Journal of Memory and Language 98 (2018) 26-44

Contents lists available at ScienceDirect

Journal of Memory and Language

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



Relative clause avoidance: Evidence for a structural parsing principle $\stackrel{\star}{\sim}$



Adrian Staub^{a,*}, Francesca Foppolo^b, Caterina Donati^c, Carlo Cecchetto^{b,d}

^a Department of Psychological and Brain Sciences, University of Massachusetts Amherst, United States

^b Department of Psychology, University of Milan-Bicocca, Italy

^c Laboratoire de Linguistique Formelle, Université Paris 7/CNRS, France

^d Structures Formelles du Langage, Université Paris 8/CNRS, France

ARTICLE INFO

Article history: Received 15 April 2017 Revision received 11 September 2017 Available online 23 September 2017

Keywords: Sentence processing Eye movements Relative clauses Long distance dependencies

ABSTRACT

Three eye movement experiments investigated the processing of the syntactic ambiguity in strings such as *the information that the health department provided*, where the *that*-clause can be either a relative clause (RC) or the start of a nominal complement clause (CC; *the information that the health department provided a cure*). The experiments tested the prediction that comprehenders should avoid the RC analysis because it involves an unforced filler-gap dependency. Readers showed difficulty upon disambiguation toward the RC analysis, and showed facilitated processing of the ambiguous material itself when the CC analysis was available; both patterns suggest rapid initial adoption of the CC analysis in preference to the RC analysis. The strength of the bias of a specific head noun (e.g., *information*) to appear with a CC did not modulate these effects, nor were these effects reliably modulated by the tendency of an ambiguous string to be completed off-line as a CC or an RC. These results add to the evidence that structural principles guide the processing of filler-gap dependencies.

© 2017 Elsevier Inc. All rights reserved.

Introduction

The comprehension of *filler-gap* dependencies has long been a central topic in sentence processing research (e.g., Clifton & Frazier, 1989; Fodor, 1978). These are structures in which an element (the *filler*, in psycholinguistic terminology) is displaced with respect to the location at which it receives its thematic role (the *gap*). In syntactic theories that posit movement, the filler is analyzed as having moved from the gap site, which contains a movement trace (e.g. Chomsky, 1981). In (1) and (2), filler-gap dependencies appear in the context of a *wh*-question and a relative clause, respectively.

- (1) Which dog did the family choose _____
- (2) The family chose the dog *that* they visited _____ on Wednesday.

E-mail address: astaub@psych.umass.edu (A. Staub).

The great majority of research on this topic has focused on processing of structures in which the filler appears to the left of the corresponding gap site (as in (1) and (2)), and has addressed questions about how, in the course of incremental processing, the parser identifies the site of the gap once a filler has been identified. One important conclusion is that incremental parsing appears to respect grammatical constraints as to the location of the gap; the parser does not posit a gap within a syntactic 'island' (Ross, 1967), namely a syntactic domain (e.g., adjunct clauses) where a gap cannot occur (e.g., Phillips, 2006; Traxler & Pickering, 1996). A second important conclusion, however, is that the parser appears to posit a gap in any grammatically licit site that it encounters, as soon as such a site becomes apparent in the course of incremental processing. This process has come to be known as 'active gap filling', and the corresponding parsing strategy as the Active Filler Strategy (Clifton & Frazier, 1989; Frazier & Clifton, 1989).

There are multiple lines of evidence for active gap filling. First, processing difficulty ensues when a potential gap site turns out to be occupied by another element. This phenomenon, known as a *filled gap effect*, was first demonstrated by Stowe (1986), and is illustrated in (3):



^{*} Thanks to Kirk Goddard, Sophia Dodge, Atreyi Mukherji, Sinthema Roy, Wesley Albright, Chhing Tiv, Jessica Tin, and Audrey O'Neill for assistance with data collection. Thanks also for helpful comments from audiences at the 2016 CUNY Conference on Human Sentence Processing at the University of Florida, and colloquium audiences at University College London and Bournemouth University. Thanks to Chuck Clifton for insightful comments on an earlier draft of this manuscript.

^{*} Corresponding author at: Department of Psychological and Brain Sciences, University of Massachusetts, 430 Tobin Hall, Amherst, MA 01003, United States.

⁽³⁾ *Which dog* did the family choose a leash for _____ at the pet store?

In this sentence, difficulty would occur when the reader encounters *a leash*. This effect appears to reflect the syntactic reanalysis that is required when the parser posits a gap after the first licensing verb that it encounters (*choose*), and then finds that this gap site is filled by other material. Second, processing difficulty results when positing a gap in a grammatically licit location yields an implausible interpretation (Traxler & Pickering, 1996), as in (4):

(4) *Which dog* did the family eat happily with _____?

This difficulty, which would arise upon encountering the gaplicensing verb (*eat*), suggests that the processor has posited a direct object gap after *eat*, and has initially computed an interpretation on which the family has eaten the dog.

Notably, such plausibility effects at the verb do not depend on that verb's subcategorization bias, i.e., the frequency with which the verb occurs with a direct object. In eyetracking and self-paced reading experiments, Pickering and Traxler (2003; Staub, 2007) found that the critical plausibility effect does not depend on whether the verb usually occurs with a noun phrase direct object, or tends to be followed by a prepositional phrase. The increase in reading time on the verb, compared to plausible controls, did not differ between (5a), where the verb *killed* tends to occur with noun phrase direct object, and (5b), where the verb *worried* tends to be followed by a prepositional phrase.

(5)	a.	That's the country that the soldier killed for during
		the war in Korea.

b. That's the car that the dog worried about after going to the vet because of an injury.

Thus, active gap filling appears to be a general parsing heuristic rather than a frequency-sensitive mechanism, as a direct object gap is actively posited in the first grammatically licit site that is encountered during incremental processing of the sentence, even when a direct object would not be expected in this location based on the licensing verb's subcategorization bias.

This work has established that the parser actively posits a gap once a filler has been identified to its left. A much smaller literature has investigated processing of sentences in which an element is extraposed to the right of its canonical position in English, so that a gap appears before the corresponding filler (Levy, Fedorenko, Breen, & Gibson, 2012; Staub, Clifton, & Frazier, 2006). Staub et al. (2006) conducted two eyetracking experiments investigating the processing of Heavy NP Shift, in which a verb's direct object is shifted to the right, over another element. The occurrence of Heavy NP Shift is conditioned by the phonological weight of the object (i.e., length) and by its discourse status (Arnold, Losongco, Wasow, & Ginstrom, 2000; Wasow, 1997). An example from Staub et al. (2006) is in (6):

(6) Jack watched _____ from the stands his daughter's attempt to shoot a basket.

Staub et al. observed a pattern of processing difficulty suggesting that whenever the verb does not categorically require a direct object (e.g., *watched*), the intransitive analysis is initially adopted in preference to the Heavy NP Shift analysis. In this case, encountering the shifted direct object resulted in disruption. Notably, this was true even for verbs that were preferentially, but not obligatorily, transitive. Only when the verb was obligatorily transitive (e.g., *praised*) was a gap readily posited in the post-verbal position. In this case, there was some disruption on the intervening prepositional phrase (e.g., *from the stands*), but the shifted direct object was processed easily. Similarly, based on a series of self-paced reading studies examining processing of extraposed relative clauses, Levy et al. (2012) also reach the conclusion that the comprehender expects an extraposed RC only when an RC is made essentially obligatory.

Thus, the literature presents a highly coherent picture of how the processor deals with ambiguity as to whether a gap is present at a given location in a string. When a filler has already been unambiguously identified, a gap is posited in the first grammatically licit location. But when no filler has been identified, the parser avoids positing a gap. In both cases, the parser's behavior appears to be more-or-less categorical, i.e., insensitive to subcategorization biases of lexical heads.

The existing evidence, then, is consistent with a general proposal by De Vincenzi (1991) governing the processing of fillergap dependencies, dubbed the *Minimal Chain Principle* (MCP):

MCP: Avoid postulating unnecessary chain members at Sstructure, but do not delay required chain members.

The second clause of the MCP describes active gap filling, and the first clause states that when there is ambiguity as to whether a filler-gap dependency is present, an analysis that does not require such a dependency should be adopted in preference to an analysis that would require one. These two clauses of the MCP can be regarded as consequences of a single principle minimizing the maintenance of dependencies: Do not form dependencies that are not required, and when they are required, make them as short as possible. The motivation for the MCP is similar to the motivation for other structural constraints that have been argued to govern the behavior of the parser, such as Minimal Attachment and Late Closure (Frazier, 1978; Frazier, 1987). This motivation is to avoid the memory burden associated with maintaining unstructured material. The two clauses of the MCP have the effect of assigning each (noun) phrase a thematic role in the sentence as rapidly as possible, and assigning each thematic role to a (noun) phrase as rapidly as possible.

In the present work, we test a further prediction arising from the first clause of the MCP. In addition to predicting that the parser will avoid positing a gap when a filler has yet to be identified, this clause predicts that the parser should avoid analyzing material as the filler in a long-distance dependency if an alternative parse in which that material is not a filler is available.

In sentences containing *wh*-questions and relative clauses, there is typically no syntactic ambiguity upon reaching the filler; in (2), for example, encountering the relative pronoun *that* unambiguously introduces a relative clause, and therefore triggers a search for the corresponding gap position. But, this is not always the case. Consider (7):

(7) The information that the health department provided

This string is ambiguous between an analysis on which *that the health department provided* is part of a nominal complement clause (CC), and an analysis on which it is a relative clause (RC). These two analyses would give rise, respectively, to continuations like in (8a–b):

- (8) a. The information that the health department provided a cure reassured the tour operators.
 - b. The information that the health department provided reassured the tour operators.

On the RC analysis, the relative pronoun (*that*) would trigger a search for the corresponding gap site, which is located after the relative clause verb (*provided*). The CC analysis does not involve a filler-gap dependency, as the clause *that the health department provided a cure* directly modifies the noun *information*, containing no gap. Notably, this ambiguity only arises for nouns that can take clausal complements. The nouns that do take CCs include nouns denoting propositional attitudes (e.g., *idea, suspicion, conclusion*), nouns denoting speech acts (e.g., *claim, announcement, proposal*), and nouns relating to epistemic modality (e.g., *evidence, possibility, truth*). For a detailed discussion of the semantics of nouns that occur with CCs, and a semantic analysis of clausal complementation more generally, the reader is referred to Moulton (2015).

The MCP makes the prediction that the CC analysis of this ambiguity is initially preferred by the parser. Indeed, it predicts that this preference should be insensitive to the frequency with which a given noun (e.g., information) occurs with a CC, compared to an RC, with the parser adopting the CC analysis whenever the noun is one that can occur with a CC. There is evidence that the argument structure biases of specific verbs can play a role in how verbal complement ambiguities are parsed (e.g., Garnsey, Pearlmutter, Myers, & Lotocky, 1997; Trueswell, Tanenhaus, & Kello, 1993; though cf. Ferreira & Henderson, 1990; Kennison, 2001; Pickering, Traxler, & Crocker, 2000). But as noted above, previous work suggests that subcategorization biases play little role in the processing of filler-gap structures; a direct object gap is posited in an immediately post-verbal gap site even when the verb is biased not to occur with a direct object (Pickering & Traxler, 2003), and a direct object gap is not initially posited in Heavy NP Shift sentences even for verbs that are transitively biased (Staub et al., 2006). In the present work, we assess the frequency with which specific CC-taking nouns occur with a CC or an RC, and explore the potential role of this factor in the resolution of the critical ambiguity. It would be of substantial theoretical interest if the processing of filler-gap structures were, across the board, insensitive to the subcategorization biases of the lexical heads that could in principle help the processor determine if a dependency chain member is present.

