I don't believe what you said before: Preschoolers retrospectively discount information from inaccurate speakers

Abstract

Children use speakers' past accuracy to make inferences about novel word meanings those individuals provide in the future. An open question is whether children can retrospectively re-evaluate information upon learning that the source was inaccurate. We addressed this question in two experiments, in which a speaker first introduced labels for novel objects and then revealed that she is either accurate or inaccurate in naming familiar objects. Experiment 1 showed that 3.5-6.5-year-olds displayed enhanced performance on a word knowledge test when they had learned novel words from a speaker who then showed herself to be an accurate labeler as opposed to an inaccurate one. Experiment 2 replicated this finding, but had a different speaker provide inaccurate label information. This manipulation did not affect learning, suggesting that children discount speakers and are not simply influenced by the demands of processing inaccurate information. Together, these results indicate that 3.5-6.5-year-olds continue to monitor the speakers' accuracy after learning new words from them, update their beliefs as accuracy data become available, and selectively retain words learned from speakers whom they deem to be epistemically competent.

I don't believe what you said before: Preschoolers retrospectively discount information from inaccurate speakers

Numerous studies have demonstrated that preschoolers learn the meaning of novel words selectively based on speakers' past accuracy (e.g., Clément, Koenig & Harris, 2004; Koenig, Clément, & Harris, 2004; Koenig & Harris, 2005; see Harris, 2012, for a review). Toddlers have selective learning capacities as well, both for the meaning of novel words (e.g., Brooker & Poulin-Dubois, 2013; Luchkina, Sobel & Morgan, 2018; Koenig & Woodward, 2010) and for nonverbal information (e.g., Over & Carpenter, 2012). By age 4, children not only selectively learn words from accurate over inaccurate speakers, but also take various social factors into consideration when making inferences about credibility (e.g., Kushnir, Wellman & Gelman, 2008; Nurmsoo & Robinson, 2009; see Sobel & Kushnir, 2013 for review).

These conclusions about children's selective learning capabilities center on procedures in which speakers' accuracy is established prior to the introduction of novel object labels. In most word learning situations, however, children rarely know speakers' reliability ahead of time. The question we ask here is whether children can discount novel linguistic information retrospectively upon learning that the source of knowledge is subsequently unreliable.

One previous study has investigated this question. Scofield and Behrend (2008, Experiment 2) showed that preschoolers could revise their judgments about novel word meanings after observing that the speaker who introduced them was subsequently inaccurate in naming familiar objects. However, only 10 of the 24 (42%) tested children in this study demonstrated this behavior. In a separate experiment (Experiment 3) in which children were retested after a delay, 22 of the 42 four-year-olds tested (52%) revised their judgments about the

label of a novel object based on subsequent information about the reliability of the informant.

These results suggest that children potentially revise their judgments about word labels based on information about a speaker's reliability, but these abilities might be nascent during the preschool or only available to a subset of the sample.

Note also that the children in Scofield and Behrend's (2008) study were presented with two informants. Both informants generated different novel labels for the novel object. One then accurately labeled familiar objects whereas the other inaccurately labeled those objects. An open question is whether children infer the reliability of each informant or whether they responded based on the contrast between these informants' information. Some research indicates that when children are only presented with a single informant, they do not show a selective preference for learning novel words from a previously accurate over a previously inaccurate informant (Vanderbilt, Heyman, & Liu, 2014). Thus, without the conflict between accurate and inaccurate labeling, children might not revise their inferences about novel labels based on the subsequent accuracy of informants.

In the present study, we built upon Scofield and Behrend's findings using a single informant paradigm. This allowed us to examine whether children could re-evaluate the novel label associated with a novel object after subsequently receiving information about the reliability of a speaker's knowledge. Our rationale for choosing a single informant design was to present children with a situation that provided minimal cues to potential informant unreliability. When two informants provide conflicting labels for the same object, children may have hypothesized that at least one of those informants was unreliable. Thus, they should pay close attention to her subsequent verbal behaviors. Our interest was to evaluate whether children spontaneously engage in accuracy-monitoring behavior and apply inferences from speakers' accuracy in

labeling to the previously taught information. Such a scenario allowed us to establish a more stringent test of children's active monitoring of the quality of their sources of information – that is, whether children keep track of informants' reliability when they have no particular reason to doubt it. This procedure was also a more realistic measure of how children might learn novel words; rarely do children encounter different informants providing different novel labels for the same object in quick succession.

