

This article was downloaded by:[University of Leipzig]
[University of Leipzig]

On: 11 July 2007

Access Details: [subscription number 738315272]

Publisher: Psychology Press

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



The Quarterly Journal of Experimental Psychology

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t716100704>

Word chunk frequencies affect the processing of pronominal object-relative clauses

Online Publication Date: 01 February 2007

To cite this Article: Reali, Florencia and Christiansen, Morten H. , (2007) 'Word chunk frequencies affect the processing of pronominal object-relative clauses', The Quarterly Journal of Experimental Psychology, 60:2, 161 - 170

To link to this article: DOI: 10.1080/17470210600971469

URL: <http://dx.doi.org/10.1080/17470210600971469>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article maybe used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

© Taylor and Francis 2007

Short article

Word chunk frequencies affect the processing of pronominal object-relative clauses

Florencia Reali and Morten H. Christiansen

Cornell University, Ithaca, NY, USA

We present experimental support for the view that fine-grained statistical information may play a crucial role in the processing of centre-embedded linguistic structure. Using both offline and online methods, we show that the processing of pronominal object-relative clauses is influenced by the frequency of co-occurrence of the word combinations (*chunks*) forming the clause. We use materials that are controlled for capacity-based factors that have been previously shown to influence comprehension of relative clauses. The results suggest that, other factors being equal, the frequency of the word chunk forming the clause affects processing difficulty. Analyses of the data indicate that the results cannot be explained by differential access to individual lexical items. Following recent constructivist approaches, we argue that frequency of co-occurrence influences the *chunking* mechanism by which multiword sequences may become fused into processing units that are easier to access.

A key question in language research pertains to the role that distributional information may play in acquisition and processing of syntactic structure. The importance of statistical information during incremental language comprehension has been primarily studied in the context of ambiguity resolution (e.g., Crocker & Corley, 2002; Desmet, De Baecke, Drieghe, Brysbaert, & Vonk, 2006; Jurafsky, 1996; MacDonald, Pearlmutter, & Seidenberg, 1994; Mitchell, Cuetos, Corley, & Brysbaert, 1995). However, much less is known about its potential role in the processing of unambiguous utterances.

Some recent studies have explored the influence of fine-grained statistics during online processing

of unambiguous sentences. For example, McDonald and Shillcock (2004) provided evidence suggesting that reading times of individual words are affected by the transitional probabilities of the lexical components. Using materials like *One way to avoid confusion/discovery is to make the changes during vacation*, they showed that transitional probabilities (high in *avoid confusion* and low in *avoid discovery*) have a measurable effect on fixation durations. They argued that the results could be explained by a Bayesian statistical model in which lexical probabilities are derived by combining transitional probabilities with the prior probability of a word's occurrence (but see Frisson, Rayner, & Pickering, 2005).

Correspondence should be addressed to Florencia Reali, Department of Psychology, Cornell University, Ithaca, NY 14853, USA.
E-mail: fr34@cornell.edu

We are grateful to Michael Spivey and Thomas Farmer for helpful discussions on this work. We also wish to thank Tessa Warren and an anonymous reviewer for providing insightful comments and suggestions regarding an earlier version of this manuscript.

Recently, Bybee (2002; Bybee & Scheibman, 1999) suggested that the representation of linguistic constituents might be affected by repetition of multiword sequences. In the spirit of constructivist approaches (e.g., Goldberg, 2006; Tomasello, 2003), they propose that when words repeatedly co-occur together in a specific order, such multiword sequences may fuse together into a single processing unit. As a consequence of this “*chunking*” process, repeated exposure to sequential stretches of words within a linguistic constituent would create a supralexicale representation of this construction, making it easier to access. That is, frequent word sequences (*chunks*) would fuse into amalgamated processing units that can be accessed and produced more easily. Additionally, this process may manifest itself as a continuum: Differences in the frequency of specific word sequences are likely to lead to different degrees of amalgamation (*chunking*), resulting in a graded process conditioned by word co-occurrence patterns.

