ed***, as in the passive sentence "The hearty meal was devour**ed** … "). The P600 was attributed to structural reprocessing (i.e. syntactic repair; Kim & Osterhout, 2005; Kim & Sikos, 2011) in response to the perceived syntactic anomaly. This conclusion was supported by the results from anomalous sentences that contained no semantic attraction, as in "The sealed envelope was devouring … ", which elicited N400 effects (and no P600 effects) relative to controls (Kim & Osterhout, 2005). These and several similar effect patterns have been labelled "semantic P600" effects (e.g. Hoeks, Stowe, & Doedens, 2004; Kim & Sikos, 2011; Kolk, Chwilla, Van Herten, & Oor, 2003; Kuperberg, Sitnikova, Caplan, & Holcomb, 2003).

Note, however, that some other semantic P600 results have appeared incompatible with the semantic

attraction account described above. Specifically, P600 effects have been observed for semantically anomalous sentences that do *not* contain thematic role reversals (Chow & Phillips, 2013; Hoeks et al., 2004; Kuperberg et al., 2003; Kuperberg, Kreher, Sitnikova, Caplan, & Holcomb, 2007; Paczynski & Kuperberg, 2011; Stroud & Phillips, 2012; Van Herten, Chwilla, & Kolk, 2006). For example, Kuperberg et al. (2007) found an increased P600 effect at the verb for sentences like "Every morning at breakfast the eggs would plant … ", relative to controls like "Every morning at breakfast the boys would eat … ", despite the fact that *eggs* does not have a strong semantic attraction to the theme role of the verb *plant*. Accordingly, several alternate explanations for semantic P600 effects have been proposed. In general, these explanations suggest that semantic P600 effects are primarily triggered by the detection of a highly implausible interpretation (e.g. animacy violation). Furthermore, they make no distinction between situations that involve semantically plausible alternatives to the grammatically supported interpretation (i.e. semantic attraction) and sentences that are simply semantically anomalous (Chow & Phillips, 2013; Kuperberg, 2007; Paczynski & Kuperberg, 2012; Stroud & Phillips, 2012). For the purposes of this paper, we will refer to these alternate explanations as the "severe anomaly" account, collapsing across certain differences between these specific proposals.

Because the semantic attraction and severe anomaly views of semantic P600 effects entail different assumptions about the functional demands imposed by semantically anomalous verbs, the existing ERP evidence does not clearly support the hypothesis that semantic information can overrule unambiguous syntax during online sentence interpretation. Without consensus on the functional antecedent conditions of the ERPs (or on the possibility that different semantic P600 effects reflect distinct functional scenarios, requiring multiple accounts), it is difficult to draw clear theoretical inferences based on the ERP results. Our understanding of the information processing dynamics at work in syntax-semantics conflict situations would be clearer if we were somehow able to tap into participants' actual interpretations precisely at the point of conflict between cues. As described in the next section, the goal of the current study is to fill this gap by extracting information about comprehenders' moment-by-moment interpretations, using a method that can bring potentially converging evidence to bear on the same question addressed in the ERP studies described above (e.g. Kim & Osterhout, 2005; Kim & Sikos, 2011).

To summarise, since the 1980s a series of findings have pushed theories of language processing to acknowledge a greater and greater role for semantics during comprehension. Many modern theories now acknowledge that semantic information can influence both online sentence interpretation when syntactic cues are indeterminate, as well as offline judgments of syntactically unambiguous but semantically anomalous sentences. However, these newer models have not made explicit predictions for whether semantic information can *overrule* unambiguous syntax during *online* sentence interpretation. Although the results of several ERP studies have led some researchers to conclude in the affirmative, their interpretations of the data have been controversial, limiting the conclusions we can draw.

Current study

The current study conducted two behavioural experiments to answer the unresolved question of whether semantic cues can overrule directly conflicting, unambiguous syntactic cues during online sentence comprehension. We will refer to this possibility as the semantic attraction hypothesis.

We tested the semantic attraction hypothesis using a speeded forced-choice sentence continuation task modified from Staub (2009). Participants were asked to read sentence fragments like "The hearty meal was devouring … " (Attraction Violation condition) and "The sealed envelope was devouring … " (No-Attraction Violation condition) and then to choose as quickly as possible which of two words (THE or BY) provided the best continuation of each fragment. Responses were made via button presses. The dependent measures were the proportions of THE and BY responses and the time required to make those responses.

Unlike studies investigating garden-path sentences, in which two competing interpretations are both compatible (temporarily) with the syntactic cues, the use of semantically implausible sentence fragments in the current study allowed us to pit semantic cues directly against unambiguous syntactic structure. The syntactic cues in both the Attraction Violation and No-Attraction Violation conditions clearly indicated an active voice construction, which would typically contain a noun phrase following the verb. This design allowed us to test the semantic attraction hypothesis at three different levels. First, if semantic cues play little to no role in participants' online interpretations of these fragments (as predicted by accounts which assume, either implicitly or explicitly, that unambiguous syntactic cues dominate online commitments), then participants should anticipate an active construction continuation. Thus, they should be more likely to select the continuation THE (which is compatible with a noun phrase like "the cookies."), rather than BY

(which is compatible with a passive interpretation like "by the boy."). Moreover, on such a syntax-dominant account the likelihood of selecting THE in either of the Violation conditions should not differ significantly from the Active Control condition. On the other hand, if participants are sensitive to semantic as well as syntactic cues during online interpretation, then we should find that the Violation conditions differ from the Active Control condition.

Second, if we find that participants are indeed sensitive to semantic information during online processing of these fragments, then we can also assess the degree to which semantics influences processing by comparing results in the Attraction Violation and No-Attraction Violation conditions. According to the severe anomaly account, participants are sensitive to highly implausible interpretations, but do not distinguish between situations that involve semantically plausible alternatives to the grammatically supported interpretation and sentences that are simply semantically anomalous (cf. Chow & Phillips, 2013; Kuperberg, 2007; Paczynski & Kuperberg, 2012; Stroud & Phillips, 2012). The severe anomaly account therefore predicts no significant difference between the likelihood of BY responses in the Attraction and No-Attraction Violation conditions. The semantic attraction hypothesis, on the other hand, predicts that the semantic cues in Attraction Violation fragments (e.g. that *hearty meals* are excellent themes for *devouring*) should influence participants' online interpretations towards a passive construction (as in "The hearty meal was devoured by … "). This outcome should increase the likelihood of BY responses. Critically, because No-Attraction Violation fragments do not contain similar semantic support for a passive construction (*envelopes* are not good themes for *devouring*), the semantic attraction hypothesis predicts important differences in the proportion of BY responses across the Violation conditions.

Third, if we do find evidence in support of the semantic attraction hypothesis over the severe anomaly account, we can also test strong and weak versions of the semantic attraction hypothesis. A strong version assumes that semantic attraction is the primary constraint in interpreting such sentences and therefore predicts that participants should strongly prefer a passive continuation for Attraction Violation fragments, leading to BY-responses on a majority of such trials.[1] A weaker version of the semantic attraction account allows other constraints to contribute to sentence interpretation and predicts that participants should prefer the passive BY continuation more often for Attraction Violation than No-Attraction Violation fragments, even if BY-continuations did not constitute a majority for either condition.

It is important to note that BY-responses in this task do not necessarily reflect a commitment to a passive structure. Sentence fragments in both the Attraction Violation and No-Attraction Violation conditions can be grammatically continued with BY, using a temporal or locative modifier (e.g. "The hearty meal was devouring … *by 12 o'clock*", "The sealed envelope was devouring … *by the window*"). However, any such BY-continuations would be semantically anomalous, in both conditions, making such interpretations relatively unlikely. Thus, there is no reason to predict that the proportion of BY-responses that are based on such temporal or locative interpretations should be affected by the presence or absence of semantic attraction. In contrast, BY-responses driven by a passive interpretation could show important differences in response preferences across conditions. In the Attraction Violation condition, a BY-response driven by a passive interpretation provides a resolution to the conflict between syntactic and semantic cues (i.e. the semantic cues override the syntactic cues). In the No-Attraction Violation condition, however, a BY-response driven by a passive interpretation would still result in a semantic anomaly.

The semantic attraction account also generates predictions for reaction times. Specifically, when choosing THE in both Attraction Violation and No-Attraction Violation conditions, participants should take longer to do so in the Attraction Violation condition. The critical assumption here is that when subjects select the syntactically licensed THE-continuation, this response should face competition from the semantically plausible passive construction, slowing down the eventual selection of THE. Thus, even in cases where participants' choices do not indicate a sensitivity to differences in semantic cues between the Attraction Violation and No-Attraction Violation conditions (and would seemingly indicate a preference for a syntactically driven interpretation), response times might reveal such sensitivity.

We did not analyse the difference in BY-response times. This is because characterising the nature of competition underlying BY-responses is much less tractable than for THE-responses. As noted above, a BY-response can be the result of multiple, qualitatively different interpretations (e.g. passive, adverbial). Consequently, any differences in BY-response times across conditions could be due to competition between multiple interpretations that lead to a BY response, as well as due to competition from the syntactically licensed THE-response. Thus, in contrast to THE-responses, the semantic attraction hypothesis cannot make clear predictions for BY-response time differences.

