Questions can answer questions about mechanisms of preschoolers' selective word learning

Abstract

This paper examined how inferences about epistemic competence and generalized labeling errors influence children's selective word learning. Preschoolers ($N=128$) learned words from informants who asked questions about objects, mentioning either correct or incorrect labels. Such questions do not convey stark differences in informants' epistemic competence. Inaccurate labels, however, generate error signals that can lead to weaker encoding of novel information. Preschoolers retained novel labels from both informants but were slower to respond in the inaccurate labeler condition. When the test procedure was not sensitive to the strength of information encoding, children performed above chance in both conditions, and their response times did not differ. These results suggest that epistemic-level inferences and error generalizations influence preschoolers' selective word learning concurrently.

Questions can answer questions about mechanisms of preschoolers' selective word learning

In order to use language effectively, children must learn correct meanings of words. Because word meanings are arbitrary and conventional (de Saussure, 1916/1966), children must use others as sources of information. Not all conversational partners, however, are equally reliable sources of lexical knowledge. Children must evaluate speakers' reliability and decide whether to accept new information from them. There is now significant evidence that preschoolers can use others' accuracy in naming to learn selectively from different sources (e.g., Clément, Koenig & Harris, 2004; Koenig, Clément, & Harris, 2004; Koenig & Harris, 2005; see Harris, 2012, for a review). An open question, however, is how children use others' accuracy in their selective learning. In this work, we examine the relative contributions of different mechanisms that employ speakers' accuracy in preschoolers' selective word learning.

One mechanism that can drive selective word learning in toddlers and young preschoolers is differential encoding of novel information, which is related to children's exposure to inaccuracy in labeling. For example, when children hear a speaker say, "This is a dog" while being shown a picture of a cat, the mismatch between the label and their expectation produces a particular neural signature (e.g., Friedrich & Friederici, 2004; Mangardich & Sabbagh, 2018; Sheehan, Namy, & Mills, 2007). If children associate enough of these error signals with an informant, they might register other information from that informant as inaccurate via association (e.g., Heyes & Pearce, 2015; Lucas & Lewis, 2010). That generalization results in weakening the encoding of novel information generated by that individual. This would result in inferior performance in learning novel words from that informant. Importantly, this mechanism

does not involve any reflection on the quality the novel information or any epistemic-level inferences about the informant.

Another mechanism that can underlie young children's selective word learning is explicit inferences about speakers' epistemic competence. Upon observing instances of mislabeling, children may integrate such an observation with other pieces of their knowledge about that speaker, such as expertise (e.g., Koenig & Jaswal, 2011; Sobel & Corriveau, 2010), knowledge base (e.g., Jaswal & Neely, 2006; Vanderborcht & Jaswal, 2009), epistemic uncertainty (e.g., Sabbagh & Baldwin, 2001), or reasoning style (e.g., Koenig, 2012). Based on such an integration, children can then make inferences about the speaker's reliability as a source of word meanings and either accept or discount novel object-label mappings provided by that speaker (e.g., Nurmsoo & Robinson, 2009).

While infants as young as 8 months can track informants' reliability and generalize their errors to future information (Tummeltshammer, Wu, Sobel, & Kirkham, 2014) and toddlers as young as 19 months were shown to register semantic errors (Friedrich & Friederici, 2004; Koenig & Echols, 2003), experimental data suggest that children begin making epistemic-level inferences only by the preschool years (e.g., Harris, Koenig, Corriveau, & Jaswal, 2018; Mills, 2013; Sobel & Kushnir, 2013). Despite that, toddlers as young as 24 months have been shown to learn words selectively from speakers who name familiar objects correctly (e.g., Koenig & Woodward, 2010; Luchkina, Sobel, & Morgan, 2018). These findings suggest that selective word learning might be driven by different mechanisms in toddlerhood and in preschool years, with epistemic-level inferences taking precedence over simpler error-tracking mechanisms as children' ability to integrate multiple pieces of information develops with age. The question we are considering is whether and how these mechanisms – (1) generalization of errors that results

in weaker encoding of novel information and (2) inferences about speakers' epistemic competence – co-exist.

One possibility is that in infancy, selective learning is governed predominantly by generalized errors. Later, as children acquire more knowledge, this mechanism is *replaced* by one that integrates such errors with other forms of knowledge. On this view, speakers' errors would have no influence on preschoolers' selective word learning if such errors couldn't inform epistemic-level inferences.