Bybee and Scheibman (1999) used evidence taken from conversations to demonstrate that repetition of multiword sequences influences the degree of phonological reduction of *don't* in American English. They showed that such reduction is more pronounced in the contexts in which *don't* occurs the most—for example, after the pronoun *I*. This effect could be explained by the chunking hypothesis favoured by the authors or by predictability effects: Accessing the next word may be easier when it is predictable, reducing production time. Additionally, Bybee and Scheibman (1999) found that vowel reduction in *don't* occurs primarily before verbs that frequently follow this expression, such as *know*, *think*, or *want*. This suggests that phonological reduction in *don't* cannot be explained as a result of simple exposure to transitional probabilities (e.g., from *I* to *don't*) because vowel reduction is also conditioned by the frequency with which the following verb occurs as part of the same construction, suggesting that the word chunks had fused together, leading to a more compact representation of constituent structure. Bybee and Scheibman (1999) argue in favour of a model according to which the frequency of phrases such as *I don't know*, *I don't think* has

“rendered them fused storage and processing units and has conditioned the loss of stress on the middle element and its consequent reduction” (p. 582).

Along similar lines of reasoning, here we present experimental data suggesting that sentences with pronominal object-relative clauses, such as *The person who I met distrusted the lawyer*, are easier to process when the embedded clause is formed by frequent pronoun–verb combinations (*I liked* or *I met*) than when it is formed by less frequent combinations (*I distrusted* or *I phoned*). We hypothesize that frequent word sequences forming object-relative clauses may fuse into more strongly amalgamated representations that are easier to process than less frequent sequences. We adhere to the view that the processing of sentence constituents (of which relative clauses are a particular case) might be affected by exposure to frequent multiword sequences (e.g., Bybee, 2002). The case of object-relatives is of special interest because of the well-established finding that *nested* (or centre-embedded) structure is more difficult to process than nonnested structure. Theories emphasizing the role of memory constraints have been proposed to account for this phenomenon, and much experimental work has been conducted to elucidate the source of this difficulty (for discussion, see Gibson, 1998).

Recent studies have shed some light on the kind of factors that may influence the production and comprehension of pronominal object-relative clauses (Race & MacDonald, 2003; Warren & Gibson, 2002). Using both complexity rating and self-paced reading tasks, Warren and Gibson (2002) examined the extent to which referential properties of the most deeply embedded subject affect comprehension of centre-embedded sentences. They found that processing difficulties depended on the degree to which the subject in the embedded clause was old or new in the discourse (e.g., pronoun *I* vs. *the scientist*). For example, they showed that the sentence *The student who the professor who I collaborated with had advised copied the article* was easier to comprehend than the sentence *The student who the professor who the scientist collaborated with had*

advised copied the article. Warren and Gibson (2002) explain these results from the perspective of the dependence locality theory (DLT; Gibson, 1998). According to DLT, the cost of syntactic integrations associated with embedded structure increases with the number of new discourse referents that are introduced between the phrasal heads that must be integrated. Recent versions of this view (Grodner & Gibson, 2005) proposed that integration cost is increased by a variety of additional factors including length of the clause (e.g., *I vs. the scientist*).

Race and MacDonald (2003) explored the use of the relativizer *that* in the production and comprehension of object-relative clauses. They found that producers less frequently insert *that* in object-relatives when the embedded subject is a pronoun. Other factors such as length-of-the-clause increased the inclusion of *that* during production, suggesting that the word *that* may be inserted to alleviate production difficulties. An additional experiment showed that comprehenders are sensitive to the observed production biases. The authors argued in favour of constraint-based interactions in production and comprehension systems: Prior comprehension experiences affect choices during production, leading to certain distributional patterns. In turn, comprehenders show sensitivity to the generated distributional patterns, finding frequent structures easier to process. This view provides an alternative explanation for the results reported in Warren and Gibson (2002): Facilitation of pronominal object-relatives could be explained, at least in part, by the frequency of the embedded subject (*I or you vs. the scientist*). Providing further support for this view, Reali and Christiansen (in press) conducted corpus analyses indicating that pronominal object-relative clauses, such as *that I liked*, occur naturally in the language with high frequency, and, in particular, these constructions are significantly more frequent than pronominal subject-relative clauses such as *that liked you*. Self-paced reading experiments indicated that the differences in processing difficulty between pronominal object-/subject-relative clauses mirrored the pattern of distribution revealed by the corpus analysis.

In sum, a growing bulk of research suggests that distributional information may influence the processing of relative clauses (see also MacDonald & Christiansen, 2002). However, a further question concerns the extent to which the frequency of token co-occurrences, such as specific *pronoun-verb* combinations in the relative clause, facilitates processing. Following the view outlined in Bybee and Scheibman (1999), here we explore two hypotheses: First, the processing of pronominal object-relative clauses may be facilitated by frequent co-occurrence of the elements forming the clause. Second, this process may manifest itself as a continuum, leading to a gradual facilitation of processing as a function of specific co-occurrence patterns.