In summary, the study presented here tests the hypothesis that semantic cues can determine online

sentence interpretations even in the face of unambiguous and contradictory syntactic information. We selected the speeded forced-choice sentence continuation task because it combines two key characteristics. First, the measure includes a categorical judgement whose linking relationship to a specific functional outcome (active or passive interpretation) is straightforward. Although it is possible that this linking hypothesis could be incorrect, its logical simplicity reduces the functional ambiguity that is sometimes present in measures like EEG and eye-tracking. Second, the task probes the state of the language processing system at the critical point of syntax-semantics conflict, thus providing an online measure—in contrast to off-line judgments of sensicality, acceptability, or plausibility, which probe at a substantial delay after this critical point (e.g. Ferreira, 2003; Ferreira & Patson, 2007).

Experiment 1

Experiment 1 tested the semantic attraction hypothesis (i.e. that semantic cues can overrule directly conflicting, unambiguous syntactic cues during online sentence comprehension), by assessing the implicit predictions that manifest immediately following an anomalous word.

Methods

Participants

Thirty-seven undergraduate students at the University of Colorado between the ages of 18 and 29 participated in Experiment 1. All participants were native speakers of English, right-handed, and all received course credit for their participation. Data from one participant were omitted due to technical error.

Materials

Ninety-six experimental items were created by modifying stimuli from Kim and Osterhout (2005). Each item had four conditions (Attraction Violation, No-Attraction Violation, Passive Control, Active Control). Examples are given in Table 1 (upper panel). Attraction Violation and No-Attraction Violation conditions paired the same verb with a different subject noun phrase in the pre-response context. Although both contexts contained inanimate subjects that were extremely unlikely agents of the verb, Attraction Violation subjects were highly plausible themes of the verb, whereas No-Attraction Violation subjects were implausible themes. Thus, Attraction Violations contained conflicting cues: semantic cues were consistent with a passive construction, while syntactic cues were consistent with an active construction. Subject noun phrases were counterbalanced across conditions by swapping the subject–verb pairings across items.

Each experimental sentence was created with two versions of the material following the participant's forced choice (Table 1A), such that one post-response sentence completion was compatible with each of the possible choices (THE or BY). THE-completions expressed direct objects consistent with an active frame, and BY-completions expressed agents as appropriate with a passive frame. The version presented during the experiment was contingent on the participant's actual response.[2]

Passive and Active Controls were created for each item by modifying the syntactic or semantic features of the items in the Attraction Violation condition to eliminate the conflict between syntactic and semantic cues (Table 1A). Passive Controls modulated the syntactic cues by modifying the verb inflection (i.e. *–ing* to *–ed*) to provide an unambiguous passive construction. BY-completions for Passive Controls expressed an animate agent as appropriate with a passive frame, while THE-completions were temporal adverbials. For Active Controls, the semantic cues of the Attraction Violation items were modified by replacing the subjects with prototypical animate agents, BY-completions were locative or manner adverbials, and THE-completions contained direct objects.

Ninety-six filler sentences of four types were created: THE-preference, BY-preference, Syntactic Anomaly, and

Table 1. Example material for Experiment 1.

A.	Experimental conditions	Pre-response context	Forced choice	Post-response completion (THE-completion/BY-completion)
	Attraction Violation	The hearty meal was devouring	THE | BY	plate of cookies/a group of energetic children
	No-Attraction Violation	The dusty tabletops were devouring	THE | BY	plate of cookies/a group of energetic children
	Passive Control	The hearty meal was devoured	THE | BY	second it was placed before them/a group of energetic children
	Active Control	The hungry boys were devouring	THE | BY	plate of cookies/mouthfuls
B.	Filler conditions	Pre-response context	Forced choice	Post-response completion (THE-completion/BY-completion)
	THE-preference	The seamstress was mending socks for	THE | BY	family/her sister
	BY-preference	The compost pile was decomposing	THE | BY	food scraps/a raised flower bed
	Syn Violation (passive)	The products were imported by a large	THE | BY	corporation/ship
	Syn Violation (active)	His father was reading yesterday's	THE | BY	newspaper/his desk
	Ambiguous	Our best employee was	THE | BY	last to retire/a water cooler

Ambiguous (Table 1B). THE-preference sentences presented the THE/BY choice immediately following a preposition, making BY-completions incompatible with the sentence. BY-preference sentences all contained intransitive verbs, making THE-completions dispreferred. Syntactic Violations presented the THE/BY choice immediately following a determiner or adjective, making either choice anomalous. Half of these sentences were in an active frame and half were in a passive frame. Finally, ambiguous sentences contained a subject and a copula, making either BY or THE a compatible choice.

Experimental items were distributed across four lists such that each list contained 24 sentences from each experimental condition but no list contained more than one condition from any single item. Experimental items were pseudo-randomly mixed with filler sentences such that no more than two items of any condition occurred in a row, and items from each condition were distributed equally across both halves of the list. Each list contained 192 sentences in total, with 62.5% well formed and 37.5% anomalous. These four lists were then reversed to create a total of eight lists.

Procedure

Participants were randomly assigned to a stimulus list and were seated in individual testing rooms in front of a PC-compatible computer running E-prime experimental software (Schneider, Eschman, & Zuccolotto, 2002). Participants were instructed to read sentences that would appear word-by-word on the screen, up until the point when a prompt would appear with the choices THE or BY. Participants were instructed to make the choice that would best continue the sentence as quickly and naturally as possible. After the choice was made, the corresponding post-response completion was presented. Participants were told that although most sentences would be normal sentences of English, some would contain unexpected words or impossible situations that may not make sense. In such cases they were instructed to choose the continuation that felt most natural. Each session began with a series of practice trials. Participants were required to respond correctly to at least six practice trials before continuing on to the experimental trials. The experimental trials were divided into 6 blocks of 32 trials each, with a break between each block. Sessions lasted approximately 30 minutes.

Trials consisted of the following sequence of events (Figure 1): trials were initiated by pressing the space bar. A fixation cross then appeared in the centre of the screen for 1000 ms, followed by a blank-screen interval of 100 ms. The pre-response context was then presented in the centre of the screen via rapid serial visual presentation (RSVP). Each word appeared for 380 ms, followed by a variable inter-stimulus interval that was proportional to word length (20 ms/character). At the point of the forced-choice task, "THE | BY" appeared in the centre of the screen, with THE always appearing on the left in order to avoid any potential confusion with the highly frequent collocation "by the". The choices remained until the participant selected one of the two words by pressing either the F key (corresponding to THE) or the J key (corresponding to BY). If response times were longer than 1200 ms, the words "TOO SLOW" appeared on the screen after the selection had been made and the trial was excluded from analyses. Following the response, the appropriate post-response completion was presented using the same RSVP parameters as above, and a screen appeared that instructed participants to push the space bar to begin the next trial. A

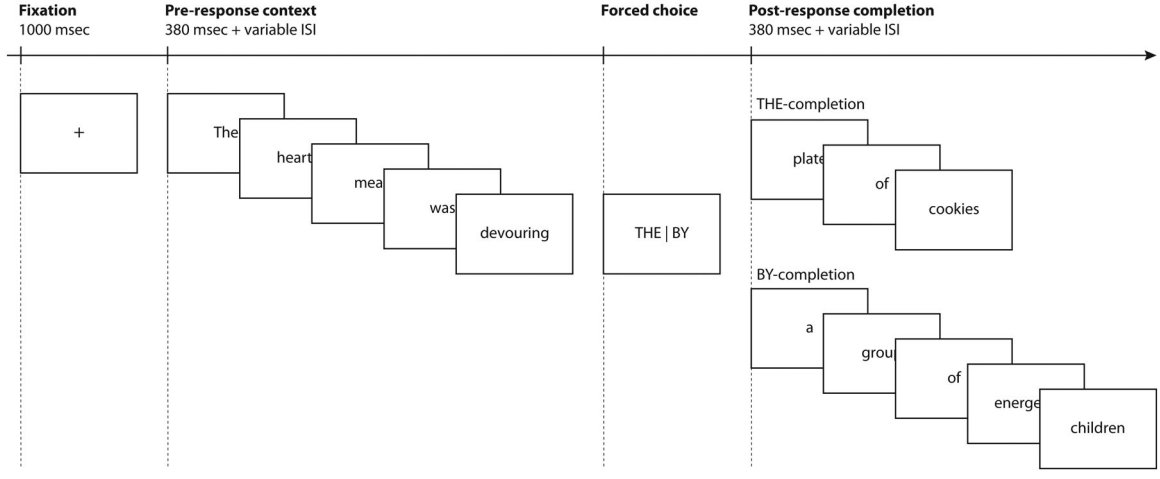

Figure 1. Trial procedure. Each trial consisted of a fixation prompt, a pre-response context, the forced-choice task, and a post-response sentence completion that was dependent on the participant's choice.

750 ms blank screen followed each trial. Response choice and reaction time were recorded for each trial.

Results and discussion

Our goal was to investigate whether semantic information can override directly conflicting, unambiguous syntactic information during online sentence interpretation, as reflected in choices about subsequent words and in reaction times to identical choices across conditions. The barplots in Figure 2 present the mean proportion of BY-responses per condition (left) and response times per condition for BY-responses (centre) and THE-responses (right). About 1.13% of experimental trials were excluded due to responses missing the 1200 ms deadline.

Here and in all subsequent analyses, response choices were analysed using generalised linear mixed-effects regression (GLMER) models with crossed random effects for subjects and items (Jaeger, 2008), and with condition as a fixed effect. These models included a logit link function and are therefore well suited for the analysis of categorical outcomes. We first present the results from the omnibus model containing all experimental conditions, and then present planned pairwise comparisons testing the individual predictions outlined above. Presentation order, response time, and interaction terms were included as fixed effects when their presence improved the fit of the model, as determined by likelihood ratio tests comparing models that differed only by the fixed effect of interest. When reaction time was included, statistical analyses were conducted using log-transformed RT to correct for skewness in the data; RTs were centred for the analysis. All models we describe here included the maximal random effects structure justified by the data based on model comparison (Jaeger, 2009).[3] Analyses were conducted using the *glmer* function (*lme4* library, version 0.999999-0; Bates & Sarkar, 2007) in the statistics software package R, version 2.15.2 (R Development Core Team, 2011). We report *p*-values as calculated from *z*-scores, Pr(>|*z*|).

