

Do children privilege phonological cues in noun class learning?

Jennifer Culbertson (jennifer.culbertson@ed.ac.uk)

Hanna Jarvinen (hanna.jarvinen@ed.ac.uk)

Frances Haggarty (frances.i.haggarty@gmail.com)

Kenny Smith (kenny.smith@ed.ac.uk)

Centre for Language Evolution, University of Edinburgh, UK

Abstract

Previous research on acquisition of noun class systems, such as grammatical gender, has shown that child learners rely disproportionately on phonological cues to class, even when competing semantic cues are more reliable. Culbertson, Gagliardi, and Smith (2017) use artificial language learning experiments with adults to argue that over-reliance on phonology may be due to the fact that phonological cues are available first; learners base early representations on surface phonological dependencies, only later integrating semantic cues from noun meanings. Here, we show that child learners (6-7 year-olds) show this same sensitivity to early availability. However, we also find intriguing evidence of developmental changes in sensitivity to semantics; when both cues are simultaneously available children are more likely to rely on a phonology cue than adults. Our results suggest that early availability and a bias in favor of phonological cues may both contribute to children's over-reliance on phonology in natural language acquisition.

Keywords: noun class; gender; language acquisition; artificial language learning; category learning

Introduction

Noun class systems, like grammatical gender, are widespread across the world's languages, and have been extensively studied from the perspective of both linguistic typology, and language acquisition. Three main types of noun classification system are often distinguished in the literature: grammatical gender, noun class systems, and numeral classifier systems. In gender and noun class systems, class is visible through agreement either with other elements in the noun phrase or with elements external to the noun phrase. For example, French has a grammatical gender system in which nouns fall into two classes, masculine and feminine; the form of the articles (e.g., *le* or *la*) varies based on noun gender; Tsez (a Nakh-Dagestanian language) has four noun classes, and exhibits verbal agreement based on noun class as well as agreement within the noun phrase. We refer to these two types of systems collectively as noun class systems. Numeral classifier systems typically feature a much larger set of classes, and do not trigger agreement; for example, there are dozens of classifiers for different types of nouns in Cantonese, and a definite noun phrase like 'one scarf' is expressed as 'one [classifier] scarf'.

Previous research on the acquisition of noun class systems in natural language has shown that child learners do not treat all potential cues to class membership equally. In particular, while most noun class systems feature a combination of noun-internal cues—i.e., phonological features of the nouns

themselves—and external semantic cues from word meanings, children appear to rely disproportionately on phonological cues to class (Karmiloff-Smith, 1981; Mills, 1985; Mariscal, 2009; Perez-Pereira, 1991; Demuth, 2000; Demuth & Ellis, 2008; Rodina & Westergaard, 2012; Gagliardi & Lidz, 2014). Following a classic study by Karmiloff-Smith (1981), for example, Gagliardi and Lidz (2014) investigated acquisition of the noun class system in Tsez by teaching children novel nouns labeling novel objects. These novel items featured phonological and semantic cues that are relevant for determining the class of real nouns in Tsez. By design, in some cases the cues were conflicting: the noun ending indicated one class, but the word meaning indicated another. They found that when Tsez-speaking children were faced with conflicting semantic and phonological cues in a novel noun, they tended to use the phonological information to determine class agreement. Importantly, this was the case even though the relevant semantic cues were shown to be *more reliable predictors of class* in Tsez. By contrast, Tsez-speaking adults did not show this effect, but rather tended to use the more reliable semantic cues.

