MANUSCRIPT ACCEPTED AT LANGUAGE LEARNING AND DEVELOPMENT

AS OF OCTOBER 13, 2021

Repetition, but not acoustic differentiation, facilitates pseudohomophone learning by children

Erin CONWELL*, Felix PICHARDO¹, Gregor HORVATH², and Amanda LOPEZ³

Department of Psychology, North Dakota State University

*Corresponding author:

Erin Conwell, Ph.D.

Department of Psychology

North Dakota State University

Fargo, ND

erin.conwell@ndsu.edu

1. Now at Minnesota Center for Twin Research, University of Minnesota, Minneapolis, MN
2. Now at Rogers Behavioral Health, Oconomowoc, WI
3. Now at Department of Speech, Language, and Hearing Sciences, University of Colorado Boulder, Boulder, CO

Abstract

Children's ability to learn words with multiple meanings may be hindered by their adherence to a one-to-one form-to-meaning mapping bias. Previous research on children's learning of a novel meaning for a familiar word (sometimes called a *pseudohomophone*) has yielded mixed results, suggesting a range of factors that may impact when children entertain a new meaning for a familiar word. One such factor is repetition of the new meaning (Storkel & Maekawa, 2005) and another is the acoustic differentiation of the two meanings (Conwell, 2017). This study asked 72 4-year-old English-learning children to assign novel meanings to familiar words and manipulated how many times they heard the words with their new referents as well as whether the productions were acoustically longer than typical productions of the words. Repetition supported the learning of a pseudohomophone, but acoustic differentiation did not.

Key words: Lexical development, homophones, frequency effects

When children learn the words of their first language, they do so in situations of tremendous referential ambiguity. Upon encountering a new word, they may need to entertain multiple possible referents for that word and to narrow down those candidate meanings. Theories of word learning typically propose that children solve this referential ambiguity problem through a set of biases or principles that guide their hypotheses about possible meanings for a novel word. One of these principles, present in one form or another in nearly every formalized word learning theory, is a one-to-one mapping between form and meaning (e.g., Markman, 1990; Slobin, 1973). Although the details of this principle vary from theory to theory, the general proposition is that children are biased to assume that every word has one meaning and every meaning has one word. Such a bias would facilitate word learning by restricting the set of possible meanings for a novel word, but also allows children to assume that they should not hypothesize new meanings for words that they already know.

However, a one-to-one mapping bias does not allow for the learning of words with more than one meaning, which are attested across the world's languages. Words with multiple meanings may be polysemes, which have multiple related meanings, or homophones, which have multiple, unrelated meanings. Polysemes tend to follow consistent patterns of meaning across languages (Srinivasan & Rabagliati, 2015), and there is evidence that children can use these consistencies to support the extension of a known word to a novel, related meaning (Srinivasan, Berner, & Rabagliati, 2019). Homophones, on the other hand, require that children posit a completely novel meaning for a familiar word, which should be more challenging under a one-to-one mapping bias.

Some research has, indeed, shown that homophones are particularly difficult for children to learn. Beveridge & Marsh (1991) presented 3- to 7-year-old children with pictures and stories

that were consistent with the lower frequency, and therefore less familiar, meaning of a familiar homophone. They found that age and higher levels of contextual information improved children's ability to select the lower frequency meaning of a homophone, but that overall, children still had difficulty using contextual cues when the more frequent meaning was also present. This shows that children struggle with competing meanings for the same word. To examine whether children have homophone pairs in their vocabularies at all, Backscheider & Gelman (1995) showed 3-year-old children a picture of the higher frequency meaning of a homophone pair and labeled it (e.g., "This is a bat"). Then they presented a set of pictures including the lower frequency meaning of the homophone and asked the children to select "another kind of" the homophone. Children were above chance, selecting about 75% of correct referents on a 3-alternative choice task. Taken together, these studies show that young children may struggle to suppress the more familiar meaning of a homophone to access its lower frequency counterpart; however, both studies indicate that children do have homophone pairs in their vocabularies. This raises the question of how children learn homophones in light of the one-to-one mapping hypothesis.