Results for the omnibus model including all conditions (Attraction Violation, No-Attraction Violation, Passive Control, Active Control) are shown in Table 2A. We used treatment coding for condition with Active Control coded as the reference level (coded as 0). Thus the intercept represents the baseline log odds of participants choosing BY at the mean Response Time when Condition is coded as zero (i.e. when Condition = Active Control) and Presentation Order is coded as zero (i.e. prior to exposure to experimental stimuli). The fact that this estimate is significant and negative indicates that the likelihood of participants choosing THE was reliably above chance (i.e. the likelihood of choosing BY was significantly below chance) in the Active Control condition, prior to any influence they might have had due to repeated exposure to stimuli. The significant parameter estimates for each condition represent the reliable increase in likelihood (in log odds) of a BY-response for that condition relative to the Active Control. Comparing the magnitude of the estimates for each condition suggests that, contra the predictions of syntax-dominant accounts, the conditions with stronger semantic cues pointing to a passive interpretation are associated with a higher proportion of BY-responses (see Figure 2, left panel). The significant log(RT) estimate represents the increase in likelihood of a BY-response for each unit increase of log(RT). This finding indicates that across all conditions, participants were more likely to choose BY as they took longer to respond (see Figure 3). The significant positive

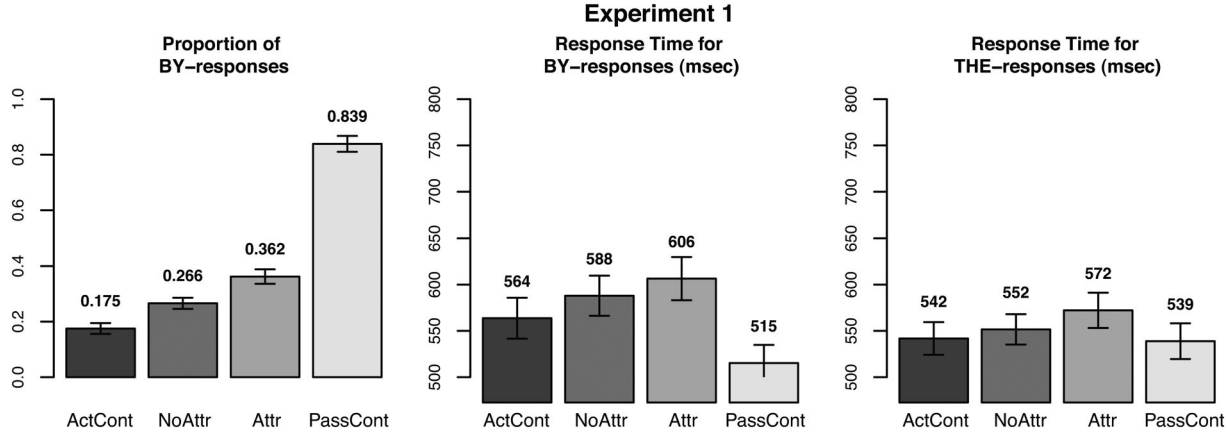

Figure 2. Experiment 1. *Left:* Proportion of BY-responses per condition: Active Control (ActCont), No-Attraction Violation (NoAttr), Attraction Violation (Attr), Passive Control (PassCont). *Centre:* Response times for BY-responses by condition. *Right:* Response times for THE-responses per condition. Error bars represent standard error of the mean.

Table 2. Experiment 1: Parameter estimates for fixed effects in GLMER and LMER models.

Parameter	Estimate	SE	Wald *z*	Pr > (|*z*|)
A. Response proportions for omnibus model containing all experimental conditions				
Intercept	−1.539	0.232	−6.638	<.001
Condition: No-Attraction Violation	0.698	0.246	2.835	<.01
Condition: Attraction Violation	1.085	0.247	4.391	<.001
Condition: Passive Control	3.172	0.401	7.905	<.001
Presentation Order	−0.002	0.002	−0.967	.333
log(RT)	0.531	0.152	3.487	<.001
No-Attraction × Presentation Order	−0.001	0.002	−0.432	.665
Attraction Violation × Presentation Order	−0.000	0.002	−0.092	.927
Passive Control × Presentation Order	0.007	0.003	2.799	<.01
B. Response proportions for planned comparison between Attraction Violation and Passive Control conditions				
Intercept	1.647	0.289	5.698	<.001
Condition: Attraction Violation	−2.057	0.350	−5.881	<.001
Presentation Order	0.005	0.002	2.595	<.01
Attraction Violation × Presentation Order	−0.007	0.002	−3.070	<.01
C. Response proportions for planned comparison between Attraction Violation and No-Attraction Violation conditions				
Intercept	−0.884	0.183	−4.818	<.001
Condition: Attraction Violation	0.435	0.119	3.650	<.001
log(RT)	0.848	0.248	3.424	<.001
Presentation Order	−0.003	0.001	−2.086	<.05
D. Response proportions for planned comparison between No-Attraction Violation and Active Control conditions				
Intercept	−1.538	0.202	−7.595	<.001
Condition: No-Attraction Violation	0.650	0.132	4.921	<.001
log(RT)	0.851	0.232	3.661	<.001
Presentation Order	−0.003	0.001	−2.125	<.05

Parameter	Estimate	SE	df	*t*-Value	*p*-Value
E. Response times for THE-responses for planned comparison between Attraction Violation and No-Attraction Violation conditions					
Intercept	581.100	17.924	1666	32.419	<.001
Condition: Attraction Violation	15.758	7.056	1666	2.233	<.05
Presentation Order	−0.201	0.098	1666	−2.045	<.05

parameter estimate for the Passive Control × Presentation Order interaction indicates a reliable increase in likelihood of participants to choose BY only in the Passive Control condition as the experiment progressed (see Figure 4).

The results for the omnibus model support an account of online interpretation in which participants are sensitive to semantic information, but it does not assess the degree to which participants are sensitive to semantic influences that contradict available syntactic cues. The semantic attraction hypothesis predicted that participants should be more likely to choose BY in the Attraction Violation condition than in the No-Attraction Violation condition as a result of participants

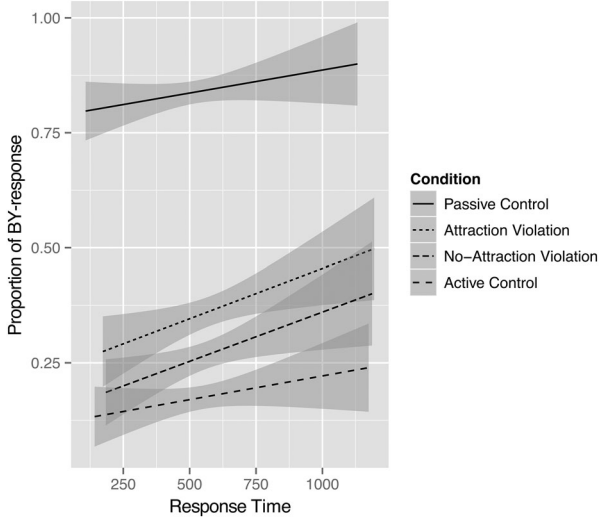

Figure 3. Experiment 1: Mean proportions of BY-responses by response times (ms) per condition. Shaded areas represent the 95% confidence interval around the smoothed continuous means.

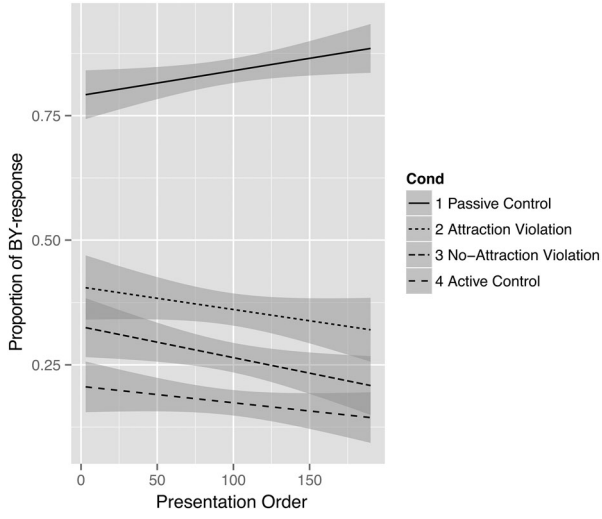

Figure 4. Experiment 1: Mean proportions of BY-responses by presentation order per condition. Shaded areas represent the 95% confidence interval around the smoothed continuous means.

distinguishing between fragments that have semantically plausible alternatives to the grammatically supported interpretation and sentences that are simply semantically anomalous. The severe anomaly account, on the other hand, predicted no difference in the proportion of BY responses between the two conditions. To test these predictions, we fit a GLMER model comparing response choices in these two conditions. The results of this analysis support the semantic attraction account (Table 2C). Participants were more likely to choose BY in the Attraction Violation condition than in the No-Attraction Violation condition—despite the fact that both conditions contained a severe anomaly (i.e. animacy violation). This finding suggests that, contra to the predictions of a severe anomaly account, semantic attraction influenced online grammatical predictions, and often led participants to interpret the anomalous sentence fragments as passive constructions containing a syntactic anomaly. Furthermore, there was a significant effect of response time, indicating that participants were more likely to choose BY when they took longer to respond. While in the omnibus model Presentation Order was not significant, in the pairwise comparison between the Attraction Violation and No-Attraction Violation conditions, it did reach significance. The small but significant effect of Presentation Order indicates that participants' preference for BY-responses decreased in both of these conditions over the course of the experiment (see Figure 4).