Another possibility is that these two systems function *in parallel* during the preschool years. The effects of errors on novel label encoding might continue to influence preschoolers' word learning implicitly as a part of a heuristic-based mechanism (i.e. System 1, e.g., Kahneman, 2011). At a more explicit level, preschoolers' selective learning is driven by inferences about speakers' epistemic competence, which are based on integrating multiple pieces of knowledge about speakers, such as their accuracy, the domain of knowledge in question, social context, etc. (i.e., System 2). The hypothesis that these two mechanisms operate in parallel is consistent with recent proposals that selective learning in general is best described by a dual-process model (Hermes, Behne, & Rakoczy, 2018). The goal of the present investigation is to tease apart these alternatives and determine how these two mechanisms contribute to preschoolers' selective word learning.

A way to dissociate between the effects of weaker encoding caused by generalized errors and inferences about speakers' epistemic competence is to have speakers produce utterances that preserve the difference in speaker-label-accuracy associations, but offer little evidence for a difference in epistemic competence. Unlike generating statements about an object's label, asking questions about an object's label can produce an error signal in labeling without conveying truth

value information that would signal differences in epistemic competence. If one speaker asks a question about an object label using the appropriate label while another speaker asks a question about an object label using an inappropriate label, hearing the latter speaker would register the mismatch error, but not produce a difference in beliefs about the epistemic competence of the two speakers. When speakers ask questions with appropriate or inappropriate labels, there is no explicit evidence that they truly believe those are the labels of the demonstrated objects.

Speakers' questions might indicate guesses about the object labels, but not a commitment to a certain object-label mapping. In each of these cases, the act of asking such questions does not provide strong evidence of a difference in the speakers' knowledge of object labels.

If children's selective word learning is only influenced by inferences about the speakers' epistemic knowledge, there should be few differences in how children learn from speakers who use appropriate and inappropriate labels in their questions. But, if there is a mismatch between the label and the object, children may register and potentially generalize an error signal and associate that error signal with the speaker. If these error signals affect children's learning beyond informing epistemic-level inferences, there should be a measurable difference in how well children learn novel labels from the two informants.

Luchkina et al. (2018) contrasted 18-month-olds' selective word learning from speakers who demonstrated their differing accuracy by making statements about familiar objects (e.g., "This is a star") and from speakers who asked questions (e.g., "Is this a star?"), mentioning either correct or incorrect labels. After familiarization, speakers made statements about the labels of novel objects. Eighteen-month-olds learned from the speaker who made accurate statements, but not the one who made inaccurate statements. Eighteen-month-olds did not learn from either question-asker. This suggests that toddlers' selective learning is not influenced exclusively by

the effects of errors in labeling on the encoding of novel labels. This work, however, did not dissociate between the contributions of different mechanisms of selective word learning. It is not clear whether the effects of errors on novel label encoding influenced toddlers' selective learning beyond informing epistemic-level inferences.

In the present study, we presented 3-4-year-olds with the same familiarization and novel label training events as in Luchkina et al. (2018), but modified the test phase to use a choice task as a response across both experiments, instead of the intermodal preferential looking paradigm (Golinkoff, Hirsh-Pasek, Cauley, & Gordon, 1987) used by Luchkina et al. with 18-month-olds. Our initial goal was to test children between the ages of 19 and 60 months in both experiments, but pilot research with children aged 19-35 months ($N = 64$), indicated that this group struggled with the procedure we used in Experiment 2. Further, this age group often refused to complete the experimental task, which resulted in a greater than 30% attrition rate. As a result, we concentrated on 3-4-year-olds in the present study.

In Experiment 1, we showed children videos of two speakers, each of whom asked questions about the labels of familiar objects. One used the appropriate labels for the objects in their questions; the other used inappropriate labels. One of these two speakers then labeled novel objects with novel labels. We then presented children with one of the novel objects from the video, and paired with a previously unseen, novel distractor object. We asked children to identify the object corresponding to one of the labels from the video. In Experiment 2, we replicated this procedure, presenting children with the same pairs of objects, but asked them to identify an object with a completely novel label (i.e., a disambiguation procedure; Merriman & Schuster, 1991). In both experiments, we recorded children's choice of object and response time.

Response time offers a more subtle measure of the process by which children learn selectively, as it has been found to reflect the effect of the strength of information encoding (e.g., Sternberg, 1969). If children make inferences about speakers' epistemic competence, then in the absence of any explicit evidence of their beliefs about object labels, they should consider the two speakers equally-accurate sources of word knowledge. However, if selective learning is also influenced by the strength of novel label encoding that results from the history of labeling errors, children who are asked to identify the novel object that the speaker labeled (Experiment 1) may be slower to respond, as their memory of the label is weaker if the speaker previously generated labeling errors.