Before describing our experiments, we review previous research that has examined processing of ambiguities similar to the one we test here. Chen, Gibson, and Wolf (2005, Experiment 2; see also Pearlmutter & Mendelsohn, 1999) examined a variant of the ambiguity we explore here in a self-paced reading experiment, comparing sentences like (9a–b):

- (9) a. The claim that the cop who the mobster attacked ignored the informant might have affected the jury.b. The claim that the cop who the mobster attacked
 - ignored might have affected the jury.

Chen et al. (2005) found that reading times were much longer on the start of the matrix verb phrase (*might have affected*) after an RC than after a CC. They interpreted this finding as reflecting a garden path effect, whereby the ambiguous material is initially processed as a CC, and encountering the matrix verb phrase in (9b) then requires the reader to engage in reanalysis. Moreover, Chen et al. found that a region consisting of the body of the *that*-clause itself (e.g., *the cop who the mobster attacked*) was read faster in (9b) than in an unambiguous RC condition in which the relative pronoun was *which* rather than *that*. This result is also consistent with the conclusion that the RC analysis. Furthermore, Chen et al. found no evidence that these results depended on the bias of the subject noun (e.g., *claim*) to occur with a CC.

These results should be regarded as merely suggestive, for several reasons. As the main theoretical question addressed by Chen et al. (2005) related to the storage costs associated with multiple embeddings, their experimental stimuli contained an additional RC (who the mobster attacked) inside the ambiguous that-clause, making the critical sentences extremely complex in both conditions; it is unclear what role is played by this added complexity. Second, and perhaps more important, the finding of an apparent garden path effect comes from comparing reading times on the matrix verb in the two critical conditions. This is the disambiguating region when the ambiguity is resolved toward the RC analysis. However, when the ambiguity is resolved toward the CC analysis, disambiguation is at the object noun phrase (e.g., the informant in 9a). Thus, the Chen et al. study compared reading times on material that serves as the disambiguating material in one condition, but comes after disambiguation in the other condition. Finally, though Chen at al. did make use of normative data regarding noun biases (from Kennison, 2000; Pearlmutter & Mendelsohn, 1999), most of the nouns that Chen et al. used were strongly CC-biased, and they included no nouns showing a clear RC bias. Thus, Chen et al. could conduct only limited analyses of the role of noun bias.

Other apparent evidence that an RC analysis is avoided comes from Altmann, Garnham, and Dennis (1992; see also Altmann, Garnham, & Henstra, 1994; Crain & Steedman, 1985; Mitchell, Corley, & Garnham, 1992). In two eyetracking experiments, Altmann, Garnham, and Dennis (1992) compared reading of sentences such as (10a-c):

- (10) a. The fireman told the woman that he had risked his life for to install a smoke detector.
 - b. The fireman told the woman that he had risked his life for many people in similar situations.
 - c. The fireman asked the woman that he had risked his life for to install a smoke detector.

In (10a–b), the string *that he had risked his life for* is initially ambiguous between an RC analysis, on which it modifies *woman*, and an analysis on which it is the initial portion of a clausal complement to the verb *told*. In conditions with no preceding context, Altmann et al. found a clear garden path effect upon disambiguation toward the RC analysis (10a), compared to disambiguation toward the CC analysis (10b), and compared to a control condition in which the CC analysis is ruled out due to the subcategorization restrictions of the main verb (*asked* in 10c). However, they also found that this garden path effect was either eliminated or reduced (see Rayner and Sereno (1994), and Altmann (1994), for discussion) when the RC was preceded by an appropriate context; e.g., (10a) was preceded by (11):

(11) An off-duty fireman was talking to two women. He was telling them how serious the situation had been when their house had caught fire. The fireman had risked his life to rescue one of the women while the other had waited outside.

Altmann et al. (1992; see also Crain & Steedman, 1985) argued that the CC preference in the null context should not be attributed to a general parsing principle that would putatively favor a CC analysis of the ambiguous material.¹ Rather, they proposed that the CC preference observed in the absence of context was the effect of a *Principle of Parsimony*, which favors the reading that carries fewer unsatisfied presuppositions. Altmann et al. argued that the RC analysis carries a presupposition that the CC analysis does not, namely that there are (at least) two possible referents for the critical noun (e.g., *the woman* in (10)) in the discourse. Thus, they argued, this analysis is disfavored if the RC is uttered out of the blue. In a context in which this presupposition is already satisfied, however, as in (11), the RC analysis is initially adopted in preference to the CC analysis.

Limited conclusions can be drawn from these experiments, too, for a number of reasons. First, interpretation of Altmann et al.'s (1992) finding of a CC preference in a null context is complicated by the fact that the verb *told*, which is the main verb in all of the items in the Altmann et al. experiments, has an extremely strong bias to appear with a CC. In the first 100 instances of the told the *INOUNI that* appearing in the Corpus of Contemporary English (Davies, 2008), there is not a single instance in which that introduces a relative clause; in all 100 cases, it introduces a complement to told. Thus, readers may adopt the CC analysis in the null context simply because they are sensitive to the overwhelming probabilistic bias for a *that*-clause in the relevant position to be a CC, rather than an RC. To the extent that the distribution of structures in a comprehender's language experience plays a role in disambiguation (e.g., Levy, 2008), a that-clause in this position should be analvzed as a CC.

Second, we dispute the characterization by Altmann et al. (1992) of the function of the context sentence in their experiments. We think that it is simply not true that a relative clause – even a restrictive relative clause – presupposes multiple possible discourse referents for the head. While it is true that one function of a relative clause is to identify the referent of the head noun in a context in which more than one potential referent has been introduced in the discourse, this is not its only function, nor even the dominant one. Consider (12a), where the relative clause simply provides additional information about the referent, much like the pre-nominal modifier in (12b).

- (12) a. I wanted to tell you that I really liked the presentation you made yesterday.b. I wanted to tell you that I really liked the keynote
 - presentation.

Clifton Jr. and Ferreira (1989) elaborate in more detail on the point that such restriction of reference, which does not presuppose a set of possible referents in the discourse, is a common function of relative clauses. They also provide the following striking example of a relative clause that cannot possibly presuppose multiple referents, which they attribute to Portner (1989):

(13) The unicorn that Richard saw was the only unicorn there ever was.

If the function of a context as in (11) is not to satisfy a presupposition of the relative clause in (10a), why does this context induce an RC analysis of the ambiguous *that*-clause? The answer builds on the fact that a singular definite description introduces a presupposition of uniqueness. For example, the definite description *the woman* cannot be used felicitously if two women are equally prominent in the context. Therefore, after the context in (11), where two women have been explicitly mentioned, the definite description *the woman* in the target sentence simply fails to refer unless it is followed by a relative clause. In other words, it is not the case that (11) *licences* a relative clause that would otherwise fail to have its presuppositions satisfied, but rather that this context *requires* a relative clause following the critical NP. Thus, it is unsurprising that a *that*-clause would be analyzed as an RC when analyzing it as a CC would lead to outright referential failure.

In sum, we regard it as not at all settled whether there is, in fact, a general tendency to analyze an ambiguous *that*-clause as a CC rather than as an RC in a null context, as the results of both Chen et al. (2005) and Altmann et al. (1992) are open to various interpretations. In particular, lexical biases supported the CC analysis in both of these studies.

In a more recent eyetracking study, Vernice, Cecchetto, Donati, and Moscati (2016) tackled the temporary ambiguity between a CC and an RC analysis in Italian and English, and found a preference for the latter, in contrast with the findings just discussed. However, only seven different nouns were used, so again the study had limited ability to investigate the role of noun bias.

Here we systematically explore the role of noun bias in the processing of the CC/RC ambiguity. We also explore the extent to which on-line processing behavior for a given item is predictable from off-line sentence completion norms provided for the same item (e.g., Taraban & McClelland, 1988). The items we use here demonstrate a high degree of variability in the tendency of the ambiguous clause to be completed as a CC or an RC in an off-line completion task. This is to be expected given the variability in the lexical content of these clauses, which gives rise to variability in the plausibility of the RC interpretation. We explore the extent to which on-line processing of this ambiguity reflects an item's completion bias, which we take to be a measure of the relative plausibility of the RC and CC interpretations of the ambiguous material.

In sum, we present three eye movement experiments examining processing of a syntactic ambiguity that constitutes a test case for the first clause of the MCP (De Vincenzi, 1991), according to which a filler-gap dependency will not be constructed when a syntactic analysis that does not involve such a dependency is available. We systematically explore the role of both noun bias and the item's completion bias in resolution of this ambiguity. To anticipate, the experiments appear to support a strong version of the MCP, holding that the CC analysis is initially constructed regardless of the head noun's bias, and regardless of a specific clause's tendency to be completed as a CC or as an RC.

Experimental stimuli: noun bias and completion bias

The three experiments presented in this study made use of sentences in which a noun was followed by a CC or by an RC. This head noun was one of 34 different CC-taking nouns. In order to provide a convincing test of the role of the noun's subcategorization bias in on-line processing of the ambiguous *that*-clause, it was necessary that these 34 nouns did indeed vary substantially in their bias. Specifically, we hoped to include both nouns that frequently occurred with a CC, and nouns that rarely did so. We also hoped that the ambiguous fragments that appeared in our experiments (e.g., *The information that the health department provided*) varied substantially in how readily they elicited CC completions in an off-line task, so that we could convincingly test the role of completion bias. We assume that a fragment's completion bias is a mea-

¹ The parsing principle that Altmann et al. (1992) and Crain and Steedman (1985) consider as potentially favoring the CC analysis over the RC analysis is Frazier's (1978) principle of Minimal Attachment. This is clearly a distinct account from the one we have put forward here, according to which the RC analysis is avoided due to the general tendency to avoid non-required filler-gap dependencies. In the General Discussion we discuss the two alternatives, and consider the question of which account better explains a preference for the CC analysis.

sure of how readily the head noun is interpreted as the object of the embedded verb; e.g., how felicitous is the situation in which a health department provides information. If this interpretation is felicitious, the fragment will often be continued as an RC, but if it is not, the fragment will tend to elicit CC completions. Fortunately, we were able to construct stimuli in which both noun bias and completion bias ranged nearly uniformly across the full possible range, and in which these two variables were only moderately correlated.

Noun bias

For each noun, the phrase *the x that* (e.g., *the information that*) was entered as a search term in Google.² For each search, one of the authors, a trained linguist, coded the first 40 unique hits (excluding the few assorted cases in which the text did not actually include one of the target structures) as either a CC, a subject relative clause (SRC), or an object relative clause (ORC). In (14) we provide examples of each of the three structures occurring from the search of the fragment *the prediction that* (conducted on July 15, 2015 by means of the Google search engine based in the U.S.):

(14)	a. The prediction that this region is overdue for a large
	earthquake (Ishibashi, 1981) is primarily based on the
	observation that sector D did not rupture in the Tokaido
	earthquake of 1944 and hence has not ruptured since
	the Ansei I event of 1854. (CC)
	(https://books.google.co.uk/)
	b. The prediction that comes forth from a hypothesis
	that one creates based on observation of nature should
	be tested through experimentation that involves a
	specific set of methods. (SRC) (http://www.geo.arizona.

edu/Antevs/nats104/wk1_lect_hyp2.html)

c. During this meeting, the prediction that the three witches make about Banquo is that his sons would be kings. (ORC)

(https://www.enotes.com/homework-help)

We note that 40 occurrences are sufficient to obtain reasonably precise estimates of the population bias for a given noun. For example, for a noun that occurs with a CC on 5 out of 40 occasions, the limits of the 95% binomial confidence interval are 0.05 and 0.26; for a noun that occurs with a CC on 20 of 40 occasions, the limits of the 95% confidence interval are 0.35 and 0.65.