In Experiment 1, we explored children's capacity to remember the referents of novel labels generated by informants who then subsequently reveal themselves to be accurate or inaccurate speakers. We also asked children to make inferences about the labels for other objects using a disambiguation paradigm (e.g., Markman & Wachtel, 1988). Children use the reliability of a speaker to make such linguistic inferences (Diesendruck, Carmel & Markson, 2010) and relate their inference to their retention of the label-object mapping (Sobel, Sedivy, Buchanan, & Hennessy, 2012).

We also considered a potential mechanism for children's retrospective capacities. Some researchers have argued that children encode but do not generalize novel information provided by unreliable speakers (e.g., Henderson, Graham, & Schell, 2015) or reject such information right away (e.g., Mangardich & Sabbagh, 2018; Sabbagh & Shafman, 2009; Sabbagh, Wdowiak, & Ottaway, 2003). On this interpretation, information is encoded but then discounted upon learning that the source is unreliable. Other researchers suggest that the increased demands of processing incorrect information affect children's ability to retain novel information provided by unreliable speakers (e.g., Barry-Anwar, 2016; Bernard, Proust, & Clément, 2014; Doebel, Rowell, & Koenig, 2016). This latter view suggests that hearing inaccurate labels for familiar objects potentially interferes with the processing of novel lexical information. This possibility

could also be extended to retrospective inferences by considering the interference in a retrospective manner (when speakers' accuracy is revealed after children learned new information from them).

We consider these alternatives in Experiment 2 by replicating our procedure from Experiment 1, but substituting a novel speaker who provides accurate or inaccurate labels for familiar objects in the place of the informant who labeled the novel objects. If information is encoded but discounted through establishing speaker reliability, children in Experiment 2 should retain the novel labels. If processing incorrect information interferes with encoding of novel labels, children should be unable to retain them.

Experiment 1

Experiment 1 extended a procedure developed by Scofield and Behrend (2008) to a group of 3.5-6.5-year-olds. Scofield and Behrend (2008) found that about half of the 4-year-olds they tested made retrospective inferences about speakers' reliability. In the current study, we included children of a broader age range to ensure the robustness of this capacity. Moreover, instead of contrasting information from accurate and inaccurate informants, we presented children with a single informant, who initially generated labels for novel objects, and then accurately or inaccurately labeled familiar objects. This ensured that children were reevaluating their belief about object labels based on subsequent information about an informant's reliability, rather than responding based on the conflict between informants. We also included two different measures of children's learning: their ability to retain the novel labels provided by the speakers, and their ability to make an inference about the labels of different novel objects, based on the principle of mutual exclusivity.

Participants

Forty-eight monolingual English-speaking children (24 per condition) aged 3.5-6.5 participated in the study ($M_{age}=60$ months; 19 girls and 29 boys; range: 42-81 months). While ethnicity data were not collected, the majority of children were Caucasian and represented middle- or high-SES backgrounds. Children were recruited and tested at a local children's museum and at the university laboratory. One additional child was tested but excluded from analyses due to the loss of interest in the task ($N=1$).

Apparatus

Novel label training and accuracy demonstration were video-recorded and presented on a MacBook 11" laptop computer screen to ensure minimal differences in children's exposure to the stimuli other than the experimental manipulation. Auditory stimuli were played through the laptop speakers set to the maximum volume. Test trials were conducted live by an experimenter who sat across the table from the child.

Materials

Six novel objects were used as test stimuli – a wooden hair pick, a folded vinyl soap holder, a white plastic pipe hook, a metal plug wrench, a yellow tarpaulin clip, and a green plastic drain cover (see Figure 1). The green drain cover and the metal plug wrench were the novel objects presented and labeled during novel label training. The remaining objects were first introduced during test trials. Three familiar objects were used during accuracy demonstration – a sneaker, a golden cardboard star, and a stuffed dog.

Procedure

The experimental procedure was designed and conducted in accordance with the approved protocol issued by the IRB at the authors' institution (Information Blinded for

Review). Children were tested in a quiet room located at the museum or at the laboratory. The experiment included three phases presented in a fixed order: *novel label training*, *accuracy demonstration*, and *test* and described in more detail below (see Figure 2).

