The effect of the number of labeled objects on novel referent selection
across short and long time delays

Erica H Wojcik

Skidmore College

ewojcik@skidmore.edu

Abstract

Children often hear many new words in one conversation, and yet word learning research overwhelmingly focuses on how children learn and retrieve the meanings of single words. The current experiment tests how the number of labeled objects affects preschoolers' novel word referent selection immediately after encoding and after a one-week delay. Seventy 3- to 6-year-olds were exposed to four novel objects. Half of the participants were given novel labels for two of the objects and half were given novel labels for all four. Label-referent mapping was tested with a four alternative forced-choice pointing task both immediately after exposure and one week later. Children performed worse overall after a week delay, replicating past work on novel word retention. While children performed significantly worse overall in the Four-Label condition, exploratory analyses revealed that this effect was driven solely by the second test trial immediately after exposure. Analyses suggest that referent selection is strongly influenced by in-the-moment constraints, such as label salience and pragmatic biases, and that these constraints are strongest immediately after novel word exposure.

Keywords: word learning; retention; referent selection; fast mapping; task validity

By three years of age, children are experts at learning novel word meanings. They can use multiple types of cues in the input, including social, perceptual, linguistic, conceptual, and pragmatic information, to successfully map a novel word to its correct referent at first exposure (Baldwin, 1993; Booth, Waxman & Huang, 2005; Smith & Samuelson, 2006; Sullivan & Barner, 2016; Naigles, 1990). These cues to word meaning, such as points, syntactic frames, or conversation topics, are typically thought of as pieces of information that help children induce a word's meaning. Investigations of word learning mechanisms, in other words, tend to focus on how children learn the meaning of an individual word. If more than one word is introduced during an experiment, it is typically done to increase statistical power and generalizability, not because of a hypothesis about the introduction of more than one novel word-label mapping affects novel word retrieval.

While the question of how children learn and retrieve an individual word's meaning is foundational to understanding word learning writ large, it ignores the fact that children do not process words one at a time. Children hear words in dialogue, and thus they may hear multiple new words in one conversation. How is word learning and retrieval affected by the number of newly labeled referents in a given moment? This question is particularly interesting because it intersects with questions about how memory processes, such as retention and retrieval, interact with word learning (see Wojcik, 2013 for a review of this intersection).

The effect of the number of novel labels on learning and retrieval cannot be gleaned from the many existing word learning studies because there is no consistency in how many novel words are presented in the training phase of an experiment. The number of novel words that children are taught in the lab ranges from one (Markson & Bloom, 1997; Vlach & Sandhofer, 2012; Waxman & Booth, 2000) to twelve (Yu & Smith, 2008). This variability not only makes it

difficult to infer how the number of novel words presented affects learning and retrieval, but it also makes many other comparisons across word learning studies meaningless. For example, as pointed out by Horst and Samuelson (2008), some studies show evidence that three year olds can retain a novel word for a month (Markson & Bloom, 1997), while others find that 2-year-olds have trouble recalling a novel word even five minutes later (Bion, Borovsky, & Fernald, 2012; Horst & Samuelson, 2008). While the age differences between these studies point to a developmental change in memory processes, it is also likely that this inconsistency is due in part to methodological differences, including the number of words presented during training: Markson and Bloom (1997) only taught children one novel word label, whereas Bion, Borovsky and Fernald (2012) taught children two.

To investigate how the number of novel labels presented during training affects retrieval performance, Axelsson and Horst (2013) exposed two year olds to four novel objects. All participants were exposed to all four objects at training. However, for some participants only one object was labeled, while for others, all four objects were labeled. Immediately after labeling, participants completed a four-alternative forced-choice task that included all four novel referents in each trial. While those who were given only one label performed near ceiling, those who were given all four labels performed at chance, even on the first test trial.

Axelsson and Horst (2013) use their results to caution that alternative forced-choice tasks may not test specific new word-meaning mappings. They argue that if only one novel referent in the test array has been labeled, then a child's point to the correct referent may only be evidence that the child knows that the object was given a name, rather than that they encoded the specific label-referent mapping. Indeed, a correct point may also be driven by low-level properties like salience; there is evidence that simply labeling an object draws more attention to it (e.g., Baldwin

& Markman, 1989; Mather & Plunkett, 2009). Given Axelsson & Horst's (2013) findings, researchers should be careful about the assumptions that they make about the mechanisms behind toddlers' successful performance in lab tasks. For example, when researchers teach children only one novel word, they may overestimate the strength of the word's representation.