A number of possible explanations for this surprising cross-linguistic tendency have been proposed, including an active bias against using external cues like semantics, when noun-internal phonological cues are available (Gagliardi, 2012; Culbertson & Wilson, 2013; Gagliardi & Lidz, 2014; Gagliardi, Feldman, & Lidz, 2017). In Culbertson et al. (2017), we presented evidence from artificial language learning experiments with adults suggesting that the over-reliance on phonology may be due instead to the fact that phonological cues are generally available earlier than semantic cues (Carroll, 1999; Polinsky & Jackson, 1999; Demuth, 2000; Culbertson & Wilson, 2013; Gagliardi et al., 2017); learners acquire early representations of phonological dependencies (e.g., between a gendered determiner and a noun) before acquiring the semantic referents of nouns. Because the system is initially built on the basis of these phonological cues, semantic cues acquired later take time to be integrated into the system. In our experiments, adults were exposed to a phonological cue and a semantic cue which were both equally reliable; during training the two cues were by design confounded (e.g., both animacy and a particular noun ending determined class membership). At test, learners had to choose the class of new labelled objects for which the cues conflicted (e.g.,

animacy suggested one class, but noun ending suggested the other). This essentially mirrors the natural language experiments with Tsez and other languages. When both cues were available, adults showed no a priori preference for attending to phonology over semantics, rather their reliance on a given cue depended on cue salience. For example, among the semantic cues tested animacy was highly salient but flexibility was not; among the phonological cues tested both a suffix and a prefix together were highly salient, but a CV ending alone was not. By contrast, when cues were *staged*—so that either the phonological or the semantic cue was available first during learning—adult learners prioritized the earlier learned cue.

We argued that these findings suggest that children’s over-reliance on phonology may not be due to an active bias against semantics. Rather, children may start building their classifications systems very early, when phonological information is available, but word meanings are not. However for this explanation to hold, children’s learning behavior must parallel that of adults. Here, we report the results of a set of experiments investigating the effect of cue type and early availability in child learners. Across a series of experiments, we show that children are indeed sensitive to the early availability of a given cue type: they are more likely to use a cue they learned earlier when determining the class of a novel noun. However, we also find evidence that children appear to prioritize the phonological cue, relative to adults. This suggests the possibility that children are in fact biased to attend to noun-internal rather than external cues when acquiring noun class systems.

Experiment 1

Before we test how children treat phonological and semantic cues when both are available, we first need to identify a set of equally-salient cues they can learn in the context of an artificial noun class system. In Experiment 1, we therefore taught children a miniature artificial noun class system in which class was deterministically cued by either semantic information about noun referents, *or* phonological information from the form of the noun itself. Evidence for the two classes in the language came from a plural marker which differed by class. The semantic cue we used was animacy-based, following our own previous work showing that animacy is a highly salient cue for adult learners in the context of artificial noun class learning (Culbertson et al., 2017). The phonological cue we used was the vowel sounds of the nouns, which were reduplicated in the stems. Since prior work has shown that reduplication aids word learning in children (Ota & Skarabela, 2016, 2017), both these cues are likely to be highly salient to children.

Participants were 6-7 year old children (N=39). We chose this age (rather than younger) because acquiring these patterns in the lab is difficult for children, and previous studies suggest they can successfully pick up simple category systems in the lab at this age. The lexicon was comprised of 12

C1VC2V words, 6 with the vowel /i/, and 6 with /a/, used to label 12 simple objects, either aliens or planets. In the phonology only condition (N=20), the cue to class came from the phonological form of the noun alone (i.e. nouns from one class were all of the form CiCi, nouns from the other class were of the form CaCa, and aliens and planets were equally frequent as referents in both classes). In the semantics only condition (N=19), the cue to class came from the semantic category of the noun (nouns from one class all referred to aliens, nouns from the other class all referred to planets, and CiCi and CaCa labels were equally frequent in both classes). Two plural markers were drawn randomly for each participant from a set of 8 CV forms (e.g., /kʌ/, /ʃæ/). During exposure (96 trials), participants were shown a series of singular and plural trials and asked to repeat what they heard (either a word alone or a word + marker). Then they were tested using a ‘wug’ test (Berko, 1958; Schuler, Yang, & Newport, 2016); participants were shown a singular picture, heard the noun stem, and were asked to produce the plural form in the language (48 trials, featuring 8 trained nouns and 4 untrained nouns, all 12 tested 4 times each).