In Experiment 1, we conducted offline rating tasks to compare complexity and plausibility ratings across doubly embedded object-relative sentences. We manipulated the frequency of word co-occurrence in the most deeply embedded clause. The pronoun *I* was used as the most deeply embedded subject in all experimental sentences, therefore providing a control for differences in referential and memory factors that had been shown to influence comprehension (Warren & Gibson, 2002). Experiment 2 was a self-paced reading task conducted on singly embedded versions of the sentences in Experiment 1. Our prediction was that the frequency of the *I-verb* combinations forming the embedded clause would facilitate its processing. In support of this view, all experiments showed a robust difference between high- and low-frequency conditions. Moreover, fine-grained analysis of the data revealed that the chunk frequency effect manifests itself as a continuum, suggesting that elements that are frequently used together may fuse into processing units as a gradual function of their specific co-occurrence patterns.

EXPERIMENT 1

Experiment 1 comprised questionnaire tasks comparing the comprehension difficulty in doubly embedded object-relative sentences in which the

pronoun *I* was the most deeply embedded subject. We manipulated the frequency of specific *I-verb* combinations forming the most deeply embedded relative clause.

Warren and Gibson (2002) used similar questionnaire experiments to show that complexity of doubly embedded sentences depends on the referential properties of the embedded noun phrase. In the present study, the type of embedded subject was not manipulated, therefore controlling for referential factors.

Method

Participants

A total of 60 native English speakers from Cornell undergraduate classes were recruited, half of which completed a questionnaire corresponding to the complexity-rating task, and the other half completed a questionnaire corresponding to the plausibility-rating task.

Materials

A total of 12 doubly nested experimental items were tested with two conditions per item. All items were object-relative sentences in which the pronoun *I* was the most deeply embedded subject. The two conditions varied in the co-occurrence patterns of the elements forming the most deeply embedded clause. We used *Google* counts (Keller & Lapata, 2003) to quantify the bigram frequency of the specific *I-verb* combinations forming the most deeply embedded clause. The materials were constructed such that the word combinations forming the embedded clause were significantly more frequent in the high-frequency condition than in the low-frequency condition ($p < .0001$). The sentences provided in (1) are examples of the stimuli (a complete list of items is included in the Appendix):

- a. The detective who the attorney who *I met* distrusted sent a letter on Monday night. (high-frequency)
- b. The detective who the attorney who *I distrusted* met sent a letter on Monday night. (low-frequency)

Crucially, across conditions sentences contained exactly the same words arranged differently. Thus, differences in complexity ratings cannot be attributed to properties of the lexical items, such as frequency of individual words.

Two types of questionnaire were created, one for the complexity-rating task and a second for a control plausibility-rating task. Following a similar paradigm to the one used in Warren and Gibson (2002), the plausibility-rating questionnaire contained a right-branching version of the experimental sentences (e.g., the right-branching version of (1a) is: *I met the attorney who distrusted the detective who sent a letter on Monday night*). Each type of questionnaire contained 52 fillers in addition to the experimental items. The two conditions were counterbalanced across lists, so each subject saw one version of each item. The lists were pseudorandomized with no two experimental items occurring back to back, and the order of the questionnaire pages was varied.

Procedure

In the complexity-rating task, participants were asked to rate the complexity of sentences on a scale from 1 to 7, 1 indicating "hard to understand" and 7 "easy to understand". The questionnaire began with a page of instructions asking participants to make their judgements based on first impressions without reading each sentence more than once. In the instructions, participants were given four practice items that varied in complexity. None of them had the same nested structure as the experimental items. Similarly, in the plausibility-rating task, participants were asked to rate the plausibility of sentences on a scale from 1 to 7, 1 denoting "not plausible" and 7 "very plausible". Additionally, the term "*plausible*" was defined as "how likely the situation described by the sentence is".