Axelsson and Horst's (2013) findings, however, potentially have implications not only for how we set up word learning experiments and interpret their results, but also for theories of how we conceptualize novel word learning and retrieval. Researchers often interpret referent selection as an indicator of label-referent mapping success. If a child can point to a word's referent correctly, then that word is assumed to be "in their lexicon." Indeed, this is the colloquial interpretation of a child's correct point. However, there is strong evidence that referent selection is influenced by many in-the-moment variables, from perceptual salience to pragmatic biases (see Samueleon, Kucker & Spencer, 2017 for a review from a dynamic systems perspective). While the typical practice of controlling retrieval context in the lab can help researchers interpret word learning studies more accurately (for example, if all objects at test have labels, then you can assume a correct point is not due to a bias towards labeled objects), it also has the untended consequence of oversimplifying the phenomenon we are trying to understand. When children are asked to point to an object, they are solving an acute communication problem using whatever information is available, including not only the label-referent mapping, but also perceptual and pragmatic constraints. By systematically re-introducing variables into a referent selection task, such as how many referents have labels, we can better understand how children communicate when their word representations are weak.

The current study builds on Axelsson and Horst (2013) by manipulating the number of labeled referents available during referent selection. In the real world, children must recall

words in the presence of a variety of objects and entities, some for which they have a word and some for which they do not. Word learning studies occasionally present a mix of unlabeled and labeled distractor objects at test (e.g. Horst & Samuelson 2008), but when Axelsson and Horst (2013) explicitly manipulated the number of novel labels presented at training, they only tested retrieval when all other distractors had been labeled or none had been labeled. Axelsson and Horst (2013) proposed that children may appear to successfully retrieve a specific newly learned word's referent, when in fact they just remember that only one object had been labeled at all. With a varied distractor retrieval context, one can test whether children are indeed strongly biased to point to any previously labeled object over any unlabeled object when asked to retrieve a newly learned word.

The effect of learning and retrieval variables, such as the number of words labeled, cannot be understood unless long-term retention is probed. While decades of research have investigated novel word encoding by testing retrieval immediately after exposure, there is a dearth of research on the contexts that aid in the preservation of a novel word mapping minutes, days, or weeks after the encoding moment (see Wojcik, 2013 for a review). Understanding the factors that contribute to novel word retention is crucial because contexts that lead to successful novel word encoding, such as the juxtaposition with an already known word in mutual exclusivity experiments, do not always result in the retention of that novel word even five minutes later (Horst & Samuelson, 2008). Axelsson and Horst (2013) only demonstrated their labeling effect immediately after training, and thus the results may reflect a transient bias. Children's attention, and therefore their pointing behavior, may have been momentarily influenced by the fact that novel object(s) had just been labeled. In other words, the feature of whether or not an object has a name may have momentarily overridden other contributions to

their behavior, including the specific label-referent mapping. If, on the other hand, the number of words labeled during training affects the strength of long-term word representations, then the observed labeling effect should still appear in a retention test one week later.

A last question is whether the number of words labeled affects older children's word learning and retrieval. Two year olds, the age group tested by Axelsson and Horst (2013), may be uniquely susceptible to the labeling effect because their lexicon is populated by new, and thus malleable and context dependent, word representations (see McMurray, Horst, & Samuelson, 2012 for a review of this perspective). Once children enter the preschool years and their lexicon is more stable, their word learning and retrieval abilities may be more consistent across both learning and retrieval contexts. On the other hand, it is possible that even preschool aged children's word learning is affected by the number of labels with which they are presented. If the labeling effect generalizes to older children, the scope of implications for both methodology and theory would widen considerably.

The current study tests 3- to 6-year-olds' novel word retrieval immediately after exposure and one week later. As in Axelsson and Horst (2013), children were all presented with four novel referents. Instead of labeling one or four of the objects, though, experimenters in the current study labeled two or four of the novel objects so that there was always another labeled distractor at test. This design allowed for the examination of how the number of labeled referents affects novel word retrieval (a) in older children and (b) across different time delays. If the effect of the number of words labeled generalizes to older children, then participants should perform better at test when fewer novel objects are labeled. If the number of words labeled affects novel word representations, and is not just a transient bias, participants who are taught fewer labels should still perform better one week later. Notably, the current design also allowed for the exploration

of whether retrieval accuracy immediately after learning affects future recall accuracy, the types of errors children make at test, and trial-by-trial performance.

Materials and Methods

The current study employed a 2 x 2 mixed-subject design, with a four-alternative forced-choice pointing task as the dependent measure. Participants were assigned to learn labels for either two novel words or four novel words (Label Condition: Two-Label vs. Four-Label). All participants were tested on novel word retrieval both immediately after training and after a week-long delay (Test Phase: Immediate vs. Delay). The study was pre-registered on the Open Science Framework [https://osf.io/7h5j8/?view_only=d0a80deb2ad844edb8c05a8c66587593] and data analysis files are available on the Open Science Framework [https://osf.io/jm8ny/?view_only=11a2996136e24cd2b3257ce3b96bcd7f].