Studies that examine homophone learning often use the *pseudohomophone* (sometimes called *pseudohomonym*) paradigm developed by Mazzocco (1997). This technique involves teaching children a new meaning for a familiar word that is not actually homophonous in their language. It allows researchers to examine how children learn multiple meanings for a word by beginning with a single-meaning word that the children are likely to know and attempting to teach them an additional meaning for that word. Learning of pseudohomophones is often compared to learning of completely novel words to ensure that children understand the word-learning task. In Mazzocco's (1997) study, 3- and 4-year-old children performed below chance when learning

pseudohomophones, but above chance when learning novel words. Second grade children were above chance at learning pseudohomophones, but still learned novel words better than second meanings for familiar words. When children in both age groups failed to select the correct referent at test, they nearly always chose the familiar referent of that word. This suggests that young children do adhere to a one-to-one mapping approach to word learning, which prevents them from adding homophones to their lexicons. One critique of Mazzocco's (1997) study is that, because the referents for the pseudohomophones were also familiar items, children were not contending merely with homophony. They also had to permit synonymy. Therefore, Doherty (2004) conducted a similar study in which the referents of the pseudohomophones were unfamiliar objects. That study found that school-aged children were still poor at assigning a new referent to a familiar word, even in the absence of pseudosynonymy.

Although these pseudohomophone studies indicate that children avoid homophony, other studies show that children will learn new meanings for a familiar word under certain circumstances. For example, Bakscheider, et al. (1999) found evidence that 3-year-old children would select the newly introduced referent of a familiar word when the familiar meaning was not an option. That study also found that children were more likely to accept a new referent for a familiar word if the new referent differed from the familiar referent in animacy. Casenhiser (2005) showed that 4-year-old children more reliably selected the new referent of a pseudohomophone if the new meaning was in a different lexical category. That is, a familiar noun was easier to map to a novel action than it was to map to a novel object. Dautriche, et al. (2018) showed similar patterns of animacy and grammatical category effects in pseudohomophone learning by toddlers. Storkel & Maekawa (2005) further found that 3- and 4-year-old children more reliably selected the novel referent of the pseudohomophone when

exposed to the new meaning multiple times during learning and when the familiar referent was unavailable. Taken together, these findings suggest that lexical category, large differences in meaning, and repeated production in the absence of the familiar referent may all indicate to learners that a particular word should be allowed to have more than one meaning.

Another cue that may help children detect homophony is acoustic differentiation. Studies of both child-directed and adult-directed speech have demonstrated that homophone meanings are acoustically distinct from one another (Conwell, 2017; Gahl, 2008; Lohmann, 2018). Specifically, the more frequent meaning of a homophone is frequently produced with a shorter duration than the less frequent meaning is. In child-directed speech (CDS), these differences can be quite variable, depending on the word pair, but the slopes reported in Conwell (2017) suggest that lower frequency words in CDS may be as much as 100 ms longer than their higher frequency counterparts. Children also produce these distinctions in their own speech, although with differences of 20-50 ms, which are somewhat smaller than those produced in CDS (Conwell, 2018). Adults can perceive differences in the range of 30 ms (Sanker, 2019). Infants are also sensitive to durational differences in noun/verb homophones (Conwell & Morgan, 2012). This is a cue to homophony that relies on neither the absence of the familiar referent nor on repeated exposure to the novel meaning. It is well-attested in children's experience (Conwell, 2017) and could support homophone learning where cues such as animacy differences, contrastive explication, and repeated exposure are unavailable. Furthermore, the use of acoustic differences to posit homophony does not require a violation of the one-to-one mapping hypothesis. Because apparent homophones have two different durations, one associated with the more frequent meaning, and one associated with the less frequent meaning, children would not need to map two meanings to a single acoustic word form. Rather, they could treat the two

durations as indicating the presence of two distinct word forms, which would allow adherence to the one-to-one mapping hypothesis.