Results and discussion

The mean complexity and plausibility ratings for each condition are presented in Figure 1. Planned comparisons across conditions indicated that when high-frequency chunks constituted the

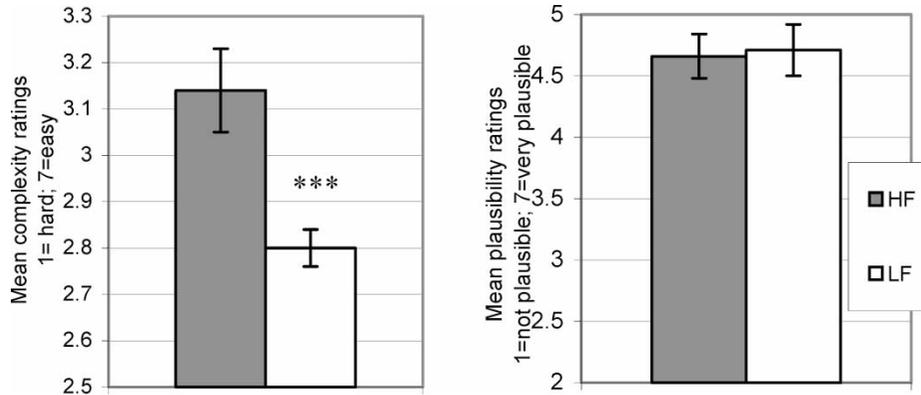


Figure 1. Results from Experiment 1: Mean complexity ratings (left) and plausibility ratings (right) for high-frequency condition (dark bars) and low-frequency condition (light bars).

most embedded clause, sentences were rated less complex ($M = 3.14$, $SD = 0.37$ in the high-frequency condition; $M = 2.80$, $SD = 0.16$ in the low-frequency condition), $t_1(29) = 11.39$, $p = .003$; $t_2(11) = 11.2$, $p = .008$. However, there was no difference between conditions in the control plausibility-rating task ($M = 4.66$, $SD = 0.65$ in the high-frequency condition; $M = 4.72$, $SD = 0.73$ in the low-frequency condition), $t_1(29) = 0.3$, $p > .5$; $t_2(11) = 0.05$, $p > .8$.

The results suggest that the frequency of the most deeply embedded clause influences complexity rating. The results cannot be due to simple lexical frequencies because in both conditions all items had the same words arranged differently. It should be noted that the frequency of the embedded clause correlates with the frequency of the verb in the most deeply embedded position. Thus, an alternative interpretation of the present findings is that sentences are easier to understand if a frequent verb occurs in the most deeply embedded position. However, the effect is observed only when the high-frequency verb appears in the internal clause and not in the external one, suggesting that statistical information must influence sentence comprehension at a deeper level than simple lexical access.

Capacity-based theories in their current form do not explain the difference in complexity ratings observed in the present study. This is because syntactic structure and embedded subjects

were identical in all items, and, therefore, capacity-related factors did not differ across conditions.

EXPERIMENT 2

In Experiment 2, we conducted a self-paced reading task to investigate the online processing of singly embedded versions of the sentences rated in Experiment 1.

Method

Participants

A total of 35 members of the Cornell community participated in this study in exchange for a \$5 payment.

Materials

The stimuli consisted of singly embedded versions of the items used in Experiment 1. The sentences provided in (2) are examples of the stimuli used in each condition (high-frequency and low-frequency, respectively):

- The attorney who *I met* distrusted the detective who sent a letter on Monday night. (high-frequency)
- The attorney who *I distrusted* met the detective who sent a letter on Monday night. (low-frequency)

Two lists were created, each containing the experimental items combined with 52 filler sentences. The two conditions were counterbalanced across lists, and the lists were randomized.

Procedure

The experimental task involved self-paced reading in a word-by-word moving window display (Just, Carpenter, & Woolley, 1982) using the Psyscope software package (Cohen, MacWhinney, Flatt, & Provost, 1993). At the start of each trial, a sentence appeared on the screen with all characters replaced by dashes. Participants pressed a key to change a string of dashes into a word. Each time the key was pressed, the next word appeared, and the previous word reverted back into dashes. The time between key-presses was recorded. After each sentence, participants answered a yes/no comprehension question. No feedback was provided for responses. Participants were asked to read at a natural pace and were given a small set of practice items in order to familiarize them with the task.