When recently taught labels are not used in test trials (Experiment 2), the quality of novel label encoding would only influence response time inasmuch as it interacts with children's recognition that one of the objects had already been labeled (e.g., Novel Name-Nameless Category principle, in which novel labels are paired with previously unnamed entities; Golinkoff, Hirsh-Pasek, Bailey, & Wenger [1992]; Mervis & Bertrand [1994]; see also Halberda [2003] for an overview of different principles that govern children's disambiguation inference in the context of novel word learning). Regardless of whether preschoolers have a clear recollection of the object-label pairing presented by a speaker, they are likely to remember the fact that the object had been labeled (Sabbagh & Shafman, 2009). On this view, children's response times should not differ between the conditions in Experiment 2, but should differ in Experiment 1.

Experiment 1

Participants

Sixty-four preschoolers (35 female, 29 male, $M_{age} = 48.25$ months, range: 35.15–60.12 months) were recruited and tested at a children's museum or through birth records and tested in a laboratory. All participants were native speakers of American English and had no developmental disabilities. Six additional participants were tested but excluded from the analyses because of loss of attention to the task ($n = 1$) or refusing to complete the test procedure ($n = 5$). While only a few families provided race, ethnicity and socioeconomic status information, most families were Caucasian and represented middle- to upper-SES backgrounds (as indicated by visitor survey data provided by the museum).

Apparatus

Children were tested individually with their caregiver present. They watched video stimuli presented on a MacBook Air 11" laptop screen placed on the table 15 cm from the edge and adjusted to be best seen from the children's seating position. The audio stimuli were played from the laptop speakers set to the maximum volume.

Materials

Familiarization and test stimuli included two sets of familiar objects (four objects per set) and two novel objects. One set included a porcelain cup, a picture book, a stuffed cat, and a rubber ball. The other set included a stuffed dog, a brown shoe, a feeding bottle, and a cardboard star painted gold. The novel objects presented during the novel label training were a green drain cover and a yellow tarpaulin clip. The unfamiliar distractor objects were a metal plug wrench and a black tripod camera mount (see Figure 1).

Insert Figure 1 here

Procedure

Children were tested in a quiet room off the museum floor or in the laboratory. Their participation was videotaped for later coding. Caregivers stayed in the room with their children for the entire duration of the procedure. Children sat at a table either on their own or accompanied by caregivers. Caregivers were instructed not to interfere with children's behavior during the experiment.

The procedure consisted of three phases: (1) *Familiarization*, (2) *Novel label training*, and (3) *Test* (Figure 2). The familiarization and novel label training phases involved children watching prerecorded videos, presented on the laptop screen. During 3-s intervals between the phases children saw a dark screen. These videos were the same as those used by Luchkina et al. (2018). Presenting children with videos instead of live speakers minimized experimenter effects and potential experimenter errors.

Before the familiarization video, the experimenter informed children that they should watch the video carefully and that they would answer questions about the video afterwards. During familiarization, children saw two female actors sit at a table with an opaque black cloth on it. Behind the table and informants was a black curtain. The actors were similar in appearance to minimize potential effects of children's subjective preferences for one speaker over the other; they were the same race and similar in age but had distinct hair color and wore college insignia sweatshirts of different colors. They took turns presenting one object, looking at it, and asking questions about its label. The *accurate labeler* always mentioned correct labels of the objects (e.g., while displaying a star, "Hmm... Is this star? A star?"). The *inaccurate labeler* always

mislabeled the objects (e.g., while displaying a book, “Hmm... Is this a cat? A cat?”). Both mentions of the target word were articulated with a rising intonation to ensure that children recognized them as questions, not statements (see the spectral analysis of the questions shown in Figure 3; see Soderstrom, Ko, & Nevzorova [2010] for a discussion on children’s use of prosody to distinguish between statements and questions). Each speaker presented and asked questions about four familiar objects. Because the speakers were asking questions, an effort was made to ensure that the questions did not appear as a pedagogical inquiry from an adult, as rhetorical, or as a joke.

Participants were randomly assigned to the accurate or inaccurate labeler conditions. In the accurate labeler condition, during the novel label training, participants saw a video of the accurate labeler introducing two novel objects and labeling each of them with a different novel label, one at a time. In the inaccurate labeler condition, participants saw a video of the inaccurate labeler do the same thing. The novel labels were nonsense words – “neem” and “lif” (“Look! A lif. This is a lif.”). These words were the same labels used by Luchkina et al. (2018) based on Werker, Cohen, Lloyd, Casasola, and Stager (1998). The novel label training video was repeated twice. During familiarization and novel label training, the locations of the object sets, the identity of the accurate labeler, and the labeler who introduced the novel objects were counterbalanced among participants.

