

Learners' bias to balance production effort against message uncertainty is independent of their native language

Lucy Hall Hartley (lucyhallhartley@email.arizona.edu)

Department of Linguistics, University of Arizona,
Tucson, AZ 85719

Masha Fedzechkina (mfedzech@email.arizona.edu)

Department of Linguistics, Graduate Interdisciplinary Program in Cognitive Science,
Graduate Interdisciplinary Program in Second Language Acquisition and Teaching, University of Arizona
Tucson, AZ 85719

Abstract

Miniature language learning is gaining increasing popularity to study biases underlying language universals. However, it is unclear whether learning preferences in these studies are influenced by learners' native language. We ask whether a previously identified bias to balance production effort against message uncertainty holds across speakers of structurally different languages. We expose English (fixed order language without case) and German (flexible order language with case) speakers to miniature languages with optional case and either fixed or flexible constituent order and study their deviations from the input. We find that English and German speakers restructure the input in the same way: They match the input constituent order proportions and use more case in the flexible order language than in the fixed order language, thus following the bias to balance production effort against message uncertainty. Our findings suggest that this bias and its specific realization are independent of learners' native language.

Keywords: Language universals; miniature language learning; L1 influences; efficient information transmission; language evolution

Introduction

Natural languages exhibit a wide amount of variation, but this variation is constrained. Unrelated languages share more features or combinations of features than expected by chance, referred to as (statistical) language universals (Greenberg, 1963). Most theories agree that the existence of language universals might be attributable to the biases in language acquisition and use (Chomsky, 1965; Hawkins, 2014). Thus, identifying the similarities between unrelated languages and understanding their causes has been a central goal of linguistics and cognitive science as this can provide insight into how language is processed and represented in the human brain.

Traditionally, language universals have been studied via historical and synchronic crosslinguistic surveys. These approaches have been instrumental in identifying many language universals (Greenberg, 1963; Dryer, 1992) but these methods also have their drawbacks. First, they often suffer from the sparsity of independent data points as many languages are genetically (common ancestor) or geographically (language contact) related to each other. Second, and more importantly, even the most robust

crosslinguistic surveys cannot directly answer questions about the underlying causes of observed language universals.

The recent decade has seen a rise in the use of a new paradigm—*miniature artificial language learning*—to study questions about language universals. In this paradigm, researchers design novel artificial languages which isolate factors of interest to test specific hypotheses about language universals. This paradigm does not suffer from data sparsity as new languages can be readily generated in the laboratory and, more importantly, can probe questions about the origins of language universals more directly than traditional approaches. To accomplish this, researchers design miniature languages with variable input, similar to a pidgin or a situation of an ongoing language change, and study whether and how learners introduce changes into the inconsistent input they are exposed to (Fedzechkina, Jaeger, & Newport, 2012; Hudson Kam & Newport, 2009; Smith & Wonnacott, 2010). A growing number of studies within this paradigm have found parallels between patterns in crosslinguistic diversity and biases operating during the acquisition of miniature languages: The changes introduced by learners into the variable miniature language input tend to bring these languages closer to crosslinguistically attested types (Culbertson, Smolensky, & Legendre, 2012; Fedzechkina, Newport, & Jaeger, 2017; Kirby, Cornish, & Smith, 2008).

However, like any paradigm, miniature language learning has its weaknesses. The main premise of this paradigm is that if the input does not bias learners in a particular direction, restructuring of the miniature languages is due to more general biases that learners might have about natural language structures, which can inform us about the causes of language universals. This requires an assumption that any observed deviations from the input cannot be attributed to the learners' native (L1) language (Goldberg, 2013).

Most miniature language learning studies investigating language universals are conducted with adult participants who are proficient in their L1 (in the overwhelming majority of cases, English). Unfortunately, L1 influences within the paradigm are not well-understood. The acquisition of miniature languages can be considered a special case of second language acquisition and while there is evidence that patterns learned from one's L1 influence learning performance in a new natural language (Flege, 1999), it remains to be seen how strongly and under which

circumstances learners' L1 influences their performance in these novel artificial languages. As the interest of applying miniature language learning to study biases underlying language universals is growing, it is pressing that we gain better insight into L1 influences in this paradigm and how they interact with more abstract universal biases.

