

# Human Information Processing Shapes Language Change



Maryia Fedzechkina<sup>1,2,3</sup>, Becky Chu<sup>4</sup>, and T. Florian Jaeger<sup>4,5,6</sup>

<sup>1</sup>Department of Linguistics, University of Arizona; <sup>2</sup>Graduate Interdisciplinary Program in Cognitive Science, University of Arizona; <sup>3</sup>Graduate Interdisciplinary Program in Second Language Acquisition and Teaching, University of Arizona; <sup>4</sup>Department of Brain & Cognitive Sciences, University of Rochester; <sup>5</sup>Department of Computer Science, University of Rochester; and <sup>6</sup>Department of Linguistics, University of Rochester

Psychological Science  
1–11  
© The Author(s) 2017  
Reprints and permissions:  
sagepub.com/journalsPermissions.nav  
DOI: 10.1177/0956797617728726  
www.psychologicalscience.org/PS  
 SAGE

## Abstract

Human languages exhibit both striking diversity and abstract commonalities. Whether these commonalities are shaped by potentially universal principles of human information processing has been of central interest in the language and psychological sciences. Research has identified one such abstract property in the domain of word order: Although sentence word-order preferences vary across languages, the superficially different orders result in short grammatical dependencies between words. Because dependencies are easier to process when they are short rather than long, these findings raise the possibility that languages are shaped by biases of human information processing. In the current study, we directly tested the hypothesized causal link. We found that learners exposed to novel miniature artificial languages that had unnecessarily long dependencies did not follow the surface preference of their native language but rather systematically restructured the input to reduce dependency lengths. These results provide direct evidence for a causal link between processing preferences in individual speakers and patterns in linguistic diversity.

## Keywords

language universals, language processing, learning biases, language structure, language evolution, dependency length, open materials

Received 4/25/16; Revision accepted 8/8/17

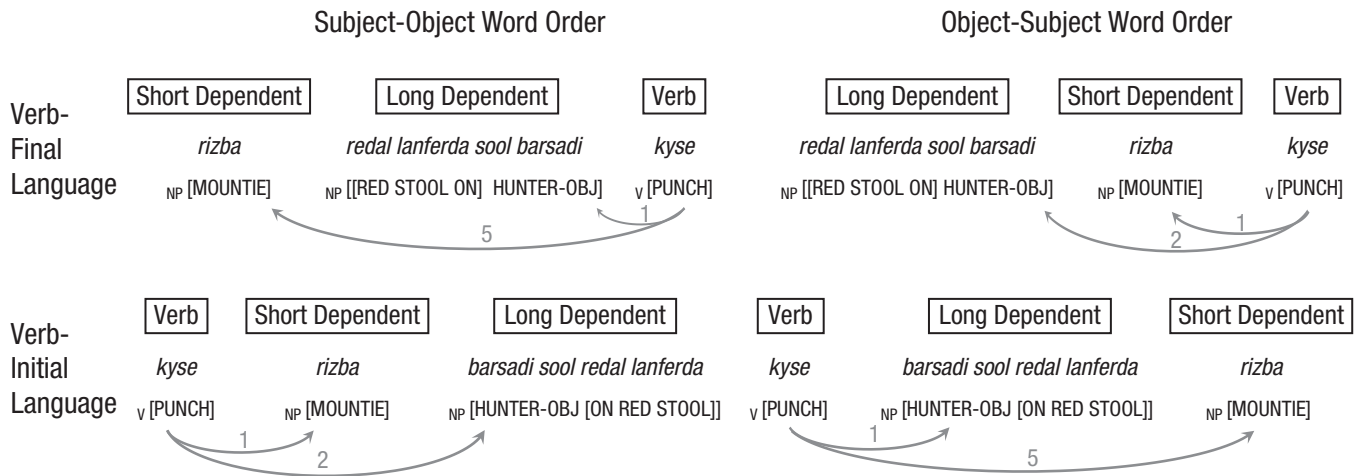
Natural languages vary along many dimensions, but this variation is not random—unrelated languages appear to share a striking number of underlying similarities. Understanding the constraints underlying these similarities has been the central question in the biological and language sciences. Most researchers agree that understanding these constraints could shed light on the mechanisms of language processing and representation in the human brain (e.g., Bates & MacWhinney, 1982; Chomsky, 1965; Christiansen & Chater, 2008; Fodor, 2001; Givón, 1991; Greenberg, 1963; Hawkins, 2014). Constraints specific to language (Chomsky, 1965; Fodor, 2001) and constraints rooted in general principles of human information processing (Christiansen & Chater, 2008; Hawkins, 2014) have been proposed.

Focusing on the latter type, we experimentally tested a hypothesized information-processing constraint operating on one of the most basic and perhaps most well-studied grammatical properties of human languages—the

way in which these languages order information in a sentence. Although the order of words in a sentence varies across languages, this variability is constrained. Across languages, some word orders are more frequent than others (Greenberg, 1963). Intriguingly, this cross-linguistic preference is sometimes also probabilistically mirrored within languages: If a language allows several word orders, the preferred orders typically correspond to the orders common across languages (Hawkins, 2014). Although it has long been hypothesized that pressures associated with human information processing influence word-order preferences across languages (Hawkins, 2014), the postulated causal link between the two has not yet been directly tested. We explored

## Corresponding Author:

Maryia Fedzechkina, Department of Linguistics, University of Arizona, 1103 E. University Blvd., Tucson, AZ 85721  
E-mail: mfedzech@email.arizona.edu



**Fig. 1.** Comparison of dependency lengths for two possible word orderings (subject-object and object-subject) in verb-final and verb-initial languages. All four sentences express the same meaning. Curved arrows represent grammatical dependencies between the verb (V) and the closest constituent boundary of its two arguments. Numbers represent dependency lengths, measured in words. For verb-final languages (top), ordering long dependents before short dependents leads to shorter total dependency length between the dependents and their head (the verb). For verb-initial languages (bottom), the relationship between total dependency length and the order of dependents relative to the head is reversed: Ordering short dependents before long dependents leads to shorter overall dependency length. The words in brackets are English translations. NP = noun phrase; OBJ = object case marking.

whether a bias toward short grammatical dependencies could explain these cross-linguistic word-order preferences.