A given noun's bias to occur with a CC was then determined in two ways: CCs as a proportion of all continuations (i.e., out of 40), and CCs as a proportion of CC or ORC continuations, i.e., not including SRCs. The former bias is potentially relevant to ambiguity resolution at *that*, where the post-nominal material can be attached as a CC or RC. The latter bias is potentially relevant to ambiguity resolution once a subsequent noun phrase has been encountered, as at this point the SRC analysis is no longer possible. We refer to the former as *inclusive CC bias*, and the latter as *exclusive CC bias*.

Fig. 1a and b shows the results for inclusive and exclusive CC bias, respectively. The mean inclusive CC bias is 0.55, indicating that across these 34 nouns, sentences including *the* x *that* are

slightly more likely to involve CCs than RCs, by a ratio of 11-to-9. The mean exclusive CC bias is 0.70; once SRCs have been eliminated, CCs outnumber ORCs by more than 2-to-1. The two bias measures are correlated r = 0.93. This level of correlation suggests that it is likely to be impossible to empirically distinguish effects of the two measures of bias. Therefore, we use only inclusive CC bias – as defined above – as a measure of noun bias in this study, and refer to it simply as *noun bias*.

Fig. 1a shows that the 34 nouns are roughly uniformly distributed in terms of their bias. It is important to note that this variability is much greater than would be expected by chance alone. Based on a binomial distribution where the probability of a CC is 0.55 (i.e., equal to the mean over all nouns), the probability of obtaining 10 or fewer CCs out of 40 is less than 0.00001, and the probability of obtaining 30 or more is about 0.01. But in fact, many nouns had 10 or fewer or 30 or more CCs. Thus, our method of assessing noun bias reveals meaningful variability among the nouns in their tendency to occur with a CC, i.e., variability that is not due to sampling variability alone. Individual nouns do reliably differ in their tendency to occur with a CC. If this had not been the case, no exploration of the role of noun bias in comprehension could possibly reveal a role of this factor.

Two nouns (*law* and *resolution*) did not occur with a CC in the 40 instances that we examined, and two nouns (*claim* and *probability*) occurred only with CCs. It may be argued, therefore, that there is no actual ambiguity when a *that*-clause follows these nouns, as the former do not allow a CC, and the latter require one. But neither of these claims would be correct. Examples (15) and (16), which we take to be clearly grammatical, show that the former nouns can indeed occur with a CC:

- (15) The law that people must stop at a red light is often violated.
- (16) The resolution that the university would divest its holdings was approved by the trustees.

Similarly, examples (17) and (18) show that the latter two nouns can be followed by ORCs:

- (17) The claim that the witness made turned out to be false.
- (18) The probability that the pollster computed was not seen
- as plausible by her colleagues.

Completion bias

We then constructed a set of sentence fragments, two for each of the 34 nouns, which were intended to be (at least to some extent) ambiguous between a CC and an RC analysis, as in (19):

b. The information that Bob released	(19)	a. b.	The information that the health department provided The information that Bob released
--------------------------------------	------	----------	---

These fragments formed the initial portion of the complete sentences that we used in our eyetracking studies, as described below. To assess the extent to which the ambiguous material in each of the 68 fragments elicits an analysis as a CC or as an ORC in an untimed, 'off-line' task, we conducted a sentence completion study using Amazon Mechanical Turk. Thirty-three subjects who selfidentified as native speakers of American English, and who responded from IP addresses in the U.S., each completed 34

² Some of these search terms occur very rarely in corpora that are standardly used in psycholinguistic studies, necessitating the use of the Google search method. For example, the phrases *the guess that* and *the notice that* occur 6 and 11 times, respectively, in the 52-million word Corpus of Contemporary American English (COCA). For phrases that do appear in COCA with some frequency, there appears to be a high degree of consistency between results from the two methods; for example, in COCA *the law that* almost never introduces a CC, *the probability that* almost always introduces a CC, and *the evidence that* introduces a CC about half of the time, mirroring the Google results for these nouns shown in Fig. 1.



Fig. 1. Distribution of inclusive (a) and exclusive (b) CC noun bias, as defined in the text.

fragments, one for each noun, in exchange for payment. Each of the 68 fragments was completed by 16 or 17 subjects. The 34 critical items in each list were intermixed with 16 unrelated fillers. The experiment was presented using Javascript on the Ibex Farm (http://spellout.net/ibexfarm/) server. Subjects were instructed to complete each sentence by typing "a sentence continuation that seems like a natural continuation of the fragment." Two subjects provided a large number of incomplete or uninterpretable responses, and were excluded from analysis. Of the 1054 responses provided by the remaining 31 subjects, all but 33 (which were removed from analysis) could be categorized as involving either a CC or ORC analysis of the ambiguous material. Each of the four (linguistically-trained) authors of this paper coded some of the responses. Spot-checking revealed no disagreements.

The distribution of CC-biases across the 68 fragments is shown in Fig. 2a. On average, the proportion of CC completions was 0.41: CCs are less common in this completion study than in the corpus results described above, where CCs and ORCs were in a 70/30 ratio. It is not possible to determine to what extent this is due to the specific content of the items we constructed, and to what extent it is due to the demands of the task; each fragment constituted a complete relative clause, while an interpretation as a CC required the subject to complete the clause. In other words, there may have been a general tendency to interpret the endpoint of the provided material as a clause boundary. Nevertheless, like the noun CC biases obtained from Google, the proportions of CC continuations obtained in this completion study are widely varying, from 0 (for six different fragments, e.g., The answer that the client provided) to 0.94 (The fear that Italy would declare). It is clear that even if the demands of the task resulted in some bias toward ORC continuations, this bias must not be very strong, as a substantial number of items did elicit mostly CCs.

The correlation was r = 0.46 between a fragment's completion bias and the noun bias for the fragment's head noun. (If the completions for both items with a given head noun are aggregated, this correlation increases to r = 0.56.) Fig.2b shows this relationship. Thus, some of the variability in completion bias across the fragments can be attributed to variability in the tendency of the fragment's head noun to take a CC, as would be expected. However, as R^2 is only 0.21, most of this variability cannot be attributed to the noun. The plausibility of the ambiguous material itself as a CC or RC likely plays an important role. In sum, the 34 nouns on which our materials are based vary widely in how frequently they occur with a CC, relative to an RC, in internet-based texts. The specific sentence fragments we use here also vary widely in their tendency to elicit a CC completion. While some of the variability in the rate of CC completion is due to the bias of the noun to appear with a CC, most of this variability cannot be attributed to the noun itself. Noun bias and completion bias are sufficiently independent that it is possible to separately assess the impact of each variable in the experiments we report here.

Experiment 1

This experiment was a first attempt to determine if there is an on-line parsing preference in favor of either an RC or a CC analysis, using eyetracking during reading. We constructed 68 sentence pairs by completing each of the fragments described above as either a CC (a) or an RC (b), as in (8) above, repeated here as (20):

(20)	a.	The information that the health department
		provided a cure reassured the tour operators.
	b.	The information that the health department
		provided reassured the tour operators.

All sentence pairs were identical except for the addition of an embedded clause object (e.g., a cure) in the CC version. The goal of the experiment was to compare the on-line effects of disambiguation toward either the CC or the RC analysis. To that end, our analyses focused on eye movement measures on the disambiguating material in the two versions (cf. Chen et al., 2005), i.e., the object noun phrase in the CC version (e.g., a cure in 20a), and the beginning of the matrix verb phrase in the ORC version (e.g., reassured in 20b); details of region definitions are provided below. Difficulty would be expected upon disambiguation in one of the conditions if there is a general preference for the other analysis. If they are equally preferred, there would be no main effect of the experimental manipulation. They may be equally preferred either because they are preferred on an equal proportion of trials, or because they are maintained with equal activation in a parallel manner. We discuss the issue of seriality vs. parallelism in the



Fig. 2. Panel (a) plots the distribution of CC completion bias across the 68 sentence fragments in the completion study and panel (b) plots the correlation between a fragment's completion CC bias and the noun bias for the fragment's head noun, with linear regression line.

General Discussion. Any preference for one analysis may also be modulated by either the head noun's bias to take a CC or an RC, or the completion bias for the sentence in question. We acknowledge at the outset that interpreting processing differences between regions containing different lexical material, and in different syntactic positions, is not straightforward. We address this issue with a different experimental design in Experiments 2 and 3, in which an identical disambiguating region is compared across conditions.

Method

Participants

Participants were 59 students at UMass Amherst who received course credit for their participation. All were speakers of English as a first language, and none reported any history of reading or language disorder. One subject was excluded due to chance performance on comprehension questions. Ten subjects were excluded based on a pre-established criterion of losing more than 20% of experimental trials to track loss or blink on the critical region, leaving 48 subjects in the analysis.

Materials

The 68 items are presented in Appendix A. These items were arranged into two experimental lists, so that each participant read two sentences with each head noun, with one continuing as a CC and one as an ORC. For example, a participant who reads (20a) would read (21b) below, and a participant who reads (20b) would read (21a).

(21)	a.	The information that Bob released the secret was
		surprising to me.

b. The information that Bob released was surprising to me.

These critical sentences were randomly intermixed with 75 other sentences, of which 24 were from an unrelated experiment on subject-verb agreement. A total of 68 sentences (17 experimental sentences and 51 fillers) were followed by two-choice comprehension questions. The 143 sentences were preceded by eight practice sentences.

Procedure

Subjects were tested individually. Eye movements were recorded using an EyeLink 1000 (SR Research, Toronto, ON, Canada) eyetracker, interfaced with a PC computer. The sampling rate was 1000 Hz. Subjects were seated 55 cm from a CRT monitor on which the sentences were displayed. At this distance the resolution of the eyetracker was substantially less than one character. Only the movement of the right eye was recorded.

All sentences were displayed on a single line in 11-point Monaco font. Before the experiment began, each subject was instructed to read for comprehension. A three-point calibration procedure was performed at the start of the experiment and as needed between trials. The subject triggered each sentence by fixating a box at the left edge of the monitor. The experiment lasted approximately 40 min. The experiment was implemented with the Eye-Track software, and initial stages of data analysis were carried out with Robodoc and EyeDry (http://blogs.umass.edu/eyelab/software/).

As noted above, one subject was excluded due to chance performance on comprehension questions; all others achieved at least 80% correct. Trials were excluded if there was a blink or track loss during first pass reading of the critical region; see below for region definitions. As noted above, ten subjects lost more than 20% of trials on this basis, and these subjects were excluded from subsequent analysis. For the remaining 48 subjects, blinks, track loss or other error resulted in deletion of 6.7% of trials. In addition, one item was removed from all analysis due to an error in the preparation of the script, leaving 3002 trials for inclusion in the analysis. Individual eye fixations less than 80 ms in duration and within one character of a previous or subsequent fixation were incorporated into this neighboring fixation. No other data trimming was carried out.

Results

We focus our analyses on the disambiguating region of each sentence, which is the object noun phrase in the CC condition (*a cure* in 20a; *the secret* in 21a) and the initial word or words of the matrix verb phrase in the ORC condition (*reassured* in 20b; *was surprising* in 21b). This latter region consisted of from one to four words, selected so as to match in length, to the extent possible, the critical region in the CC condition. This region always included the matrix verb. The disambiguating verbal material in the RC condition was on average slightly longer (mean of 13.04 characters) than the disambiguating NP in the CC condition (mean of 11.37 characters); region length is included as a predictor in the models we report below.