Figure 1: Accuracy in plural form production for old (trained) and new (untrained) words in each condition in Experiment 1. Colored circles represent individual participant means, bars give mean of those by-participant means, error bars represent 95% confidence intervals.

Figure 1 shows average accuracy in plural marker production across conditions for old (trained) and new (untrained) words. Data were analyzed using logistic mixed-effects regression models.¹ Overall, children performed above chance in this task for both trial types (as indicated by a significant model intercept; old: $\beta = 0.85 \pm 0.15, p < 0.001$; new: $\beta = 0.79 \pm 0.19, p < 0.001$). To compare across conditions, we ran a model predicting responses by trial type (old vs.

¹Unless otherwise noted, all models reported here were run in R (R Development Core Team, 2010) using the lme4 package (Bates, 2010), with by-participant random intercepts and by-participants slopes for trial type. All fixed-effects were sum coded.

new) and condition (Phonology Only vs. Semantics Only). This model revealed no significant effect of trial type ($\beta = -0.01 \pm 0.07, p = 0.89$), condition ($\beta = 0.12 \pm 0.16, p = 0.47$), or their interaction ($\beta = 0.06 \pm 0.06, p = 0.37$). In other words, no difference in behavior was found across conditions, suggesting that children are similarly sensitive to both of these cues in isolation.

Experiment 2

Having established these two cues as learnable and equally salient for children in the context of this task, in Experiment 2, we follow Culbertson et al. (2017) in asking whether a preference for one cue over the other will emerge when the cues are in conflict. We do this by deliberately confounding the cues during training, such that both animacy and the reduplicative vowel cue predict noun class equally well. We then test learners on new items where the two cues conflict (i.e. the semantic cue points to one marker, and the phonological cue the other). By design, the system learners are trained on is ambiguous as to which cue the learner should use when they are in conflict. If the learner has a preference for the phonological cue, then they should use this cue to determine the class of a new item, regardless of its animacy. By contrast, if the learner has a preference for the semantic cue, they should use animacy to determine the class of a new noun, and disregard the form of the word itself. Given that we are specifically interested in uncovering whether children show a bias in favor of phonology compared to adults, in Experiment 2 we will also test adult learners under these same conditions. If both adult and child learners are equally likely to go with the semantic or phonological cue, this would suggest that no meaningful bias exists. This is what we might expect for two equal salience cues based on the results reported in Culbertson et al. (2017). By contrast, if children are significantly more likely to use the phonological cue compared to adults in the same task, this will point to the possibility of a bias on the part of child learners.

Participants were 6-7 year old children (N=20), and adults (university undergraduates, N=20). The exposure phase (96 trials) was identical to Experiment 1, but objects seen during exposure have both the phonological and semantic cues which point to one of the two markers deterministically (e.g. all nouns from one class were of the form CiCi and had alien referents, all nouns from the other class were of the form CaCa and had planet referents). This was followed by a practice production phase (24 trials) with feedback, and finally the critical testing phase involving production without feedback (48 trials). During the testing phase, participants encountered objects seen during exposure (aligned trials, where both the phonological and semantic cues agree in which marker they suggest) and new untrained objects; those untrained objects have conflicting cues, with the phonological cue used for one class of nouns during exposure, but the semantic cue used for the other (i.e. returning to the example above, a conflicting item might have a label of the form CiCi but a planet refer-

ent). Responses on conflicting trials can therefore align with either the semantic cue *or* the phonological cue.




Figure 2: Average proportion of trials on which children and adults choose the semantic cue for aligned vs. conflicting trials in Experiment 2. For aligned trials, the semantic and phonological cues are aligned, thus these trials indicate how well participants have learned the language they were trained on. Conflicting trials indicate whether participants chose the plural marker on the basis of the semantic or phonological cue on novel trials in which the two cues conflict. Error bars represent 95% CIs. Colored circles represent individual participant means.