Participants

Seventy monolingual children aged 3 to 6 years (21 3-year-olds; 27 4-year-olds, and 12 5-year-olds) were recruited from a predominantly white, middle-class preschool and the surrounding community in the Northeast United States. G*Power 3 (Faul et al., 2007) was used to calculate the sample size needed to have 80% power. Based on previous work (Axelsson & Horst, 2013; Booth, 2009), a predicted medium-large effect size (Cohen's $f = 0.30$) was used for calculations. For a 2 x 2 mixed design ANOVA, 68 total participants are needed (34 in each condition). Note that the pre-registered, planned analysis was not an ANOVA, but rather a mixed-effects logistic regression. At the time of pre-registration, there was no consensus as to the correct way to run a priori power analysis for mixed-effects models. However, mixed-effects models are statistically more powerful than ANOVA models (Baayen, Davidson & Bates, 2008), and thus the pre-planned sample size is justified (note that two additional participants were run

and included for completeness; excluding the last two participants does not change the results). Based on our pre-registered exclusion criteria, an additional five children were tested but excluded for not completing every test trial. Note that while we planned on excluding kids who did not respond correctly on the practice trials, all children were accurate on these trials. An additional 16 children were excluded because they did not complete the second testing session either due to experimenter error (11) or absence from school during the testing window (5).

Materials

The stimuli were four novel words, each consisting of a referent and label chosen from the NOUN database (Horst & Hout, 2016). The novel labels were “coodle”, “tulver”, “manu”, and “bosa”, chosen to be phonetically legal in English as well as to have distinct onsets and offsets from each other. The novel referents (see Figure 1) were chosen to have nameability and familiarity scores of under 50% (Horst & Hout, 2016). The pairing between the label and the referent was counterbalanced between participants. Additionally, four familiar objects and labels were used for practice trials. All learning and test trials were presented on a 13” Macbook Pro using Microsoft Powerpoint.

Figure 1. The novel referents that participants were exposed to. Label-referent mappings were counterbalanced across participants.

Procedure

The study consisted of a learning phase, an immediate test phase, and a delay test phase.

Learning phase. Children were seated in a quiet space in front of a laptop next to Experimenter 1 (E1). Experimenter 2 (E2) sat in the same room, but was positioned so that s/he was not able to see the laptop screen. After getting the child's assent, E1 said, "You're going to learn some new words! I'm going to show you a bunch of pictures of funny toys, and I'll tell you what some of them are called. Are you ready?" Using Powerpoint, E1 then presented the child with four novel objects on the laptop screen, one at a time, twice through, for a total of eight trials. On each trial, the novel referent was centered on the screen.

Half of the participants heard novel labels for two of the objects (Two-Label condition), and half of the participants heard labels for all four (Four-Label condition). The assignment of the labels to the objects was counterbalanced between participants. The objects' labels were

introduced with two phrases on each trial: “Look at the (name). Wow, check out the (name).” In the two object condition, the unlabeled objects were accompanied by two attention-getting phrases, such as “Look at this one! Wow!” Order of object presentation was quasi-randomized between participants, with the same orders used for the Two- and Four-label condition so that object presentation order was maintained across conditions.

Immediate test phase. To test novel word encoding, immediately after the learning phase, E2 ran a four-alternative forced choice task (a second experimenter was used so that the tester was blind to the correct answer). E2 first introduced the test phase: “Hi! Let’s play a game. I’m going to say a word, and your job is to point to a picture of that word.” Then, there were two practice trials in which the child saw four familiar objects on the screen. E2 asked “Where’s the ball?”. A second trial with the same four objects in new positions followed with the question “Where’s the shoe”?

The two test trials followed. For both the Two- and Four-Label conditions, each trial presented all four novel objects on the screen, and the experimenter asked, “Where’s the (name)”? The experimenter did not give performance feedback. If the child did not respond, they were asked to make their best guess. Note that only 2 of the 4 labeled objects were tested in the Four-Label condition to equate testing parameters across conditions. The two items were tested in the same order that they were presented in the learning phase (as in Axelsson & Horst, 2013), skipping the objects that were not labeled in the Two-label condition.

During both the practice and test trials, E2 recorded the child’s responses on a score sheet. The child was given a sticker after the test phase, regardless of their performance and whether they completed the experiment.

Delay test phase. To test novel word retention, participants' novel word retrieval was tested again after one week (range: 7-10 days, Mean = 7.35 days, Median = 7) by E2. The delay test was the same as the immediate encoding test, but for the participants in the Four-Label condition, the other two novel words were tested after the first two.

Results

Data were coded such that for each test trial, participants were given a 1 (correct) or 0 (incorrect). Data were analyzed in R (version 3.5.1; R Core Team, 2018) using the lme4 package (version 1.1-18-1; Bates, Maechler, Bolker, & Walker, 2015). Data analysis files are available on the Open Science Framework

[https://osf.io/jm8ny/?view_only=11a2996136e24cd2b3257ce3b96bcd7f].