Grammatical dependencies are asymmetric relations between the *head* (a word that licenses the presence of other words) and a *dependent* (a word that modifies the head). For example, in the sentence *The boy is kicking the ball*, the head (the verb *kicking*) forms two grammatical dependencies—one with the subject (*the boy*) and one with the direct object (*the ball*). Psycholinguistic evidence shows that the length of the dependency (i.e., the distance between the head and its dependent) affects comprehension efficiency: Longer dependencies are associated with greater processing difficulty (Grodner & Gibson, 2005), an effect that is presumably due to memory retrieval (Bartek, Smith, Lewis, & Vasishth, 2011). Likewise, language production also exhibits a preference for shorter dependencies. When several word-order choices are available to convey the same message, speakers of verb-initial languages (i.e., languages that place the verb before its dependents) and verb-medial languages (i.e., languages that place the verb after the subject and before the object), such as English, tend to place short postverbal constituents before long postverbal constituents (Arnold, Wasow, Losongco, & Ginstrom, 2000; Wasow, 2002). In contrast, speakers of verb-final languages (i.e., languages that place the verb after the dependents), such as Japanese, typically prefer long preverbal constituents before short preverbal constituents (Ros, Santesteban, Fukumura, & Laka, 2015; Yamashita & Chang, 2001). The respective verb-dependent orderings

reduce the average dependency length in a sentence (Fig. 1).

Although the processing advantage of shorter dependencies is well established, its contribution to historical word-order change is still under debate. Recent large-scale computational studies have provided some support for the processing account: Among the languages studied so far (almost 40), average dependency lengths are significantly shorter than would be expected by chance (Ferrer i Cancho, 2004; Futrell, Mahowald, & Gibson, 2015; Gildea & Temperley, 2010), and some languages are close to the theoretical minimum (Gildea & Temperley, 2010). Although these studies suggest a correlation between the preference for shorter dependencies in online processing and cross-linguistic word-order constraints, they also face two critical limitations. First, typological (i.e., cross-linguistic) data are sparse, which makes it difficult to convincingly test the validity of cross-linguistic correlations (see debates in Croft, Bhattacharya, Kleinschmidt, Smith, & Jaeger, 2011; Dryer, 2011; Dunn, Greenhill, Levinson, & Gray, 2011). Second, and more crucially, typological data cannot directly address questions about the underlying causes of this hypothesized correlation. Thus, although it has been widely assumed that dependency-length minimization (DLM) influences one of the fundamental abstract properties of language—grammatical constraints on word order—the causal link between the two has not yet been directly tested. This leaves open the possibility that the word-order patterns consistent with DLM that have been observed in previous correlational studies are spurious.

In the current study, we used a miniature-artificial-language learning paradigm (Hudson Kam & Newport, 2009; Kirby, Tamariz, Cornish, & Smith, 2015) to directly probe the causal link between processing biases in individual learners and the DLM preference observed across languages. Learning of miniature languages has been successfully used to study mechanisms of first- and second-language acquisition (Pajak & Levy, 2014; Saffran, Aslin, & Newport, 1996). Recent work has adapted this paradigm to explore the underlying causes of cross-linguistic patterns by creating situations of atypical highly variable input (reminiscent of situations of pidgin or of language change) in the laboratory and studying how learners deviate from the atypical input they receive (Culbertson, Smolensky, & Legendre, 2012; Fedzechkina, Jaeger, & Newport, 2012; Hudson Kam & Newport, 2009; Kirby et al., 2015; Smith & Wonnacott, 2010).

We tested whether the cross-linguistic bias toward shorter dependencies originates in the limitations of the human processing system. We directly assessed one specific pathway—biases operating during language learning—by which DLM could come to cause the observed cross-linguistic patterns. We presented learners with input languages that had inefficient (unnecessarily long) dependencies and tested whether learners shifted the language toward more efficient (shorter) dependencies. If DLM causes learners, on average, to produce languages that deviate slightly from the original input toward word orders with shorter dependencies, the input for the next generation of learners would, on average, contain shorter dependencies. This would allow biases toward shorter dependencies to accumulate over generations of learners. We note that this account does not predict that all languages converge on the same word orders, because it is plausible to assume that there are trade-offs between DLM and other learning and processing biases.<sup>1</sup>

If learners are indeed biased toward shorter dependencies, this would support proposals that attribute certain cross-linguistic word-order patterns to DLM (e.g., Hawkins, 2014). On the other hand, if learners exhibit no preference toward languages with shorter dependencies, this would constitute a serious challenge for such accounts. Testing whether DLM cause learners to deviate from the input can show whether and how a specific processing preference can contribute to patterns in cross-linguistic word-order variation.

## Method

### *Participants*

The Research Subjects Review Board at the University of Rochester approved the recruitment of participants


and the execution of this study. Participants in the experiment were monolingual native English speakers between the ages of 18 and 30 recruited from the University of Rochester and the surrounding community. Each participant was exposed to only one language and received \$30 for participation. To reduce the researcher degrees of freedom, we continued recruitment until 20 participants successfully learned each language (as in our earlier work; Fedzechkina et al., 2012; Fedzechkina, Newport, & Jaeger, 2017). Most participants successfully learned the assigned language (to have 20 successful learners for each of the two languages, we recruited 45 participants; for details, see the Scoring section).

### *Design and materials*

The participants learned miniature artificial languages by watching short videos describing simple transitive events performed by two human actors (e.g., “chef punch referee”) and hearing their descriptions in the novel language. Both languages had flexible word order, so that subject-object (SO) and object-subject (OS) orders occurred equally frequently in the input. Like many languages with flexible word order (Blake, 2001), our languages had consistent case marking—a noun suffix that disambiguated who was doing what to whom in the scene. The case marker was always “di” and occurred on all direct objects. The languages shared the same lexicon of four transitive verbs, eight nouns (six animate and two inanimate), three adpositions (“with,” “next to,” and “on”), and two color adjectives (“blue” and “red”).<sup>2</sup> For more details about the languages, see the Supplemental Material available online. Both languages contained adpositional phrases (e.g., “chef next to blue skateboard”; see Fig. 2). The order of the adposition (e.g., “next to”) relative to its dependent (“blue skateboard”) and head (“chef”) followed patterns that are common across languages (Dryer, 2013), as shown in Figure 1.