Results and discussion

Comprehension accuracy in the high-frequency and low-frequency conditions was 90% and 91%, respectively, and did not differ significantly across conditions ($p > .5$). Figure 2 shows mean reading times (RTs) per word. RTs were removed if they exceeded 3,000 ms. A 2 (high-frequency vs. low-frequency) \times 2 (Verb 1 vs. Verb 2) analysis of variance (ANOVA) revealed an effect of frequency condition in the region consisting of the two verbs following the pronoun (e.g., *met distrusted* vs. *distrusted met* in Example 2), $F_1(1, 34) = 6.22$, $MSE = 21,604$, $p = .018$; $F_2(1, 11) = 9.16$, $MSE = 5,189$, $p = .012$. The advantage of comparing this region is that averaging across the two verbs controls for differences in frequency and length of individual words. As shown in Figure 3, planned comparisons between the RTs averaged across the two-verb region revealed lower means in the high-frequency condition ($M = 443$ ms, $SD = 44$ ms) than in the low-frequency condition ($M = 507$ ms,

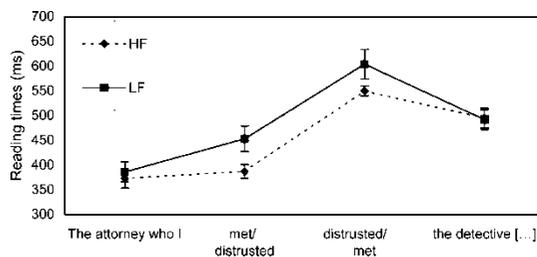


Figure 2. Results from Experiment 2: Mean reaction times across regions for high-frequency condition (dashed line) and low-frequency condition (solid line).

$SD = 64$ ms), $t_1(34) = 2.82$, $p = .004$; $t_2(11) = 2.06$, $p = .032$. The two-verb region contained the same words arranged differently across conditions (e.g., *met distrusted* in 2a, and *distrusted met* in 2b), and therefore the results cannot be explained by the frequency of individual words. Note, however, that the less frequent verb (e.g., *distrusted* in 2) is read first in the low-frequency condition and second in the high-frequency condition. Thus, processing spillover from the harder verb would remain within the target region in the low-frequency condition but could spill over to the following noun-phrase region in the high-frequency condition. However, RT

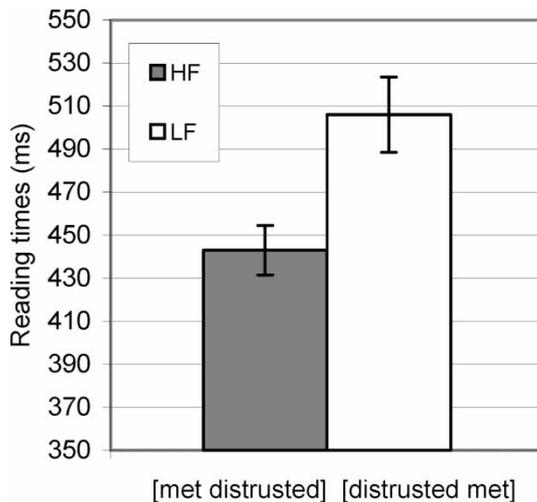


Figure 3. Mean reading times averaged across the two-word critical region for high-frequency condition (dark bar) and low-frequency condition (light bar).

comparisons in the region following the second verb (e.g., *the detective* in 2) revealed no measurable effect of spillover, $F_1(1, 34) < 0.5$; $F_2(1, 11) < 0.5$, $ps > .5$. This indicates that, if present, spillover effects produced by individual verbs are indistinguishable across conditions.

These findings suggest that the online processing of object-relative sentences is affected by the frequency of the embedded clause. A further question concerns the extent to which RTs are predicted by word-chunk frequencies across individual items. To explore this issue, we conducted a series of regression analyses to investigate the predictive power of the co-occurrence frequency of individual *I-verb* combinations forming the embedded clause.

In Regression 1 we explored whether the RTs recorded in the target region were predicted by the individual frequencies of the *I-verb* combinations forming the relative clause. The dependent variable consisted of the RTs averaged across the two-verb target region (*met distrusted* and *distrusted met* in 2), while the independent variable was the \log_{10} transform of the frequency (henceforth log-frequency) of the *I-verb* combinations in the object-relative clause (*I met* and *I distrusted* in 2). RTs were collapsed across high-frequency and low-frequency conditions into a single regression analysis, leading to a total of 24 data points (two conditions per item). As shown in Figure 4, the log-frequency of the *I-verb* combinations significantly predicted RTs across the two-verb target region, accounting for more than 55% of the variance, $\beta = -.74$, $R^2 = .556$, $F(1, 22) = 27.59$, $p < .0001$. This analysis provides strong evidence that the frequency of the embedded *I-verb* chunk facilitates overall object-relative processing. However, there is a significant correlation between the log-frequency of the *I-verb* combination (*I met* in 2a) and the log-frequency of the individual verb in the embedded clause (*met* in 2a), $R^2 = .54$, $p < .005$. Thus, a possible objection to our interpretation could be that the facilitation of object-relative processing is caused by the frequency of the individual verb appearing in the embedded position rather than by the frequency of the *I-verb* combination. To explore this