The effect of label condition on encoding and retention

The primary, pre-registered analysis investigated whether participants' performance depended on how many words they were taught (Label Condition, between-subjects) and time delay (Test Phase, within-subjects). All data except for the last two retention trials for the Four-Label condition were used in this analysis so that data were comparable across the Label Condition levels. Trial-by-trial accuracy was predicted in a logistic mixed-effects model with Label Condition and Test Phase as fixed effects and by-subject and by-item random intercepts. An offset corresponding to the logit of chance performance (0.25) was applied. Results showed that children's overall accuracy ($M = 0.49$, $SD = 0.30$), differed significantly from chance performance ($b = 1.08$, Wald 95% CI = [0.75, 1.40], Wald $X^2(1) = 42.81$, $p < 0.001$). There was a main effect of Label Condition ($b=0.69$, Wald 95% CI = [0.035, 1.34], Wald $X^2(1) = 4.27$, $p = 0.039$), such that participants in the Two-Label condition ($M = 0.56$, $SD = 0.31$) performed significantly better than in the Four-Label condition ($M = 0.42$, $SD = 0.28$). There was also a

main effect of Test Phase ($b = -0.83$, Wald 95% CI = $[-1.36, -0.29]$, Wald $X^2(1) = 9.30$, $p = 0.002$), such that participants performed better in the Immediate test ($M = 0.58$, $SD = 0.37$) than the Delay test ($M = 0.40$, $SD = 0.41$). There was no significant interaction between Label Condition and Test Phase ($b = -0.70$, Wald 95% CI = $[-1.74, 0.34]$, Wald $X^2(1) = 1.74$, $p = 0.19$; see Figure 2 for a graph of accuracy means and Figures 3 & 4 for an alternate form of data presentation). A model was run with age as an additional fixed effect to explore potential developmental differences. There was no significant main effect of age or any interaction of age with other variables in the model.

Figure 2. Performance on the four-alternative forced-choice task. The Immediate test was administered right after exposure, and the Delay test one week later. Error bars represent one standard error. The dashed line represents chance performance.

Figure 3. Performance in the Immediate test phase. Each bar represents the number of participants for a given Label Condition that got a certain number of trials correct (max = 2).

Figure 4. Performance in the Delay test phase (one week after exposure) for the words that were previously tested in the Immediate phase. Each bar represents the number of participants for a given Label Condition that got a certain number of trials correct (max = 2).

Exploratory follow-up tests were conducted via dummy coding to compare each cell to chance. Participants in the Two Label condition performed significantly above chance (0.25) for both the Immediate ($M = 0.69$, $SD = 0.39$; $b = 2.01$, Wald 95% CI = [1.39, 2.63], Wald $X^2(1) = 40.42$, $p < 0.001$) and Delay test phases ($M = 0.44$, $SD = 0.43$; $b = 0.83$, Wald 95% CI = [0.25, 1.41], Wald $X^2(1) = 7.93$, $p = 0.005$). Participants in the Four-Label condition, however, performed significantly above chance in the Immediate phase ($M = 0.47$, $SD = 0.32$; $b = 0.97$, Wald 95% CI = [0.40, 1.54], Wald $X^2(1) = 11.06$, $p < 0.001$), but not in the Delay phase ($M = 0.37$, $SD = 0.39$; $b = 0.50$, Wald 95% CI = [-0.09, 1.09], Wald $X^2(1) = 2.74$, $p = 0.10$).

The testing effect: Performance on previously tested versus untested items

A pre-registered analysis investigated a general testing effect on retention in the Four-Label condition. Specifically, a model was run to test whether Delay test performance in the Four-Label condition depended on whether the words has been tested during the Immediate test phase (within-subjects; recall that only two of the four words were tested in the Immediate phase so that the procedure matched that of the Two-Label condition.) Trial-by-trial accuracy was predicted in a logistic mixed-effects model with Retrieval History (tested vs. not tested) as a fixed effect, applying an offset corresponding to the logit of chance performance (0.25). The model included by-subject and by-item random intercepts. There was no significant effect of Retrieval History ($b = -0.44$, Wald 95% CI = [-1.19, 0.31], Wald $X^2(1) = 1.31$, $p = 0.25$). However, while trials for items that were not previously tested differed significantly from chance ($M = 0.46$, $SD = 0.41$; $b = 0.89$, Wald 95% CI = [0.27, 1.51], Wald $X^2(1) = 7.83$, $p = 0.005$), those that were previously tested did not ($M = 0.35$, $SD = 0.38$; $b = 0.44$, Wald 95% CI = [-0.31, 1.19], Wald $X^2(1) = 1.31$, $p = 0.25$).

Exploratory error analyses

We ran two exploratory analyses that were not pre-registered to further explore the errors that participants made in the study. First, we tested whether Delay test performance was related to Immediate test accuracy. In other words, we asked whether accuracy after the one-week delay was predicted by whether participants were accurate for that item at the immediate test. If this were the case, then it would suggest that word retention is strengthened by immediate recall. Note that for this analysis, only the first two retention trials could be used for participants in the Four Label condition since the second two trials tested items that were not tested before. A logistic mixed-effects model was run with trial accuracy as the outcome variable, Immediate test accuracy (yes vs. no) as a fixed effect, and subject and item as random intercepts (as in all models reported above). Immediate test accuracy was not a significant predictor of Delay accuracy, ($b = 0.29$, Wald 95% CI = $[-0.66, 1.24]$, Wald $X^2(1) = 0.36$, $p = 0.55$; see Table 1). Thus, there is no clear evidence that immediate test performance influenced delay test performance.