In *novel label training*, children were first seated across the table in front of an experimenter who explained the task to them and started the video. While children watched the video, the Experimenter remained in the back of the room, so that the children could not see her. In each trial, a female actor on the video introduced two novel objects and labeled them “*lif*” and “*wug*” (“Look, a wug! Wug!”). Each object was labeled twice and manipulated twice, so that label-object pairs were each heard four times total. Label-object pairings and the order of object presentation were counterbalanced among participants.

Following novel label training, children were assigned to one of two conditions for the accuracy demonstration trials. Children in the *accurate speaker* condition saw the same actor from the novel label training phase introduce three familiar objects and label them accurately. Children in the *inaccurate speaker* condition saw the actor label the same objects inaccurately with familiar inaccurate labels (i.e., “cat”, “book”, “cup”). Each familiar label was mentioned twice (e.g., “Look, a cat! Cat!”). During each trial, the actor alternated between looking at the object and at the camera to emulate eye contact and to establish joint attention.

Immediately following the accuracy demonstration, the experimenter removed the laptop and sat across from the child to administer four *test* trials – two *retention* (one per novel label) and two *disambiguation* trials, presented in alternation. On each test trial, three objects – the two novel objects from the video and one unfamiliar novel object – were placed in front of the child. The unfamiliar novel object was different on every trial. On retention trials, the Experimenter invited children to choose the referent associated with the label used during the novel label

training trials (e.g., “Which one is the lif/wug?”). On disambiguation trials, the Experimenter invited children to choose the referent associated with an unfamiliar novel label (e.g., “Which one is the dax?”; “*dax*” and “*neem*” were used for these trials).

We included both retention and disambiguation trials to ensure that children’s test performance did not reflect a potential tendency to select visually familiar objects. Chance-level performance with three possible object choices is achieved when children select the novel distractor object 33.33% of the time. If children’s choices are driven by avoiding that object, then there is a possibility that even if children only select the correct object slightly more than 50% of the time, their performance can be evaluated as above chance. Such a conclusion may overestimate children’s ability to retrospectively evaluate the quality of information provided to them in the past. To avoid making such an error, we included disambiguation trials, on which children had to suppress a potential impetus to select previously demonstrated objects and infer that the novel label must stand for the unnamed object. Children were not provided any feedback on their responses other than a neutral acknowledgement, “Ok, this one.” The location of the target object was counterbalanced across trials.

Results

Responses on test trials were coded as correct (children chose the previously-labeled object on retention trials or the novel object on disambiguation trials) or incorrect. We conducted preliminary analyses of correct responses by constructing a Generalized Estimating Equation (GEE) with a binary logistic distribution, to explore effects of age and gender. The overall model was not significant, nor were these factors, Wald $\chi^2(1, N = 48) = 2.74$ and Wald $\chi^2(1, N = 48) = 0.70$, p -values $> .05$, respectively. These variables were excluded from further analyses.

Our main analysis focused on constructing a GEE on correct responding as the dependent measure, assuming a binary logistic distribution. Condition (accurate speaker, inaccurate speaker) was a between-subject factor. Trial type (disambiguation, retention) and trial number (first, second) were within-subject factors. This analysis revealed a significant effect of Condition, Wald $\chi^2(1, N = 48) = 7.45, p = .006$. Whereas children in the accurate speaker condition responded correctly on 78% of the trials, children in the inaccurate speaker condition responded correctly on only 55% of the trials. There was a significant effect of trial type, Wald $\chi^2(1, N = 48) = 4.58, p = .032$, with participants performing better on disambiguation trials (74% correct) than on retention trials (60%). The effect of trial number was also significant, Wald $\chi^2(1, N = 48) = 4.26, p = .039$, with children responding correctly on 63% of first test trials and 72% of second test trials.

Because the significant effect of trial number suggested that performance on the second retention and disambiguation trials was affected by children's inferences on the first trials, we re-ran all analyses only for the first trials. In this analysis, the effect of condition remained significant, Wald $\chi^2(1, N = 48) = 5.61, p = .018$, but the effect of trial type was no longer significant.