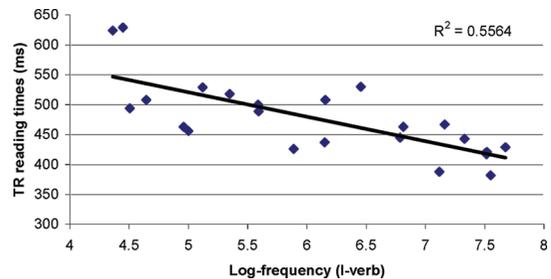


Figure 4. Results from Regression 1. The y-axis represents the averaged RTs across the target region (TR) comprising the two verbs following the pronoun *I*. The x-axis represents the log-frequency of *I-verb* combinations that form the relative clause.

possibility we conducted a hierarchical regression analysis (Regression 2) in which the dependent variable was the same as that in Regression 1, but in which two predictors were included in the analysis: The first variable was the log-frequency of the *I-verb* combination (*I met* in 2a), while the second variable was the log-frequency of the individual embedded verb (*met* in 2a). When both variables were entered, the model accounted for a significant amount of the variance in RTs, $R^2 = .556$, $F(2, 23) = 14.43$, $p = .0001$. However, analyses of individual contributions revealed that only the log-frequency of the *I-verb* combination was a significant predictor when the other factor was controlled for. That is, after the log-frequency of the *I-verb* chunk had been taken into account, the inclusion of the log-frequency of the embedded verb did not significantly improve prediction, $\beta = -.08$, $t(23) = 0.4$, $p = .68$. However, after the log-frequency of the embedded verb had been taken into account, the log-frequency of the *I-verb* combination still accounted for a significant amount of the variance in RTs, $\beta = -.68$, $t(33) = 3.31$, $p = .003$. This indicates that the facilitation effect is not explained by the frequency of individual verbs in the embedded position, but rather by co-occurrence patterns of the word sequence forming the relative clause.

In Regressions 1 and 2, the RTs recorded from both the high-frequency and the low-frequency conditions were collapsed in the regression analyses. Thus, the results might be partly due to categorical differences between RTs in the

high-frequency vs. low-frequency conditions. To explore this possibility we conducted a third regression analysis (Regression 3) in which the dependent variable was the across-condition *difference* in RTs in the target region (e.g., the RTs for *met distrusted* minus the RTs for *distrusted met* in 2), while the independent variable was the across-condition *difference* in the log-frequency of the *I-verb* combinations forming the clause—for example (log-frequency of *I met*) minus (log-frequency of *I distrusted*) in 2. Regression 3 revealed that the across-condition differences in log-frequencies significantly predicted the across-condition differences in RTs, $\beta = .72$, $R^2 = .52$, $F(1, 10) = 11.02$, $p = .007$.

Finally, we investigated whether the frequency of the *I-verb* combinations affected the RTs of the upcoming verb—that is, the main verb of the sentence. To do that, we conducted a regression analysis (Regression 4) in which the independent variable was the log-frequency of the *I-verb* combinations (*I met* in 2a), while the dependent variable consisted of the RTs of the main verb (*distrusted* in 2a). As shown in Figure 5, main-verb RTs were significantly predicted by the log-frequency of the *I-verb* combinations forming the preceding clause, $\beta = -.54$, $R^2 = .30$, $F(1, 22) = 9.44$, $p = .005$. Because there is no overlap between the *predictive* and *predicted* regions, these results cannot be explained by transitional probabilities of the type explored in MacDonald and Shillcock (2004). Rather, not only are word chunks more easily processed by themselves but also, as a by-product, they lead to further processing facilitation downstream when integrating the main verb into the ongoing interpretation. This account is further supported by the absence of a significant correlation between main-verb RTs and the log-frequency of the main verb itself ($p > .3$).