The difference in BY-responses in the Violations conditions was robust. However, it did not test for an outright preference for a passive interpretation in the Attraction Violation condition (i.e. BY-response), as predicted by the strong version of the semantic attraction account. The strong version predicts that Attraction Violation stimuli should be treated similarly to passive constructions. To test the strong version, we fit a GLMER model that compared the likelihood of participants choosing BY in the Attraction Violation condition to the likelihood of a BY-response in the Passive Control condition. Passive Control was coded as the reference level. Table 2B presents the results of this analysis. The prediction was not borne out. Participants were significantly less likely to choose BY in the Attraction Violation condition (36% of responses) than in the Passive Control condition (84% of responses). In addition, there was an effect of Presentation Order in both conditions, indicating that participants responded differently to both conditions as the experiment progressed. The significant simple effect of Presentation Order indicates that over the course of the experiment, participants became more likely to choose BY in the Passive Control condition. In contrast, the magnitude of the significant negative effect of the Attraction Violation × Presentation Order interaction indicates that participants were more likely to choose THE (i.e. less likely to choose BY) in the Attraction Violation condition as the study progressed (see Figure 4).

Additional evidence suggesting that semantic cues can influence online processing even in the face of contradictory syntactic cues may be found in comparing response choices in the No-Attraction Violation condition to the Active Control condition. In the No-Attraction Violation items all subject nouns were inanimate, while in the Active Control condition subject nouns were animate (see Table 1). Participants' response choices in each of these conditions may reflect sensitivity to implicit knowledge about the animacy of subject nouns in passive and active constructions. Inanimate subject nouns are a common feature of passive constructions in English, compared to prototypical transitive active constructions, which contain animate subject nouns. In other words, the animacy of the subject noun may serve as a weak semantic cue. This may lead some participants, on some portion of No-Attraction trials, to commit to a passive structural analysis even without the (additional) influence of semantic attraction between the subject noun and the theme role of the verb. To test this prediction, we fit a GLMER model comparing response choices in the Active Control and No-Attraction conditions. Results of this analysis are presented in Table 2D. The prediction was borne out: participants were significantly more likely to choose BY in the No-Attraction condition than in the Active Control. As above, both response time and Presentation Order were also significant, indicating that participants were more likely to choose BY as they took longer to respond and that participants' preference for BY-responses decreased for both of these conditions over the course of the experiment (see Figure 4).

Finally, we analysed response times when participants selected the syntactically licensed THE-continuation (Figure 2, right panel) to test for evidence of conflict between selected response and an alternative interpretation based on semantic plausibility. If Attraction Violation fragments cause participants to consider a semantically plausible alternative to the syntactically supported continuation, then this alternative might compete with and slow down the selection of the syntactically licensed THE-continuation. If No-Attraction Violation fragments do not support such alternative interpretations—or only provide weaker cues to them—they should engender less competition for the THE-continuation. Here and in subsequent analyses, response times were analysed with linear mixed-effects regression (LMER) with crossed random effects for subjects and

items, and with condition and presentation order as fixed effects. All models included the maximal random effects structure justified by the data based on model comparison. Analyses were conducted using the *lme* function (*nlme* library, version 3.1–105; Pinheiro, Bates, DebRoy, Sarkar, & the R Development Core Team, 2013) in R. Consistent with competition due to semantic attraction, participants were slower to make THE-responses in the Attraction Violation condition than in the No-Attraction condition (Table 2E).

We should note that there could be other explanations for this difference in response times.[4] Perhaps the unrelatedness of the verb and subject in No-Attraction items allows participants to more quickly identify these fragments as anomalous, thereby speeding responses. In other words, participants might immediately recognise that fragments in the No-Attraction Violation condition cannot be resolved and therefore "abort" further processing without weighing the degree of conflict between semantic and syntactic cues. However, such an account by itself only explains the response time differences—it cannot explain why participants are more likely to choose THE than BY in No-Attraction relative to Attraction conditions. On the other hand, a general bias for THE-responses (see below) could explain this difference in response proportions, but not the difference in response times. Thus, combining a THE bias account with an "abort" account might explain the overall pattern of results observed in the data. While this combination of explanations is indeed plausible, we believe that semantic attraction provides a more parsimonious account for the overall pattern of results: because the Attraction Violation condition generates more competition than the No-Attraction Violation condition, not only does it decrease the likelihood of choosing THE, it also slows down response times for choosing THE.

Presentation order was also significant, suggesting that there was an overall task adaptation effect such that participants' response times steadily decreased over the course of the study.

Filler responses

We present the responses to the filler conditions to affirm that participants performed as expected in the speeded forced-choice sentence continuation task used here. Figure 5 (left) presents the mean proportion of BY-responses per filler condition. These proportions show that participants were more likely to choose BY than THE in the BY-preference filler condition, more likely to choose THE than BY in the THE-preference filler condition, and were at chance in the Ambiguous filler condition. Moreover, participants were also at chance in both the Syntactic Violation (active) and Syntactic Violation (passive) conditions. These responses are consistent with the hypothesis that participants made their choices based on the available cues. Moreover, the responses to the Syntactic Violation filler conditions suggest that anomaly alone does not lead to an increase in BY-responses, and this helps constrain interpretation of the experimental conditions.

Response times were also informative in the filler conditions. Figure 5 presents response times for BY-responses (centre) and THE-responses (right) per condition. Participants were fastest in the BY- and THE-preference conditions, which were both grammatical and had strong cues biasing towards one response or the other. Participants were slowest in the Syntactic Violation conditions, where neither response was grammatical.

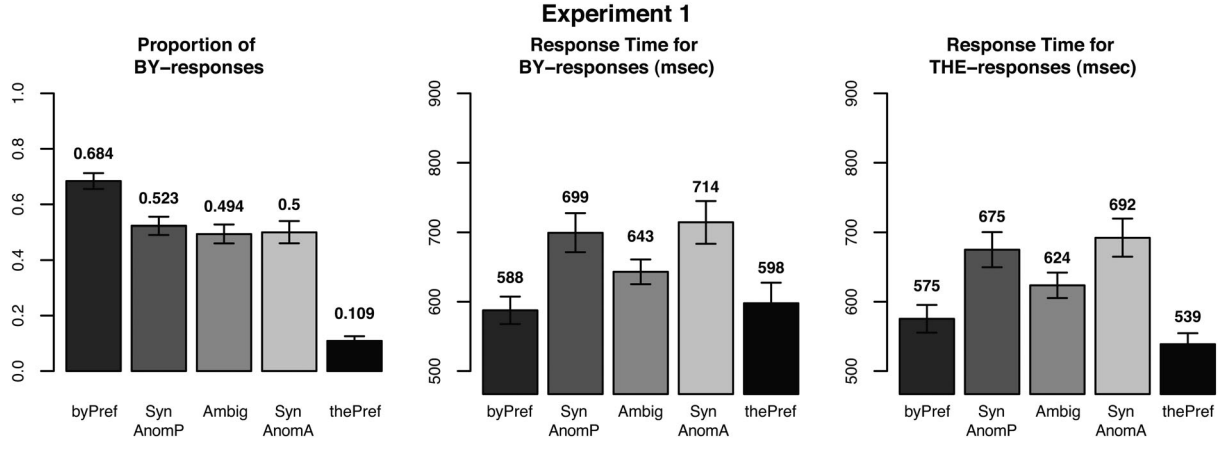

Figure 5. Experiment 1. *Left:* Proportion of BY-responses per filler condition: BY-preference (byPref), Syntactic Anomaly-passive (SynAnomP), Ambiguous (Ambig), Syntactic Anomaly-active (SynAnomA), THE-preference (thePref). *Centre:* Response times for BY-responses by condition. *Right:* Response times for THE-responses per condition. Error bars represent standard error of the mean.

Finally, response times were intermediate for the Ambiguous condition which was grammatical but had no strong cues to bias responses towards either BY or THE.

Potential effects of sentence completions

Recall that on each trial a sentence completion compatible with the participant's choice was presented following their response. Although these completions were intended to affirm the participant's choice, it could be argued that the completions may have influenced participants' response preferences as the experiment progressed. It is possible that participants had a pre-experimental bias for THE-responses (active interpretations) for both Attraction and No-Attraction conditions due to the distributional patterns of English, in which active sentences are more frequent overall than passive sentences (Roland, Dick, & Elman, 2007). If so, the completions could have reinforced this bias, resulting in an artificially inflated proportion of THE-responses in the Attraction Violation condition, increasing the difference between the Attraction Violation and Passive Control conditions and masking evidence in support of the strong version of the semantic attraction hypothesis. This possibility is supported by a significant effect of presentation order in the pairwise GLMER analyses reported above, which demonstrated that participants' preferences for THE-responses increased slightly during the course of the experiment for the Attraction Violation, No-Attraction Violation, and Active Control conditions (see Figure 4). As such, removing the effect of completions could result in a majority of BY-responses in the Attraction Violation condition, thereby providing support for the strong version of the semantic attraction account. To test for sentence completion effects, we conducted a replication study with the modification that no overt completions were provided. In Experiment 2, responses were followed by ellipses, during which time participants were encouraged to simply imagine how they might complete the sentence on their own. Our hypothesis was that while participants might imagine a completion that would be compatible with active frames (e.g. THE-continuations), the lack of overt input would be less likely to bias responses than when completions were provided.