We also compared performance on the first retention and disambiguation trial to chance responding (33.33%, see Figure 3). In the inaccurate speaker condition, performance on the disambiguation trial was not significantly different from chance (43%), Binomial Test, $p = .21$ whereas on the retention trial, children chose the target objects significantly more often than chance (57%), Binomial Test, $p = .02$. In the accurate speaker condition, children chose the target object significantly more often than chance both on the disambiguation (68%) and retention (80%) trials, Binomial Tests, both p -values $< .001$.

Discussion

By the preschool years, children use speaker accuracy information to predict the reliability of future information provided by the same speakers. Experiment 1 suggests that they also evaluate previously learned information retrospectively. Consistent with Scofield and Behrend (2008), these results suggest that children track speakers' reliability continuously, recognize the relevance of such information to their beliefs about word meanings, and update such beliefs upon receiving information about speakers' reliability.

However, to justify this conclusion, we must consider alternative explanations of performance in the inaccurate speaker condition. Instead of retrospectively discounting the information provided by the unreliable source, children might have accepted novel labels provided by accurate and inaccurate speaker equivalently but were then affected by memory demands associated with processing inaccurate labels (Koenig & Doebel, 2013); hearing incorrect labels could have interfered with children's retention of novel words.

One piece of evidence is inconsistent with this possibility: we observed a difference in the types of errors children made on retention trials. Participants in the accurate speaker condition chose the distractor object 20% of the time when they made an error (4% of all responses), whereas the proportion of distractor choices was 70% in the inaccurate speaker condition (30% of all responses; see Figure 4). This difference suggests that most children thought that the inaccurate speaker had not simply mixed up the object-label pairings but that she had provided completely invalid information. Experiment 2, however, directly tests this possibility.

Experiment 2

Experiment 2 replicated the procedure used in Experiment 1 with one change. Novel label training was presented by one speaker (Speaker 1) and accuracy demonstration was presented by a different speaker (Speaker 2) instead of by the original speaker (as in Experiment 1). If the difference between the accurate and inaccurate conditions reflected children's retrospective inferences, then this manipulation would eliminate the effect of condition because Speaker 2's accuracy has no implications for Speaker 1's reliability. Alternatively, if the effect of condition in Experiment 1 was driven by the increased memory demands, the effect of condition should remain as children must process the equivalent inaccurate information.

Participants

Fifty (25 per condition) monolingual English-speaking 3.5-6.5-year-olds participated in the experiment ($M_{\text{age}} = 62.37$ months, 27 girls and 23 boys, age range: 41-83 months). As in Experiment 1, the majority of children were Caucasian and represented middle- or high-SES backgrounds. Children were recruited and tested at a local children's museum and the laboratory.

Apparatus and Materials

The apparatus and materials were the same as in Experiment 1, with the exception of the materials used in the accuracy demonstration phase. The original accuracy demonstration video-clip was replaced with a clip featuring a different female actor labeling familiar objects accurately or inaccurately. The objects and the labels used in the video were the same as in Experiment 1.

Procedure

The procedure was identical to the procedure in Experiment 1 (see Figure 2), with two exceptions. First, as mentioned above, two actors were used: one in the novel object training

phase, and a second actor in the accuracy demonstration phase. Second, to ensure that children attended to the change of actor between the two phases, we included the following control question: “Do you remember the video you just saw? How many ladies did you see?”

Results and Discussion

Eight out of 50 children answered the control question about the number of speakers in the video incorrectly. Their data was excluded from analyses¹.

As in Experiment 1, preliminary analyses via GEE showed no effects of gender and age on children’s choice of object, Wald $\chi^2(1, N = 42) = 0.79$ and Wald $\chi^2(1, N = 42) = 1.68$, both p -values $> .05$. These variables were not considered further. For the main analysis, we constructed a Generalized Estimating Equation with a binary logistic distribution, with choice of object as a dependent variable. Condition (accurate speaker, inaccurate speaker) was a between-subject factor. Trial type (disambiguation, retention) and Trial number (first, second) were within-subject factors. This analysis revealed no significant effects of Trial type or Condition, Wald $\chi^2(1, N = 42) = 0.43$ and Wald $\chi^2(1, N = 42) = 0.04$, both p -values $> .05$. As in Experiment 1, there was a significant effect of Trial number, Wald $\chi^2(1, N = 42) = 3.84$, $p = .05$, with children responding correctly 63% of the time on first trials and 75% on the time on second trials.