The miniature languages differed in whether they were verb-final or verb-initial. As is common cross-linguistically, the verb-final language used prenominal postpositional phrases (as in Japanese or Hindi), ordering the adposition after its dependent and before its head (e.g., “blue skateboard next to chef”). The verb-initial language used postnominal preposition phrases (as in English), ordering the adposition after its head and before its dependent (e.g., “chef next to blue skateboard”).

In training, participants were exposed to sentences that contained either two short constituents (i.e., neither subject nor object had adpositional-phrase modification; 50% of training scenes) or two long constituents (i.e., both subject and object had adpositional-phrase modification; 50% of training scenes). Sentences in



**Fig. 2.** Illustration of the constituent-length manipulation in the experiment. Visual scenes like those on the left tend to elicit short descriptions, whereas visual scenes like those on the right tend to elicit more complex, long descriptions. Example descriptions (provided only auditorily in the experiment) are shown for the verb-final miniature language; participants did not hear or see the English glosses.

which subject and object phrases differed in length were not part of the input. Word order was thus independent of phrase length in the input, and both short-short and long-long scenes occurred equally frequently with OS and SO orders. During the production test, participants described previously unobserved scenes that contained either one long subject constituent, one long object constituent, or no modification of either constituent. Each of these three possibilities occurred equally often across scenes.

### Procedure

The experiment was conducted in a 1-hr session on each of three consecutive days (in some cases, a day was skipped between two of the sessions). All three sessions involved similar combinations of exposure and test blocks (see Fig. 3); there was more intensive vocabulary exposure on Day 1 and more intensive sentence exposure on Days 2 and 3.

**Noun exposure and tests.** Participants saw pictures of characters or objects one at a time, accompanied by their names in the novel language, and were instructed to repeat the names out loud to facilitate learning. After noun exposure, participants completed noun comprehension and production tests. In the comprehension test, participants were shown a set of four character pictures accompanied by a name in the novel language and asked to choose the character matching the name. In the production test, participants were asked to name the character shown on the screen. Feedback on performance was provided after each trial in both tests.

**Phrase exposure and tests.** Participants were explicitly informed that they would learn phrases in the new language. These phrases contained a character modified by a description (for more details, see Fig. 2). The same procedure used in the vocabulary training and tests was used here.

**Sentence exposure and comprehension test.** Participants learned the grammar by watching short videos and hearing descriptions of the videos in the novel language. Participants were instructed to repeat the sentences aloud to facilitate learning. On Day 1, participants could replay the videos and the sound as many times as they wished; no repetitions were allowed on subsequent days.

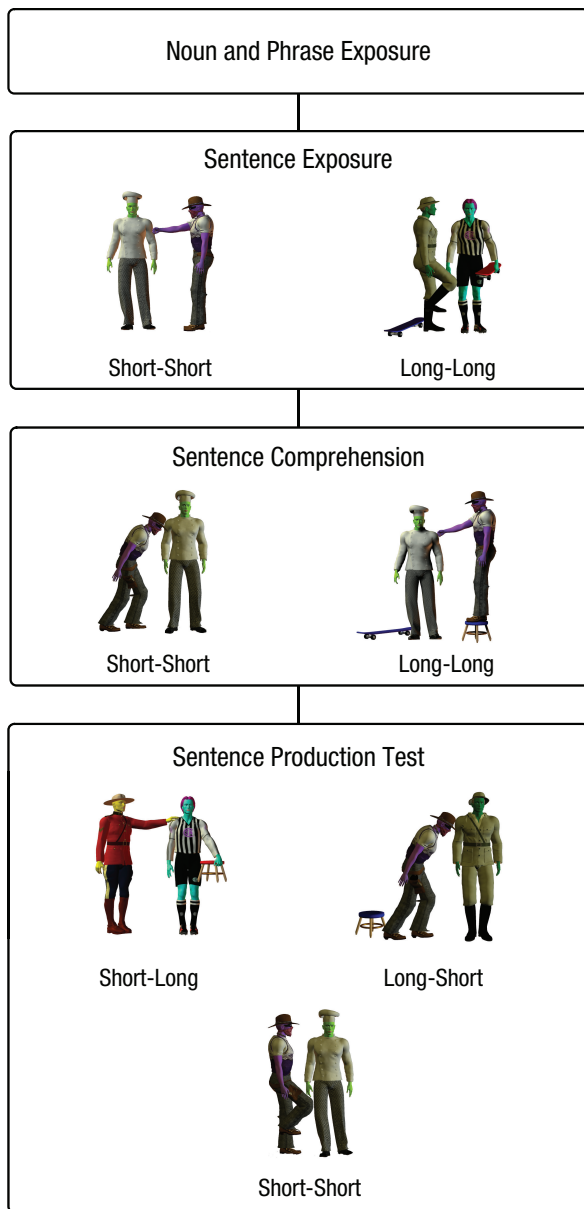
After sentence exposure, participants performed a sentence-comprehension test. Participants were presented with two side-by-side videos accompanied by an auditory description. The two videos showed the same action and characters, but the order of the actor and patient of the action reversed. Participants were asked to choose the video that matched the description. Feedback on performance was provided on each trial.

**Production test.** Participants were shown two previously unseen videos, side by side, depicting the same action, but with the characters' grammatical roles switched (i.e., in one video, the character was the subject of the sentence, and in the other video, the character was the object of the sentence). One of the videos was highlighted. The videos disappeared from the screen after 1,200 ms and were replaced by a crosshair in the center of the screen. Participants were instructed to describe the highlighted video after seeing the crosshair. A verb prompt was provided to facilitate the descriptions. No feedback on performance was provided during this test.