Consider one such abstract bias – a bias to balance message uncertainty against production effort--previously observed by Fedzechkina et al., 2017. In their study, the researchers presented adult English speakers with miniature languages that had optional case marking and either fixed or flexible constituent order. They found that learners in the experiment restructured the input distributions of case marking in a manner consistent with the crosslinguistic trade-off between constituent order flexibility and the presence of case in a language (Blake, 2001): Learners maintained case in the flexible order language but dropped case in the fixed order language. Interestingly, learners matched the constituent order distributions in the input. Fedzechkina et al. (2017) interpreted these findings as evidence for learners' bias to efficiently balance message uncertainty against production effort as learners dropped case marking when it did not contribute to uncertainty reduction but required effort to produce (i.e., in the fixed order language) but retained it despite increased production effort when it reduced message uncertainty (i.e., in the flexible order language).

While the findings of Fedzechkina et al. (2017) are consistent with a bias to balance message uncertainty against production effort, this study was conducted on English speakers only, thus leaving open a possibility that some aspects of learners' performance could be attributed to their experience specific to English and not to a more general bias. For instance, it is possible that learners drop case marking in the fixed order language not because it is uninformative and requires effort to produce but rather in an attempt to bring the language closer to English (a fixed order language with no case). Or perhaps because constituent order is a strong cue to grammatical function assignment in their L1, English speakers pay more attention to it than the unfamiliar cue of case, thus matching the input distributions of constituent order while failing to perfectly learn (i.e., match) the case input distribution.

Here we ask whether the bias to balance message uncertainty against production effort found in Fedzechkina et al. (2017) holds for speakers of structurally different languages. Specifically, we compare learning preferences in native speakers of English and German. These L1's were chosen for our study as they differ in the relative strength they attribute to constituent order and case as cues to grammatical function assignment. English strongly relies on constituent order – it has fixed order and no case. German, on the other hand, has flexible order and obligatory case marking. If the bias to balance message uncertainty against production effort is abstract and independent of learners' L1 structure, we expect to replicate the findings of Fedzechkina et al. (2017) in both English and German speakers. However, if learners' preferences observed in Fedzechkina et al. (2017) are driven

by their experience with English, we expect German speakers to show different preferences. For example, German speakers might not follow a bias to balance message uncertainty against production effort at all, which would suggest that this bias is not an abstract universal bias. Alternatively, German speakers might follow this bias but realize it in their productions in a different way (e.g., by fixing constituent order in the flexible order language or by consistently using case), which would suggest that the observed bias is potentially universal but its realization is dependent on learners' L1.

Methods

Participants

72 participants completed the experiment via Prolific Academic, a crowdsourcing platform. Based on the built-in prescreening system, participants were native speakers of either English or German with no language-related disorders and had a 95% approval rate based on 10 or more submissions on Prolific Academic. The experiment was administered via FindingFive, an online study administration platform for behavioral research (FindingFive Team, 2019). Participants received \$12.50 for completing two sessions, each lasting on average 40-45min.

Miniature input languages

Participants were instructed that they would be learning a miniature language by watching a series of videos accompanied by their descriptions. The lexicon of both languages contained six nouns that depicted animate actors ('chef', 'referee', 'mountie', 'hunter', 'conductor', and 'bandit'), two verbs that depicted transitive actions ('hug' and 'kick'), and a case suffix '-dak' that marked the object of the action. All words were phonotactically legal non-words of English and German. They were individually synthesized using the Greek-accented Melina voice on the Mac speech synthesizer. This voice was chosen to keep the accent in the miniature language neutral for both English and German speaking participants. Words were concatenated into sentences using Praat with 35 ms of silence between words. This procedure ensured the miniature languages contained no prosodic cues to grammatical function assignment. Each scene in the language was accompanied with an auditory and written description. The decision to present written descriptions was motivated by the need to familiarize learners with novel orthography as they produced the language by typing.

Participants were randomly assigned to one of the two languages: fixed order or flexible order. Both languages consisted of simple transitive sentences that described short videos created using Poser Pro software. The languages differed in the amount of constituent order flexibility they allowed. The fixed order language allowed one constituent order—all utterances were 100% VSO. The flexible order language allowed VSO and VOS orders in free variation (50/50%). These constituent orders were chosen as they were

not present in either English or German (the native languages of our participants). Both languages had optional case marking: The marker ‘-dak’ was present on the object in 67% of sentences and absent in 33% of sentences. The case marker, thus, did not provide much information about sentence meaning above and beyond that already expressed by constituent order in the fixed order language, but it was informative in the flexible order language (when present).