Given this result, as in Experiment 1, we reran our analyses on the first trials only. These analyses revealed a significant effect of Trial type, Wald $\chi^2(1, N = 42) = 9.09$, $p = .003$, with children responding correctly on 78% of retention trials and 47% of disambiguation trials. There

¹ Children who responded to the control question incorrectly performed on the word knowledge test like children in Experiment 1 and displayed a significant effect of condition, Wald $\chi^2(1, N=8) = 69.802$, $p < .001$, and a significant effect of trial type, Wald $\chi^2(1, N = 8) = 133.854$, $p < .001$. Children in the accurate speaker condition chose the target object 100% of the time on retention trials and 80% of the time on disambiguation trials. For children in the inaccurate speaker condition, the numbers were 67% and 33% for retention and disambiguation trials respectively. Although these analyses are on a small sample, they speak to the necessity of recognizing that the second speaker was different from the first, and that children are tracking informants’ reliability over time.

was no effect of Condition on children's choice of objects during test, Wald $\chi^2(1, N = 42) = 0.60$, $p > .05$.

We also compared performance on the first retention and disambiguation trials to chance. On the retention trial, children responded differently from chance on both the accurate speaker (70% correct) and inaccurate speaker conditions (86% correct), Binomial tests, both p -values $< .001$ (see Figure 5). On the disambiguation trial, children in both conditions did not select the target object more often chance (children responded at 50% and 45% correct for the accurate and inaccurate conditions respectively) Binomial tests, $p = .07$ and $p = .16$. The difference in performance on the two trial types suggests that disambiguation inferences were substantially more difficult for children than the retrieval of the novel word-object pairings.

General Discussion

We investigated children's ability to make retrospective inferences about word meanings learned from speakers who subsequently revealed their accuracy or inaccuracy in naming familiar objects. Experiment 1 replicated and extended Scofield and Behrend's (2008) findings using a single informant. When children learned novel labels from speakers who later mislabeled familiar objects, children were less likely to retain the object-label pairing used by that speaker on test trials than they were when they learned novel labels from speakers who named familiar objects correctly.

Two patterns of results are worth noting. First, although in Experiment 1 children in the inaccurate speaker condition performed significantly above chance on retention trials, there was a notable difference between the conditions (57% vs. 80%). It is possible that in the absence of inaccurate labeling episodes, performance on these questions was supported both by episodic and

semantic memory, so that children remember the labeling episodes and form semantic links between novel objects and novel labels. When children witnessed inaccurate labeling following novel label training, they might engage in semantic blocking (e.g., Sabbagh & Shafman, 2009) and their above chance performance would be driven by a residual event-based memory of a labeling event (Henderson, Graham, & Schnell, 2015).

Second, unlike in Experiment 1, performance on disambiguation trials in Experiment 2 did not deviate from chance levels in both conditions (50% and 45% for the accurate and the inaccurate condition respectively). We hypothesize that this effect was driven by the novelty of information presented during familiarization in Experiment 2. This result also may be related to the distinction between episodic and semantic memory. When children were first exposed to novel labels introduced by one person and then observed a different person label familiar objects, the novelty of the social context may have interfered with the formation of episodic memory (due to the need to store a new event), more than the formation of a novel semantic memory. In this case, when children engaged in the disambiguation inference, they might not have had access to the heuristic that requires remembering a labeling event (i.e., these objects have been labeled, therefore, the unfamiliar label must refer to the unfamiliar object). If so, in making disambiguation inferences, children must rely on the process of elimination that is based on comparing the novel label with the labels of the previously labeled objects. Such a process may have required substantially more working memory resources and thus resulted in a higher proportion of errors.

Overall, the results of Experiments 1 and 2 establish that by 3.5-6.5 years, children are not just using speakers' history of accuracy to make predictions about their future information.

They are flexibly using accuracy information to re-evaluate their beliefs about the past information provided by those speakers.