Figure 2 shows the average proportion of trials on which children and adults chose the semantic cue for each of these trial types. A logistic mixed-effects model was run, predicting semantic choice based on trial type (aligned vs. conflicted) and age (adults vs. children). The model revealed a significant effect of age ($\beta = 2.39 \pm 0.36, p < 0.001$), indicating that children differ from adults across both types of trials—on aligned trials, that simply indicates that children learned the language less accurately. There was also a significant effect of trial type ($\beta = 1.71 \pm 0.49, p < 0.001$): both age groups rely less on the semantic cue on conflicting trials compared to aligned trials. Finally, Figure 2 suggests that the difference between trial types is not the same across age groups; rather it appears that adults strongly tend to use the semantic cue for conflicting trials while children are much more likely to use the phonological cue. However, the interaction term in this model was not significant ($\beta = -0.10 \pm 0.39, p = 0.79$). Although this might seem surprising at first glance, it is due to the confluence of two issues: variability in the child data, and near-ceiling scores for adults. First, there is substantial individual variation in the children’s data, with 5 children tending to use the semantic cue, one child alternating between the two cues, and the remaining 14 tending to choose the phonological cue, some very strongly. This limits the precision with which the fixed

effects in the model can be estimated.² Secondly, because the adults are very close to ceiling, the relatively small difference between how often they use the semantic cue across trial types in probability space is actually quite large in logit space. More concretely, while the average difference between trial types is much *larger* for children than adults in probability space (by a factor of 3.5), it is actually *smaller* in log-odds space. Because there is no in principle reason for evaluating log-odds rather than comparing behavior in probability space, we ran several additional analyses to investigate the interaction further. A standard by-participants ANOVA revealed a main effect of age ($F(1, 74) = 19.00, p < 0.001$), trial type ($F(1, 74) = 1.20, p = 0.001$), and their interaction ($F(1, 74) = 7.10, p = 0.009$). We also computed a difference score for each participant by subtracting the average proportion of semantic choices on aligned and conflicting trials. We subjected the difference scores for each age group to a two-sample t-test, which again confirmed a significant difference between adults and children ($t = -3.25, p = 0.002$).

To summarize, in Experiment 1 we showed that children were able to learn a novel noun class system based on either a semantic cue (animacy) or a noun internal phonological cue (vowel reduplication). In fact, these two cues were learned equally well when they provided the only cue to class. In Experiment 2 however, the data suggest that these cues are not treated the same by adults and children when they are both present and equally reliable. In particular, when the two cues conflict, adults almost categorically use the semantic cue to determine class (17/20 participants). By contrast, children tend to use the semantic cue much less (5/20). While this apparent effect—which resembles the findings from acquisition of natural language noun classes—is intriguing, the nature of the data, and the method of analysis in this case necessitate caution in drawing overly strong conclusions. To address this issue, we are currently testing younger children (e.g., 4 year olds); our data suggest a developmental effect, and in fact most of the children in our dataset who used the semantic cue in conflicting trials were in the older half of our sample, therefore younger children may show a stronger effect.

Experiment 3

Results from Experiment 2 suggests that when both semantic and phonological cues to class are present and conflicting, children may rely more on phonological cues than adults. In our final experiment, we ask whether this effect could be further amplified by the early availability of phonological cues. Culbertson et al. (2017) hypothesized that the explanation for children’s over-reliance on phonology might come from early access to phonological cues, prior to learning noun mean-

²It is worth noting that although children were more variable in their responses to conflicting trials than adults, they were not responding randomly. To confirm this, we calculated how often each child used their preferred cue type (for example, if a child chose based on semantics in 25% of trials, then they used their preferred cue type—the phonological cue—75% of the time). Children were well above chance on this measure ($\beta = 1.79 \pm 0.27, p < 0.001$).