We report four eye movement measures on this region. *First fix-ation duration* is the duration of the first eye fixation on the region, on the first pass through the sentence. *First pass time* is the sum of the duration of all first pass fixations on the region, before the eyes first leave the word to either the left or right. If the reader made

only a single first pass fixation before leaving the region, first fixation duration and first pass time are identical for that trial. *Go-past time* is the sum of all fixation durations beginning with the first on the region, and including all fixations until the reader moves past the region to the right; this measure includes the durations of any fixations made after a regressive eye movement to the left. Finally, *regression proportion* is the proportion of trials on which first pass inspection of a region ended with an eye movement to the left rather than the right. For all of these measures, a trial is excluded if the region was skipped on first pass reading; this was rare due to the length of the critical region.

Statistical analysis was carried out by means of linear mixed effects models of reading times and logistic mixed effects models of regression probability, implemented using the lme4 package (version 1.1-7: Bates, Maechler, Bolker, & Walker, 2015) for the R statistical programming environment (version 3.1.2; R Core Team, 2014). The modeling strategy was as follows. For each measure, we began with a model that included condition as a fixed effect, with the CC condition receiving a value of -0.5, and the RC condition receiving a value of 0.5. Centered region length was also entered as a fixed effect. Random factors included intercepts for subjects and items, as well as random subject and item slopes for the condition factor. We then constructed two models in which an additional factor was included: one in which centered noun bias was included as a fixed effect, along with its interaction with condition; and one in which the fragment's centered completion bias was included, along with its interaction with condition. We compared model fits for each of these models to the initial model, by means of log likelihood comparison, and report the larger model when the improvement in fit was significant at p < 0.05, based on the Chi-Square distribution with df equal to the difference in number of parameters, where the value of the statistic is equal to two times the log likelihood difference.

For the three reading time models, adding completion bias and its interaction with condition improved model fit, and the results of these models are shown in Table 1. Including noun bias and its interaction with condition did not improve model fit except in the case of first pass time, where the interaction term just reached significance (t = 2.06). We constructed an additional model of first pass time with both completion bias and noun bias as predictors, and this model did not fit better than the model with completion bias only, though it did fit better than the model with noun bias only. Moreover, only the completion bias interaction was significant in this larger model. This suggests that the significance of noun bias as a predictor is likely due to its correlation with completion bias, so as in the case of first fixation duration and go-past time, we report the model that includes completion bias only. We note that models of log-transformed reading times did not reveal any differences in patterns of significance.

For all three reading time measures, the results are similar. There is a large and significant effect of condition (treating |t| > 2 as significant) in first pass time and go-past time, with slower reading times upon disambiguation toward an RC. In first fixation duration this effect does not reach significance but trends in the same direction. This effect is estimated at 8 ms in first fixation duration, 45 ms in first pass time, and 138 ms in go past time. There are also significant length effects on all measures, with longer disambiguating regions eliciting shorter first fixation durations, and longer first pass and go-past times. Finally, there is a significant interaction in each measure between completion bias and condition, in the form of slower reading time when a sentence with a strong bias toward CC completion is disambiguated as an RC.

For the regressions out measure, neither of the more complex models including noun bias or completion bias improved on the fit of the simplest model. We therefore report the results of this simple model in Table 1. The effect of condition is significant; overall there were 22.2% regressions in the RC condition and 16.7% in the CC condition. The effect of region length was not significant.

The pattern in the reading time measures is shown in Fig. 3. For visualization, we have adjusted for region length by computing residual go-past time on each trial, based on a linear model in which region length is the sole predictor. Each point on the plot represents an item's go-past time penalty upon disambiguation toward an RC, i.e., the mean residual go-past time in the RC condition minus the mean residual go-past time in the CC condition. Fig. 3a illustrates that this penalty increases linearly as a function of completion bias. The penalty tends to be large for items that are biased toward a CC completion, while it is neutralized for items that are biased toward an RC completion. It is notable, however, that there is no reversal: RC disambiguation does not actually become easier than CC disambiguation. An extreme RC bias in the completion study merely eliminated the difficulty associated with disambiguation toward an RC. Fig. 3b illustrates that there is little apparent effect of noun bias, though the trend does go in the same direction.

Discussion

Readers showed substantial difficulty when the critical ambiguity was resolved as an RC rather than as a CC. This difficulty appeared in the form of increased first pass time and go-past time on the disambiguating material in the RC condition, as well as an increase in the probability of regressing from this material. These results suggest that readers initially adopted the CC analysis, and were forced to reanalyze upon encountering the main clause verb.

This difficulty was modulated by an individual item's tendency to elicit CC continuations in the completion study. The difficulty was most severe for items that elicited mostly CC continuations, and was essentially eliminated for items that elicited few, or no, CC continuations. Interestingly, however, the effect was never reversed, as there was no subgroup of items for which disambiguation toward a CC analysis was more difficult.

This pattern suggests that completion bias may interact with the RC penalty by means of an effect on reanalysis difficulty, rather than an effect on the probability of initially adopting one analysis or the other. When the fragment is one that is not easily interpreted as an RC, reanalysis is difficult, but when the fragment is one that is very easily interpreted as an RC, reanalysis is very easy. Because the RC analysis is not initially adopted even when the clause tends to elicit RC continuations, there is no penalty on disambiguation toward the CC analysis. The fact that noun bias did not significantly interact with condition is also consistent with a general initial preference for the CC analysis. Noun bias, if it were to have an effect, would presumably influence the reader's initial parsing decision, as lexically-specific subcategorization information is available early in incremental processing of the sentence. On the other hand, information regarding the plausibility of the ambiguous material as a CC or an RC can only be available late, after this material has been read. However, we postpone further discussion of these effects because, to anticipate, Experiments 2 and 3 do not replicate the interaction with completion bias observed in Experiment 1.

Although the results obtained in Experiment 1 suggest that a CC analysis is initially adopted, independent of noun bias, the finding of difficulty upon disambiguation toward an RC may be open to a range of interpretations due to the nature of the materials used in the first experiment. Specifically, Experiment 1 compared very different disambiguating material in the two conditions, an object noun phrase that disambiguates toward a CC analysis and the beginning of a verb phrase that disambiguates toward an RC analysis. The two regions did not necessarily contain the same number of words for a given item; the object noun phrase that

Table 1

Free of the sector of the		14 - 6	···· ··· · · · · · · · · · · · · · · ·		1.1		L	41	
EXDELILIELL	STATISTICAL	results from	i iiiixeu-enecis	s models, wo	iel specificatioi	I IS DESCH	Ded III	une i	lexi

Measure		Estimate	SE	<i>t/z</i> -value
First fixation duration	Condition	8.37	4.67	1.79
	Completion bias	17.41	9.27	1.88
	Length	-1.74	0.68	-2.58
	Cond \times Comp bias	40.71	13.83	2.94
First pass time	Condition	45.28	13.86	3.27
	Completion bias	33.78	18.19	1.86
	Length	23.99	1.62	14.77
	Cond × Comp bias	160.41	40.32	3.97
Go past time	Condition	138.49	23.78	5.82
	Completion bias	141.42	47.87	2.95
	Length	27.68	3.15	8.78
	Cond \times Comp bias	258.18	75.92	3.40
Regression probability	Condition	0.45	0.15	2.98
-	Length	-0.00	0.02	-0.02



Fig. 3. Experiment 1 go-past time penalty for RC disambiguation, by item, as a function of completion CC bias (a) and noun CC bias (b). Points reflect differences in lengthresidualized go-past time, as described in the text. Dashed line is linear regression line.

disambiguated toward the CC analysis ranged from one word (e.g., *animals*) to three (e.g., *a new galaxy*), while the region that disambiguated toward the RC analysis ranged from one word (e.g., *reassured*) to four (e.g., *is not actually stated*). Moreover, lexical frequency was not controlled in these regions. Thus, the processing disadvantage on the matrix verb phrase could, in principle, be due to uncontrolled lexical differences between the critical regions.

More concerning, however, is the possibility that some or all of the difficulty on the matrix verb in the RC condition is due to the difficulty of processing the object relative clause itself, rather than syntactic garden-pathing. Indeed, Staub, Dillon, and Clifton (2017) have shown that there is difficulty on the matrix verb when it immediately follows an unambiguous object gap. Though the present effect is larger than the effects observed by Staub et al., it is important to rule out this alternative explanation. Experiments 2 and 3 were designed to achieve this.

Experiment 2

To address the issue raised above, Experiment 2 re-used 34 of the RC sentences from Experiment 1 – one with each subject noun. However, instead of comparing processing of these sentences to matched CC sentences, this experiment compared each one to another RC sentence, differing only in that the subject noun does not allow a clausal complement, e.g., *vaccination* in (22a):

(22)	a.	The vaccination that the health department
		provided reassured the tour operators.
	1	

b. The information that the health department provided reassured the tour operators.

The design of Experiment 2 allows an even stronger test of the claim that the parser initially adopts the CC analysis when it is available, as suggested by Experiment 1. In this design, any difficulty at the start of the matrix verb phrase in the (b) condition is more clearly attributable to syntactic garden-pathing (i.e., initial adoption of a CC analysis) and subsequent reanalysis.

Previous studies have established that the processing of unambiguous ORCs causes incremental difficulty in the eye movement record (Gordon, Hendrick, Johnson, & Lee, 2006; Staub, 2010; Staub et al., 2017; Traxler, Morris, & Seely, 2002). This difficulty, which appears most strongly on the embedded subject (e.g., the health department in 22), but also appears on the embedded verb (e.g., provided in 22), has been argued to reflect both violation of syntactic expectations (e.g., Hale, 2001; Levy, 2008) and memory retrieval difficulty (Gibson, 1998; Grodner & Gibson, 2005). In previous experiments, this difficulty has been demonstrated by comparing processing of ORCs to both SRCs (Gordon et al., 2006; Staub, 2010; Staub et al., 2017; Traxler et al., 2002) and verb complement clauses (Staub, 2010). This difficulty during the incremental processing of unambiguous ORCs gives rise to a novel prediction for the present experiment. If the ambiguous material is analyzed as a CC when the noun allows this analysis (e.g., information), first pass processing of the that-clause itself (e.g., that the health department provided) should be less difficult in this case than when the noun does not allow the CC analysis (e.g. vaccination). Only in the latter case would the ORC analysis be constructed, as it is the only available analysis, and only in this case would the traditional ORC penalty be evident in the processing of the that-clause.

In short, the claim that the ambiguous material is parsed as a CC when the head noun allows this analysis predicts, in the present experiment, a full 'crossover' pattern of processing difficulty. The ORC itself should be more difficult to process when the noun does

not allow a CC analysis (e.g., *vaccination*), but there should be difficulty at the matrix verb, upon disambiguation toward the ORC analysis, when the noun does allow the CC analysis (e.g., *information*).

Method

Participants

Participants were 46 subjects from the same pool as Experiment 1, but who did not participate in the previous study. One subject was excluded due to poor performance on comprehension questions (72% accuracy; no other subject was below 80%), and four were excluded based on losing more than 20% of experimental trials to track loss or blink on the critical region, leaving 41 subjects in the analysis.

Materials

Thirty-four of the items from Experiment 1, one for each head noun, were adapted for this experiment as described above. These items were arranged into two experimental lists, so that each subject read 17 items in each of the two experimental conditions. These 34 items were randomly intermixed with 86 other sentences, of which 36 were from an unrelated experiment on subject-verb agreement. The 120 sentences were preceded by eight practice sentences. A total of 50 items were followed by twochoice comprehension questions.

Procedure

The procedure was identical to Experiment 1. As noted above, four subjects lost more than 20% of trials on the basis of blinks or track losses on the critical region (the same region at the start of the matrix verb phrase used in Experiment 1), and these subjects were excluded from subsequent analysis. For the remaining 41 participants, blinks, track loss or other error resulted in deletion of 7.5% of trials, leaving 1290 trials for inclusion in the analysis.