GENERAL DISCUSSION

Distributional properties of language are often described without considering differences between type and token frequencies. In Experiments 1 and 2, we showed that offline and online comprehension of pronominal

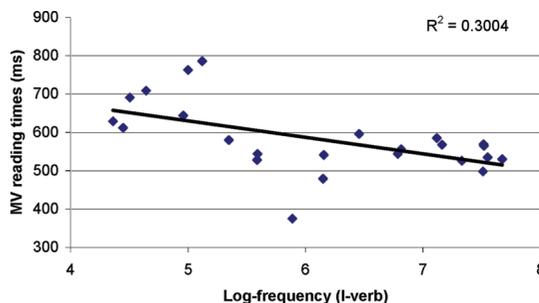


Figure 5. Results from Regression 4. The y-axis represents the averaged RTs recorded at the main verb (MV) region. The x-axis represents log-frequency of *I-verb* combinations that form the relative clause.

object-relative clauses is facilitated when the tokens forming the clause tend to co-occur frequently in the language. Importantly, the results cannot be explained by capacity-based theories in their current form. This is because the syntactic structure and the subject type in the mostly embedded position were identical in all items, and, therefore, integration and memory costs associated with these factors did not differ across conditions. However, it should be noted that capacity-based theories could be revised to accommodate these results, provided that they incorporate chunk-frequency as a factor capable of affecting memory demands during comprehension. It is also worth noting that the pronoun *I* was the only type of embedded subject in the materials used here. Thus, the question remains whether these results would generalize to other types of pronoun-verb combinations. Consistent with experience-based approaches, we expect generalization of these findings. However, it is hard to anticipate the nature of the possible interactions between fine-grained statistics and other probabilistic factors, such as, for example, contextual constraints defined at the discourse level.

The results suggest that, other factors being equal, the frequency of word chunks forming a relative clause influences its comprehension. The series of regression analyses conducted in Experiment 2 provided a way to explore some fine-grained aspects of the chunk frequency effect. Hierarchical regression analyses indicated

that log-frequency of the embedded *I-verb* combination significantly predicted RTs after controlling for frequency of the embedded verb. In contrast, *verb* frequency was not a significant predictor after controlling for the frequency of the *I-verb* chunk, suggesting that the effect on RTs was not due to differences in access to individual lexical items. Rather, access to word chunk representations may become easier as a function of the sequential co-occurrence patterns of their components. This interpretation is further supported by the results of Regression 4: Main-verb RTs were significantly predicted by the frequency of the relative clause, suggesting that the integration of the main verb into the unfolding interpretation may be facilitated by easier processing of the preceding clause. Additionally, the results of Regression 3 indicate that the frequency of the embedded *I-verb* combinations facilitates sentence processing in a *gradual* fashion. Elements that are frequently used together may be fused into processing units as a continuous function of their specific co-occurrence patterns. The gradual nature of the chunk frequency effect is consistent with sentence-processing approaches that advocate the existence of a continuity between language experience and comprehension.

In sum, these findings point toward a model of sentence processing and constituent representation in which language use and repetition play a crucial role. In the spirit of constructivist approaches, we have provided experimental support for the view that statistical tracking occurring at multiple levels of utterance representation affects the way we understand and represent linguistic structure, implicating a deep continuity between learning and comprehension processes over the course of development.

Original manuscript received 9 June 2006
Accepted revision received 1 August 2006
First published online 30 October 2006

REFERENCES

- Bybee, J. (2002). Sequentiality as the basis of constituent structure. In T. Givón & B. Malle (Eds.), *The evolution of language out of pre-language* (pp. 107–132). Philadelphia: John Benjamins.
- Bybee, J., & Scheibman, J. (1999). The effect of usage on degrees of constituency: The reduction of *don't* in English. *Linguistics*, *37*, 575–596.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavioral Research Methods, Instruments & Computers*, *25*, 257–271.
- Crocker, M. W., & Corley, S. (2002). Modular architectures and statistical mechanisms. In P. Merlo & S. Stevenson (Eds.), *The lexical basis of sentence processing* (pp. 157–180). Amsterdam: John Benjamins Publishing Company.
- Desmet, T., De Baecke, C., Drieghe, D., Brysbaert, M., & Vonk, W. (2006). Relative clause attachment in Dutch: On-line comprehension corresponds to corpus frequencies when lexical variables are taken into account. *Language and Cognitive Processes*, *21*, 453–485.
- Frison, S., Rayner, K., & Pickering, M. J. (2005). Effects of contextual predictability and transitional probability on eye movements during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 862–877.
- Gibson, E. (1998). Linguistic complexity: Locality and syntactic dependencies. *Cognition*, *68*, 1–76.
- Goldberg, A. (2006). *Constructions at work: The nature of generalizations in language*. New York: Oxford University Press.
- Grodner, D., & Gibson, E. (2005). Consequences of the serial nature of linguistic input. *Cognitive Science*, *29*, 261–291.
- Jurafsky, D. (1996). A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science*, *20*, 137–194.
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes and in reading comprehension. *Journal of Experimental Psychology: General*, *3*, 228–238.
- Keller, F., & Lapata, M. (2003). Using the web to obtain frequencies for unseen bigrams. *Computational Linguistics*, *29*, 459–484.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: A comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, *109*, 35–54.
- MacDonald, M., Pearlmutter, N., & Seidenberg, M. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*, 676–703.