Previous research has demonstrated that children will use acoustic information during online lexical processing. Paquette-Smith, Fecher, & Johnson (2016) found that English-learning two-year-old children show better word recognition when the coarticulatory cues in a monosyllabic word match the target meaning (e.g., when the coda of the word *boat* comes from a token of that word) than when the coarticulatory cues are inconsistent with the target meaning (e.g., when the coda of the word *boat* has been cross-spliced from a token of the word *bone*). Zamuner, Moore, & Desmeules-Trudel (2016) reported similar findings with two-year-olds, three-year-olds, and adults. The use of acoustic information, such as coarticulation, in these studies indicates a possibility that young children might also use this information during word learning. On the other hand, children should certainly not treat all acoustic differences as evidence of lexical contrast, and prior work has shown that they may not use even phonological changes to motivate the assignment of a novel meaning to a familiar form. When presented with a novel word that differs in terms of a single segment from a familiar word, two-year-old children tend to select the referent of the familiar word as the target, both by touching it and by looking more to that object (Merriman & Schuster, 1991; Swingley, 2016). This indicates that perceptible phonological differences may not be treated as lexically contrastive for learning purposes, which suggests that more difficult to perceive acoustic differences might also not be used. The conflict between children's sensitivity to acoustic information in processing, but their apparent non-use of larger distinctions for word learning further motivates the question of whether acoustic differences support the learning of a second meaning for a known word.

The present study asks whether children can, in fact, recruit durational differences in homophone productions to learn that a familiar word has a second meaning. In a pseudohomophone task, modeled on that used by Doherty (2004), children were asked to learn a second meaning for a familiar word. To determine whether children are better able to learn a second meaning for a known word when uses associated with that meaning have longer durations than usual, the lengths of some of the familiar words were extended during exposure to the novel meaning. This study also varied the number of exposures to the novel meaning, which has been previously shown to improve pseudohomophone learning (Storkel & Maekawa, 2005) and examined how repetition interacted with the availability of durational cues to homophony in pseudohomophone learning.

Method

Participants

Participants were 72 four-year-old English-learning children (range: 3;10-4;11; mean age: 4;5), 38 male and 34 female. Although specific racial demographic information was not collected, nearly all participants in the study were white. Four additional children completed the study but were excluded from the analysis due to inattentiveness (1), age outside of the target range (2) and significant (>25%) exposure to a language other than English (1). Parents or guardians gave written, informed consent for their children to participate, and the children gave verbal assent for participation. All children received a book for participating, and their parents received a small monetary reimbursement.

Design

This study used a pseudohomophone learning paradigm similar to that used by Doherty (2004). In this paradigm, children are exposed to familiar words used with novel meanings to

simulate the experience of encountering a second meaning of a homophone. However, because the familiar words used in this study do not have homophone counterparts and were paired with novel meanings, they are considered *pseudohomophones*. On each trial, children heard a brief story (3-4 sentences) that referred to unfamiliar objects using either a familiar, but unrelated, word (pseudohomophone) or a nonsense word. The nonsense words were included to ensure that children understood the word-learning task. After each story, children were asked to select the referent of the target word from a grid of four choices. All children completed eight pseudohomophone trials and two nonsense word trials. To test whether homophone duration affected children's willingness to assign a novel meaning to a familiar word, the duration of the pseudohomophones was manipulated as a within-subjects variable. On half of the pseudohomophone trials for each participant, the target word was modified in both the story and the question to be 25ms longer in duration than its typical pronunciation. This length was selected based on acoustic differences between homophones in CDS and those produced by children in elicitation studies (Conwell, 2017; 2018). (See the Stimuli section for details on this process.) For the other half of the pseudohomophone trials, the target word was unmodified from its baseline pronunciation. All children heard four modified pseudohomophones, four unmodified pseudohomophones, and two nonsense words. Which pseudohomophones were modified was counterbalanced across participants.

To test whether frequency of use impacted children's willingness to assign a novel meaning to a familiar word, the frequency with which children heard the target words in the stories was manipulated as a between-subjects variable. Half of participants (36) heard each target word only once, while the other half (36) heard each target word three times.

Stimuli

Eight familiar monosyllabic nouns were selected from the MacArthur-Bates CDI (Fenson, et al., 1994) to use as pseudohomophones: *key*, *sky*, *bib*, *pen*, *cake*, *cow*, *egg*, and *cup*. Additionally, *vik* (/vIk/) and *kesk* (/kɛsk/). were used as nonsense words. Each pseudohomophone and nonsense word referred to an inanimate object that was likely to be unfamiliar to 4-year-old children from the American Midwest. These unfamiliar objects included specialized equipment (e.g., pet dryer), post-modern furniture, tropical fruits and vegetables (e.g., hala fruit), and novel objects created for research purposes (i.e., *fribbles*; Barry, Griffith, De Rossi & Hermans, 2014). An additional 30 unfamiliar objects were presented as distractor objects during the identification phase of each trial.