Experiment 2

The goal of Experiment 2 was to replicate the key findings of Experiment 1, while minimising the potentially biasing influence of sentence completions that were consistent with participant responses.

Methods

Participants

Thirty-six undergraduate students at the University of Colorado between the ages of 18 and 22 participated in Experiment 2, none of whom participated in Experiment 1. All participants were native speakers of English, and all received course credit for their participation. Data from one participant were omitted due to technical error.

Materials

Materials for Experiment 2 were identical to those of Experiment 1, with the exception that stimulus sentences did not contain post-response completions.

Procedure

The procedure for trials in Experiment 2 was identical to that of Experiment 1 up to the point of the forced-choice task. As in Experiment 1, trials with response times greater than 1200 ms were excluded from analyses. Upon selecting THE or BY, participants were presented with ellipses in the centre of the screen for 1400 ms. Participants were instructed to imagine during this time how they might complete the sentence on their own. After 1400 ms, a screen appeared that instructed participants to push the space bar to proceed to the next trial.

Results and discussion

The barplots in Figure 6 present the mean proportion of BY-responses per condition (left) and response times per condition for BY-responses (centre) and THE-responses (right). About 1.70% of experimental trials were excluded due to responses missing the 1200 ms deadline.

In Experiment 1, we observed an effect of presentation order in some pairwise comparisons such that the likelihood of a THE-response increased over the course of the experiment. We suggested that the sentence completions provided to participants following their choice may have reinforced a pre-experimental bias for THE-responses, and influenced subsequent choices, artificially inflating the proportion of THE-responses in the Attraction Violation condition. If this were the case, removing the sentence completions from the experimental procedure should have ameliorated the reinforcement of any bias and resulted in an overall increase of BY-responses. A comparison of the barplots in Figures 2 and 6 suggests that removing the sentence completions did not substantially alter the qualitative pattern of results. Although the proportion of BY-responses increased somewhat in the Passive Control (83.9–95%) and Attraction Violation (36.2–

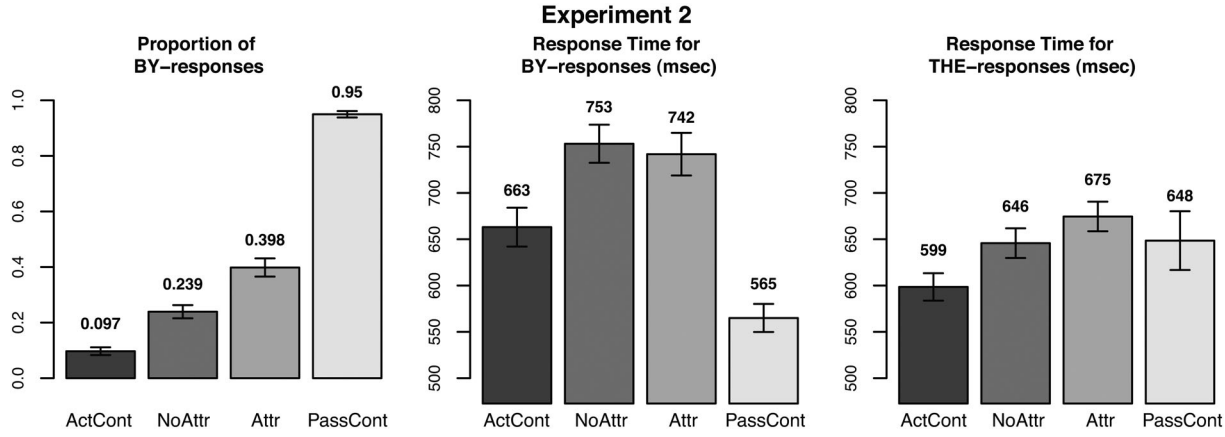

Figure 6. Experiment 2: *Left:* Proportion of BY-responses per condition: Active Control (ActCont), No-Attraction Violation (NoAttr), Attraction Violation (Attr), Passive Control (PassCont). *Centre:* Response times for BY-responses by condition. *Right:* Response times for THE-responses per condition. Error bars represent standard error of the mean.

38.9%) conditions, there was also a small decrease in the No-Attraction Violation (26.6–23.9%) and Active Control (17.5–9.7%) conditions. Furthermore, the increase in BY-responses in the Attraction Violation condition does not appear to reflect participants choosing BY-responses a majority of the time, nor does it appear to show participants treating the Attraction Violation condition similarly to the Passive Control. Thus, the sentence completions provided in Experiment 1 do not appear to have masked any evidence for the strong version of the semantic attraction hypothesis. Instead, the pattern of results found in Experiment 2 is similar to that found in Experiment 1, supporting the weak version of the hypothesis. We therefore analysed response choices and response times in Experiment 2, fitting the models as specified above.

Results for the omnibus model for Experiment 2 are presented in Table 3A. Participants were again reliably

Table 3. Experiment 2: Parameter estimates for fixed effects in GLMER and LMER models.

Parameter	Estimate	SE	Wald *z*	Pr > (|*z*|)	
A. Response proportions for omnibus model containing all experimental conditions					
Intercept	−2.410	0.231	−10.466	<.001	
Condition: No-Attraction Violation	0.922	0.213	4.320	<.001	
Condition: Attraction Violation	1.880	0.228	8.240	<.001	
Condition: Passive Control	29.430	5.225	5.634	<.001	
log(RT)	1.997	0.474	4.215	<.001	
Presentation Order	−0.001	0.001	−0.604	.546	
No-Attraction Violation × log(RT)	0.4822	0.580	0.832	.406	
Attraction Violation × log(RT)	−0.717	0.548	−1.307	.191	
Passive Control × log(RT)	−3.681	0.810	−4.545	<.001	
B. Response proportions for planned comparison between Attraction Violation and Passive Control conditions					
Intercept	3.415	0.258	13.258	<.001	
Condition: Attraction Violation	−4.203	0.348	−12.084	<.001	
log(RT)	−2.464	0.685	−3.598	<.001	
Attraction Violation × log(RT)	4.188	0.915	4.579	<.001	
C. Response proportions for planned comparison between Attraction Violation and No-Attraction Violation conditions					
Intercept	−1.455	0.151	−9.635	<.001	
Condition: Attraction Violation	0.909	0.124	7.308	<.001	
log(RT)	2.978	0.450	6.612	<.001	
Attraction Violation × log(RT)	−1.393	0.434	−3.207	<.01	

Parameter	Estimate	SE	df	*t*-Value	*p*-Value
D. Response times for THE-responses for planned comparison between Attraction Violation and No-Attraction Violation conditions					
Intercept	691.180	18.357	1608	37.652	<.001
Condition: Attraction Violation	18.418	11.281	1608	1.633	.098
Presentation Order	−0.227	0.162	1608	−1.398	.162

Parameter	Estimate	SE	Wald *z*	Pr > (|*z*|)	
E. Response proportions for planned comparison between No-Attraction Violation and Active Control conditions					
Intercept	−2.551	0.167	−15.262	<.001	
Condition: No-Attraction Violation	1.007	0.191	5.259	<.001	
log(RT)	2.425	0.293	8.278	<.001	

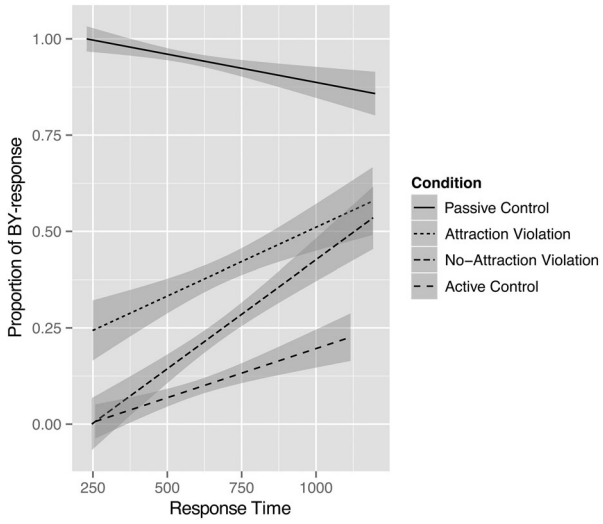

Figure 7. Experiment 2: Mean proportions of BY-responses by response times (ms) per condition. Shaded areas represent the 95% confidence interval around the smoothed continuous means.