ings. We trained adults on a system with confounded semantic and phonological cues of varying salience. A high salience semantic cue (animacy) was paired with an equally salient phonological cue (a paired noun beginning and ending similar to a prefix and suffix combination), or a lower salience phonological cue (a noun ending only). In order to simulate the hypothesized early availability of phonology cues, an additional exposure phase was added in which participants were trained on nouns and class markers alone, without pictured referents (i.e. the phonological cue was available, but not the semantic cue). We found that when the two cues were of similar salience, adults tended to rely more on the earlier available phonological cue in later conflicting trials. Even the relatively weak phonological cue (noun ending only) could be strengthened via early availability. Importantly, staging the input such that the semantic cue was available early instead (by training participants on pictures and markers without the noun itself present) had a similar effect; when paired with a phonological cue of similar salience, the early available semantic cue was used more often. In Experiment 3, we test whether the effect of cue availability is similar in children. One possibility, in light the findings of Experiment 2, is that children may simply prefer to use the phonological cue regardless of its availability. Relatedly, early availability may further strengthen the phonological cue, but not the semantic cue. However, if cue availability and any a priori preference for phonology are independent of one another, they we expect to see both cues strengthened by our staging manipulation.

Participants were 6-7 year old children (N=40). Half of the children were first exposed to the phonological information alone, and the other half to the semantic information alone. In the single cue exposure phase (48 trials) all trials were plural. In the semantics-first condition, on each trial, a set of the same objects appeared on the screen, and the plural marker was presented auditorily (without the any label). In the phonology-first condition, on each trial, a label and plural marker were presented auditorily (without any picture). Following single-cue exposure, the remainder of the experiment was identical across conditions, and the same as Experiment 2 (the second exposure phase was 48 trials total).

Figure 3 shows the average proportion of trials on which children chose the semantic cue for aligned and conflicting trial types in each of the two staged conditions, with the results from the unstaged Experiment 2 repeated for comparison. A logistic mixed-effects model was run predicting semantic choice based on condition (Phonology First vs. Semantics First) and trial type (aligned vs. conflicted). The model revealed a significant effect of condition ($\beta = -0.56 \pm 0.28, p = 0.05$), trial type ($\beta = -1.56 \pm 0.24, p < 0.001$), and a significant interaction ($\beta = 0.77 \pm 0.23, p < 0.001$). This indicates that in both conditions, conflicting trials generated fewer semantic choices compared to aligned trials, but as predicted the two conditions differ on how large this difference is. In particular, in the Phonology First condition, participants were more likely to based their choice on phonology

when cues conflict, while participants in the Semantics First condition we more likely to go with semantics.

Comparing these results to those of Experiment 2, we can see that in both cases, staging has an effect on children’s behavior in conflicting trials. Their reliance on phonology is strengthened by having the phonological cue available first, and weakened by having the semantic cue available first. This is confirmed by logistic mixed-effects models comparing the Unstaged data to the Phonology First condition (no effect of condition: $\beta = 0.13 \pm 0.39, p = 0.74$; significant effect of trial type: $\beta = 1.85 \pm 0.11, p < 0.001$; significant interaction: $\beta = -0.38 \pm 0.15, p = 0.01$, note that this interaction is driven by a small change in *both* trial types) and to the Semantics First condition (no effect of condition: $\beta = -1.0 \pm 0.58, p = 0.10$; significant effect of trial type: $\beta = 0.73 \pm 0.10, p < 0.001$; significant interaction: $\beta = 0.79 \pm 0.14, p < 0.001$).




Figure 3: Average proportion of trials on which children choose the semantic cue for aligned vs. conflicting trials in Unstaged (Experiment 2 results, repeated from above), Phonology First, and Semantics First conditions (Experiment 3). Error bars represent 95% CIs. Colored circles represent individual participant means.