The use of two videos was meant to encourage participants to produce adpositional phrases (e.g., “with skateboard”) for highlighted videos when the visual scene required an adpositional phrase. Arguably, a better way to elicit adpositional phrases might have been to present two videos, one for which an adpositional phrase was required and one for which it was not (e.g., “chef” in Video 1 and “chef with skateboard” in Video 2), rather than two videos with switched subject and object roles. However, participants overwhelmingly produced adpositional phrases as required by the scene. This was reflected in the high production accuracies reported in the next section.

### Results

Before turning to the predictions and central findings of our work, we describe how our data were scored and discuss learners' acquisition accuracy. We then outline our predictions and present the analyses of learners' language production.



**Fig. 3.** Schematic of the phases of exposure to novel miniature languages and testing of sentence comprehension and sentence production at all sessions. Participants were first shown pictures of characters or objects one at a time, accompanied by their names in the novel language, and then learned how to put the words together into phrases (not shown here). They then watched short videos of simple transitive events performed by two human actors and heard one-sentence descriptions of those events in the novel language; each sentence contained either two short constituents or two long constituents. During the comprehension test, participants watched two videos and heard a description and were asked which video best fit the description. In the sentence-production test, participants were shown two videos depicting the same action but with the characters' grammatical roles switched and were then asked to describe one of the videos. The measure of interest was the word order that learners used in this test, which could have two short constituents or one long constituent (either a long subject or a long object). The images shown represent still frames of the video stimuli used in the experiment.

## Scoring

We first examined the accuracy of acquisition of both languages. We used, as a measure of comprehension accuracy, whether participants chose the correct video to match the sentence they heard. Because all sentences were disambiguated by case marking, this measure allowed us to assess how well learners acquired the grammar of the novel language. Recruitment continued until the number of participants who achieved 70% accuracy on sentence-comprehension tests on the final day of training reached 20 in each language. Participants who failed to pass this accuracy requirement (3 participants in the verb-final language and 2 participants in the verb-initial language) were removed from further analyses. The pattern of results reported below does not depend on this exclusion.

The 40 learners submitted for further analysis achieved a high level of comprehension accuracy on the final day of training (97% accuracy for both languages). Production performance showed a similarly high degree of accuracy, which suggests that the task was feasible (for details on production scoring, see the Supplemental Material). For the verb-final language, participants made 8.2% lexical mistakes and 3.5% grammatical mistakes on the final day of training. For the verb-initial language, participants made 12.5% lexical mistakes and 2.7% grammatical mistakes. All analyses reported here are based only on grammatically correct sentences from the production test. We follow our previous work in not removing lexical mistakes from the analysis. The results reported below do not depend on this decision.

Given the high accuracy of acquisition of both languages, any observed word-order preferences are unlikely to be due to insufficient knowledge of the lexicon and syntactic structure of the novel language.

## Predictions

The central hypothesis of our study was that learners are biased toward shorter grammatical dependencies. We assessed whether such bias existed in learners' language production in two ways. We first explored whether learners ordered constituents within each language in a manner predicted by DLM accounts. Second, we tested whether the DLM preference caused learners to deviate from the input toward significantly shorter overall dependencies when we accounted for the amount of word-order flexibility in the language. These two tests complemented each other: *Whether* and *how much* participants shortened dependency lengths compared with the input language also depended on participants' overall word-order preferences.

### **Relative constituent length predicts learners' word-order choices in production**

We began by testing whether learners' production preferences were affected only by the surface ordering preferences in their native language or whether they were also driven by a deeper underlying principle of DLM. If learners' word-order preferences were affected only by the surface ordering biases of their native language, we would expect learners to follow English-like short-before-long ordering (Arnold et al., 2000; Wasow, 2002). If, on the other hand, learners' word-order preferences were driven by a deeper underlying principle of DLM, we would expect learners to introduce a preference for shorter dependencies in their scene descriptions. This preference should result in opposite surface orderings for the two languages: long-before-short ordering in the verb-final language and short-before-long ordering in the verb-initial language (see Fig. 1).

To assess learners' preferences in length-based ordering, we conducted a mixed-effects logistic regression analysis. We predicted learners' SO word-order frequency from constituent length (all constituents short vs. object long, subject long vs. all other cases; Helmert coded), day of training (2 vs. 1, 3 vs. all other cases; Helmert coded), and their interactions. This analysis thus assessed learners' ordering preferences on the basis of the relative order of constituents within a language, regardless of what other biases might affect overall word-order preferences. The model contained the maximal random-effects structure justified by the data according to backward model comparison (by-participant random intercept, by-participant random slopes of day and constituent length). The same results were obtained in the model with the maximal random effects structure that still converged.

**Verb-final miniature language.** As predicted by the DLM hypothesis, learners' word-order preferences in the verb-final miniature language revealed a bias for shorter dependencies. Despite receiving an unbiased input and having the opposite short-before-long preference in their native language, learners of the verb-final language introduced a long-before-short ordering in their own scene descriptions (see Fig. 4). Across all 3 days of training, learners were significantly more likely to use SO order for sentences with long subject constituents and short object constituents compared with other sentences types,  $\hat{\beta} = 1.36$ ,  $z = 5.56$ ,  $p < .001$ . Likewise, learners were significantly more likely to use SO order for sentences in which both subject and object constituents were short, compared with sentences with short subject constituents and long object constituents,  $\hat{\beta} = 0.66$ ,  $z = 2.54$ ,  $p = .011$ . There

was no main effect of day of training ( $ps > .408$ ), but day of training did interact with the effects of constituent length. On Day 2, the difference in SO word-order frequency between sentences with long subjects and all other sentence types was significantly smaller than on Day 1,  $\hat{\beta} = -0.56$ ,  $z = -4.02$ ,  $p < .001$ . The difference in SO word-order frequency for sentences with two short constituents compared with the sentences with long objects was significantly greater on Day 3 than on Day 2,  $\hat{\beta} = 0.18$ ,  $z = 2.12$ ,  $p = .034$ .