Twelve stories, each 3-4 sentences long, introduced the unfamiliar objects to the children. Each story was followed by the question “Can you point to the [TARGET] from the story?” The stories were initially recorded with a nonsense word in the target position which was later replaced with the target word. The pseudohomophones were recorded in sentence-medial contexts using their familiar meaning and then spliced into the stories. All audio stimuli were produced by a female adult native speaker of American English (North Plains dialect) naïve to the purpose of the study.

Following recording, the pseudohomophones were excised from the recorded passages to be used as unmodified tokens. A modified token of each pseudohomophone was also created by adding 25ms to the vowel duration in PRAAT (Boersma & Weenink, 2020). The nonsense words and the words used in the practice trials were not modified. The target words were then spliced into the story recordings in the target position. Two versions of each story were created: one with the modified pseudohomophone and one with the unmodified pseudohomophone. Modified and unmodified pseudohomophones were also spliced into the question that followed each story.

Each story was illustrated as a series of 3-4 scenes, each described by 1-2 sentences. Cartoon-style illustrations of the target and distracter objects were also created. Illustrations and audio were presented on a laptop computer using automatic presentation in PowerPoint. Following each story, children saw a 2x2 grid containing the target object and three distracter objects. On pseudohomophone trials, the familiar referent of the target word was one of the distracter objects, following Doherty (2004). Upon presentation of this grid, children heard the question “can you point to the *target* from the story?”¹ An example of one of these grids is shown in Figure 1. The location of the target object in the grid was randomized across trials. All experimental materials for this study are available on the Open Science Framework (<https://osf.io/6q45m/>).

Procedure

Before the study began, parents completed a survey regarding their child’s knowledge of the objects and words used in the study. The list included all of the objects and words the child would see and hear in the experiment as well as a few unrelated terms (41 in total). If an object had more than one common name, all names were included in the survey. This survey was intended to ensure that children were familiar with the common meanings of the words that would be used as pseudohomophones and that they did not already have labels for the novel referents. Surveys from two of the 72 participants were unavailable due to experimenter error.

During the study, children sat in front of a laptop computer with the experimenter beside them. The experimenter instructed the children to listen closely to the stories. They were also

¹ We modeled this phrasing after that used in Doherty (2004), but the use of this question opens the possibility of children using the strategy of simply selecting the object that was in the story on every trial. If they used this strategy, we would expect no differences between conditions. However, future studies may wish to include a baseline condition with no word learning, but that asks children to select the item “from the story” to compare with word learning conditions.

told that after each story they would be asked a question and that they should point to the answer on the screen. Two practice trials involving only familiar words and their familiar meanings were completed without headphones, and the children received feedback on their responses. After the practice trials, children put on headphones and adjusted the volume to a comfortable level. The stories played automatically and then paused on the response screen to allow for a response. The location to which the child pointed was recorded with pencil and paper during the study and coded offline as target meaning, familiar meaning, or distractor.

Results

Parental reports on the vocabulary survey indicated that children knew, on average, the familiar meanings for 7.94 of the 8 pseudohomophones. They knew, on average, 0.214 of the real-world names of the 8 new referents of these words. This establishes that the children in this study were mapping unfamiliar referents to familiar words, making these cases all true instances of pseudohomophony without pseudosynonymy.

The results of this study are shown in Figure 2. The data were analyzed using a mixed multinomial logistic regression implemented using the lme4 package (Bates, Maechler, Bolker & Walker, 2015) in R (R Core Team, 2019). The outcome variable was response type and had three possible values: target, familiar, distractor. Familiar was the reference category. The predictor variables were pseudohomophone type (modified or unmodified; manipulated within subjects) and repetition condition (single use or multiple uses; manipulated between subjects). Participant and item were both included as random effects. The fixed factors were centered using contrast matrices generated in R. The complete results of this analysis are shown in Table 1.