Discussion and conclusions

Previous research on the acquisition of noun classification systems in natural language has found that child learners rely disproportionately on phonological cues to class, even in cases where competing semantic cues are more reliable (Karmiloff-Smith, 1981; Mills, 1985; Mariscal, 2009; Perez-Pereira, 1991; Demuth, 2000; Demuth & Ellis, 2008; Rodina & Westergaard, 2012; Gagliardi & Lidz, 2014). In our previous experiments with adults (Culbertson et al., 2017), we found evidence that the early availability of phonological cues might explain this otherwise puzzling result. We argued that because phonological information is available to learners before word meanings are acquired, children may begin by building noun class systems on the basis of surface phonological cues and only later integrate semantic information. This may lead to a developmental stage during which phonological cues are weighted more heavily than semantics. By staging

the early availability of these different types of cues in an artificial language learning paradigm with adults, we were able to show that even relatively low salience phonological cues (e.g., a single syllable noun ending) were used more than semantic cues (e.g., animacy) if they were learned first. Importantly, we found that adults would also rely more on early available semantic cues. Based on our findings, we therefore argued that early availability rather than an active bias against semantics (or phonology) determined adult learners’ use of cues to noun class. However, we left open the possibility that children might behave differently. Here we reported a series of experiments aimed at exploring this.

In Experiment 1, we trained children on a novel noun class system in which a class marker was perfectly predicted by *either* a semantic *or* a phonological cue. We chose two cues—one based on animacy, and the other on vowel reduplication—which we predicted would be highly salient to child learners. Indeed, participants were highly accurate at learning and generalizing both these two cues, with no significant differences in the level of learning. Having shown that these cues were both similarly learnable for children, in Experiment 2, we then trained both children and adults on a system in which the cues were confounded as in Culbertson et al. (2017). During training, class markers for a given noun were determined by both the semantic and phonological cue. At test, we presented learners with trials in which the cues conflicted (i.e., the semantic cue for one class was present, along with the phonological cue to the other class). Here, whereas adults were highly likely to rely on the semantic cue, children were instead more likely to rely on the phonological cue. The analysis of this effect was complicated by the fact that adult participants were near-ceiling in their preference for semantic cues—while the crucial interaction was not statistically significant when analyzed using logistic regression, analyses in probability space suggest a highly significant effect. Finally, in Experiment 3 we explored whether children are sensitive to the early availability of a particular cue. Like adults, we found that learning either the phonological *or* the semantic cue first led child to use that cue more in trials where the two cues conflict.

These results provide tentative support for two proposals put forward in the literature to explain why children acquiring noun class systems in natural language appear to over-rely on phonological cues. First, there is an effect of early availability of cues on learning which appears to hold across development; both adult and child learners tend to rely more on early available cues, regardless of whether the cues come from noun-internal phonological features or word meanings. Our findings are also consistent with the possibility that children have an active bias against semantic cues; when both a noun-internal and external cue are available, children prefer to use the noun-internal cue (Gagliardi, 2012; Culbertson & Wilson, 2013; Gagliardi & Lidz, 2014; Gagliardi et al., 2017).

Why might this be the case? One possibility is that phonological cues may be preferred in this context since both the

class marker and the noun are within the same local domain. Indeed there is evidence from language comprehension which points to a similar over-reliance on local or “bottom-up” information. For example, five-year-old children readily use bottom-up cues such as lexical statistics, prosodic rhythm and stress to resolve syntactic ambiguities, but do not use contextual cues from the global environment in an adult-like way (Snedeker & Trueswell, 2004; Snedeker & Yuan, 2008). Similarly, compared to adults, 4-year-old children are more likely to use local lexical statistics rather than global contextual cues to disambiguate word senses (Rabagliati, Pytkänen, & Marcus, 2013). If a bias to use local cues *and* a persistent reliance on earlier learned cues are both at work during the acquisition of noun classes, this would lead child learners to rely heavily on phonology, fully integrating semantic cues only later in development.

Importantly, further work is required to show convincingly that children have a bias to use local, noun-internal information when forming grammatical dependencies like noun class. As mentioned above, if this bias changes across development as our results suggest, we predict it will be stronger in even younger learners. Of course, here we have also restricted our investigation to two particular cues—an animacy-based cue, and reduplicated stem vowel cue. Given previous work suggesting the possible importance of reduplication in children’s early word learning (Ota & Skarabela, 2016, 2017), investigating this apparent bias using alternative phonological cues will help establish the generalizability of our findings.