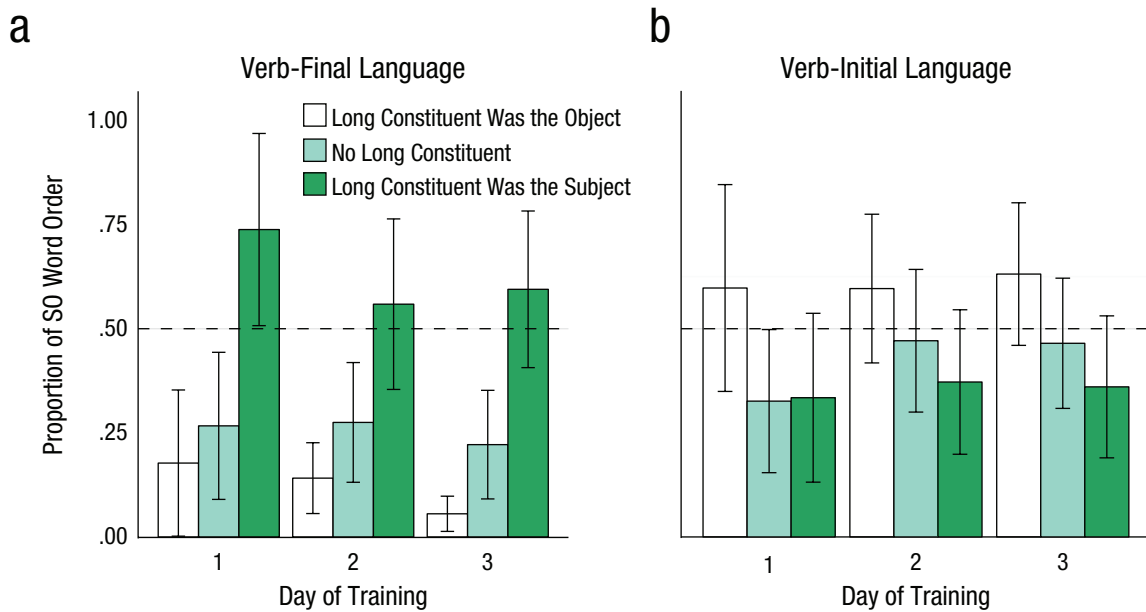
The analysis of simple effects revealed that learners used SO word order significantly more frequently in sentences with long subjects than in all other cases on all days of training—Day 1:  $\hat{\beta} = 1.96$ ,  $z = 5.66$ ,  $p < .001$ ; Day 2:  $\hat{\beta} = 0.84$ ,  $z = 3.56$ ,  $p < .001$ ; Day 3:  $\hat{\beta} = 1.27$ ,  $z = 5.21$ ,  $p < .001$ . The difference in SO word-order frequency for sentences with long objects compared with sentences with two short constituents reached significance only on the final day of training, after participants became sufficiently fluent in the novel language,  $\hat{\beta} = 1.02$ ,  $z = 3.36$ ,  $p < .001$ .

**Verb-initial miniature language.** The verb-initial language was analyzed according to the same statistical procedure and variable coding as the verb-final language.

As predicted by the DLM hypothesis, learners of the verb-initial language introduced a short-before-long ordering preference in their utterances—the opposite preference of that observed in the verb-final language (see Fig. 4). Across all days, learners were significantly less likely to use SO word order in sentences with long subject constituents and short object constituents than in all other sentence types,  $\hat{\beta} = -0.44$ ,  $z = -2.5$ ,  $p = .012$ . Likewise, learners were significantly less likely to use SO order in sentences with short subject constituents and short object constituents than in sentences with long object constituents and short subject constituents,  $\hat{\beta} = -0.47$ ,  $z = -2.01$ ,  $p = .044$ . This preference did not interact with day of training ( $ps > .2$ ), and there was no main effect of day of training ( $ps > .6$ ).

Simple-effects analyses showed that the bias against SO word order in sentences with long subjects compared with all other sentence types was significant on Day 2,  $\hat{\beta} = -0.42$ ,  $z = -2.27$ ,  $p = .023$ , and Day 3,  $\hat{\beta} = -0.52$ ,  $z = -2.9$ ,  $p = .003$ , and marginally significant on Day 1,  $\hat{\beta} = -0.37$ ,  $z = -1.76$ ,  $p = .078$ . The difference in SO word-order frequency for sentences with two short constituents compared with sentences with long objects became significant with sufficient proficiency in the novel language—on the final day of training,  $\hat{\beta} = -0.6$ ,  $z = -2.47$ ,  $p = .014$ .

As predicted by the DLM hypothesis, learners preferred opposite length-based constituent orders for verb-initial and verb-final languages. This suggests that



**Fig. 4.** Subject-object (SO) word-order frequency in the sentences produced by participants who learned (a) the verb-final language and (b) the verb-initial language. Results are shown separately for each day of training and for sentences with long object constituents, sentences with no long constituents, and sentences with long subject constituents. The dashed line indicates the proportion of SO order in the training input, which was equal across all sentence types and languages. The error bars represent 95% confidence intervals.

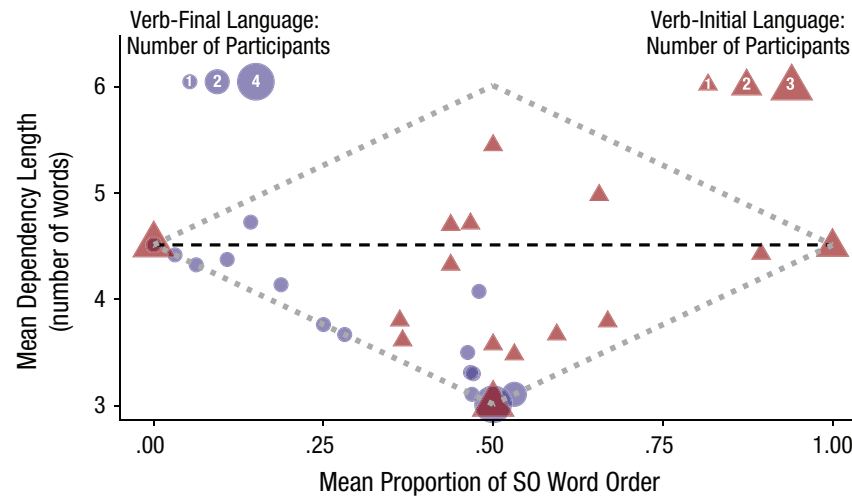
their word-order choices in production are driven by a deeper underlying preference for DLM.