Table 1

Relation between the number of trials answered accurately versus inaccurately by test phase

		Delay test performance	
		Accurate	Inaccurate
Immediate test	Accurate	36 (.63)	45 (.54)
	Inaccurate	21 (.37)	38 (.46)
Total		57	83

Note. There were 140 trials total (2 for each of the 70 participants). Each cell reports the total number of trials (proportion, where the denominator is total number of delay trials at each accuracy level).

Second, we tested whether when participants were inaccurate in the Two-Label condition, they were more likely to have chosen the other labeled referent (this test was impossible in the Four-Label condition because all distractors had been labeled). If this is the case, it would suggest that there is a bias for participants to select any labeled object (as suggested by Axelsson & Horst, 2013). Of the 140 test trials (two trials from each testing phase), participants were inaccurate on 61. Of those 61, they chose the labeled distractor on 27, and an unlabeled distractor on 34. To test whether participants' distractor choice (labeled vs. unlabeled) on inaccurate trials differed from chance, a logistic mixed-effects model was run on the 61 inaccurate trials, with the outcome variable as which type of distractor they chose (labeled vs. unlabeled). Test Phase (Immediate vs. Delay) was included as a fixed effect, and subject and item as random intercepts (as in all models reported above). Participants' distractor choice did not differ from chance ($b = -0.11$, Wald 95% CI = $[-0.64, 0.43]$, Wald $X^2(1) = 0.15$, $p = 0.70$), and Test Phase did not predict distractor choice, ($b = -0.95$, Wald 95% CI = $[-2.02, 0.13]$, Wald $X^2(1) = 3.00$, $p = 0.08$). Thus, there is no clear evidence for a bias towards labeled objects at test.

Exploratory analyses of trial-by-trial performance

Lastly, we ran two analyses to explore children's performance across test trials in order to investigate contextual biases on referent selection¹. While the design of the study controlled for the familiarity of objects at test, it by definition did not control for how many objects at test had labels: in the Two-Label condition, two of the four choices at test had been labeled, while in the Four-Label condition, all choices had been labeled. Additionally, to control for recall history in the Delay condition, there were only two test trials in the Immediate test phase, even for those in the Four-Label condition. Together, this leads to a different testing context across the label conditions that potentially explains the Two-Label advantage reported above. Specifically, for

¹ Thank you to an anonymous reviewer for this suggestion.

children in the Two-Label condition, since only two referents had been given labels, as long as they knew one mapping, they could determine the other mapping by process of elimination during test. This was not the case for children in the Four-Label condition: knowing one mapping would still lead to three potential referents for any of the remaining labels.

To explore the possibility that children in the Two-Label condition received a boost from this reduced problem space, the same model that was run on the complete dataset was run again, but only using the first trial in each phase (Axelsson & Horst, 2013 ran this first-trial only analysis as well, and it is common in the word learning literature, although it was not pre-registered for the current study). Results showed that children's overall accuracy ($M = 0.58$, $SD = 0.31$), differed significantly from chance performance ($b = 1.47$, Wald 95% CI = [1.13, 1.83], Wald $X^2(1) = 67.46$, $p < 0.001$). There was a main effect of Test Phase ($b = -1.11$, Wald 95% CI = [-1.81, -0.40], Wald $X^2(1) = 9.41$, $p = 0.002$), such that participants performed better in the Immediate test ($M = 0.71$, $SD = 0.46$) than the Delay test ($M = 0.46$, $SD = 0.50$). However, unlike in the model with both test trials, there was no main effect of Label Condition ($b = 0.28$, Wald 95% CI = [-0.42, 0.99], Wald $X^2(1) = 0.62$, $p = 0.43$). There was no significant interaction between Label Condition and Test Phase ($b = -0.56$, Wald 95% CI = [-1.98, 0.85], Wald $X^2(1) = 0.62$, $p = 0.43$; see Figure 5 for a graph of accuracy means and Figure 6 for the distribution of choices.)

As with the full model, follow-up tests were conducted via dummy coding to compare each cell to chance. Participants in the Two-Label condition performed significantly above chance (0.25) for both the Immediate ($M = 0.77$, $SD = 0.43$; $b = 2.31$, Wald 95% CI = [1.53, 3.10], Wald $X^2(1) = 33.07$, $p < 0.001$) and Delay test phases ($M = 0.46$, $SD = 0.51$; $b = 0.93$, Wald 95% CI = [0.26, 1.59], Wald $X^2(1) = 7.46$, $p = 0.006$). Participants in the Four-Label

condition also performed significantly above chance in both the Immediate ($M = 0.66$, $SD = 0.48$; $b = 1.75$, Wald 95% CI = [1.05, 2.45], Wald $X^2(1) = 24.13$, $p < 0.001$), and Delay phases ($M = 0.46$, $SD = 0.51$; $b = 0.93$, Wald 95% CI = [0.26, 1.59], Wald $X^2(1) = 7.46$, $p = 0.006$).