Each noun occurred equally frequently as the subject and object of sentences. Each verb occurred equally frequently in each constituent order allowed in each language. To ensure there were no unwanted associations between a word form and constituent order or case marking, the assignment of word forms to meanings was rotated across two lists.

Procedure

The procedure was based on Fedzechkina & Jaeger (2020) with the following changes. Participants in the current study produced utterances by typing rather than choosing the lexical items from the provided lexicon as in Fedzechkina & Jaeger (2020). The lexicon in the current study was slightly smaller to compensate for the increased difficulty of the production test (as participants were not provided the lexicon during production in our set-up). The study was administered over two sessions spread over two days with at least 12 hours in between. The procedure was identical for both sessions. Participants were instructed that they would be learning an alien language with the help of the alien informant named Lumi. The experiment was broken up into 12 blocks (described below) focusing on different aspects of lexical and grammar learning.

Noun Exposure The experiment started with noun exposure: Participants were shown pictures of each of the six characters paired with the corresponding name in the novel language (12 trials, two for each noun). Noun comprehension followed noun exposure.

Noun Comprehension In this block, participants were presented with four pictures of characters and heard a name of a character in the novel language. They were instructed to click on the correct character. Participants received feedback indicating whether or not they had chosen the correct character on each trial (12 trials, two for each noun). Noun comprehension was followed by noun production.

Noun Production Participants were presented with a picture of a character and asked to type its name into a textbox below the picture. They then received feedback on the correct name for the character (six trials, one for each noun).

A series of one noun exposure, one noun comprehension, and one noun production block was repeated twice before moving to sentence exposure.

Sentence Exposure Participants were shown videos accompanied by simple transitive sentences that described the scene. Participants could replay the video and sentence as many times as they liked (24 trials). After completing two sentence exposure blocks, participants moved to sentence comprehension.

Sentence Comprehension Participants were provided with a sentence in the novel language and asked to match it to one of the two videos shown on the screen. The videos differed in which character was the subject and which was the object of the same action. No feedback was provided (24 trials).

Participants were presented with two sets of two sentence exposure and one sentence comprehension block before moving to sentence production (our critical test).

Sentence Production Participants were shown previously unseen videos and asked to type a sentence in the miniature language that described what was happening in the video. Participants were provided with a verb prompt to facilitate production. No feedback was provided (24 trials).

Results

Before we turn to the analysis of our main prediction, we describe how our data was scored, list exclusions, and report learning accuracy.

Scoring

Sentence comprehension accuracy was calculated over unambiguous (i.e., case-marked) trials. To assess lexical accuracy in noun and sentence production trials, we employed soft-string matching with a custom Python script using the SequenceMatcher algorithm to calculate Levenshtein distance (number of insertions, deletions, and substitutions) from the target word. Any word where the Levenshtein distance between the word produced by a learner and the target word was larger than two was coded as a lexical error. For example, the production *flc* for *flaki* was coded as a lexical error as it involved two deletions and a substitution but *floki* was coded as lexically correct since it only contained one substitution. For a small number of sentences, constituent order could not be determined because both nouns contained lexical errors. These productions (2% of the data across both days) were coded as containing both a lexical error and a grammatical mistake. We also coded constituent order used in production, the presence of case marking, and the presence of grammatical mistakes (using a constituent order not present in the input or marking case on a word other than the object). Productions that contained grammatical mistakes were removed from all analyses; productions containing only one lexical mistake were included in the analysis.

Exclusions

Following prior work (Fedzechkina et al., 2017; Fedzechkina and Jaeger, 2020), seven participants were excluded from the analysis for failure to achieve 70% accuracy for unambiguous (case-marked) trials on the final sentence comprehension block on Day 2. All participants were learners of the flexible order language (six English speakers and one German speaker). One participant was excluded from the analysis for failing to follow the instructions (German speaker in the flexible order language). This left 64 participants for analysis (19 English speakers in the fixed order language, 17 English speakers in the flexible order language, 14 German speakers

in the fixed order language, and 14 German speakers in the flexible order language).