The results also reveal some differences in learners' preferences across the two languages. First, the effect appears stronger in the verb-final language than in the verb-initial language: Learners of the verb-final language introduced more pronounced changes into the input word order than learners of the verb-initial language. Comparisons with the input discussed in the next section confirm this observation.

Second, Figure 4 reveals that learners of the two languages differed in their overall preference for SO order. This is evident when considering only baseline (short-short) trials, for which DLM makes no ordering predictions. For baseline trials, learners of the verb-initial language matched the input on their final day of training, with 46% SO production (Wilcoxon signed-rank test over by-participant proportions:  $V = 64.5$ ,  $z = -0.39$ ,  $p = .693$ ). Learners of the verb-final language used SO word order significantly less often than in the input (22% SO order;  $V = 15$ ,  $z = -3.26$ ,  $p = .001$ ). These word-order preferences speak against direct native-language influences on learners' performance. If learners transferred surface-based ordering preferences from their native language into our experiment, we should have found a preference for the subject-first SO order (as in English), compared with the input. However, this was not the case.

The bias against SO in the verb-final language was probably due to a strong preference to provide case marking at the beginning of the sentence (in the miniature languages in the experiment, case marking occurred only on the object)—a bias we have repeatedly observed in previous work (Fedzechkina et al., 2012, 2017). One possible cause for this effect is a processing preference for providing informative cues at sentence onset in parsing (Hawkins, 2014; for independent evidence from artificial languages, see Fedzechkina, Jaeger, & Trueswell, 2015). Given the incremental nature of sentence processing, placing a case-marked constituent at sentence onset would allow comprehenders to converge on the correct interpretation early on and avoid costly revisions. This explanation would leave open the reason for the smaller bias against SO in the verb-initial language compared with the verb-final language. One possible explanation—left to future work—is that verbs in verb-initial languages tend to be highly informative about the correct sentence interpretation and comprehenders are sensitive to this information (Garnsey, Perlmutter, Meyers, & Lotocky, 1997). Thus, it is possible that the perceived utility of case marking is reduced in verb-initial languages.

Regardless of the overall difference in their preference for SO order, learners of both languages ordered longer constituents further away from the verb, as predicted by the DLM hypothesis. We then analyzed



**Fig. 5.** Average dependency length in participants' utterances with one long constituent on the final day of training as a function of the overall proportion of subject-object (SO) word order. The size of each data point is proportional to the number of participants represented (1–4). The bold dashed line shows the per-sentence dependency length expected if learners exhibited no length-based ordering preferences. Points for participants who reduced dependency length compared with the input fall below the dashed line. The dotted gray lines encompass the area into which all points had to fall. Data points for learners whose dependency length was as short as possible given their overall SO preference fall on the lower gray lines.

whether the respective length-based ordering preferences in learners' utterances resulted in shorter dependency lengths compared with the input.

### ***Learners deviate from the input toward shorter dependency lengths***

The analyses conducted so far show that word-order preferences within each language followed the DLM prediction when overall biases in word-order frequency were ignored. This leaves open the question of whether length-based orderings introduced by learners result in shorter average dependency lengths compared with the input when learners' overall word-order preferences in the language are taken into account, as would be expected if DLM strongly affects word-order preferences in learners' utterances.

To address this question, we compared average per-sentence dependency length (measured in words) on the final day of training with the expected average per-sentence dependency length in the input (which did not contain length-based ordering preferences). Baseline trials (short-short) were excluded from this analysis because they are uninformative with regard to the evaluation of DLM accounts.

As predicted by the DLM hypothesis, the output languages produced by learners had significantly shorter dependency lengths than the input dependency length of 4.5. The verb-final language had an average dependency length of 3.64 on Day 3, Wilcoxon signed-rank

test over by-participant proportions:  $V = 4$ ,  $z = -3.06$ ,  $p < .001$ . The verb-initial language had an average dependency length of 4.09 on Day 3,  $V = 22$ ,  $z = -2.09$ ,  $p = .033$ . Overall, all but 5 of 40 learners either matched (6 learners, 15%) or reduced (28 learners, 70%) dependency length compared with the input.

Thus, the reduction in dependency length in our experiment was driven by a clear majority of the learners. As a final assessment, we quantified the degree of DLM compared with the theoretically possible minimization. As shown in Figure 5, the amount of DLM that can be achieved is conditional on learners' overall word-order preference. Minimal theoretically possible dependency lengths are attainable only if learners maintain perfectly flexible (SO vs. OS) word-order frequency. It is thus worth determining whether learners minimize dependency length given their overall SO-versus-OS preference.

Twenty of the 40 learners (50%) achieved the minimal theoretically possible dependency length, given their overall preference for SO word order (i.e., their utterances fall on the lower gray lines in Fig. 5). These 20 included 7 of 9 learners who maintained perfect word-order flexibility in their utterances and produced output languages with dependency lengths indistinguishable from the absolute minimal dependency length possible. In addition, 6 learners (15%) used completely fixed word order, thereby trivially producing the minimal possible dependency length for their overall word-order preference.



Figure 5 also reveals that learners of the verb-final language followed the DLM principle more strongly than learners of the verb-initial language. One possible explanation is that the DLM preference is enhanced when it favors a word-order variant that is preferred in the language for other reasons. Recall that learners of the verb-final language produced significantly more OS order, which is consistent with a preference to provide informative cues at sentence onset. When DLM favored OS order, learners of the verb-final language followed this preference significantly more strongly than when it favored SO order, Wilcoxon signed-rank test over by-participant proportions:  $W = 97.5$ ,  $z = -2.99$ ,  $p = .012$ . Learners of the verb-initial language, who used SO and OS orders equally frequently, followed DLM equally strongly for both orders ( $W = 206$ ,  $z = 0.13$ ,  $p = .880$ ). Note that learners of both languages in our experiment showed a preference to reduce dependency lengths compared with the input, which suggests that the observed learning outcomes in the two languages cannot be fully explained by learners' baseline word-order preferences.