Figure 5. Performance on the four-alternative forced-choice task for the first test trial only. The Immediate test was administered right after exposure, and the Delay test one week later. Error bars represent one standard error. The dashed line represents chance performance.

Figure 6. Distribution of choices in the first test trial. In the Two-Label condition, “Tested” refers to the other labeled object, while “Untested” refers to either of the two distractors. In the Four-Label condition, because all objects were labeled, “Tested” refers to the other object that they were tested on during the Immediate test phase; there is no difference between “Tested” and “Untested” objects in this condition.

To further explore the effects of testing context, second trial performance was analyzed. To allow for comparisons across the Two and Four-Label conditions, in this analysis, choices on the second trial were coded as “Correct”, “Tested”, or “Untested”. “Correct” refers to the correct answer; “Tested” refers to incorrect answers where the selection had been the correct answer on the first trial; and “Untested” refers to either of the two distractors. Note that in the Two-Label condition, the “Untested” object had not been given labels, while in the Four-label condition, they had. This coding scheme allows for the investigation of whether the observed differences between the Two- and Four-label condition performance can be attributed to the constrained

testing context created by the fact that a) only two objects had been labeled, and b) only two words were tested.

The data was split by first trial performance (Correct vs. Incorrect). Two chi-square tests were run to investigate whether second trial choice (Correct, Tested, or Untested), differed by label condition (Two vs. Four) for each part of the data. If participants' second trial choices were distributed equally, regardless of label condition and first trial performance, this would suggest that there was no advantage in the Two-Label condition provided by the constrained testing context. However, if participants were more likely to choose the Tested or Correct object on the second trial in the Two-Label condition only, this would suggest that the overall higher accuracy in this condition is not due (or not entirely due) to stronger word representations, but to the fact that the testing task was made easier by the constraints intrinsic to the experimental design.

When the participants were incorrect on the first trial, there was no significant difference between Label conditions, $\chi^2(2, N=58) = 0.16, p = 0.92, V=0.05$, suggesting that participants' second trial choices did not differ across label conditions when they were incorrect on the first trial. However, when participants were correct on the first trial, their performance was better in the Two-Label condition than in the Four-Label condition, $\chi^2(2, N=82) = 11.79, p = 0.003, V=0.38$. Participants in the Two-Label condition were 2.69 times more likely than those in the Four-Label condition to choose the correct referent on the second trial if they were correct on the first. Although there were not enough observations per cell to run a chi-squared test on the Immediate and Delay conditions separately, Figure 7 shows that this result was primarily driven by performance in the Immediate test phase. Together, the results provide some evidence that

having only two labeled objects makes the second trial easier because only participants in this condition can use their first trial choice to constrain their second.

Figure 7. Distribution of choices in the second trial separated by first trial accuracy. “Tested” refers to incorrect answers where the selection had been the correct answer on the first trial, while “Untested” refers to either of the two distractors.

Discussion

The current study investigated whether the number of labeled novel objects affects novel word retrieval, both immediately following exposure and after a one-week retention delay. In a four-alternative forced-choice (4AFC) task, participants (aged 3 to 6 years) performed significantly worse after a one-week delay regardless of how many objects had been labeled, in

line with previous work demonstrating a decay in memory for novel words across time (Vlach & Sandhofer, 2012). Participants additionally performed significantly worse when all four objects had been labeled (Four-Label condition) compared to when only two of the four objects had been labeled (Two-Label condition), and there was no interaction with Test Phase (Immediate vs. Delay). Analyses also revealed that while performance after a one-week delay was significantly greater than chance in the Two-Label condition, showing strong evidence of retention of the newly learned words, performance was not above chance after one week in the Four-Label condition.

Exploratory analyses of the patterns of choices across test trials, however, revealed that on the first trial in both the Immediate and Delay test phases, children's performance was statistically identical across the Two- and Four-Label. Moreover, participants in the Two-Label condition were more likely to choose the correct referent on the second trial if they were correct on the first trial, while participants in the Four-Label condition were not. Together, the results suggest that the reason that participants performed worse overall in the Four-Label condition was not due to differences in label-referent mapping strength, but rather due to the fact that participants in the Two-Label condition benefitted from contextual constraints in the testing phase.

Past work has found that two-year-olds perform better on a 4AFC task immediately after learning when one object is labeled compared to all four (Axelsson & Horst, 2013). This previous finding was mainly discussed as a methodological lesson, demonstrating that when word learning studies are designed such that only one novel object in a test set is labeled (e.g., Markson & Bloom, 1997), children's strong performance may only show that they can pick out the object that had been labeled. In other words, they may not have learned the specific label-

referent mapping at all, and instead simply learned with object has a name. Indeed, Axelsson and Horst (2013) found that when all four objects were labeled, children performed at chance during the 4AFC retrieval task.