Results

We analyzed eye movements on the relative clause (e.g., *that the health department provided*) and the subsequent disambiguating region (e.g., *reassured*). We analyze the entire clause as a single region (cf. Staub, 2010; Staub et al., 2017), because the present focus is not on the details of incremental relative clause processing. Moreover, we did not have specific predictions for how the additional difficulty on this region when the head noun does not allow a CC should be manifested on each of the constituents within the RC.

For each of the two critical regions, we computed the same four eyetracking measures as in Experiment 1. We follow the same modeling approach as in Experiment 1, with the exception that, because identical material is now being compared across the two conditions of each item, we remove region length as a predictor. The fit of the simplest model, with only condition as a fixed effect, was never improved upon by the more complex models that included either completion bias or noun bias. The results of the statistical models are shown in Table 2.

On the relative clause region, the effect of condition was not significant in first fixation duration or first pass time, but was significant in go-past time and regressions out. These effects are in the direction of shorter go-past time and fewer regressions when the noun does permit a CC. The regressions effect corresponds to overall regression rates of 24.8% when the noun does not permit a CC, and 17.8% when it does.

On the disambiguating region, the pattern is reversed, as shown in Fig. 4. The main effect of condition was not significant in first fixation duration or first pass time, but was fully significant in go-past

Table 2

Experiment 2 statistical results from mixed-effects models. Estimates are for effect of experimental condition; as described in text, models with additional fixed effects were not justified by model comparison.

Measure	Estimate	SE	<i>t/z</i> -value
Relative clause region			
First fixation duration	-8.09	5.11	-1.58
First pass time	46.13	23.37	1.63
Go past time	-67.74	30.45	-2.23
Regression probability	-0.83	0.25	-3.24
Disambiguating region			
First fixation duration	2.79	5.30	0.53
First pass time	13.25	11.28	1.18
Go past time	79.35	25.30	3.14
Regression probability	0.47	0.20	2.35

time and the regressions out measure. In this case the significant effects are in the direction of longer go-past time and more regressions when the noun does permit a CC. The regressions effect corresponds to 13.7% regressions when the noun does not permit a CC, and 20.0% when it does.

Neither completion bias nor noun bias interacted significantly with condition in the models that included these factors, and, as noted above, these models were not justified by model comparison. However, the trends for the interaction of completion bias and condition were in the expected direction on both regions. On the relative clause region, the go-past advantage for the condition in which the noun does permit a CC was numerically more pronounced when the fragment was biased toward a CC completion $(\beta = -66.91, SE = 109.20, z = -0.61)$. On the disambiguating region, the penalty was larger for the condition with a CC-permitting noun when the clause tended to be completed as a CC (β = 83.42, SE = 86.21, z = 0.97). For the noun bias factor, on the other hand, the interactive trend on the relative clause region was negligible $(\beta = -25.82, SE = 90.68, z = -0.29)$, and on the disambiguating region, it was actually in the unpredicted direction ($\beta = -67.86$, SE = 70.91, z = -0.96). Fig. 5 illustrates the go-past effects on each of the two regions, as a function of completion bias and noun bias. Each point represents the mean go-past time in the condition with a CC-permitting noun minus the mean go-past time in the condition with a non-CC-permitting noun, for a specific item. These differences are mostly negative on the relative clause, and are mostly positive on the disambiguation region, and do not seem to depend on either completion bias or noun bias.

Finally, we conducted an additional exploratory analysis motivated by the observation that because all 34 of the critical sentences were disambiguated as relative clauses, the experiment provides an ideal circumstance to test for the presence of syntactic adaptation (e.g., Fine, Jaeger, Farmer, & Qian, 2013). Fine et al. found that a garden path effect can be eliminated over the course of an experiment, or even reversed, which they attributed to comprehenders adapting their expectations to the statistics of the recent input. This account predicts attenuation, over the course of the experiment, of both the advantage on the relative clause region when the noun allows a CC and the advantage on the disambiguating region when the noun does not allow a CC.

To test this prediction statistically, we computed mixed effects models of go-past time on each region, in which (centered) condition, (centered) trial order (1 through 34), and their interaction were treated as fixed effects. We employed the maximal random effect structure. On the relative clause region, this analysis found the expected main effect of condition ($\beta = -66.21$, SE = 32.19, t = -2.06), a main effect of order, indicating faster reading as the subject proceeds through the experiment ($\beta = -6.67$, SE = 1.69, t = -3.94), and a nonsignificant interaction ($\beta = -4.24$, SE = 3.56, t = 1.19). Though this interaction does not approach significance,



Fig. 4. Mean go-past time with standard errors (a) and regression proportions (b) on each region, by noun status as CC-permitting or not CC-permitting.



Fig. 5. Experiment 2 go-past time penalty in the CC-permitting condition, by item, on relative clause regions (panels a and b) and disambiguation region (panels c and d) as a function of completion CC bias and noun CC bias. Dashed line is linear regression line. Negative values represent an advantage in the CC-permitting condition.

the fact that the trend conforms with the adaptation account can be seen clearly in Fig. 6. On the disambiguating region, on the other hand, there was the expected main effect of condition (β = 79.29, SE = 26.29, *t* = 3.02), but neither the effect of order (β = -1.71, SE = 1.44, t = -1.19), nor the interaction effect ($\beta = -0.87$, SE = 2.57, t = -0.34) neared significance. In Fig. 6, the garden path effect is as clearly in evidence at the end of the experiment as at the beginning.



Fig. 6. Experiment 2 mean go-past time by condition and trial order, for the relative clause region and disambiguation region. Dashed lines are linear regression lines.

Discussion

When the head noun preceding a relative clause does not allow a CC, there is difficulty on the relative clause itself, compared to when the head noun does allow a CC. This difference is expected if the relative clause is parsed as a relative clause only when the head noun does not allow a CC, as only in this case does the relative clause cause the incremental processing difficulty that is seen in unambiguous ORCs. The difficulty on the relative clause reached significance only in regression-based measures, i.e., go-past time and regressions out. This result is consistent with previous studies, which have shown that most of the difficulty induced by an ORC is due to the tendency of the ORC subject to induce regressive eve movements (Staub, 2010; Staub et al., 2017). Indeed, the presence of a nonsignificant but reversed trend in first pass time is also consistent with some previous experiments (Staub, 2010, Experiment 1; Staub et al., 2017, Experiment 2). First pass reading measures on an ORC can be especially short, as first pass reading is often truncated by a regression out of the region.

This difference between conditions is reversed when the reader reaches the disambiguating main clause verb; now there is difficulty when the head noun does allow a CC, compared to when it does not. Like on the relative clause, the difference between conditions reached significance only in the go-past and regressions measures. In Experiment 1, there was difficulty on disambiguation toward the RC analysis in first pass time as well. However, the more salient difference between the two experiments is that the reading time effects on the disambiguating region in Experiment 2 are smaller across the board. The estimate of the go-past time effect, which was significant in both experiments, was reduced from 138 ms to 79 ms. A natural inference from this difference is that while a large part of the disambiguation effect in Experiment 1 was due to syntactic garden-pathing, some of the effect was also due simply to the difficulty that occurs on the matrix verb following an object gap, even in unambiguous structures (Staub et al., 2017). In Experiment 2, the entire effect can be attributed to syntactic garden-pathing, as the matrix verb would have induced this difficulty in both conditions.

Another salient difference between the two experiments is in the role of completion bias. In Experiment 1, this factor interacted significantly with condition, as shown in Fig. 3. In Experiment 2, this interaction never neared significance, though there were trends in the expected directions on both the relative clause region and the disambiguating region. However, Experiment 2 reinforces the conclusion from Experiment 1 that noun bias does not modulate the processing of these sentences.

In sum, Experiment 2 confirmed the detailed predictions of an account according to which the CC analysis of a *that*-clause is first

adopted if the head noun allows it. When the noun does not allow a CC, the RC analysis is adopted, and difficulty occurs on the *that*clause itself. On the other hand, when the noun does allow a CC, difficulty occurs on the matrix verb phrase that disambiguates toward the RC analysis. The effects on the disambiguating material were smaller in Experiment 2 than in Experiment 1, suggesting that some part of the effect in Experiment 1 may be attributable to a general difficulty on the matrix verb after an ORC. However, a substantial garden path effect still survives. Experiment 2 failed to replicate the significant interaction effect with completion bias on the disambiguating region, though the nonsignificant trends did go in the expected direction.

Finally, Experiment 2 also suggested that the reader's preference for the CC analysis of this ambiguity is resistant to syntactic adaptation. Though the ambiguous string in this experiment was resolved as an RC on every single occasion, and no CCs occurred in the experiment, the garden path effect was clearly still in evidence at the end of the experiment.

The crossover pattern that was obtained in this experiment demonstrates a very high degree of incrementality in the processing of these sentences, with the pattern on the disambiguating region emerging despite the potential for an opposing 'spillover' effect from the preceding relative clause region. A main goal of Experiment 3 was to assess the replicability of this striking pattern. In addition, Experiment 3 provided an opportunity to further clarify the role of completion bias.

Experiment 3

Experiment 3 used the same sentences with CC-permitting nouns that were used in Experiment 2 (e.g., 22b). In this experiment, however, the control condition eliminated the possibility of the CC analysis in a different way. Instead of using a head noun that did not allow a CC, as in Experiment 2, in this third experiment the control sentences used *which* rather than *that* as the relative pronoun, as shown in (23a).

- (23) a. The information which the health department provided reassured the tour operators.b. The information that the health department provided reassured the tour operators.
- Because the relative pronoun *which* in (23a) rules out a CC analysis of the relative clause, the predictions for this experiment are identical to Experiment 2: We expected increased difficulty on the ORC itself in the sentences with *which*, but increased difficulty on the disambiguating material in the sentences with *that*.

Method

Participants

Participants were 60 subjects from the same pool as Experiment 1, who did not participate in either of the previous experiments. One subject was excluded due to poor performance on comprehension questions (68.5% accuracy; no other subject was below 80%), and eleven were excluded based on losing more than 20% of experimental trials to track loss or blink on the critical region, leaving 48 subjects in the analysis.

Materials

The thirty-four items used in Experiment 2 were modified for this experiment as described above. These items were arranged into two experimental lists, so that each subject read 17 items in each of the two experimental conditions. These 34 items were randomly intermixed with 98 other sentences, of which 48 were from an unrelated experiment on subject-verb agreement. The 132 sentences were preceded by eight practice sentences. A total of 50 items were followed by two-choice comprehension questions.

Procedure

The procedure was identical to Experiment 1. As noted above, 11 subjects lost more than 20% of trials on the basis of blinks or track losses on the critical disambiguating region, and these subjects were excluded from subsequent analysis. For the remaining 48 subjects, blinks, track loss or other error resulted in deletion of 6.6% of trials, leaving 1525 trials for inclusion in the analysis.

Results

We analyzed eye movements as in Experiment 2. Results from the largest justified models of each measure (described below) are shown in Table 3.

On the relative clause region, only the simplest model, with condition as a fixed effect, was justified by model comparison. A significant effect of condition appeared in all four measures, with shorter reading times and fewer regressions when the relative pronoun was that. The regressions effect corresponds to 27.2% regressions from the relative clause region in the which condition, and 19.3% in the *that* condition. Though the effect of completion bias and the interaction between completion bias and condition did not reach significance for any measure in the models that included these factors, the interaction did trend in the predicted direction. In go-past time, the parameter estimate for the interaction was $(\beta = -116.18, SE = 112.90, z = -1.03)$, indicating that the advantage for that compared to which was somewhat larger when the clause tended to elicit CC completions. On the other hand, the nonsignificant interaction involving noun bias went in the opposite direction (β = 109.45, SE = 91.63, *z* = 1.20). The opposing signs on these nonsignificant effects are preserved in a model that includes both factors and their interactions with condition, despite the fact that completion bias and noun bias are themselves positively correlated.