Learning Accuracy

Of the remaining 64 participants, learning accuracy was high. The accuracy for unambiguous trials on the final sentence comprehension block was above 95% for all input languages and all L1's. The number of sentence production trials that contained lexical errors on Day 2 was low (lowest lexical error rates [2%] were observed for English speakers in the flexible order language, highest lexical error rates [7%] were found for German speakers in the flexible order language). Interestingly, the rate of case errors on Day 2 was comparable across L1's (0% for English and 0.6% for German speakers in the fixed order language; 1% for both English and German speakers in the flexible order language), potentially suggesting that the presence of case marking in an L1 might not have a strong influence on the accuracy of its acquisition in a miniature artificial language. The high degree of accuracy in sentence comprehension and low error rates in sentence production suggest that the grammars of the miniature languages were acquired by both English and German speaking participants and that any differences between the groups in case and constituent order use are not due to poor learning.

Constituent Order Use in Production

As the first step in our comparison of learning preferences across English and German, we analyzed learners' constituent order use in production. Given that English and German differ in the amount of constituent order flexibility they allow (English has relatively fixed order and German has relatively flexible order), we asked whether speakers of English and German will have different preferences in constituent order use in the miniature language. We tested this by comparing constituent order use across learners' L1's and by comparing constituent order use to the input in each L1 group. We limited this analysis to the flexible order language as given our scoring (see above), all participants trivially matched the input VSO proportion in the fixed order language.

To investigate whether English and German speakers had different constituent order preferences in the flexible order language, we fit a generalized linear mixed effects model that predicted VSO use from L1 (sum coded, 1=German, -1=English), day of training (sum coded, 1=Day 2, -1=Day 1), and their interactions. The model contained the fullest converging random effects structure (by-participant random intercept). There was no main effect of L1 ($\hat{\beta} = -0.01, z = -0.07, p = 0.93$): English and German speakers used the same amount of VSO in their productions across both days of training. There was a main effect of day of training ($\hat{\beta} = -0.13, z = -2.32, p < 0.05$): Overall, participants used significantly less VSO on Day 2 than on Day 1. Day of training interacted with L1 ($\hat{\beta} = -0.13, z = -2.36, p < 0.05$). A simple effects test revealed that German speakers used significantly less VSO on Day 2 compared to Day 1 ($\hat{\beta} = -0.27, z = -3.23, p < 0.01$)

but VSO use did not differ across days for English speakers ($\hat{\beta} = 0.002, z = 0.02, p = 0.98$, see Figure 1).

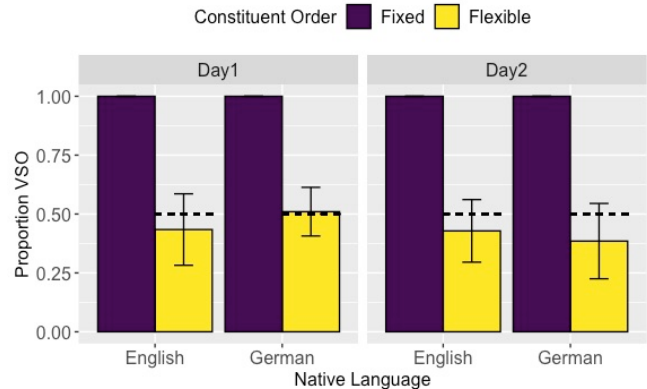


Figure 1: VSO use in production by day of training and L1 background. The dashed line represents the input proportion for the flexible order language (VSO input for the fixed order language is 1.0). Error bars are bootstrapped 95% confidence intervals.

We next asked whether English and German learners of the flexible order language had different preferences in VSO use compared to the 0.5 input proportion on Day 2 of training. We used a generalized linear mixed effects model that predicted VSO use from L1 (treatment coded) and the fullest converging random effects structure (by-participant random intercept). An offset term corresponding to the input language proportion of VSO (0.5, 0 log-odds) was added to the model. The intercept of this model captures whether the L1 coded as the reference level differs from the input proportion. We ran this model twice, first with English coded as the reference level and then with German coded as the reference level. The model revealed that VSO use did not significantly differ from the input proportion for either English speakers ($\hat{\beta} = -0.37, z = -0.95, p = 0.34$) or for German speakers ($\hat{\beta} = -0.82, z = -1.85, p = 0.06$).