Thus, the DLM hypothesis is supported by both the length-based ordering preferences within each language and the reduction of dependency length compared with the input. As predicted by the DLM hypothesis, learners of the verb-initial and verb-final languages introduced opposite length-based orders into their utterances. Learners did so in ways that resulted in a significant reduction of the average dependency length compared with the input.

## Discussion

The current study presents the first direct test of the hypothesized causal link between a processing bias for shorter grammatical dependencies and cross-linguistic word-order distributions (as predicted in Hawkins, 2014). Our learners shared the same language background and received input languages with the same statistics but had different word-order preferences depending on the verb (head) position in the language. As predicted by DLM, learners preferred short-before-long ordering in the verb-initial language and long-before-short ordering in the verb-final language, which resulted in shorter dependencies in the two languages. This lends credibility to the hypothesis that the cross-linguistic preference for short dependencies originates in constraints on human information processing.

Our work adds to the debate on the role of linguistic-specific versus domain-general constraints on word-order distributions. Traditionally, grammatical constraints on word order have been explained without a reference to processing by postulating linguistic-specific generalizations such as harmony universals (e.g., a preference

to place heads either consistently before or after their dependents; Baker, 2001; Travis, 1984) or basic word-order universals (e.g., a cross-linguistic preference for SOV order; Coopmans, 1984). Later researchers, drawing on cross-linguistic correlational data, have proposed alternative explanations of these universals in terms of DLM—and thus, as widely assumed, in terms of human information processing (Hawkins, 2014). We find that DLM indeed influences word-order distributions (at least when the input language allows two orders): Learners consistently produce output languages that have shorter dependency lengths. This suggests that DLM-based explanations of harmony and basic word-order universals are plausible, which makes DLM a potential unifying cause behind several types of cross-linguistic word-order generalizations.

Learners' preferences in our experiment were driven by an underlying DLM preference. Learners, however, did not produce languages that have optimal dependency lengths. Instead, DLM introduced small shifts into learners' utterances, thus providing a seed for this cross-linguistic preference. An important open question for future research is whether these changes accumulate as the language is transmitted over generations of speakers (as we assume), thereby causing gradual language change over historical time (e.g., Christiansen & Chater, 2008; Kirby et al., 2015).

Can our findings be accounted for by learners' native-language preferences? Native-language transfer effects are widely attested in second-language acquisition (for a review, see Pajak, Fine, Kleinschmidt, & Jaeger, 2016) and thus present a serious consideration when interpreting our results. Our participants' native language (English) has an overall short-before-long preference. This could explain the result for the verb-initial miniature language, but not the inverse long-before-short preference in the verb-final language. This rules out direct surface-based transfer from English to the miniature languages as a source of the observed effects.

A related possibility is that learners transfer some form of context-specific ordering bias from English. For example, English allows topicalization (e.g., "Cheese, John already bought") and left dislocation (e.g., "Cheese, John already bought it"). These structures realize phrases that would otherwise occur after the verb—here the direct object—at sentence onset. There is suggestive evidence that long phrases are more likely to be topicalized or left dislocated than short phrases—a preference that is itself predicted by DLM (Snider & Zaenen, 2006). This raises the possibility that the long-before-short preference in the verb-final language is explained by a native language preference to topicalize or left-dislocate long phrases. But several properties of these structures in English make this possibility rather unlikely. First, they are licensed only in specific

discourse contexts (Prince, 1995) that differ from those in our experiment. Second, both structures are extremely rare in English (< 0.7% of all sentences; Gregory & Michaelis, 2001). Low-frequency native-language structures might give rise to transfer effects in miniature-language studies (Goldberg, 2013). However, this would still raise the question of why we found no evidence of a more direct transfer from English, such as an overall preference for SO order.

One important question that is left open pertains to the origin of the DLM preference that learners exhibit. Is this preference based on an innate cognitive principle or on an abstract principle acquired from the statistics of the learners' native language (Culbertson & Adger, 2014)? English exhibits DLM particularly strongly—its average dependency lengths are close to the theoretical *minimum* (Gildea & Temperley, 2010). Thus, it is possible that native speakers of English are especially attuned to DLM and are readily extending this abstract preference to the novel miniature languages. Future extensions of our work to languages with weaker DLM preferences (e.g., German or Japanese) could address this possibility. If the preference observed in our experiment is indeed learned from the statistics of English, it raises the question of why English expresses this preference. For now, we note that DLM provides a unifying explanation for the existence of these biases both in English and in the novel miniature language. Another potentially appealing aspect of this hypothesis is that it is part of a more general proposal suggesting that the human information-processing system prefers certain structures and thus can provide a parsimonious domain-general account of constraints on language structure.

### Action Editor

Matthew A. Goldrick served as action editor for this article.

### Author Contributions

M. Fedzechkina and T. F. Jaeger developed the study concept and contributed to the design. B. Chu conducted the study and collected the data. M. Fedzechkina analyzed the data with input from T. F. Jaeger. All the authors contributed to writing the manuscript and approved the final version of the manuscript for submission.

### Acknowledgments

We thank Madeline Clark, Irene Minkina, Andy Wood, and Cassandra Donatelli for help with data coding and stimuli creation.

### Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

### Funding

This work was supported in part by National Science Foundation CAREER Grant IIS-1150028 (to T. F. Jaeger).

### Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797617728726>

### Open Practices



All materials have been made publicly available via the Open Science Framework and can be accessed at <https://osf.io/dbf2k/>. The complete Open Practices Disclosure for this article can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797617728726>. This article has received the badge for Open Materials. More information about the Open Practices badges can be found at <https://www.psychologicalscience.org/publications/badges>.

### Notes

1. Although this was not our primary goal, our experiment allowed us to assess trade-offs between DLM and other learning and processing biases against one specific, already-documented learning bias—a bias toward simplifying the grammar by fixing previously variable word order (Hudson Kam & Newport, 2009).
2. The term *adposition* refers to both prepositions and postpositions. The words used in this section (e.g., “red,” “next to”) are the English translations of the words used in the miniature artificial language.