The data again show the crossover pattern seen in Experiment 2, with increased difficulty on the disambiguating region in the *that* condition, as shown in Fig. 7. On the disambiguating region, the fit of the simplest model of the first fixation and first pass time data was not improved by adding completion bias or noun bias into the model. For first fixation duration the effect of condition was near significance, while for first pass time it was not. For gopast time, the model that included completion bias and its interaction with condition provided a significant improvement in fit. In this model, the effect of condition was significant, as was the main effect of completion bias; go-past time on the disambiguating

Table 3

Experiment 3 statistical results from mixed-effects models. With the exception of gopast time on the disambiguating region, all estimates are for main effect of experimental condition. As described in text, a model of go-past time that included effects of completion bias and the interaction between completion bias and condition was justified by model comparison.

Measure	Estimate	SE	t/z-value
Relative clause region First fixation duration First pass time Go past time Regression probability	-11.13 -71.19 -233.48 -0.36	5.20 25.29 33.61 0.17	-2.14 -2.82 -6.95 -2.12
Disambiguating region First fixation duration First pass time Go past time (condition) Go past time (completion bias) Go past time (cond × comp bias) Regression probability	9.12 8.35 91.06 195.85 -56.50 0.43	4.58 10.95 31.76 86.08 97.89 0.16	1.99 0.76 2.87 2.28 -0.58 2.79

region was longer when the relative clause was introduced with *that*, and was longer when the fragment elicited mostly CC completions in the norming study. However, the interaction effect was not significant, and was in fact in the opposite direction from the effect in the previous experiments. The model that included noun bias was not justified by model comparison, and showed little hint of either a main effect of noun bias ($\beta = -6.31$, SE = 76.18, *t* = -0.08) or an interaction ($\beta = 21.74$, SE = 80.43, *t* = 0.27). Finally, only the simplest model of the regressions data was justified, and in this model there was a significant main effect of condition. This corresponds to 15.6% regressions when the relative pronoun was *which*, and 22.8% when the relative pronoun was *that*.

The four plots in Fig. 8 illustrate the go-past effects on each of the two regions, as a function of completion bias and noun bias. Each point represents the mean go-past time in the *that* condition minus the mean go-past time in the *which* condition, for a specific item. These differences are consistently negative on the relative clause, and are mostly positive on the disambiguation region, and do not show clear influence of either completion bias or noun bias.

Finally, we again performed an order analysis to test for adaptation to the distribution of syntactic structures in the experiment. As in Experiment 2, subjects in this experiment encountered only RCs. As Fig. 9 indicates, there was no hint of such adaptation, on either the relative clause region or the disambiguating region; the penalty on the relative clause region in the *which* condition was as pronounced at the end of the experiment as at the beginning, and the penalty on the disambiguating region in the that condition was a pronounced at the end as the beginning. The statistical models confirm this. For go-past time on the relative clause, there was of course a large main effect of condition ($\beta = -238.24$, SE = 33.68, t = -7.07), and there was also a main effect of order $(\beta = -4.86, SE = 1.90, t = -2.56)$. However, unlike in Experiment 1, the nonsignificant interactive trend was actually in the opposite direction from the prediction of an adaptation hypothesis, with a somewhat larger effect at the end of the experiment ($\beta = -2.81$, SE = 3.39, t = 0.83). For go-past time on the disambiguating region, there was the expected main effect of condition ($\beta = 86.77$, SE = 31.80, *t* = 2.73), but no main effect of trial order (β = -2.27, SE = 1.56, t = -1.46), and no hint of an interaction ($\beta = -0.71$, SE = 4.46, t = -0.21).

Discussion

Experiment 3 fully replicated the crossover pattern observed in Experiment 2. Processing of the relative clause itself was substantially easier when this clause was introduced by *that* than when it



Fig. 7. Mean go past time with standard errors (a) and regression proportions (b) on each region, by relative pronoun type.



Fig. 8. Experiment 3 go-past time penalty in the *that* condition, by item, on relative clause regions (panels a and b) and disambiguation region (panels c and d) as a function of completion CC bias and noun CC bias. Dashed line is linear regression line. Negative values represent an advantage in the *that* condition.

was introduced by *which*. This difficulty difference appeared in all measures, as opposed to only go-past time and regressions out in Experiment 2. Moreover, the go-past effect was much larger in the present experiment, 233 ms vs. 68 ms. We attribute the size of this effect, and the fact that it also appeared in early eye movement measures, to additional difficulty that may have been introduced by *which*. For some speakers, *which* is used primarily to introduce non-restrictive relative clauses, and corresponds to the

use of a prosodic contour that would normally be indicated with a comma. On the disambiguating region, the condition differences reversed, with difficulty appearing in the *that* condition. This effect is very similar in size across the two experiments: 79 ms in go-past time in Experiment 2, and 91 ms in Experiment 3.

In this experiment there were no significant interactions between either completion bias or noun bias and condition, on either region. The effects of the manipulation of the relative pro-



Fig. 9. Experiment 3 mean go-past time by condition and trial order, for the relative clause region and disambiguation region. Dashed lines are linear regression lines.

noun on both the relative clause region and the disambiguating region were not significantly modulated by either the head noun's tendency to occur with a CC, or the clause's tendency to be completed with a CC. The results of Experiment 1 suggested that there is a garden path effect when an ambiguous *that*-clause is completed as an RC, but that this effect is reduced – or indeed, eliminated – when the clause is one that is strongly biased to elicit an RC completion. This modulation of the garden path effect by completion bias did not reach significance in Experiment 2, and appears to be completely absent in Experiment 3.

There was, however, a large and significant main effect of completion bias on go-past time on the disambiguating region in Experiment 3. One interpretation of this pattern is as follows. Because only the relative pronoun differed between conditions, whatever factors make a particular fragment relatively more or less felicitous as an RC - and therefore induce more or fewer RC completions - would apply equally to both conditions. We suspect that the main effect simply reflects the fact that some of these RC sentences are more felicitous than others, and that the level of felicity is about the same in the two conditions. It is now clear, in any case, that the modulation of the garden path effect by completion bias that was observed in Experiment 1 does not appear in experiments where the disambiguating material is constant across conditions. The frequency with which a fragment elicits a CC completion in an untimed production task does not reliably modulate on-line comprehension patterns.

Finally, Experiment 3 confirmed the lack of syntactic adaptation. Though there was a non-significant trend on the relative clause region in Experiment 2 that was consistent with the presence of an adaptation effect, that trend was completely absent in Experiment 3. Moreover, in neither experiment was the garden path effect on the disambiguating region modulated by trial order. The present null effects may be regarded as a challenge to the notion that comprehenders routinely adapt to the recent syntactic environment, given that only RCs were present in these experiments. In any case, it is clear that the tendency to analyze the present ambiguity as a CC is resistant to adaptation.

General discussion

We presented three eyetracking experiments investigating the on-line processing of the ambiguity that arises in a string such as *The information that the health department provided*, between an analysis in which the *that*-clause is an CC and one on which it is an RC. These experiments present a test of the first clause of the MCP ("Avoid postulating unnecessary chain members at Sstructure"). This principle predicts that comprehenders should adopt the CC analysis as part of a more general strategy of minimizing the maintenance of long-distance filler-gap dependencies. The MCP predicts that an analysis that does not involve positing a such a dependency should always be adopted in preference to one that does.

The central results from these experiments are as follows. First, all three experiments revealed processing difficulty when an ambiguous that-clause is disambiguated toward the RC analysis. In Experiment 1, this result was obtained in a comparison of material that disambiguated toward an RC or toward a CC, which differed in both lexical content and syntactic role. In Experiments 2 and 3, difficulty upon disambiguation toward the RC analysis appeared in a design that enabled comparison of identical material in two conditions. In Experiment 2, the ambiguous condition was compared to one in which the CC analysis was ruled out by the subcategorization restrictions of the head noun (e.g. vaccination as opposed to *information*); in Experiment 3, the ambiguous condition was compared to one in which the CC analysis was ruled out by the use of the pronoun *which*, which unambiguously introduces an RC. It is clear across the three experiments that the penalty on the disambiguating material is, in fact, a garden path effect.

The second clear finding, emerging in both Experiments 2 and 3, is that when the CC analysis of the ambiguous *that*-clause is available, the *that*-clause itself is read faster. This is expected if readers do initially adopt the CC analysis when it is available, and if adopting this analysis eliminates the difficulty associated with processing the *that*-clause as an object relative. We obtained a full crossover pattern in both Experiments 2 and 3, in which fully significant effects in the expected (opposite) directions were obtained on the *that*-clause and on the disambiguating material.

The third important finding from these experiments is a consistent null effect: A noun's bias to occur with a CC never significantly modulated either of the two critical effects, i.e., the garden path effect upon disambiguation toward an RC, or the facilitation on the that-clause itself when a CC analysis was available. Though Experiment 1 showed a numerical trend consistent with a reduction of the garden path effect when the noun was RC-biased, this trend was absent in Experiments 2 and 3. Moreover, the trend on the *that*-clause itself was either absent (Experiment 2) or in the opposite direction from the predicted effect (Experiment 3). Thus, across five regions in three experiments, the effect was never significant, and does not show a consistent trend. While any one of these null effects may be dismissed as resulting from insufficient power, five null effects are not easily explained on these grounds, especially because these experiments did detect substantial main effects of the experimental manipulation on all five of these regions. (Note that even if the power to detect this interaction were as low as 0.4, the probability of five independent null effects, if an interaction does exist, is 0.6⁵, or less than 0.08.) It is also important

to note that the manipulation of noun bias was a strong one, with a nearly uniform distribution of noun biases across the range from 0 to 1. We conclude that a noun's bias to appear with a CC plays little if any role in the parser's initial analysis of an ambiguous *that*-clause; if the noun allows a CC, a CC analysis is preferred, whether the noun occurs with a CC frequently or only rarely. It is worth not-ing, finally, that the limited analysis of the role of noun bias conducted by Chen et al. (2005) came to the same conclusion.

There was also no consistent effect of a specific sentence fragment's off-line tendency to be completed as a CC or an RC. In Experiment 1, we found a significant modulation of the garden path effect by completion bias, with this effect being substantially reduced, or even eliminated, in the most RC-biased sentences. In Experiments 2 and 3, however, this interaction with completion bias never reached significance. In Experiment 2, there were nonsignificant effects of completion bias on both regions, in the expected directions. In Experiment 3, the trend also went in the expected direction on the relative clause itself, but not on the disambiguating region. Taken as a whole, the present results are consistent with the possibility that clauses that are more likely to be interpreted as RCs in an off-line task are very slightly more likely to elicit an RC analysis online, and/or are easier to reanalyze toward this analysis if they initially elicit a CC analysis. But if so, this must be a small effect that does not appear reliably in experiments of the size reported here.

Finally, there was no evidence that readers adapted to the distribution of syntactic structures in the experiments. In Experiments 2 and 3, the critical ambiguity was always resolved toward the RC analysis. Nevertheless, garden path effects were as pronounced at the end of the experiments as they were at the beginning.