Participants correctly selected the target meaning at rates above chance (25%) in all conditions. Across all conditions, children selected the target image on a majority of homophone

trials (mean 4.625 out of 8 possible, s.d.=2.165; range 0-8). Children selected the target meaning from among the four options on 56.6% of modified trials and on 58.7% of unmodified trials. Children who heard the novel meaning of the pseudohomophone only once selected the target meaning on 54.7% of trials, and those in the repetition condition selected the target item on 72.5% of trials. The analysis revealed no evidence of an effect of token modification on target selection ($z=1.047$, $p=0.295$), but a significant effect of repetition condition on target selection ($z=2.28$, $p=0.022$). There was no evidence of an interaction of the two predictors with regard to target selection ($z=1.414$, $p=0.157$).

Nonce trials were not included in the multinomial logistic regression because they were only included in the design to ensure that children understood the word learning task, and because they did not have a meaningful “familiar” item for children to select. Instead, children’s performance on nonce trials was analyzed using a binomial test (target selected vs. other item selected), using 25% correct as the chance level. Children in all conditions selected the target object on nonce trials 87.5% of the time, a rate well above chance (126/144, binomial test $p<.001$). This indicates that children understood the word learning task and had no difficulty assigning a meaning to a novel word. All data from this study are available on OSF (<https://osf.io/6q45m/>).

Discussion

This study asked whether acoustic differentiation and repetition affect children’s ability to assign a novel referent to a familiar word. The results showed that repeated exposures to a novel meaning of a familiar word improved children’s ability to select the novel meaning at test, while acoustic differentiation did not appear to support the mapping of the novel meaning to the familiar word. This suggests that, although acoustic differentiation of homophone meanings is

robustly attested in child-directed English (Conwell, 2017), children do not use that information to support homophone learning. However, multiple exposures to the new meaning can improve the mapping of a familiar word to a novel referent.

The facilitative effect of repetition for learning pseudohomophone is not surprising. Repetition in sequential sentences generally supports word learning (Schwab & Lew-Williams, 2016), and Storkel and Maekawa (2005) demonstrated a repetition effect for pseudohomophone learning specifically. Repetition of a homophone may support learning by providing multiple opportunities for the learner to notice the absence of the typical referent and to determine the intended novel referent. These findings are also consistent with discriminative learning accounts of language learning (e.g., Nixon, 2020), as multiple repetitions may provide learners with the opportunity to learn from prediction errors that arise from the absence of the familiar referent. The present data cannot differentiate between a discriminative learning account and a purely statistical mechanism, however. Additionally, whether unfamiliar meanings for homophonous words are introduced to children in this repetitious way in natural speech environments is an open question.

On the other hand, the lack of a facilitative effect of acoustic differentiation suggests that children do not use a cue to homophony that is readily available in their natural environments (Conwell, 2017) and that both infants and adults are able to perceive (Conwell & Morgan, 2012; Sanker, 2019). One possible reason that children did not use acoustic differences in the context of this study is that these differences were not consistent across trials. The length of the pseudohomophones was manipulated within subjects, meaning that this cue was available to children on some trials, but not others. The speaker identity, however, was consistent across

trials. This may have indicated to children that token length was not a reliable cue to novel meaning for this speaker.

Alternatively, children may have been unable to detect the contrast in token length in this study because no uses of the familiar words with the familiar meaning were available in the context of the study. Because previous work demonstrated that children struggle to learn pseudohomophones when the familiar meaning is present in the context (Mazzocco, 1997; Doherty, 2004), we designed this study to avoid the presence of the familiar referent for each word in the story context. We also did not use any of the words with their familiar meanings in the study. This means that children may have been unable to detect the acoustic differences in the manipulated tokens because they had no information about how this speaker produces those words with the familiar meanings. Evidence from both adults (Nygaard & Pisoni, 1998) and infants (Weatherhead & White, 2016) has shown that learners are sensitive to speaker variability and that speaker-specific production information may be learned and used for subsequent speech processing. In the present study, learners did not have access to the speaker's productions of the familiar word meaning, and so may not have been aware that the lengthened tokens were, in fact, different from that speaker's usual production. In a natural learning context, speakers may be more familiar to children and so children might have a better representation of their interlocutors' usual productions of words, meaning that they may be better able to detect the differences between homophone meanings that are present in natural speech environments.