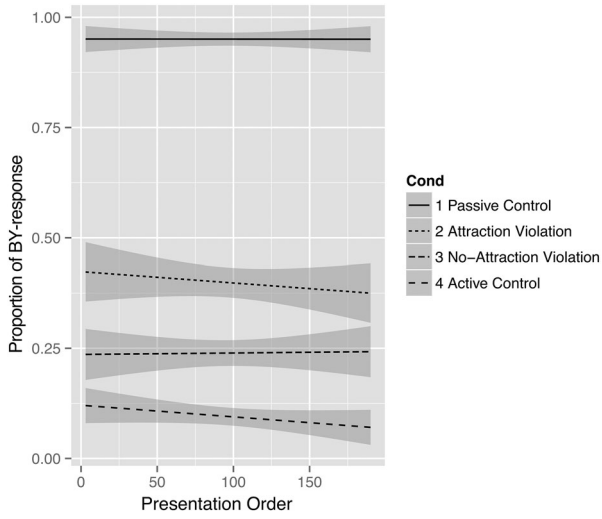

Figure 8. Experiment 2: Mean proportions of BY-responses by presentation order per condition. Shaded areas represent the 95% confidence interval around the smoothed continuous means.

above chance in the likelihood of choosing THE in the Active Control condition. Contra the predictions of syntax-dominant accounts, the significant condition parameter estimates indicate an increase in likelihood of a BY-response for each condition relative to the Active Control, suggesting that participants are sensitive to semantic influences during online interpretation. As in Experiment 1, there was a significant effect of response time, indicating that participants were more likely to choose BY as they took longer to respond (see Figure 7). In contrast to Experiment 1, however, a Passive Control × Response Time interaction revealed that longer response times in the Passive Control condition led to a decrease in BY-responses rather than an increase. Moreover, no interaction between Condition and Presentation Order was found in Experiment 2. Figure 8 reveals that this is because participants' response preferences did not change substantially over the course of the study. This finding lends support to the speculation that the explicit completions provided in Experiment 1 did have some influence on participants' response preferences as the experiment progressed, but not enough to change the overall pattern of results.

The results again support a semantic attraction account of the response patterns over a severe anomaly account. Table 3C presents the results of the analysis. Participants were significantly more likely to choose BY in the Attraction Violation condition than in the No-Attraction Violation condition, independent of response time. There was also a simple effect of response time, as well as a Response Time × Condition (Attraction Violation, No-Attraction Violation) interaction. Taken together, these results indicate that participants were overall more likely to choose BY when taking longer to respond but that this effect was stronger in No-Attraction Violation condition relative to the Attraction Violation condition (see Figure 7).

The lack of explicit sentence completions in Experiment 2 did not result in a majority of BY-responses in the Attraction Violation condition. As in Experiment 1, the strong version of the semantic attraction account was not supported (Table 3B). Participants were significantly less likely to choose BY in the Attraction Violation condition (40% of responses) than in the Passive Control condition (95% of responses). As in the omnibus model, there was a significant effect of response time, affirming that participants were more likely to choose BY as they took longer to respond. There was also an Attraction Violation × Response Time interaction, capturing the opposite effect of increased response time on response choices in the Passive Control and Attraction Violation conditions (see Figure 7).

As in Experiment 1, we analysed response times in trials when participants chose the syntactically licensed THE-continuation in Attraction Violation and No-Attraction Violation conditions (Figure 6, right panel). This was done to test for evidence of conflict between the syntactically licensed choice and a semantically plausible alternative. Again, we predicted that if participants considered a semantically plausible alternative to the syntactically supported continuation, then this alternative might compete with and slow down the selection of the syntactically licensed THE-continuation. The results are presented in Table 3D. Although participants were

numerically slower to choose THE in the Attraction Violation condition than in the No-Attraction Violation condition, this difference was not statistically reliable, unlike Experiment 1. There was no effect of Presentation Order.

As in Experiment 1, participants were significantly more likely to choose BY in the No-Attraction condition than in the Active Control condition. Results are presented in Table 3E. This finding provides further evidence supporting the speculation that implicit statistical knowledge about inanimate subject nouns being more frequent in passive than active constructions may serve as a weak semantic cue, leading to a passive interpretation even in the absence of semantic attraction. This idea comports well with the overall pattern of BY-responses found across all experimental conditions in both experiments (Figures 2 and 6, left panel). As can be seen, conditions containing more semantic cues pointing to a passive interpretation are associated with greater BY-responses.

Filler responses

Figure 9 presents the mean proportion of BY-responses per filler condition (left), as well as response times per condition for BY-responses (centre) and THE-responses (right). The same pattern of effects can be seen as in Experiment 1: participants were more likely to choose BY than THE in the BY-preference condition, more likely to choose THE than BY in the THE-preference condition, and were at chance in the Ambiguous condition, Syntactic Violation (active) and Syntactic Violation (passive) conditions. The response times also patterned like Experiment 1. Participants were fastest in the BY- and THE-preference conditions, slowest in the Syntactic Violation conditions, and intermediate in the Ambiguous condition.

General discussion

Our findings demonstrate that semantic knowledge can control comprehenders' online interpretations of sentences, even in the face of unambiguous but contradictory syntactic cues. In two speeded forced-choice sentence continuation studies, participants rapidly selected between two words, BY and THE, as continuations for anomalous sentence fragments. Participants were more likely to select BY for fragments in which the subject noun was semantically attracted to the theme role of the verb (e.g. "The hearty meal was devouring … ", Attraction Violation), than for fragments that were equally anomalous but contained a subject noun was an implausible theme of the verb (e.g. "The sealed envelope was devouring … ", No-Attraction Violation). This evidence supports the weak version of the semantic attraction hypothesis. The strong version of the hypothesis, however, was not supported: participants did not choose BY on a majority of Attraction Violation trials.

When participants continued anomalous fragments with THE, they were slower to do so in the Attraction Violation condition than the No-Attraction Violation condition (this effect was reliable in Experiment 1 but not reach significance in Experiment 2). This result suggests that the semantically plausible alternative was entertained—and was therefore in competition with the syntactically supported continuation—thus contributing to a slow down in the selection of the syntactically licensed THE-continuation.[5]

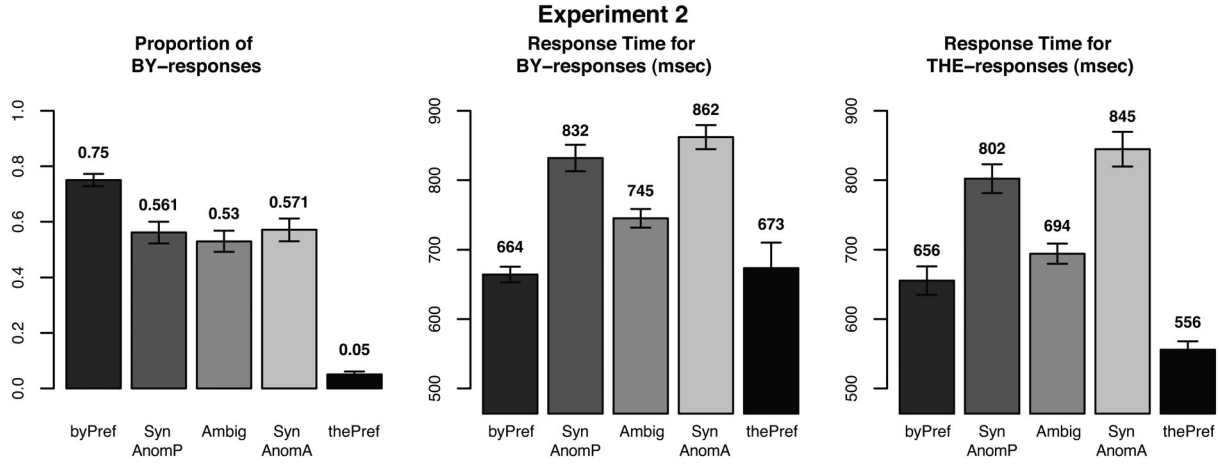

Figure 9. Experiment 2. *Left:* Proportion of BY-responses per filler condition: BY-preference (byPref), Syntactic Anomaly-passive (SynAnomP), Ambiguous (Ambig), Syntactic Anomaly-active (SynAnomA), THE-preference (thePref). *Centre:* Response times for BY-responses by condition. *Right:* Response times for THE-responses per condition. Error bars represent standard error of the mean.

Although BY responses were less common in No-Attraction Violation than Attraction Violation conditions, they were more common in the No-Attraction Violation than in the Active Control condition. We speculate that comprehenders have knowledge about inanimate subject nouns being more frequent in passive than active constructions, and that this knowledge serves as a weak semantic cue that sometimes drives commitments to a passive structural analysis. While animacy alone is not semantic attraction *per se*, this finding comports well with the claim that semantic cues should be able to influence online processing even in the face of unambiguous syntax.

Taken together, these results contradict influential older comprehension models that only allow syntactically licensed interpretations to be pursued (e.g. Ferreira & Clifton, 1986; Frazier & Fodor, 1978). Indeed, the strongest version of a syntax-dominant account would predict that responses to Attraction Violation and No-Attraction Violation sentences should not differ from Active Control sentences (e.g. "The hungry boys were devouring … "). Instead, the current results are consistent with the spirit of more modern comprehension models that allow semantics to have a greater role in determining online sentence comprehension than prior syntactocentric models (e.g. Ferreira, 2003; Kamide et al., 2003; Trueswell et al., 1994)—however, the studies cited here do not directly support the hypothesis that semantic information can influence *online* sentence interpretation.

Thus, the present study fills critical gaps in the current theoretical and empirical landscape. While most previous demonstrations of semantic influences on sentence interpretation have either involved situations where syntactic cues are indeterminate, or have relied on off-line measures of interpretation, the current findings offer some of the clearest available behavioural evidence that semantic representations can control real-time interpretation without syntactic support. The results point to online semantic processing mechanisms that may underlie a number of off-line results whose temporal dynamics have not been clear, including "good enough" interpretation (e.g. Ferreira, 2003), pragmatic inference of a speech error or typographical error (Grice, 1975), and semantic illusion (Erickson & Mattson, 1981).