Following novel label training, participants were tested using a forced choice procedure administered by a live female experimenter who did not appear in the video. On each of two trials, the experimenter presented a novel object from the video paired with an unfamiliar distractor object (also a novel object, but one that was unlabeled). The experimenter placed both objects in front of the child and asked children to give her the object corresponding to a novel

label from the video. The location of the target object (left/right) was randomized among participants and counterbalanced across the test trials. If children did not make a choice, the experimenter repeated her request and encouraged them to pick up the object corresponding to the label. To control for potential bias on the part of the experimenter, the experimenter made an explicit effort to not look at either object and tried to look between children's eyes to minimize the chance of her looking behavior influencing children's responses.

Insert Figure 2 here

Data Coding and Analyses

All data were coded offline by a naïve coder. The dependent measures were (1) children's choice of object and (2) their response time, calculated as the time elapsed from the last frame of the experimenter's utterance ("Give me the lif!") to the first frame when the object contacted the experimenter's hand. The experimenter had her hand extended from the moment she completed the utterance and kept it extended until children handed the object. The labeler (accurate, inaccurate) who introduced the novel objects was a between-subject independent variable, and trial number (first, second) was a within-subject variable, as we anticipated potential effects of trial number (e.g., responding faster on the second trial upon figuring out the goal of the task after the first trial).

Ten percent of videos were coded by a second coder. There was 100% intercoder agreement about children's choice of object. Coders' judgments of the response time differed on average by 38 ms (0.79% of the average response time). The average value of the coders' judgments was taken to resolve the discrepancies.

Results and Discussion

Figure 4 shows the percentage of trials on which children chose the target object, as well as their response times. Children generally retained the labels learned from the video and demonstrated strong performance in both conditions. In the accurate labeler condition, children chose the target object 90% of the time with 24 children (out of 30) selecting the target object on both trials and 6 children selecting the target object on one trial. Performance in the inaccurate labeler condition was similar; children chose the target object on 90.3% of the trials with 22 children (out of 28) selecting the target object on both trials, 6 children selecting the target object on one trial. Children never selected the distractor object on both trials. Chi-square goodness of fit tests revealed that children in both conditions chose the target object significantly more often than chance, $\chi^2(2, N = 30) = 13.20, p = .001$ and $\chi^2(2, N = 28) = 11.71, p = .003$ for the accurate and inaccurate labeler conditions respectively. These findings were also significant when the full dataset was considered (including the outliers, described below).

We fitted Generalized Linear Mixed Models (GLMMs) to children's test performance (their choice of object and response time in separate GLMMs) to evaluate the effects of experimental condition and trial number. We used GLMMs because of their flexibility in accounting for fixed and random effects, as well as repeated measures, and its suitability for the analyses of binary (choice of object) and continuous (response time) measures. Because every child was accompanied by different number of family members, had visited different museum exhibits before participating in the experiment, and expressed different level of enthusiasm about participating, we entered subject ID as a random effect (for the intercept) in all analyses. In the analyses of children's choice of object we used a logit link function (children could either choose the target or the distractor object). In the analyses of children's response time we used an identity

link function applied to the log-transformed response time (before the log-transformation the distribution of the response time data significantly deviated from normal, as measured by the Kolmogorov-Smirnov tests, Smirnov [1948]). Three additional participants – two in the accurate labeler condition and one in the inaccurate labeler condition – were excluded from these analyses as outliers because their response times were more than two standard deviations away from the mean on one or both trials.

Preliminary analyses (GLMM with a logit link function for the choice of objects and GLMM with an identity link function for response time) of the effects of age, gender, and the counterbalancing variables did not indicate any significant effects on the choice of object or response time, all p -values $> .36$. These variables were therefore excluded from subsequent analyses. A GLMM fitted to the choice of object revealed no significant effects of any independent variables, all p -values $> .95$. The analyses of response time revealed a significant effect of condition, $z = -2.15$, $p = .036$, with children in the inaccurate labeler condition taking longer to respond than children in the accurate labeler condition, $M_{\text{response time}} = 6567$ ms and 4002 ms in the inaccurate and accurate labeler conditions respectively (see Figure 5). There was also a significant effect of trial number, $z = 14.39$, $p < .001$, with children responding faster on the second trial than the first, $M_{\text{response time}} = 6163$ ms and 2213 ms of the first and second trial respectively.