Another explanation for children's failure to use acoustic information to learn pseudohomophones in this study is that they may interpret the length differences between homophone meanings as uninformative variability, even though that variability does contain information. Speakers robustly produce these differences, but learners must contend with a high

degree of acoustic and phonetic variability in the speech that they hear. If a particular kind of variability, such as the duration of a word, is only related to word meaning for some word types, then learners might ignore that cue altogether. Prior work on infant word learning shows that babies are willing to ignore phoneme substitutions when the substituted speech sound is similar to the target sound (White & Morgan, 2008) or when only one segment has been changed (Swingley, 2016). Toddlers also have difficulty learning a novel word if it is the phonological neighbor of a known word (Dautriche, Swingley, & Christophe, 2015). If learners do not use phoneme-level differences to learn a new word, then it is quite possible that they similarly would not map a novel meaning to a word on the basis of subphonemic differences such as duration.

Overall, this study has shown that homophone learning may be best supported by multiple, clustered exposures to novel meanings for familiar words, rather than by acoustic differences between meanings. Although the literature is clear that acoustic differences between meanings are attested in children's experience, there is not, to our knowledge, evidence on the extent to which homophone meanings are clustered or contrasted in CDS. This will be an important research direction to improve our understanding of children's acquisition of homophones, specifically, but also of lexical acquisition in general.

Acknowledgements

This research was supported by Grant 1R15HD077519-01 to the first author from the National Institute of Child Health and Human Development, part of the National Institutes of Health (NIH). The contents of this article are the sole responsibility of the authors and do not necessarily represent the official views of NICHD or NIH. The authors also thank Shanda Lauer for recording the audio stimuli, and all participants and their families for their support of this research. Comments from two anonymous reviewers and the Editor improved this article and we thank them for their suggestions.

Declaration of Interest Statement

The authors have no conflicts of interest, financial or otherwise, to disclose.

Data Availability Statement

The complete dataset for this study is available on the Open Science Framework at <https://osf.io/6q45m/>.

References

- Backscheider, A. G., Gelman, S. A., Martinez, I., & Kowieski, J. (1999). Children's use of different information types when learning homophones and nonce words. *Cognitive Development, 14*, 515-530.
- Backscheider, A. G. & Gelman, S. A. (1995). Children's understanding of homonyms. *Journal of Child Language, 22*, 107-127.
- Barry, T. J., Griffith, J. W., De Rossi, S., and Hermans, D. (2014). Meet the Fribbles: Novel stimuli for use within behavioral research. *Frontiers in Psychology*.
<https://doi.org/10.3389/fpsyg.2014.00103>
- Beveridge, M. & Marsh, L. (1991). The influence of linguistic context on young children's understanding of homophonic words. *Journal of Child Language, 18*, 459-467.
- Boersma, P. & Weenink, D. (2020). Praat: Doing phonetics by computer. [Computer Program] <http://www.praat.org>.
- Casenhiser, D. (2005). Children's resistance to homonymy: An experimental study of pseudohomonyms. *Journal of Child Language, 32*, 319-343.
- Conwell, E. (2017). Are homophones acoustically distinguished in child-directed speech? *Language Learning and Development, 13*, 262-273.
- Conwell, E. (2018). Token frequency effects in homophone productions: An elicitation study. *Language and Speech, 61*, 466-479.
- Conwell, E. & Morgan, J. L. (2012). Is it a noun or is it a verb? Resolving the ambicategoricity problem. *Language Learning and Development, 8*, 87-112.