The lack of difference between the accurate and inaccurate speaker conditions in Experiment 2 suggests that the effect of condition we observed in Experiment 1 was not explained by the increased memory demands associated with processing incorrect labels in the inaccurate speaker condition. Rather, we suggest that children made retrospective inferences about the novel information provided by those speakers upon receiving subsequent information about the speaker's reliability.

The pattern of errors in Experiment 1 also supports this interpretation. If incorrect responses were related to increased memory demands for processing incorrect information, children should have been equally likely to confuse object-label pairings and to choose the distractor object during test. Children, however, were significantly more likely to choose the distractor. It is likely that their choice of object was informed by retrospective inferences (the speaker is unreliable and must have provided wrong information, thus a completely unfamiliar object is more likely to be the correct referent of the label than either of the previously labeled objects).

One concern with this conclusion is that children might not have made retrospective inferences about the quality of the novel labels provided by the speaker, but instead form domain-general retrospective distrust in an informant's knowledge. That is, instead of making an inference about the speaker's reliability in the verbal domain based on her accuracy or inaccuracy in labeling, children may have made a broader valence inference about the speaker's competence and concluded that she cannot be relied on in more general terms. While our experimental design does not distinguish between these possibilities, there is limited evidence

that children do make sweeping inference about informants' competence (e.g., "halo" or "pitchfork" effects). But in the cases where children do demonstrate halo effects (preferring stronger, informants to weaker ones, kinder to meaner ones, etc.), the informants were puppets rather than adult human beings (e.g., Brosseau-Liard & Birch, 2010; Fusaro, Corriveau & Harris, 2011). When human adults acted as informants, preschoolers did not demonstrate halo effects (e.g., Koenig & Jaswal, 2011; Sobel & Corriveau, 2010). Given that in our studies, children observed adult human actors and, most likely, had previously formed an association between adults and competence, it is unlikely that by observing three instances of mislabeling, they inferred domain-general incompetence or unreliability of the speaker. Future research should investigate this theoretical possibility.

Our findings from Experiment 1 are broadly consistent with the idea that selective word learning in preschool years might be driven by semantic blocking (Mangardich & Sabbagh, 2018). This process describes the suppression of establishing semantic mapping between a label and its referent, if the label was provided by an unreliable informant, while retaining episodic memory of the labeling event. Children's depressed test performance in the inaccurate speaker condition suggests that despite equivalent conditions of the initial encoding of novel labels, subsequent inaccuracy in labeling interfered with semantic consolidation. Chance-level performance on the disambiguation trials, however, suggests that semantic blocking alone cannot fully explain our findings. If children's retrospective learning was driven by a process by which semantic consolidation but not episodic memory was interfered by the exposure to inaccurate labeling, we would expect to observe above chance performance on disambiguation trials in the inaccurate speaker condition. Despite failed semantic consolidation, children should still remember that two out of three objects have already been labeled and, therefore, would have

selected the object corresponding to the unfamiliar label. The observed chance-level performance on these trials suggests that children did not simply block recently learned semantic mappings upon observing their source mislabel familiar objects, but instead allowed for the possibility that the previously labeled objects may have different names.

The present results are also consistent with children's general capacity to retrospectively revise their beliefs. When presented with initially ambiguous data, preschoolers can integrate new data to come to appropriate causal conclusions (e.g., Bonawitz et al., 2008; Griffiths et al., 2011; Sobel et al., 2004). Moreover, the present results are consistent with the "wait and see" strategy for mutual exclusivity (e.g., Savage & Au, 1996). Children potentially keep in mind possible labels for an object for a period of time, until the information that allows them to revise their hypotheses becomes available. As Scofield and Behrend (2008) correctly point out, children often make these inferences from speakers with whom they have no experience (e.g., Markson & Bloom, 1997) and such a strategy is consistent with previous findings suggesting that children are influenced by pragmatic cues generated by the speakers (Diesendruck et al., 2010; Sobel et al., 2012).

To conclude, our results highlight the role of children's active involvement in their selective learning and add to the growing body of literature on the development of its underlying mechanisms. Earlier work shows that children as young as 18 months already incorporate social cues to their selective learning (e.g., Luchkina, Sobel, & Morgan, 2018). By 3-4 years, children's selective learning is predominantly driven by inferences about speakers' epistemic competence but there is still a residual effect of associative generalizations (Hermes, Behne, & Rakoczy, 2018; Luchkina, Morgan, & Sobel, 2019). The present findings suggest that by the time children

are 3.5-6.5 years old, they make inferences about epistemic competence both prospectively and retrospectively.