Learners in our study matched the input proportion of constituent order, thus replicating and extending the work by Fedzechkina et al. (2017). Interestingly, this input-matching behavior does not depend on the amount of constituent order flexibility allowed in the learners' L1. German speakers, whose L1 has substantial constituent order flexibility, show the same preference in constituent order use as English speakers, whose L1 has little constituent order flexibility. This suggests that learners' preferences in matching input constituent order distributions are independent of the amount of constituent order flexibility in their L1. We now turn to the analysis of case use in learners' production.

Case Marker Use in Production

If a bias to efficiently balance uncertainty about the intended message and production effort observed in prior work (Fedzechkina et al., 2017) is independent of learners' L1 structure, we expect to find the same bias in both English and German speakers. Specifically, we expect that participants

exposed to the fixed order language would drop case marking as it is a redundant cue that requires effort to produce, while participants exposed to the flexible order language would retain case marking because it reduces uncertainty about grammatical function assignment.

To address this question, we employed a generalized linear mixed effects model to predict case use from L1 (sum coded, 1=German, -1=English), constituent order flexibility (sum coded, 1=flexible order, -1=fixed order), and day of training (sum coded, 1=Day 2, -1=Day 1) with all interactions. The model contained the fullest converging random effects structure (by-participant and by-object-noun random intercepts).




Figure 2: Case marker use in production by day of training and L1 backgrounds. The dashed line represents the input proportion (same across languages). Error bars are bootstrapped 95% confidence intervals.

There was a main effect of day of training on case use: Participants used significantly more case on Day 2 than on Day 1 ($\hat{\beta} = 0.33, z = 6.92, p < 0.001$), suggesting that case use increased as learners became more proficient in the novel language, consistent with prior work (Fedzechkina et al., 2017). L1 interacted with day of training ($\hat{\beta} = -0.14, z = -2.84, p < 0.01$), suggesting that German speakers showed less increase in case use from Day 1 to Day 2 compared to English speakers. Case use increased more from Day 1 to Day 2 in the flexible order language as suggested by a significant constituent order by day of training interaction ($\hat{\beta} = 0.17, z = 3.53, p < 0.001$). There was a significant three-way interaction between and constituent order flexibility, L1, and day of training ($\hat{\beta} = -0.10, z = -2.234, p < 0.05$), suggesting that for English speakers, there was a greater difference in case use between the fixed and flexible order conditions on Day 2 compared to Day 1 than there was for German speakers. There was a main effect of constituent order flexibility on case use ($\hat{\beta} = 0.94, z = 4.13, p < 0.001$, see Figure 2): Learners of the flexible order language used significantly more case than learners of the fixed order language across both days of training. Critically, there was no significant main effect of L1 on case use ($\hat{\beta} = 0.13, z = 0.61, p = 0.54$): English and German speakers used the same amount of case in production. L1 did not interact with

constituent order flexibility ($\hat{\beta} = 0.07, z = 0.33, p = 0.73$), suggesting that English and German speakers used case in a similar manner across the input languages.

Participants in our study used more case in the flexible order language compared to the fixed order language, thus replicating the findings of Fedzechkina et al. (2017). Learners' preferences in case use did not differ depending on their L1: German speakers, whose L1 has a rich case system, used case marking in the miniature languages in the same way as English speakers whose L1 does not have case at all. These findings suggest that learners' preferences in case use were driven by a bias to balance production effort and message uncertainty and that this bias was realized in the same way by participants of different L1 backgrounds.

Discussion

We asked whether a bias to balance message uncertainty against production effort is expressed by learners from structurally different L1 backgrounds in the same way. To test this, we compared the behavior of English and German speakers exposed to miniature languages with optional case marking and either fixed or flexible constituent order. English and German employ case and constituent order as cues to grammatical function assignment in different ways: English has relatively fixed constituent order and no case marking; German has relatively flexible constituent order and a rich case system.