### References

- Arnold, J. E., Wasow, T., Losongco, T., & Ginstrom, R. (2000). Heaviness vs. newness: The effects of structural complexity and discourse status on constituent ordering. *Language*, 76, 28–55.
- Baker, M. (2001). *The atoms of language: The mind's hidden rules of grammar*. New York, NY: Basic Books.
- Bartek, B., Smith, M., Lewis, R., & Vasishth, S. (2011). In search of on-line locality effects in sentence comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 1178–1198.
- Bates, E., & MacWhinney, B. (1982). Functionalist approaches to grammar. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 173–218). Cambridge, England: Cambridge University Press.
- Blake, B. J. (2001). *Case*. Cambridge, England: Cambridge University Press.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Christiansen, M. H., & Chater, N. (2008). Language as shaped by the brain. *Behavioral & Brain Sciences*, 31, 489–509. doi:10.1017/S0140525X08004998
- Coopmans, P. (1984). Surface word-order typology and universal grammar. *Language*, 60, 55–69.

- Croft, W., Bhattacharya, T., Kleinschmidt, D., Smith, D. E., & Jaeger, T. F. (2011). Greenbergian universals, diachrony, and statistical analyses. *Linguistic Typology*, *15*, 433–453. doi:10.1515/lity.2011.029
- Culbertson, J., & Adger, D. (2014). Language learners privilege structured meaning over surface frequency. *Proceedings of the National Academy of Sciences, USA*, *111*, 5842–5847.
- Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*, 306–329. doi:10.1016/j.cognition.2011.10.017
- Dryer, M. S. (2011). Evidence for word order correlation. *Linguistic Typology*, *15*, 335–380.
- Dryer, M. S. (2013). Relationship between the order of object and verb and the order of adposition and noun phrase. In M. S. Dryer & M. Haspelmath (Eds.), *The world atlas of language structures online*. Retrieved from <http://wals.info/chapter/95>
- Dunn, M., Greenhill, S. J., Levinson, S. C., & Gray, R. D. (2011). Evolved structure of language shows lineage-specific trends in word-order universals. *Nature*, *473*, 79–82. doi:10.1038/nature09923
- Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences, USA*, *109*, 17897–17902. doi:10.1073/pnas.1215776109
- Fedzechkina, M., Jaeger, T. F., & Trueswell, J. (2015). Production is biased to provide informative cues early: Evidence from miniature artificial languages. In D. Noelle, A. Waelaumont, J. Yoshimi, T. Matlock, C. Jennings, & P. Maglio (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 674–679). Austin, TX: Cognitive Science Society.
- Fedzechkina, M., Newport, E. L., & Jaeger, T. F. (2017). Balancing effort and information transmission during language acquisition: Evidence from word order and case marking. *Cognitive Science*, *41*, 416–446. doi:10.1111/cogs.12346
- Ferrer i Cancho, R. (2004). Euclidean distance between syntactically linked words. *Physical Review E*, *70*, 056135.
- Fodor, J. D. (2001). Setting syntactic parameters. In M. Baltin & C. Collins (Eds.), *The handbook of contemporary syntactic theory* (pp. 730–767). Oxford, England: Blackwell. doi:10.1002/9780470756416.ch23.
- Futrell, R., Mahowald, K., & Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences, USA*, *112*, 10336–10341.
- Garnsey, S. M., Perlmutter, N. J., Meyers, 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.
- Gildea, D., & Temperley, D. (2010). Do grammars minimize dependency length? *Cognitive Science*, *34*, 286–310.
- Givón, T. (1991). Markedness in grammar: Distributional, communicative and cognitive correlates of syntactic structure. *Studies in Language*, *15*, 335–370. doi:10.1075/sl.15.2.05giv
- Goldberg, A. E. (2013). Substantive learning bias or an effect of similarity? Comment on Culbertson, Smolensky, & Legendre (2012). *Cognition*, *127*, 420–426.
- Greenberg, J. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (Ed.), *Universals of human language* (pp. 73–113). Cambridge, MA: MIT Press.
- Gregory, M., & Michaelis, L. (2001). Topicalization and left-dislocation: A functional opposition revisited. *Journal of Pragmatics*, *33*, 1665–1706.
- Grodner, D., & Gibson, E. (2005). Consequences of the serial nature of linguistic input. *Cognitive Science*, *29*, 261–290.
- Hawkins, J. A. (2014). *Cross-linguistic variation and efficiency*. Oxford, England: Oxford University Press.
- Hudson Kam, C., & Newport, E. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, *59*, 30–66.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, *141*, 87–102.
- Pajak, B., Fine, A. B., Kleinschmidt, D. F., & Jaeger, T. F. (2016). Learning additional languages as hierarchical inference: Insights from L1 processing. *Language Learning*, *66*, 900–944.
- Pajak, B., & Levy, R. (2014). The role of abstraction in non-native speech perception. *Journal of Phonetics*, *46*, 147–160.
- Prince, E. F. (1995). On the limits of syntax, with reference to left-dislocation and topicalization. In P. W. Culicover & L. McNally (Eds.), *The limits of syntax* (pp. 281–302). San Diego, CA: Academic Press.
- Ros, I., Santesteban, M., Fukumura, K., & Laka, I. (2015). Aiming at shorter dependencies: The role of agreement morphology. *Language, Cognition, and Neuroscience*, *30*, 1156–1174.
- Saffran, J., Aslin, R., & Newport, E. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.
- Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, *116*, 444–449. doi:10.1016/j.cognition.2010.06.004
- Snider, N., & Zaenen, A. (2006). Animacy and syntactic structure: Fronted NPs in English. In M. Butt, M. Dalrymple, & T. H. King (Eds.), *Intelligent linguistic architectures: Variations on themes by Ronald M. Kaplan* (pp. 323–338). Stanford, CA: CSLI.
- Travis, L. (1984). *Parameters and effects of word order variation* (Doctoral dissertation, Massachusetts Institute of Technology). Retrieved from <http://www.ai.mit.edu/projects/dm/theses/travis84.pdf>
- Wasow, T. (2002). *Postverbal behavior*. Stanford, CA: CSLI.
- Yamashita, H., & Chang, F. (2001). “Long before short” preference in the production of a head-final language. *Cognition*, *81*, B45–B55.