The results indicate, then, that when a reader encounters a *that*clause after a noun that can occur with a CC, this clause is initially parsed as a CC, regardless of the frequency of CCs with that particular noun, and even if the clause is highly felicitous as an RC. We note that the data do not rule out the possibility that this parsing preference is applied on merely most trials rather than all, as on a variable-choice parsing model (Traxler, Pickering, & Clifton, 1998; Van Gompel, Pickering, & Traxler, 2001). However, if this is the case, neither the frequency with which the noun occurs with a CC nor the plausibility of the clause as an RC has much influence on the probability that the CC analysis is adopted. The statistics of the recent input also do not seem to play a role.

Parsing the that-clause as a CC enables the reader to avoid the processing difficulty that would result from parsing this clause as an ORC. One possible source of the difficulty in ORCs is the violation of expectations when the relative pronoun that is followed by a noun phrase as opposed to a verb, arising because SRCs are more frequent than ORCs (e.g., Roland, Dick, & Elman, 2007), or because ORCs have specific discourse licensing conditions (Roland, Mauner, O'Meara, & Yun, 2012). Notably, the ORC penalty was present in our experiments despite the fact that only inanimate head nouns were used, which has been shown to decrease the ORC penalty (e.g., Traxler, Williams, Blozis, & Morris, 2005). Another likely source of difficulty is the memory retrieval of the relative clause head that is required upon reaching the ORC gap site (e.g., Gordon, Hendrick, & Johnson, 2001). Still another explanation for the ORC penalty is the syntactic intervention effect created by the subject noun phrase, which intervenes between the RC head and its gap (Friedmann, Belletti, & Rizzi, 2009). These sources of difficulty would not play a role when the clause is parsed as a CC: If that is treated as a complementizer, then a subject noun phrase is expected as the next constituent, and there is no gap site within the clause at which a long-distance retrieval is required. But upon disambiguation at the matrix verb, difficulty arises, as the CC analysis must be abandoned and an ORC analysis constructed.

The preceding account of processing of the critical ambiguity assumes that parsing is serial, with only one analysis adopted at the point of ambiguity, and with this analysis being abandoned only when inconsistent input is encountered (e.g., Frazier, 1978). Could the present data be accounted for by a parallel model (e.g., Levy, 2008; McRae, Spivey-Knowlton, & Tanenhaus, 1998), on which CC and ORC parses are both activated to some degree, and are maintained simultaneously while processing the ambiguous material? The pattern on the relative clause region in Experiments 2 and 3, which was read faster when it was fully ambiguous between two analyses than when it was unambiguous, seems to present an argument against parallelism. Many parallel accounts predict that ambiguity should result in processing difficulty (Clifton & Staub, 2008). However, this ambiguity advantage could be accounted for by a specific version of parallelism, implemented in Levy's (2008) Surprisal model. Levy accounts for such an ambiguity advantage in other experiments (e.g., Van Gompel et al., 2001) by proposing that processing difficulty is a function of the input's combined conditional probability on all syntactic parses (though see Clifton and Staub (2008), for a counterargument). If ambiguous input has non-zero probability on two active parses, rather than on just one, the probability of the input will be greater overall in the ambiguous condition.

We see two problems for this account of the present data. First, it does not predict that the difficulty in the unambiguous ORC condition should take the specific form that it does, both here and in other experiments, i.e., increased regressive eye movements. Like Staub (2010), we regard these regressions as arising from difficulty associated with a specific parse of the input. Second, parallel models propose that the activations of competing parses are weighted according to their syntactic probability (Levy, 2008) or factors such as plausibility and subcategorization bias (McRae et al., 1998). This view predicts that disambiguation difficulty, which is conceptualized in terms of a change in the activation level of the multiple parses, should depend on noun bias, completion bias, or both; disambiguation toward the RC analysis should be easy when the noun's bias and/or the plausibility of the clause as an RC favor this analysis, and difficult when they do not. These patterns are not present in the data. We conclude, then, that the data pattern as a whole is most consistent with serial parsing guided by a preference for the CC analysis.

On this view, the experiments provide evidence of more-or-less categorical avoidance of an analysis where *that* introduces an RC, when an alternate analysis is available on which *that* introduces a CC. Thus, the parser is then able to avoid positing a longdistance dependency, and to avoid a search for a downstream gap site. This conclusion is complementary to that of Staub et al. (2006), who found that the parser categorically avoids positing the gap associated with a Heavy NP Shift analysis when another analysis is available. It is consistent with the first clause of the MCP (De Vincenzi, 1991), which states that filler-gap dependencies should not be posited except when necessary. Even more broadly, it is consistent to lexical biases; when a gap is necessary, it will be posited in an eager, predictive manner, but when a filler-gap analysis can be avoided, the parser will do so.

However, it is also important to consider the possibility that some other structural principle favors the CC analysis over the RC analysis. As noted above, Crain and Steedman (1985; Altmann et al., 1992) suggested that Frazier's (1978) principle of Minimal Attachment favors a CC analysis of the ambiguity that they explored. The principle of Minimal Attachment states that the parser will attach incoming material into the phrase marker using the fewest possible syntactic nodes, i.e., the parser will construct the structurally simplest analysis, where simplicity is defined in terms of number of nodes. Given contemporary representational assumptions, the CC analysis of the ambiguous *that*-clause requires attaching this clause as sister to the preceding noun, while the RC analysis requires attaching it as an adjunct, i.e., as sister to a higher projection. So it is indeed reasonable to conclude that the CC analysis is favored by Minimal Attachment. (We note that argument analyses may also generally be favored over adjunct analyses; Schütze & Gibson, 1999; Liversedge, Pickering, Branigan, & van Gompel, 1998; cf. Clifton, Speer, & Abney, 1991).

We cannot entirely rule out the possibility that structural simplicity, in the sense of Minimal Attachment, may play some role in the preference for the CC analysis, and we also cannot rule out the possibility that Minimal Attachment and avoidance of longdistance dependencies effectively conspire to make the RC analysis especially unlikely to be adopted. However, the categoricity of this preference suggests that the critical factor may be avoidance of long distance dependencies, rather than Minimal Attachment. At least some studies (e.g., Garnsey et al., 1997; Trueswell et al., 1993), have suggested that lexical biases can overcome Minimal Attachment-based parsing preferences, while the lack of a role for lexical biases, as observed here, is consistent with studies of filler-gap processing (e.g., Pickering & Traxler, 2003; Staub, 2007; Staub et al., 2006). Other studies have shown that Minimal Attachment-based preferences may be overcome when the nonminimal parse leads to a more plausible interpretation (e.g., Trueswell, Tanenhaus, & Garnsey, 1994). Again, filler-gap processing is different, as active gap-filling appears to be indifferent to the plausibility of the resulting analysis (e.g., Traxler & Pickering, 1996). Our results pattern like the latter, as the plausibility of an ambiguous string as an RC or as a CC, as indicated by the off-line completion preferences, did not reliably modulate the effects in these experiments.

The results might also be relevant for a recent debate in syntactic theory. Some authors (Arsenijević, 2009; Kayne, 2010) have claimed that, despite the superficial differences between them (i.e. the presence/absence of a gap), at a more abstract level CCs should be analyzed as a special type of RCs. Our results speak against this hypothesis, unless it is further specified how to account for the clear differences in processing between them (see De Cuba (2017) for a criticism of this hypothesis based on crosslinguistic data).

An interesting remaining question is whether prior context that introduces multiple possible referents for the head noun could eliminate, or even reverse, the strong CC preference. For example (7), repeated here, could be preceded by (24), as follows:

(24)	Two pieces of information were given to the public.
(7)	The information that the health department provided

Unlike the contexts used by Altmann et al. (1992), this context does not render an RC analysis of the *that*-clause in (7) completely obligatory. This is because the uniqueness presupposition of the definite description can also be satisfied by a CC, e.g.:

(25) The information that the health department provided a cure was contradicted by the information that the disease was spreading across the southern part of the country.

We leave it as an open question whether such a context would indeed reduce the CC preference, or even induce an RC preference.

We believe that the results reported here confirm and extend our understanding of the principles that govern the processing of filler-gap structures. The results suggest that active gap filling is actually one manifestation of a more general dependency minimization principle, as originally suggested by the MCP. They also add to previous evidence that the processing of long-distance dependencies is one domain in which structural heuristics override the influence of lexical biases.

Appendix A

Materials used in Experiment 1, comprising 68 sentences, two with each of 34 head nouns. The embedded clause object in the CC condition is in parentheses. The 34 items used in Experiment 2 are denoted with an asterisk, and the corresponding non-CCtaking head noun used in that experiment is provided in parentheses after each of these items. In Experiment 3 each sentence used in Experiment 2 was matched with a corresponding sentence in which *that* was substituted by *which*.

- 1. The announcement that Massachusetts expected (a new governor) was released this morning.
- *2. The announcement that the president is considering (an intervention) will alarm the citizens. (intervention)
- 3. The answer that the client provided (the material) was not accepted by anyone.
- *4. The answer that the colony transmitted (the virus) was regarded with suspicion by the medical staff. (virus)
- 5. The assumption that the researcher tested (volunteers) is not actually stated in the paper.
- ^{*}6. The assumption that the witness entertained (the defendant) was criticized by the judge. (suspect)
- ^{*}7. The bet that the player will cancel (the match) was for a small amount of money. (donation)
- 8. The bet that the siblings would announce (their father's identity) was shocking to some family members.
- *9. The claim that the copywriter invented (the brand) increased his standing in the agency. (product)
- 10. The claim that the committee will favor (a candidate) was disputed among the participants.
- 11. The concern that the other countries are fostering (terrorism) has caused the negotiations to break down.
- *12. The concern that the pediatrician suggested (a specialist) worried the parents about a possible allergy. (specialist)
- ^{*}13. The conclusion that the author condemns (religion) appeared in the review. (chapter)
- 14. The conclusion that the editor published (junk) was widespread among the readers.
- 15. The desire that women would show (their elegance) inspired the whole collection.
- ^{*}16. The desire that their child wrote (poetry) was shared by both parents. (language)
- 17. The doubt that he had raised (the child) was now widespread among the jurors.
- ^{*}18. The doubt that the priest expressed (disapproval) was spreading in the parish. (love)
- 19. The dream that his wife personifies (his mother) made things hard for their relationship.
- *20. The dream that the painter was chasing (his shadow) inspired the best piece of the collection. (bird)
- 21. The evidence that the witness discovered (the gun) was not admitted by the judge.
- ^{*}22. The evidence that the husband found (the weapon) gave new life to the investigation. (device)

- ^{*}23. The fear that Italy would declare (bankruptcy) came as no surprise in Europe. (bankruptcy)
- 24. The fear that his words revealed (the truth) kept her awake all night.
- 25. The guess that Sally made (the brownies) was disputed by everyone at the party.
- ^{*}26. The guess that the gambler would make (a million dollars) excited all his friends. (salary)
- ^{*}27. The hope that the local administration transmitted (funding) boosted the community. (funding)
- 28. The hope that this marriage represents (stability) is worth all the trouble and sacrifice.
- ^{*}29. The idea that the movie represents (reality) is quite shocking to most audience members. (world)
- The idea that the laboratory tested (animals) was never known to the ethics committee.
- ^{*}31. The implication that the project revealed (a new galaxy) was disputed by most experts. (planet)
- 32. The implication that the student condemns (equal rights) was never intended by the professor.
- 33. The information that Bob released (the secret) was surprising to me.
- ^{*}34. The information that the health department provided (a cure) reassured the tour operators. (vaccination)
- 35. The insinuation that Howard suggested (mischief) was denied by all the other children.
- ^{*}36. The insinuation that the director made (a blunder) was understood by the members of the board. (film)
- 37. The law that all objects follow (gravity) has been known for many centuries.
- ^{*}38. The law that the European Community will cancel (the public debt) was not appreciated in the US. (ceasefire)
- 39. The notice that the woman released (a book) appeared in the Sunday paper.
- *40. The notice that the committee expected (a new chairman) influenced the choices made by all the members. (backing)
- 41. The notion that the politician entertained (the diplomats) was widespread on TV.
- *42. The notion that the students proposed (less homework) was taken into account in preparing his course. (syllabus)
- ^{*}43. The possibility that the manager is considering (a new job) has to be monitored. (candidate)
- 44. The possibility that the job offered (a good salary) was the main reason for him to apply.
- 45. The prediction that the coach would make (a mistake) was very controversial.
- ^{*}46. The prediction that the software will compute (the result) has been verified by the technician. (sum)
- 47. The probability that the system will compute (the answer) is not known to the physicists.
- ^{*}48. The probability that the scientist uncovered (the allegations) was not revealed in court. (disease)
- 49. The proof that the linguist found (the dialect) was presented at the conference.
- ^{*}50. The proof that the witness falsified (his deposition) was the best hope of the defense. (passport)
- ^{*}51. The remark that the candidate offered (money) had devastating consequences in the campaign. (funding)
- 52. The remark that the teacher uncovered (the vandalism) was overlooked by the student's parents.
- 53. The report that the journalist falsified (the story) turned out to be old news.