Despite these L1 differences, we found that both English and German speakers follow the bias to efficiently balance message uncertainty against production effort. Interestingly, speakers of both languages employ the same strategies in case and constituent use to restructure the miniature language to better express this bias. Specifically, both English and German speakers matched the input proportion of constituent order. They also dropped case marking when it was a redundant cue that took effort to produce (i.e., in the fixed order language) and retained case in production when it was informative about grammatical function assignment (i.e., in the flexible order language), mirroring the crosslinguistically common pattern in case and constituent flexibility trade-off. Thus, learning behavior in our study replicates and extends the findings of Fedzechkina et al. (2017) suggesting that the bias to balance message uncertainty and production effort observed in their work is not specific to English but is rather a general bias that spans across structurally different L1's.

In addition, our findings suggest that the strategies learners adopt to realize this bias are not modulated by their L1's. L1 influences on learning outcomes are an issue in miniature language experiments (Culbertson & Adger, 2014; Fedzechkina, Chu, & Jaeger, 2018) and it is somewhat unexpected that we did not observe greater L1 influences on learning performance in our study. Why might this be the case? One possibility is that L1 influences did not arise because the structures of our miniature languages were different enough from both English and German (the native languages of our participants). The flexible order language in the current study was similar to German at an abstract level –

it had flexible constituent order and case; while our fixed order language was similar to English in that it had no constituent order flexibility. But this is where the similarities ended. The words in the miniature languages were intentionally synthesized with a Greek accent to be distinct from English and German accents. The constituent order variants used in our grammar (VSO and VOS) are not present in German or English. Furthermore, the miniature language in our experiment realized case with a suffix, while in German, case is realized on the article in combination with a suffix or on the article alone. Thus, it is possible that the similarities that were present between our miniature language and learners' L1's were too abstract to give rise to L1-driven influences. This raises an intriguing question that we leave to future work—how high does L1-to-miniature-language similarity need to be to give rise to L1-driven influences? Perhaps if German speakers were exposed to miniature languages with article case marking or with more familiar constituent orders L1 influences would play a bigger role in their learning preferences.

To our knowledge, the current study is the first one to test L1 influences on miniature language performance by asking whether a bias that's assumed to be L1-independent (a bias to efficiently balance message uncertainty against production effort) is realized differently by speakers of structurally different L1's. We find no evidence for L1 influences on this bias: Speakers of English and German realized this bias in the same way. Our study shows that by collecting crosslinguistic data from speakers of structurally different languages we can begin addressing questions about the precise circumstances of L1 influences and their interactions with more abstract universal biases in miniature language learning.

References

- Blake, B. J. (2001). *Case*. Cambridge: Cambridge University Press.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Culbertson, J., & Adger, D. (2014). Language learners prefer structure to frequency. *Proceedings of the National Academy of Sciences*, *111*(16), 5842-5847.
- Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*(3), 306-329.
- Dryer, M. S. (1992). The Greenbergian word order correlations. *Language*, *68*, 81-138.
- Fedzechkina, M., Chu, B., & Jaeger, T. F. (2018). Human information processing shapes language change. *Psychological Science*, *29*(1), 72-82.
- Fedzechkina, M., & Jaeger, T. F. (2020). Production efficiency can cause grammatical change: Learners deviate from the input to better balance efficiency against robust message transmission. *Cognition*, *196*, 104115.
- 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.
- 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*(2), 416-446.
- FindingFive Team (2019). FindingFive: A web platform for creating, running, and managing your studies in one place. FindingFive Corporation (nonprofit), NJ, USA. <https://www.findingfive.com>
- Flege, J. E. (1999). Age of learning and second- language speech. In D. P. Birdsong (Ed.), *Second language acquisition and the critical period hypothesis*. Hillsdale, NJ: Erlbaum.
- Goldberg, A. (2013). Substantive Learning Bias or an Effect of Similarity? Comment on Culbertson, Smolensky, Legendre (2012). *Cognition*, *127*, 420-26.
- Greenberg, J. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (Ed.), *Universals of human language*. Cambridge, MA: MIT Press.
- Han, Z.- H. (2004). *Fossilization in adult second language acquisition*. Clevedon, UK: Multilingual Matters.
- Hawkins, J. A. (2014). *Cross-linguistic variation and efficiency*. Oxford, England: Oxford University Press.
- Hudson Kam, C., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, *59*, 30-66.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, *105*(31), 10681-10686.
- Smith, K., Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, *116*, 444-449.