- McDonald, S. A., & Shillcock, R. C. (2004). Eye-movements reveal the on-line computation of lexical probabilities during reading. *Psychological Science, 14*, 648–652.
- Mitchell, D. C., Cuetos, F., Corley, M. M. B., & Brysbaert, M. (1995). Exposure-based models of human parsing: Evidence for the use of coarse-grained (non-lexical) statistical records. *Journal of Psycholinguistic Research, 24*, 469–488.
- Race, D. S., & MacDonald, M. C. (2003). The use of “that” in the production and comprehension of object relative clauses. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Meeting of the Cognitive Science Society* (pp. 946–951). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Real, F., & Christiansen, M. H. (in press). Processing of relative clauses is made easier by frequency of occurrence. *Journal of Memory and Language*.
- Tomasello, M. (2003). *Constructing a language: A usage based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- Warren, T., & Gibson, E. (2002). The influence of referential processing on sentence complexity. *Cognition, 85*, 79–112.

APPENDIX

Experiment 1

Experimental items across conditions: high-frequency (HF) relative clause/low-frequency (LF) relative clause:

1. The comedian who the actress who [I loved] (HF)/[I phoned] (LF) phoned/loved presented his new show yesterday.
2. The kid who the teacher who [I hated] (HF)/[I praised] (LF) praised/hated transferred to a better school.
3. The scientist who the biologist who [I visited] (HF)/[I assisted] (LF) assisted/visited criticized the article about cloning.
4. The man who the woman who [I called] (HF)/[I bothered] (LF) bothered/called signed the petition for better schools.
5. The detective who the attorney who [I met] (HF)/[I distrusted] (LF) distrusted/met sent a letter on Monday night.
6. The boy who the girl who [I liked] (HF)/[I annoyed] (LF) annoyed/liked organized the school meeting last month.
7. The judge who the lawyer who [I asked] (HF)/[I contradicted] (LF) contradicted/asked rescheduled the audience for next week.
8. The assistant who the director who [I remembered] (HF)/[I promoted] (LF) promoted/remembered put the coupons into the locker.
9. The freshman who the student who [I helped] (HF)/[I advised] (LF) advised/helped improved his grades significantly.
10. The detective who the inspector who [I recognized] (HF)/[I evaded] (LF) evaded/recognized took a break in the hotel lobby.
11. The employee who the workman who [I recommended] (HF)/[I designated] (LF) designated/recommended fixed the problem easily.

12. The candidate who the politician who [I mentioned] (HF)/[I denounced] (LF) denounced/mentioned emphasized the need of funding.

Experiment 2

1. The actress who [I loved] (HF)/[I phoned] (LF) phoned/loved the comedian who presented his new show yesterday.
2. The teacher who [I hated] (HF)/[I praised] (LF) praised/hated the kid who transferred to a better school.
3. The biologist who [I visited] (HF)/[I assisted] (LF) assisted/visited the scientist who criticized the article about cloning.
4. The woman who [I called] (HF)/[I bothered] (LF) bothered/called the man who signed the petition for better schools.
5. The attorney who [I met] (HF)/[I distrusted] (LF) distrusted/met the detective who sent a letter on Monday night.
6. The girl who [I liked] (HF)/[I annoyed] (LF) annoyed/liked the boy who organized the school meeting last month.
7. The lawyer who [I asked] (HF)/[I contradicted] (LF) contradicted/asked the judge who rescheduled the audience for next week.
8. The director who [I remembered] (HF)/[I promoted] (LF) promoted/remembered the assistant who put the coupons into the locker.
9. The student who [I helped] (HF)/[I advised] (LF) advised/helped the freshman who improved his grades significantly.
10. The inspector who [I recognized] (HF)/[I evaded] (LF) evaded/recognized the detective who took a break in the hotel lobby.
11. The workman who [I recommended] (HF)/[I designated] (LF) designated/recommended the employee who fixed the problem easily.
12. The politician who [I mentioned] (HF)/[I denounced] (LF) denounced/mentioned the candidate who emphasized the need of funding.