References

- Barry, R. A. (2016). *Selective word learning in toddlers: An eye tracking investigation of the mechanism* (Doctoral dissertation). Retrieved from ProQuest Dissertations Publishing. (Accession No. 10165848).
- Bernard, S., Proust, J., & Clément, F. (2014). The medium helps the message: Early sensitivity to auditory fluency in children's endorsement of statements. *Frontiers in Psychology, 5*, 1412. <https://doi.org/10.3389/fpsyg.2014.01412>
- Bonawitz, E. B., Fischer, A., & Schulz, L. (2008, January). Training a Bayesian: Three-and-a-half-year-olds' Reasoning about ambiguous evidence. *In Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 30, No. 30). <https://escholarship.org/uc/item/5mx3g61x>
- Brooker, I., & Poulin-Dubois, D. (2013). Is a bird an apple? The effect of speaker labeling accuracy on infants' word learning, imitation, and helping behaviors. *Infancy, 18*(s1), E46–E68. <https://doi.org/10.1111/infa.12027>
- Brosseau - Liard, P. E., & Birch, S. A. (2010). 'I bet you know more and are nicer too!': what children infer from others' accuracy. *Developmental Science, 13*(5), 772-778. <https://doi.org/10.1111/j.1467-7687.2009.00932.x>
- Clément, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*(4), 360-379. <https://doi.org/10.1111/j.0268-1064.2004.00263.x>
- Diesendruck, G., Carmel, N., & Markson, L. (2010). Children's sensitivity to the conventionality of sources. *Child Development, 81*(2), 652-668. <https://doi.org/10.1111/j.1467-8624.2009.01421.x>

- Doebel, S., Rowell, S. F., & Koenig, M. A. (2016). Young children detect and avoid logically inconsistent sources: The importance of communicative context and executive function. *Child Development*, 87(6), 1956-1970. <https://doi.org/10.1111/cdev.12563>
- Fusaro, M., Corriveau, K. H., & Harris, P. L. (2011). The good, the strong, and the accurate: Preschoolers' evaluations of informant attributes. *Journal of Experimental Child Psychology*, 110(4), 561-574. <https://doi.org/10.1016/j.jecp.2011.06.008>
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*, 35(8), 1407-1455. <https://doi.org/10.1111/j.1551-6709.2011.01203.x>
- Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge, MA: Harvard University Press. <https://doi.org/10.4159/harvard.9780674065192>
- Henderson, A. M., Graham, S. A., & Schell, V. (2015). 24-month-olds' selective learning is not an all-or-none phenomenon. *PloS one*, 10(6), e0131215. <https://doi.org/10.1371/journal.pone.0131215>
- Hermes, J., Behne, T., & Rakoczy, H. (2018). The Development of Selective Trust: Prospects for a Dual-Process Account. *Child Development Perspectives*, 12(2), 134-138. <https://doi.org/10.1111/cdep.12274>
- Heyes, C., & Pearce, J. M. (2015). Not-so-social learning strategies. *Proceedings of the Royal Society of London B: Biological Sciences*, 282(1802), 20141709. <https://doi.org/10.1098/rspb.2014.1709>
- Jaswal, V. K., Pérez-Edgar, K., Kondrad, R. L., Palmquist, C. M., Cole, C. A., & Cole, C. E. (2014). Can't stop believing: Inhibitory control and resistance to misleading testimony. *Developmental Science*, 17(6), 965-976. <https://doi.org/10.1111/desc.12187>

- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76(6), 1261-1277. <https://doi.org/10.1111/j.1467-8624.2005.00849.x>
- Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, 15(10), 693.5-6.598. <https://doi.org/10.1111/j.0956-7976.2004.00742.x>
- Koenig, M. A., & Doebel, S. (2013). Children's understanding of unreliability. *Navigating the social world: What infants, children, and other species can teach us*, 235-240.
- Koenig, M. A., & Jaswal, V. K. (2011). Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Development*, 82(5), 1634-1647. <https://doi.org/10.1111/j.1467-8624.2011.01618.x>
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition*, 107(3), 1084-1092. <https://doi.org/10.1016/j.cognition.2007.10.004>
- Luchkina, E., Sobel, D. M, Morgan, J. L. (2018). Eighteen-month-olds selectively generalize words from accurate speakers to novel contexts. *Developmental Science*, e12663 <https://doi.org/10.1111/desc.12663>
- Luchkina, E., Morgan, J. L., Sobel, D. M. (2018). *Questions answer questions about the mechanism of selective word learning*. Manuscript submitted for publication. Brown University.
- Mangardich, H., & Sabbagh, M. A. (2018). Mechanisms of Selective Word Learning: Evidence and Implications. In *Active Learning from Infancy to Childhood* (pp. 147-166). Springer, Cham. https://doi.org/10.1007/978-3-319-77182-3_9

- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, 20(2), 121-157. [https://doi.org/10.1016/0010-0285\(88\)90017-5](https://doi.org/10.1016/0010-0285(88)90017-5)
- Markson, L., & Bloom, P. (1997). Evidence against a dedicated system for word learning in children. *Nature*, 385(6619), 813. <https://doi.org/10.1038/385813a0>
- Nurmsoo, E., & Robinson, E. J. (2009). Children's trust in previously inaccurate informants who were well or poorly informed: When past errors can be excused. *Child Development*, 80(1), 23-27. <https://doi.org/10.1111/j.1467-8624.2008.01243.x>
- Over, H., & Carpenter, M. (2012). Putting the social into social learning: explaining both selectivity and fidelity in children's copying behavior. *Journal of Comparative Psychology*, 126(2), 182. <http://psycnet.apa.org/doi/10.1037/a0024555>
- Sabbagh, M. A., & Shafman, D. (2009). How children block learning from ignorant speakers. *Cognition*, 112(3), 415-422. <https://doi.org/10.1016/j.cognition.2009.06.005>
- Sabbagh, M. A., Wdowiak, S. D., & Ottaway, J. M. (2003). Do word learners ignore ignorant speakers? *Journal of Child Language*, 30(4), 905-924. <https://doi.org/10.1017/S0305000903005828>
- Savage, S. L., & Au, T. K. F. (1996). What word learners do when input contradicts the mutual exclusivity assumption. *Child Development*, 67(6), 3120-3134. <http://dx.doi.org/10.2307/1131770>
- Scofield, J., & Behrend, D. A. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development*, 23(2), 278-290. <https://doi.org/10.1016/j.cogdev.2008.01.003>

- Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, 120(4), 779.
<http://psycnet.apa.org/doi/10.1037/a0034191>
- Sobel, D. M., Sedivy, J., Buchanan, D. W., & Hennessy, R. (2012). Speaker reliability in preschoolers' inferences about the meanings of novel words. *Journal of child language*, 39(1), 90-104. <https://doi.org/10.1017/S0305000911000018>
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28(3), 303-333. https://doi.org/10.1207/s15516709cog2803_1
- Vanderbilt, K. E., Heyman, G. D., & Liu, D. (2014). In the absence of conflicting testimony young children trust inaccurate informants. *Developmental Science*, 17(3), 443-451.
<https://doi.org/10.1111/desc.12134>

Figure 1. Novel objects used in novel label training and on test trials in Experiments 1–2

Figure 2. Procedure in Experiments 1–2 (depicted here is the inaccurate speaker condition)

Experiment 1

Novel Label Training (video)

Look! A lif!
Lif!

Look! A wug!
Wug!

Accuracy Demonstration (video)

Look! A
cup! Cup!

Look! A book!
Book!

Test (live)

Disambiguation trial:
Which one is the dax?

Retention trial:
Which one is the wug?

Experiment 2

Novel Label Training (video)

Look! A lif!
Lif!

Look! A wug!
Wug!

Accuracy Demonstration (video)

Look! A
cup! Cup!

Look! A book!
Book!

Test (live)

Disambiguation trial:
Which one is the dax?

Retention trial:
Which one is the wug?

Different
person

Figure 3. Children's test performance on first trials in Experiment 1.

Figure 4. Proportions of responses on the first retention trial in Experiment 1.

Figure 5. Performance on first test trial in Experiment 2.