- Dautriche, I., Fibla, L., Fievet, A.-C. & Christophe, A. (2018). Learning homophones in context: Easy cases of favored in the lexicon of natural languages. *Cognitive Psychology*, 104, 83-105.
- Dautriche, I., Swingle, D. & Christophe, A. (2015). Learning novel phonological neighbors: Syntactic category matters. *Cognition*, 143, 77-86.
- Doherty, M. J. (2004). Children's difficulty in learning homonyms. *Journal of Child Language*, 31, 203-214.
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., Tomasello, M., Mervis, C. B. & Stiles, J. (1994). Variability in early communication development. *Monographs of the Society for Research in Child Development*, 59, 1-185.
- Gahl, S. (2008). Time and thyme are not homophones: The effect of lemma frequency on word durations in spontaneous speech. *Language*, 84, 474-496.
- Lohmann, A. (2018). Time and thyme are NOT homophones: A closer look at Gahl's work on the lemma-frequency effect, including a reanalysis. *Language*, 94, e180-e190.
- Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14, 57-77.
- Mazzocco, M. M. M. (1997). Children's interpretations of homonyms: A developmental study. *Journal of Child Language*, 24, 441-467.
- Merriman, W. E. & Schuster, J. M. (1991). Young children's differentiation of object name reference. *Child Development*, 62, 1288-1301.
- Nixon, J. S. (2020). Of mice and men: Speech sound acquisition as discriminative learning from prediction error, not just statistical tracking. *Cognition*, 197, 104081.

- Nygaard, L. C. & Pisoni, D. B. (1998). Talker-specific learning in speech perception. *Perception & Psychophysics*, 60, 355-376.
- Paquette-Smith, M., Fecher, N., & Johnson, E. K. (2016). Two-year-olds' sensitivity to subphonemic mismatch during online spoken word recognition. *Attention Perception and Psychophysics*, 78, 2329-2340.
- R Core Team (2019). R: A language and environment for statistical computing. [Computer Program] R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Sanker, C. (2019). Effects of lexical ambiguity, frequency, and acoustic details in auditory perception. *Attention, Perception, & Psychophysics*, 81, 323-343.
- Schwab, J. F. & Lew-Williams, C. (2016). Repetition across successive sentences facilitates young children's word learning. *Developmental Psychology*, 52, 879-886.
- Slobin, D. I. (1973). Cognitive prerequisites for the development of grammar. In C. A. Ferguson & D. I. Slobin (eds.), *Studies in Child Language Development*. New York: Holt, Reinhart, Winston.
- Srinivasan, M. & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124-152.
- Srinivasan, M., Berner, C. & Rabagliati, H. (2019). Children use polysemy to structure new word meanings. *Journal of Experimental Psychology: General*, 148, 926-942.
- Storkel, H. L. & Maekawa, J. (2005). A comparison of homonym and novel word learning: The role of phonotactic probability and word frequency. *Journal of Child Language*, 32, 827-853.
- Swingle, D. (2016). Two-year-olds interpret novel phonological neighbors as familiar words. *Developmental Psychology*, 52, 1011-1023.

- Venables, W. N. & Ripley, B. D. (2002). *Modern Applied Statistics with S, Fourth Edition*. New York: Springer.
- Weatherhead, D. & White, K. S. (2016). He says potato, she says potahto: Young infants track talker-specific accents. *Language Learning and Development*, 12, 92-103.
- White, K. S. & Morgan, J. L. (2008). Sub-segmental detail in early lexical representations. *Journal of Memory and Language*, 59, 114-132.
- Zamuner, T. S., Moore, C., & Desmeules-Trudel, F. (2016). Toddlers' sensitivity to within-word coarticulation during spoken word recognition: Developmental differences in lexical competition. *Journal of Experimental Child Psychology*, 152, 136-148.

Table 1. Complete results of the multinomial logistic regression. The dependent measure was the selected image; familiar was the reference category. * $p < 0.05$

	Coefficient	Standard Error	z
Intercept	0.508	0.369	1.378
Trial Type	0.212	0.202	0.295
Repetition Condition	0.774	0.339	0.022*
Trial Type x Repetition Condition	0.572	1.414	0.157

Figure 1.

Figure 2.

Figure Captions

Figure 1. A sample test grid. In this example, the audio asked children to select “the pen from the story.” The object in the upper right quadrant was referred to as a “pen” in the story they had just heard, and is therefore the target. The actual pen in the lower left is the familiar meaning of this word and the other two objects are distractors.

Figure 2. Proportion of each image selected by children across conditions. Children were asked to select one image of four, which included the target meaning, the familiar meaning, and two distractors. Modified tokens were lengthened by 25 ms from their original production, while unmodified tokens were unchanged.