Acknowledgments

The authors would like to thank all the families who participated in our studies, as well as South Morningside Primary, Culter Primary, Skene Square Primary, Aberdeen and Edinburgh Councils, and the Dalkeith Rainbows. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement 681942, held by KS).

References

Bates, D. (2010). lme4: Mixed-effects modeling with R. <http://lme4.r-forge.r-project.org/book>.

Berko, J. (1958). *The child’s learning of english morphology*. Unpublished doctoral dissertation, Radcliffe College.

Carroll, S. E. (1999). Input and SLA: Adults’ sensitivity to different sorts of cues to French gender. *Language Learning*, 49(1), 37.

Culbertson, J., Gagliardi, A., & Smith, K. (2017). Competition between phonology and semantics in noun class learning. *Journal of Memory and Language*, 92, 343–358.

Culbertson, J., & Wilson, C. (2013). Artificial grammar learning of shape-based noun classification. In *Proceedings of Cogsci 35* (pp. 2118–2123).

Demuth, K. (2000). Bantu noun class systems: loanword and acquisition evidence of semantic productivity. *Systems of Nominal Classification*. Cambridge University Press, Cambridge, UK, 270–292.

Demuth, K., & Ellis, D. (2008). *Revisiting the acquisition of Sesotho noun class prefixes*. Mahwah, NJ: Lawrence Erlbaum.

Gagliardi, A. (2012). *Input and intake in language acquisition*. Unpublished doctoral dissertation, University of Maryland.

Gagliardi, A., Feldman, N., & Lidz, J. (2017). Modeling statistical insensitivity: Sources of suboptimal behavior. *Cognitive Science*, 41(1), 188–217.

Gagliardi, A., & Lidz, J. (2014). Statistical insensitivity in the acquisition of Tsez noun classes. *Language*, 90(1), 58–89.

Karmiloff-Smith, A. (1981). *A functional approach to child language: A study of determiners and reference* (Vol. 24). New York, NY: Cambridge University Press.

Mariscal, S. (2009). Early acquisition of gender agreement in the Spanish noun phrase: starting small. *Journal of Child Language*, 36(1), 143–171.

Mills, A. E. (1985). The acquisition of German. In D. Slobin (Ed.), *The crosslinguistic study of language acquisition* (Vol. 1). Hillsdale, NJ: Lawrence Erlbaum.

Ota, M., & Skarabela, B. (2016). Reduplicated words are easier to learn. *Language Learning and Development*, 12(4), 380–397.

Ota, M., & Skarabela, B. (2017). Reduplication facilitates early word segmentation. *Journal of Child Language*, 1–15.

Perez-Pereira, M. (1991). The acquisition of gender: What Spanish children tell us. *Journal of Child Language*, 18(03), 571–590.

Polinsky, M., & Jackson, D. (1999). Noun classes: Language change and learning. In B. A. Fox, D. Jurafsky, & L. Michaelis (Eds.), *Cognition and function in language* (pp. 29–50). Stanford, CA: CSLI.

R Development Core Team. (2010). R: A language and environment for statistical computing [Computer software manual]. Vienna.

Rabagliati, H., Pytkänen, L., & Marcus, G. F. (2013). Top-down influence in young children’s linguistic ambiguity resolution. *Developmental psychology*, 49(6), 1076.

Rodina, Y., & Westergaard, M. (2012, 11). A cue-based approach to the acquisition of grammatical gender in Russian. *Journal of Child Language*, 39, 1077–1106.

Schuler, K., Yang, C., & Newport, E. (2016). Testing the tolerance principle: Children form productive rules when it is more computationally efficient to do so. In *Proceedings of Cogsci 38*.

Snedeker, J., & Trueswell, J. C. (2004). The developing constraints on parsing decisions: The role of lexical-biases and referential scenes in child and adult sentence processing. *Cognitive psychology*, 49(3), 238–299.

Snedeker, J., & Yuan, S. (2008). Effects of prosodic and lexical constraints on parsing in young children (and adults). *Journal of memory and language*, 58(2), 574–608.