The current study expands on this lesson, first finding that across all test trials, pre-school aged children also perform better at test when fewer objects have been labeled, demonstrating that older children are also affected by perceptual and pragmatic testing constraints. Notably, though, the children performed above chance even in the Four-Label condition, unlike the younger children in Axelsson and Horst (2013), suggesting that novel label-referent mappings are more robust for older children. This finding is analogous previous studies that have found stronger novel word encoding as children develop (e.g., Bion et al., 2013).

Additionally, while further studies are needed to replicate the exploratory analyses with more power, the current study provides new evidence for several mechanisms underlying the Two-Label condition boost. First, the findings suggest that participants in the Two-Label condition are not simply eliminating the un-labeled referents at test, as suggested by Axelsson and Horst (2013). In the exploratory error analysis, when participants in the Two-Label condition erred, they were not more likely to select the other labeled referent. Additionally, when only first trial performance was examined, there was no difference in performance between those in the Two- and Four-Label conditions; indeed, the distribution of choices on the first trial is remarkably similar (see Figure 6). If participants in the Two-Label condition could eliminate two of the four choices, they would perform better even on the first trial.

Second, there is evidence that participants have the pragmatic bias to not select the same referent twice (akin to a mutual exclusivity bias, but at the process rather than representation level). When participants in the Two-Label condition were correct on trial one, they were more

likely to select the correct (other) referent on trial two, and when they were incorrect on trial one, they were numerically more likely to choose the other labeled object than the same referent again. Third, there is evidence that this pragmatic bias interacts with the labeled object bias. If participants' choices were only influenced by a pragmatic bias not to select the same referent twice, participants across the label conditions should show a similar distribution of choices on the second trial; after all, in both conditions, there were three unchosen referents available in trial two. However, instead we see that participants in the Two-Label condition were more likely than those in the Four-Label condition to choose correctly on trial two. This must be due to the fact that they were biased to choose the other labeled referent (or a bias to not select an unlabeled referent). However, as noted above, this labeled object bias cannot be the only force at play, or the accuracy in the Two-Label condition would have been higher than in the Four-Label condition on the first trial, which was not evident. Future work should include the systematic manipulation of contextual variables, as well as computational modeling, to unpack how these biases interact.

The current study also expands on past work by testing novel word retrieval after a one-week delay. Testing novel word learning after a delay is crucial to our understanding of word learning. Most studies on word learning mechanisms only test novel word knowledge immediately after exposure, and yet in the real world, children must be able to retrieve their newly learned word days and weeks later, and studies that test retention often find changes in novel word representations across time (see Wojcik, 2013 for a review). Thus, our understanding of any word learning mechanism is incomplete without a retention test. In the current study, adding a retention phase broadly supports past experiments demonstrating the rapid decay of

novel word representations (e.g., Bion, Borovsky & Fernald, 2013; Gordon et al. 2016; Vlach & Sandhofer, 2012).

The results from the retention delay test present an intriguing potential mechanism behind the Two-Label boost in the current study. The distribution of responses on the first test trial of the retention delay test were nearly identical across the label conditions, showing a similar decrease in the number of correct choices from the immediate test phase, and a similar split of incorrect trials across tested and untested distractors (Figure 6). Responses in the retention delay test are also similarly distributed on the second trial (Figure 7). This similar performance in the second trial is notable, though, because in the immediate phase, performance across the label conditions diverged in this trial. The lack of a Two-Label performance boost at retention suggests that participants benefit from a pull towards labeled objects only immediately after training, when the labels were made salient. The current study does not have enough power to make strong conclusions based on the second trial performance by test phase, but the findings suggest that individual label-referent mappings are stronger than the more general encoding of which objects have a label at all. Unlike in the immediate test phase, after a week delay, children's performance at retention is remarkably similar across the label conditions on all testing trials, with no evidence for a label bias. There is analogous evidence of this type of memory specificity effect in phonology. Preschool-aged children's phonological representations of novel words are either remembered with specificity after a delay or forgotten; children do not retain a general representation of the word form over time (Gordon et al., 2016). These exploratory findings highlight how including a retention trial is crucial to teasing out the mechanisms behind word learning and retrieval, as they adds to previous work showing that

memory processes interact with contextual biases during word learning (Liittschwager & Markman, 1994).

By demonstrating that the number of novel labels presented during a learning moment affects later retrieval, the current results suggest that theories of word learning mechanisms must move beyond explaining how children map meaning onto individual words and instead take into account the full context surrounding word retrieval, including the number of labeled referents and time from first exposure. The results show that word representation and word processing are intimately connected. Both developmental scientists and people who interact with children in the real world routinely assume that when children can point to a ball, they have learned that word. It is simply in their lexicon. However, the present findings demonstrate that children's ability to identify the correct referent of a word does not necessarily show that their label-referent mapping is strong. If asked to recall a word, they can narrow the problem space using contextual factors, such as which objects have been labeled in the past, and these contextual factors with time. The current findings suggest that researchers should not study word learning in a vacuum, assuming that children are just accumulating a lexicon of strong representations. Instead, researchers should embrace the fact that young children's primary goal is to solve acute, in the moment communication problems; they rely on contextual features of a situation (perceptual or pragmatic) to efficiently, and seemingly correctly, use language even when their word representations are weak.