Relating current results to semantic P600 findings

In this section we discuss the implications of our findings for on-going debates surrounding the interpretation of semantic P600 effects. As mentioned in the introduction, ERP studies report that P600 rather than N400 effects are reliably elicited by anomalous words in sentences similar to those in the current study's Attraction Violation condition (Hoeks et al., 2004; Kim & Osterhout, 2005; Kim & Sikos, 2011; Kolk et al., 2003; Kuperberg et al., 2003). Some accounts of semantic P600 effects have emphasised the role of severe anomaly and make no distinction during online processing between situations that involve semantically plausible alternatives to the grammatically supported interpretation and sentences that are simply semantically anomalous (e.g. Chow & Phillips, 2013; Kuperberg, 2007; Paczynski & Kuperberg, 2012; Stroud & Phillips, 2012).

The primary goal of the present study was to test similar questions about semantic influences on sentence comprehension as those investigated by the aforementioned ERP studies, but using a different and potentially converging measure. Thus, our results can provide insight into the functional demands of stimuli that elicit semantic P600 effects, and may help constrain the functional interpretation of these ERPs.

Contra the severe anomaly account, our findings show an important difference in the processing of words that are simply difficult to integrate into the previous context (No-Attraction Violation) and words that create a conflict between semantic and syntactic cues due to semantic attraction (Attraction Violation). Despite both fragment types being highly implausible and both containing animacy violations, participants were more likely in the latter condition to entertain an interpretation that overruled the syntactic cues. This finding is compatible with accounts that attribute semantic P600 effects to attempts to resolve conflict between a syntactically licensed interpretation and semantically plausible alternative (Kim & Osterhout, 2005; Kolk et al., 2003; Kuperberg et al., 2003). In contrast, this result is problematic for the severe anomaly account, which attributes semantic P600 effects to the detection of highly implausible interpretation (e.g. animacy violation), regardless of plausibility (e.g. Chow & Phillips, 2013; Kuperberg, 2007; Paczynski & Kuperberg, 2012; Stroud & Phillips, 2012).[6]

Our results also raise new questions for interpretations of the previous ERP findings. What should we make of the current results indicating that people interpret Attraction Violation-type sentences as passives only 36–40% of the time? Why then do the anomalous words in ERP studies elicit P600s and not elicit N400s, if the active interpretation is more frequent (60–64%)?[7] One explanation for the presence of P600 effects is that semantic attraction may interact with other constraints such that a subset of items within the Attraction Violation condition drives the grand average semantic P600 effect. For example, it may be that P600s are largest for items with both a high degree of semantic attraction and a

verb that is frequently used in passive constructions, while P600s might be smallest for items with a weak semantic attraction and a verb that is rarely used in passive constructions. Future work combining corpus analyses, item-wise measures of semantic attraction, and item-wise ERP analyses could shed light on this hypothesis.

Another possible explanation for this divergence between the presence of P600 effects in ERP studies and the overall bias for THE-responses in the study presented here is as follows. Although the semantic information in the Attraction Violation condition reliably affects participants' online processing, some of the current participants, on some portion of trials, may not have fully completed the processes that would result in a P600 effect and may have instead defaulted to a prepotent THE-response. This explanation is consistent with the finding that participants were more likely to choose BY when they took longer to respond. Future work using the forced-choice sentence continuation paradigm could test this hypothesis by manipulating the latency of the response prompt in order to probe the state of the system at different points during processing of the anomalous word (e.g. before P600 processing begins vs. after P600 processing is complete).

Finally, regarding the absence of an N400 response to such stimuli, one possible contribution has to do with the close semantic relationship between the subject noun phrase and the anomalous verb in Attraction Violation-like sentences. On the common interpretation that N400 amplitude is modulated by the ease or difficulty of retrieving word-associated semantic knowledge within a given context (Kutas & Hillyard, 1980; Kutas, Van Petten, & Kluender, 2006), the N400 is likely to be attenuated to some degree due to lexical priming from the related subject noun. Crucially, this attenuation might be large enough to offset any increase in N400 amplitude due to semantic anomaly, even if the majority of participants, on the majority of trials, entertain an active (syntactically driven) interpretation—consistent with the high proportion of THE-response in the current study. Although we agree that lexical priming needs to be taken into account as at least as part of the explanation for the absence of N400 effects within the overall semantic P600 effect pattern, there is reason to believe that priming is not sufficient to completely offset an N400 effect driven by semantic anomaly. In a recent ERP study, Oines and Kim (2014) presented participants with Attraction Violation and No-Attraction Violation sentences and included a between-subjects instruction manipulation. Half of the participants were told that the sentences would contain occasional semantic anomalies and were shown a visualisation of a cartoon hamburger eating a meal. The other half were told that the sentences would contain occasional syntactic anomalies (i.e. "there will be lots of typos"). In the former group, anomalous words in the Attraction Violation condition elicited an N400 effect (plus a left anterior negativity: LAN) but no P600 effect, relative to the No-Attraction condition. In the latter group, anomalous words in the Attraction Violation condition elicited a P600 effect without N400 or LAN effects, relative to the No-Attraction condition. Thus, it appears that when participants actually interpret Attraction Violation sentences as being semantically anomalous, lexical priming does not sufficiently offset N400 amplitude.

Methodological contributions

We would also like to highlight several advantageous features of the current speeded forced-choice sentence continuation paradigm (adapted from Staub, 2009) for the study of conflict during online language processing. First, because the probe elicits responses precisely when direct conflict between syntactic and semantic cues becomes apparent, this method provides a window into participants' interpretations at the critical point in time. That is, the measure is not only online, it also taps into participants' interpretations in a way that alternative online methods like EEG and eye-tracking cannot. Second, participants provide a behavioural response whose functional interpretation is relatively straightforward. Thus, the method combines functional clarity with reasonably high temporal resolution. These properties allow this method to complement other online measures in which functional ambiguity can arise due to theoretical disagreement about the demands imposed by the stimuli. Finally, given the amount of emphasis on predictive mechanisms in current sentence processing research (e.g. Altmann & Mirkovic, 2009; Brothers, Swaab, & Traxler, 2015; Farmer, Brown, & Tanenhaus, 2013; Federmeier, 2007; Fine, Jaeger, Farmer, & Qian, 2013; Garrod & Pickering, 2015; Kim & Lai, 2012; Levy, 2008; Lewis & Bastiaansen, 2015; Van Petten & Luka, 2012; see Kuperberg & Jaeger, 2016 for a review), we believe that this method may be valuable in addressing a variety of research questions regarding how language users leverage predictive processes to help resolve conflicts online (e.g. in cases of co-reference resolution or syntactic ambiguity resolution).

A potential objection is that the inclusion of an interpretive decision might add an artificial component to this task. However, we believe that online comprehension in the wild requires *repeated sequences* of such interpretive decisions as participants incrementally update their mental representation (their "interpretation") of an

unfolding utterance. As information comes into the system, comprehenders must integrate that information with the current mental representation, and make implicit predictions (or have implicit expectations) for subsequent input. The current task taps into this natural process.

Conclusion

Over the last 30 years, psycholinguistic findings have gradually pushed theories of language processing to acknowledge an ever-increasing role for semantics during comprehension. We have presented evidence from two experiments that tested whether participants' online interpretations were influenced by semantic cues when reading sentences that were syntactically unambiguous but semantically anomalous due to an animacy violation. The results from both experiments demonstrated that semantically attractive interpretations can be pursued in real time even in the face of contradictory syntactic cues. These findings suggest that semantic attraction plays a reliable role in the interpretation of such sentences and provide an important contribution relative to previous behavioural studies (e.g. Ferreira, 2003; Kamide et al., 2003) by demonstrating that semantic information can influence *online* sentence interpretation even when syntactic cues are *un*ambiguous. These results provide converging evidence for conclusions from prior ERP studies which suggest that semantic cues can impact comprehenders' interpretations of semantically anomalous sentences (Kim & Osterhout, 2005; Kim & Sikos, 2011). Thus, the current study extends theories of language comprehension by acknowledging a greater role for semantics during sentence interpretation.

Funding information

This work was supported by the NICHD grant to Albert Kim [grant number 1R03HD071094-01A1].

Notes

1. A maximally strong version of the semantic attraction hypothesis is not tenable. If the comprehension system was completely dominated by semantic cues, then comprehenders should fail to notice any anomaly in the Attraction Violation condition, since semantic cues are the only thing that matters. However, it is clear that such sentences will be perceived as anomalous, and sentence-level acceptability judgments about similar stimuli show that people do indeed perceive them as such (e.g., Kim & Osterhout, 2005).
2. Note that this response-dependent sentence completion differs from the standard method used in ERP studies, where no mid-sentence behavioural response is collected. For example, the post-critical-word continuations in Kim and Osterhout (2005) were equally distributed across by-phrase, noun-phrase, and adverbial phrase completions in order to avoid biasing either the passive or active interpretations. In the present paradigm, however, such a strategy would be problematic because any sentence completion that is inconsistent with the participant's response would lead to a double anomaly in the violation conditions (i.e. first at the critical word and then at the word immediately following the response). In contrast, response-compatible completions can serve to affirm the participant's choice.
3. The appropriate random effects structure for each model was determined via a forward step-wise model comparison procedure. All models started with only random intercepts for subject and items. Random effects were then added iteratively and compared to the previous model using log-likelihood ratio tests. If the result of the test was significant, the random effect was kept in the model for the subsequent iteration.
4. We thank an anonymous reviewer for this suggestion.
5. The idea that longer response times reflect competition is also supported by the finding that when participants took longer to respond, they were more likely to choose BY—within all fragment types and across both experiments for all but the Passive Control condition. We speculate that these correlations indicate that the selection of a BY-response (whether the result of a passive or adverbial interpretation) was under strong competition from a preferred, syntactically licensed THE-response. This may be the case even in the Attraction Violation condition, which provides stronger cues to a passive interpretation than does either the No-Attraction Violation or Active Control conditions.
6. We do not intend to suggest that severe anomaly (or other factors) play *no* role in semantic P600 effects. Severe anomaly may very well be a key factor in situations where P600 effects have been observed for semantically anomalous sentences that do not contain thematic role reversals.
7. For instance, an active interpretation of "The hearty meal was devouring … " indicates that the meal is the agent of devouring, which is semantically anomalous, and thus should lead to N400 effects.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