- *54. The report that the FBI invented (the affair) was believed by the Russian agents. (spy)
- 55. The resolution that the lawyer will favor the plea surprised the court.
- ^{*}56. The resolution that the unions are fostering (dissent) was regarded with suspicion by the media. (hostility)
- ^{*}57. The rule that the meter maid wrote (the tickets) was discussed by the council. (ticket)
- 58. The rule that the policeman personifies (authority) is not always followed.
- *59. The statement that the journal published (fake data) was reported worldwide. (figure)
- 60. The statement that the miners would declare (a strike) caused unwarranted reactions.
- ^{*}61. The suggestion that I proposed (the new band) was slightly offensive. (name)
- 62. The suggestion that the members follow (the leader) was put forth by the majority of the group.
- 63. The suspicion that the patient had (lyme disease) was verified by the doctor.
- 64. The suspicion that the company would announce (a takeover) was shaking Wall Street. (firing)
- ^{*}65. The truth that Jane was chasing (a fugitive) would shake up the political establishment. (senator)
- 66. The truth that her daughter discovered (his origins) was slammed in his face with no pity.
- 67. The wish that my mother expressed (her emotions) would not be fulfilled anytime soon.
- 68. The wish that his diaries would show (his love) was thrilling to her. (passion)

References

- Altmann, G. T. (1994). Regression-contingent analyses of eye movements during sentence processing: Reply to Rayner and Sereno. *Memory & Cognition*, 22, 286–290.
- Altmann, G. T., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language*, 31, 685–712.
- Altmann, G., Garnham, A., & Henstra, J. A. (1994). Effects of syntax in human sentence parsing: Evidence against a structure-based proposal mechanism. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 209–216.
- Arnold, J. E., Losongco, A., Wasow, T., & Ginstrom, R. (2000). Heaviness vs. newness: The effects of structural complexity and discourse status on constituent ordering. *Language*, 76, 28–55.
- Arsenijević, B. (2009). Clausal complementation as relativization. *Lingua*, 119, 39–50.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1–48.
- Chen, E., Gibson, E., & Wolf, F. (2005). Online syntactic storage costs in sentence comprehension. Journal of Memory and Language, 52, 144–169.
- Chomsky, N. (1981). Lectures on government and binding. Dordrecht: Foris.
- Clifton Jr., C., & Ferreira, F. (1989). Ambiguity in context. Language and Cognitive Processes, 4, SI77-SI103.
- Clifton, C., Jr., & Frazier, L. (1989). Comprehending sentences with long distance dependencies. In G. Carlson & M. Tanenhaus (Eds.), *Linguistic structure in language processing* (pp. 273–317). Dordrecht, The Netherlands: Kluwer Academic.
- Clifton, C., Speer, S., & Abney, S. P. (1991). Parsing arguments: Phrase structure and argument structure as determinants of initial parsing decisions. *Journal of Memory and Language*, 30, 251–271.
 Clifton, C., & Staub, A. (2008). Parallelism and competition in syntactic ambiguity
- Clifton, C., & Staub, A. (2008). Parallelism and competition in syntactic ambiguity resolution. Language and Linguistics Compass, 2, 234–250.
- Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. *Natural Language Parsing*, 320–358.
- Davies, M. (2008). The corpus of contemporary American English (COCA): 520 million words, 1990-present http://corpus.byu.edu/coca/.
- De Cuba, C. (2017). Noun complement clauses as referential modifiers. *Glossa, 2*, 3.
- De Vincenzi, M. (1991). Syntactic parsing strategies in Italian. Dordrecht, The Netherlands: Kluwer Academic.

- Ferreira, F., & Henderson, J. M. (1990). Use of verb information in syntactic parsing: Evidence from eye movements and word-by-word self-paced reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 555–568.
- Fine, A. B., Jaeger, T. F., Farmer, T. A., & Qian, T. (2013). Rapid expectation adaptation during syntactic comprehension. *PLoS ONE*, 8(10), e77661.
- Fodor, J. D. (1978). Parsing strategies and constraints on transformations. *Linguistic Inquiry*, 9, 427–473.
- Frazier, L. (1978). On comprehending sentences: Syntactic parsing strategies, Unpublished doctoral dissertation. University of Connecticut, Storrs.
- Frazier, L., & Clifton, C. Jr., (1989). Successive cyclicity in the grammar and the parser. Language and Cognitive Processes, 28, 331–344.
- Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), Attention and performance XII: The psychology of reading (pp. 559–586). Hillsdale, NJ: Earlbaum.
- Friedmann, N., Belletti, A., & Rizzi, L. (2009). Relativized relatives: Types of intervention in the acquisition of A-bar dependencies. *Lingua*, 119, 67–88.
- Garnsey, S. M., Pearlmutter, N. J., Myers, E., & Lotocky, M. A. (1997). The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. *Journal of Memory and Language*, 37, 58–93.
- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68, 1–76.
- Gordon, P. C., Hendrick, R., & Johnson, M. (2001). Memory interference during language processing. Journal of Experimental Psychology: Learning, Memory and Cognition, 27, 1411–1423.
- Gordon, P. C., Hendrick, R., Johnson, M., & Lee, Y. (2006). Similarity-based interference during language comprehension: Evidence from eye tracking during reading. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 32, 1304–1321.
- Grodner, D. J., & Gibson, E. A. F. (2005). Consequences of the serial nature of linguistic input for sentential complexity. *Cognitive Science*, 29, 261–291.
- Hale, J. (2001). A probabilistic early parser as a psycholinguistic model. In Proceedings of NAACL (Vol. 2, pp. 159–166).
- Kayne, R. (2010). Why isn't this a complementizer? In R. Kayne (Ed.), Comparisons and contrasts (pp. 190–227). Oxford: Oxford University Press.
- Kennison, S. M. (2000). Two assessments of American usage frequencies for ninetyseven sentence complement-taking nouns. *Journal of Psycholinguistic Research*, 29, 249–264.
- Kennison, S. M. (2001). Limitations on the use of verb information during sentence comprehension. Psychonomic Bulletin & Review, 8, 132–137.
- Levy, R. (2008). Expectation-based syntactic comprehension. Cognition, 106, 1126–1177.
- Levy, R., Fedorenko, E., Breen, M., & Gibson, E. (2012). The processing of extraposed structures in English. Cognition, 122, 12–36.
- Liversedge, S. P., Pickering, M. J., Branigan, H. P., & van Gompel, R. P. (1998). Processing arguments and adjuncts in isolation and context: The case of byphrase ambiguities in passives. *Journal of Experimental Psychology: Learning*, *Memory, and Cognition*, 24, 461–475.
- McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, 38, 283–312.
- Mitchell, D. C., Corley, M. M., & Garnham, A. (1992). Effects of context in human sentence parsing: Evidence against a discourse-based proposal mechanism. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 69–88.
- Moulton, K. (2015). CPs: Copies and compositionality. Linguistic Inquiry, 46, 305–342.
- Pearlmutter, N. J., & Mendelsohn, A. A. (1999). Serial versus parallel sentence comprehension. Unpublished manuscript.
- Phillips, C. (2006). The real-time status of island phenomena. Language, 82, 795-823.

- Pickering, M. J., & Traxler, M. J. (2003). Evidence against the use of subcategorisation frequencies in the processing of unbounded dependencies. *Language and Cognitive Processes*, *18*, 469–503.
- Pickering, M. J., Traxler, M. J., & Crocker, M. W. (2000). Ambiguity resolution in sentence processing: Evidence against frequency-based accounts. *Journal of Memory and Language*, 43, 447–475.
- Portner, P. (1989). Processing indefinite noun phrases in quantified sentences. Unpublished manuscript.
- R Core Team (2014). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing http://www.R-project.org/>.
- Rayner, K., & Sereno, S. C. (1994). Regressive eye movements and sentence parsing: On the use of regression-contingent analyses. *Memory & Cognition*, 22, 281–285.
 Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of English grammatical
- structures: A corpus analysis. *Journal of Memory and Language*, 57, 348–379. Roland, D., Mauner, G., O'Meara, C., & Yun, H. (2012). Discourse expectations and
- relative clause processing. Journal of Memory and Language, 66, 479–508. Ross, J. R. (1967). Constraints on variables in syntax. Bloomington, IN: Indiana
- University Linguistics Club.
- Schütze, C. T., & Gibson, E. (1999). Argumenthood and English prepositional phrase attachment. *Journal of Memory and Language*, 40, 409–431.
- Staub, A. (2007). The parser doesn't ignore intransitivity, after all. Journal of Experimental Psychology: Learning, Memory, and Cognition, 33, 550–569.
- Staub, A. (2010). Eye movements and processing difficulty in object relative clauses. Cognition, 116, 71–86.
- Staub, A., Clifton, C., Jr., & Frazier, L. (2006). Heavy NP shift is the parser's last resort: Evidence from eye movements. *Journal of Memory and Language*, 54, 389–406.
- Staub, A., Dillon, B., & Clifton, C. (2017). The matrix verb as a source of comprehension difficulty in object relative sentences. *Cognitive Science*, 41, 1353–1376.
- Stowe, L. A. (1986). Parsing WH-constructions: Evidence for on-line gap location. Language and Cognitive Processes, 1, 227–245.
- Taraban, R., & McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language*, 27, 597–632.
- Traxler, M., Morris, R. K., & Seely, R. E. (2002). Processing subject and object relative clauses: Evidence from eye movements. *Journal of Memory and Language*, 47, 69–90.
- 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.
- Traxler, M. J., Pickering, M. J., & Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, 39, 558–592.
- Traxler, M. J., Williams, R. S., Blozis, S. A., & Morris, R. K. (2005). Working memory, animacy, and verb class in the processing of relative clauses. *Journal of Memory* and Language, 53, 204–224.
- Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language*, 33, 285–318.
- Trueswell, J., Tanenhaus, M. K., & Kello, C. (1993). Verb specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 528–533.
- Van Gompel, R. P., Pickering, M. J., & Traxler, M. J. (2001). Reanalysis in sentence processing: Evidence against current constraint-based and two-stage models. *Journal of Memory and Language*, 45, 225–258.
- Vernice, M., Cecchetto, C., Donati, C., & Moscati, V. (2016). Relative clauses are not adjuncts: An experimental investigation of a corollary of the raising analyses. *Linguistische Berichte*, 246, 139–169.
- Wasow, T. (1997). End-weight from the speaker's perspective. Journal of Psycholinguistic Research, 26, 347–361.