References

- Axelsson, E. L., & Horst, J. S. (2013). Testing a word is not a test of word learning. *Acta Psychologica*, 144(2), 264-268. <https://doi.org/10.1016/j.actpsy.2013.07.002>
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390-412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Baldwin, D. A. (1993). Infants' ability to consult the speaker for clues to word reference. *Journal of Child Language*, 20(2), 395-418. <https://doi.org/10.1017/S0305000900008345>
- Baldwin, D. A., & Markman, E. M. (1989). Establishing word-object relations: A first step. *Child Development*, 381-398. <https://doi.org/10.2307/1130984>
- Bates, D., Maechler, M., Bolker, B. Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1-48. doi:10.18637/jss.v067.i01.
- Bion, R. A., Borovsky, A., & Fernald, A. (2013). Fast mapping, slow learning: Disambiguation of novel word-object mappings in relation to vocabulary learning at 18, 24, and 30 months. *Cognition*, 126(1), 39-53. <https://doi.org/10.1016/j.cognition.2012.08.008>
- Booth, A. E. (2009). Causal supports for early word learning. *Child Development*, 80(4), 1243-1250. <https://doi.org/10.1111/j.1467-8624.2009.01328.x>
- Booth, A. E., Waxman, S. R., & Huang, Y. T. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology*, 41(3), 491. <https://doi.org/10.1037/0012-1649.41.3.491>
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175-191. <https://doi.org/10.3758/BF03193146>

- Gordon, K. R., McGregor, K. K., Waldier, B., Curran, M. K., Gomez, R. L., & Samuelson, L. K. (2016). Preschool children's memory for word forms remains stable over several days, but gradually decreases after 6 months. *Frontiers in Psychology*, 7, 1439. <https://doi.org/10.3389/fpsyg.2016.01439>
- Horst, J. S., & Hout, M. C. (2016). The Novel Object and Unusual Name (NOUN) Database: A collection of novel images for use in experimental research. *Behavior Research Methods*, 48(4), 1393-1409. <https://doi.org/10.3758/s13428-015-0647-3>
- Horst, J. S., & Samuelson, L. K. (2008). Fast mapping but poor retention by 24-month-old infants. *Infancy*, 13(2), 128-157. <https://doi.org/10.1080/15250000701795598>
- Liittschwager, J.C., & Markman, E. M. (1994). Sixteen- and 24-month-olds' use of mutual exclusivity as a default assumption in second-label learning. *Developmental Psychology*, 30(6), 955-968. doi:10.1037/0012-1649.30.6.955
- Markson, L., & Bloom, P. (1997). Evidence against a dedicated system for word learning in children. *Nature*, 385(6619), 813-815. <https://doi.org/10.1038/385813a0>
- Mather, E., & Plunkett, K. (2009). Learning words over time: The role of stimulus repetition in mutual exclusivity. *Infancy*, 14(1), 60-76. <https://doi.org/10.1080/15250000802569702>
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, 119(4), 831-877. <https://doi.org/10.1037/a0029872>
- Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, 17(2), 357-374. <https://doi.org/10.1017/S0305000900013817>
- R Core Team (2018). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

- Samuelson, L. K., Kucker, S. C., & Spencer, J. P. (2017). Moving word learning to a novel space: A dynamic systems view of referent selection and retention. *Cognitive Science*, 41, 52-72. <https://doi.org/10.1111/cogs.12369>
- Smith, L. B., & Samuelson, L. (2006). An attentional learning account of the shape bias: Reply to Cimpian and Markman (2005) and Booth, Waxman, and Huang (2005). *Developmental Psychology*, 42(6), 1339–1343. <https://doi.org/10.1037/0012-1649.42.6.1339>
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106(3), 1558-1568. <https://doi.org/10.1016/j.cognition.2007.06.010>
- Sullivan, J., & Barner, D. (2016). Discourse bootstrapping: Preschoolers use linguistic discourse to learn new words. *Developmental Science*, 19(1), 63-75. <https://doi.org/10.1111/desc.12289>
- Vlach, H., & Sandhofer, C. M. (2012). Fast mapping across time: memory processes support children's retention of learned words. *Frontiers in Psychology*, 3, 46. <https://doi.org/10.3389/fpsyg.2012.00046>
- Waxman, S. R., & Booth, A. E. (2000). Principles that are invoked in the acquisition of words, but not facts. *Cognition*, 77(2), B33-B43. [https://doi.org/10.1016/S0010-0277\(00\)00103-7](https://doi.org/10.1016/S0010-0277(00)00103-7)
- Wojcik, E. H. (2013). Remembering new words: integrating early memory development into word learning. *Frontiers in Psychology*, 4. <https://doi.org/10.3389/fpsyg.2013.00151>