Altmann, G. T., & Mirkovic, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, *33*(4), 583–609. doi:10.1111/j.1551-6709.2009.01022.x

Bates, D., & Sarkar, D. (2007). *lme4: Linear mixed-effects models using S4 classes*. 'R' package. Version 0.9975-12. http://CRAN. R-project.org

Brothers, T., Swaab, T. Y., & Traxler, M. J. (2015). Effects of prediction and contextual support on lexical processing: Prediction

takes precedence. *Cognition*, *136*, 135–149. doi:10.1016/j.cognition.2014.10.017

Caramazza, A., & Zurif, E. B. (1976). Dissociation of algorithmic and heuristic processes in language comprehension: Evidence from aphasia. *Brain and Language*, *3*(4), 572–582. doi:10.1016/0093-934X(76)90048-1

Chow, W. Y., & Phillips, C. (2013). No semantic illusions in the "Semantic P600" phenomenon: ERP evidence from Mandarin Chinese. *Brain Research*, *1506*, 76–93. doi:10.1016/j.brainres.2013.02.016

Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, *42*(4), 368–407. doi:10.1006/cogp.2001.0752

Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, *20*(5), 540–551. doi:10.1016/S0022-5371(81)90165-1

Farmer, T. A., Brown, M., & Tanenhaus, M. K. (2013). Prediction, explanation, and the role of generative models in language processing. *Behavioral and Brain Sciences*, *36*(3), 211–212. doi:10.1017/S0140525X12002312

Federmeier, K. D. (2007). Thinking ahead: The role and roots of prediction in language comprehension. *Psychophysiology*, *44*(4), 491–505. doi:10.1111/j.1469-8986.2007.00531.x

Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive Psychology*, *47*, 164–203. doi:10.1016/S0010-0285(03)00005-7

Ferreira, F., & Clifton Jr., C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*(3), 348–368. doi:10.1016/0749-596X(86)90006-9

Ferreira, F., & Patson, N. D. (2007). The 'good enough' approach to language comprehension. *Language and Linguistics Compass*, *1*(1-2), 71–83. doi:10.1111/j.1749-818X.2007.00007.x

Fine, A. B., Jaeger, T. F., Farmer, T. A., & Qian, T. (2013). Rapid expectation adaptation during syntactic comprehension. *PLoS ONE*, *8*(10): e77661. doi:10.1371/journal.pone.0077661

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press/Bradford Books.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*(4), 291–325. doi:10.1016/0010-0277(78)90002-1

Garrod, S., & Pickering, M. J. (2015). The use of content and timing to predict turn transitions. *Frontiers in Psychology, 6*, 751. doi:10.3389/fpsyg.2015.00751

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics, Vol.3: Speech acts* (pp. 225–242). New York, NY: Seminar Press.

Hagoort, P., Brown, C., & Groothusen, J. (1993). The syntactic positive shift (SPS) as an ERP measure of syntactic processing. *Language and Cognitive Processes*, *8*(4), 439–483. doi:10.1080/01690969308407585

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*(1), 59–73. doi:10.1016/j.cogbrainres.2003.10.022

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, *59*(4), 434–446. doi:10.1016/j.jml.2007.11.007

Jaeger, T. F. (2009). *Random effect: Should I stay or should I go*. HLP/Jaeger lab blog.

Kamide, Y., Altmann, G. T., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, *49*(1), 133–156. doi:10.1016/S0749-596X(03)00023-8

Kim, A., & Lai, V. (2012). Rapid interactions between lexical semantic and word form analysis during word recognition in context: Evidence from ERPs. *Journal of Cognitive Neuroscience*, *24*(5), 1104–1112. doi:10.1162/jocn_a_00148

Kim, A., & Osterhout, L. (2005). The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language*, *52*(2), 205–225. doi:10.1016/j.jml.2004.10.002

Kim, A., & Sikos, L. (2011). Conflict and surrender during sentence processing: An ERP study of syntax-semantics interaction. *Brain and Language*, *118*(1), 15–22. doi:10.1016/j.bandl.2011.03.002

Kolk, H. H. J., Chwilla, D. J., Van Herten, M., & Oor, P. J. W. (2003). Structure and limited capacity in verbal working memory: A study with event-related potentials. *Brain and Language*, *85*(1), 1–36. doi:10.1016/S0093-934X(02)00548-5

Kuperberg, G. R. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research*, *1146*, 23–49. doi:10.1016/j.brainres.2006.12.063

Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, *31*(1), 32–59. doi:10.1080/23273798.2015.1102299

Kuperberg, G. R., Kreher, D. A., Sitnikova, T., Caplan, D. N., & Holcomb, P. J. (2007). The role of animacy and thematic relationships in processing active English sentences: Evidence from event-related potentials. *Brain and Language*, *100*(3), 223–237. doi:10.1016/j.bandl.2005.12.006

Kuperberg, G. R., Sitnikova, T., Caplan, D., & Holcomb, P. J. (2003). Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research*, *17*(1), 117–129. doi:10.1016/S0926-6410(03)00086-7

Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, *207*(4427), 203–205. doi:10.1126/science.7350657

Kutas, M., & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, *307*, 161–163. doi:10.1038/307161a0

Kutas, M., Van Petten, C., & Kluender, R. (2006). *Psycholinguistics electrified II: 1994e2005. Handbook of psycholinguistics* (2nd ed.). New York, NY: Elsevier.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*(3), 1126–1177. doi:10.1016/j.cognition.2007.05.006

Lewis, A., & Bastiaansen, M. (2015). A predictive coding framework for rapid neural dynamics during sentence-level language comprehension. *Cortex*, *68*, 155–168. doi:10.1016/j.cortex.2015.02.014

Oines, L., & Kim, A. E. (2014). *Integrate or repair? ERP responses to semantic anomalies depend on choice of processing strategy*. Talk presented at Architectures and Mechanisms of Language Processing, University of Edinburgh.

Osterhout, L., Holcomb, P. J., & Swinney, D. A. (1994). Brain potentials elicited by garden-path sentences: Evidence of the application of verb information during parsing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(4), 786. doi:10.1037/0278-7393.20.4.786

Paczynski, M., & Kuperberg, G. R. (2011). Electrophysiological evidence for use of the animacy hierarchy, but not thematic role assignment, during verb-argument processing. Language and Cognitive Processes, 26(9), 1402–1456. doi:10.1080/01690965.2011.580143

Paczynski, M., & Kuperberg, G. R. (2012). Multiple influences of semantic memory on sentence processing: Distinct effects of semantic relatedness on violations of real-world event/state knowledge and animacy selection restrictions. *Journal of Memory and Language*, *67*(4), 426–448. doi:10.1016/j.jml.2012.07.003

Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., & the R Development Core Team. (2013). *nlme: Linear and Nonlinear Mixed Effects Models. R package version 3.1-105*. The R Foundation for Statistical Computing. Vienna, Austria.

R Development Core Team (2011). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. http://www.R-project.org

Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of basic English grammatical structures: A corpus analysis. *Journal of Memory and Language*, *57*(3), 348–379. doi:10.1016/j.jml.2007.03.002

Saffran, E. M., Schwartz, M. F., & Linebarger, M. C. (1998). Semantic influences on thematic role assignment: Evidence from normals and aphasics. *Brain and Language*, *62*(2), 255–297. doi:10.1006/brln.1997.1918

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-prime user's guide*. Pittsburgh, PA: Psychology Software Tools.

Slobin, D. I. (1966). Grammatical transformations and sentence comprehension in childhood and adulthood. *Journal of Verbal Learning and Verbal Behavior*, *5*(3), 219–227. doi:10.1016/S0022-5371(66)80023-3

Staub, A. (2009). On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language*, *60*(2), 308–327. doi:10.1016/j.jml.2008.11.002

Stroud, C., & Phillips, C. (2012). Examining the evidence for an independent semantic analyzer: An ERP study in Spanish. *Brain and language*, *120*(2), 108–126. doi:10.1016/j.bandl.2011.02.001

Tabossi, P., Spivey-Knowlton, M. J., McRae, K., & Tanenhaus, M. K. (1994). Semantic effects on syntactic ambiguity resolution: Evidence for a constraint-based resolution process. In C. Umilta & M. Moscovitch (Eds.), *Attention and performance: Vol. 15. Conscious and nonconscious information processing* (pp. 589–615). Cambridge, MA: MIT Press.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268* (5217), 1632–1634. doi:10.1126/science.7777863

Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, *33*(3), 285–318. doi:10.1006/jmla.1994.1014

Van Herten, M., Chwilla, D. J., & Kolk, H. H. (2006). When heuristics clash with parsing routines: ERP evidence for conflict monitoring in sentence perception. *Journal of Cognitive Neuroscience*, *18*, 1181–1197. doi:10.1162/jocn.2006.18.7.1181

Van Petten, C., & Luka, B. J. (2012). Prediction during language comprehension: Benefits, costs, and ERP components. *International Journal of Psychophysiology*, *83*(2), 176–190. doi:10.1016/j.ijpsycho.2011.09.015