Upon completing our planned analyses of the effects of condition and trial number on children's choice of object and their response time, we conducted exploratory analyses of the relation between these two measures. We examined whether longer response times predicted the choice of non-target objects. If our hypothesis that children's response time was driven by the strength of encoding is correct, it is possible that children who took a longer time to respond

were also less likely to select the target object. To test this hypothesis, we added log-transformed response time to the GLMM analyzing children's object choice, and re-ran the analyses.

Consistent with our conjecture, there was a significant effect of log-transformed response time on children's choice of object, $z = -4.65, p < .001$. This result is not surprising given that both response time and the choice of object were measures of children's retention of novel words. The effect of response time on children's choice of object was present in both conditions; separate GLMMs fitted to the choice of object data revealed significant effects of log-transformed response time in the accurate labeler condition, $z = -9.67, p < .001$, and the inaccurate labeler condition, $z = -5.50, p < .001$ (see Figure 5).

Insert Figure 4 here

Insert Figure 5 here

In sum, preschoolers judged both labelers who asked questions about familiar objects as reliable sources of word knowledge. Unlike toddlers in Luchkina et al. (2018), the preschoolers in this experiment did not treat labelers who asked questions as epistemically incompetent, regardless of whether they uttered accurate or inaccurate labels. Response time, however, reflected the accuracy of the labels in those questions. Children were faster when novel labels were provided by the accurate labeler than by the inaccurate labeler. Children also showed an overall difference in response time between the two test trials. The latter result is potentially due to higher familiarity of the task on the second trial, resulting in faster response times overall. Because the locations of target objects were randomized across participants and counterbalanced across trials, there is little chance that the effect of trial could be attributed to other factors.

These data suggest that children's selective word learning might be influenced by two learning mechanisms – one that evaluates epistemic competence and another that registers semantic errors in informants' speech and potentially weakens children's encoding of novel object-label pairs provided by the same informants. However, a challenge with this conclusion is that the accurate labeler might have been seen as more epistemically competent than the inaccurate labeler. Mentioning the accurate label in a question indicates that the speaker has that word in her lexicon. The inaccurate labeler does not indicate that she possesses this knowledge. This might suggest that the accurate labeler has more linguistic knowledge than the inaccurate labeler.

To understand whether children could have perceived one of the speakers as less competent than the other (i.e., the speaker who mentions incorrect labels might be viewed as less knowledgeable), we coded children's behavior during the familiarization trials. We measured how often children looked to the back of the room to their caregivers or the experimenter (i.e. social referencing, Feinman, 1982; see Warreyn, Roeyers, & De Groote, 2005 for examples of social referencing measures in preschoolers). If children considered the situations in the video to be socially incongruous – adult speakers asking questions about the labels of highly familiar items – they would attempt to solicit other adults' help in interpreting the situation. If one speaker was considered less competent than the other, children's social referencing behaviors might differ depending on the speakers' labeling accuracy.

Spearman's rank correlation coefficient did not reveal a significant relation between label accuracy during familiarization and the number of times children looked to caregivers or the experimenter (summed across the four trials of each type), $r_s(114) = .05, p = .58$ (see Table 2). To test whether social referencing behaviors during familiarization related to children's test

performance and response times, we considered the number of times children looked to their caregivers or to the experimenter as an independent variable in a new set of GLMMs. The results of these analyses did not reveal any significant effects of social referencing behaviors during familiarization on children's choice of object on test trials, $z = 0.26$, $p = .80$, or on their response time, $z = -0.33$, $p = .74$. These results suggest that children did not treat the two speakers differently, at least as measured by the frequency of their social referencing behaviors during the eight familiarization trials. The low frequency of such behaviors ($M = 0.31$ and 0.32 times for the accurate and the inaccurate labeler trials, respectively, for the entire familiarization phase; see Table 1) also suggests that children did not consider familiarization trials to be incongruous or difficult to interpret and did not often solicit competent adults' input.

In addition to analyzing the frequency of children's looks toward caregivers on the experimenter, we analyzed the measures of children's engagement with the video – laughing, frowning, and responding to the labelers' questions (see Table 1). Spearman's correlation analyses did not reveal any significant relations between children's laughing or frowning at the labelers' questions or responding to them depending on labelers' accuracy, all $|r_s(114)\text{-values}| < .09$, all $p\text{-values} > .32$ (see Table 2). In sum, children's behavior during familiarization suggest that they did not consider either labeler less competent than the other and there were no differences in their engagement with the video depending on whether the accurate or the inaccurate labeler presented an object. Based on these results, we do not think that the observed difference in children's response time was driven by a difference in their assessment of labelers' epistemic competence.

To test this conclusion more explicitly, and to explore the mechanisms underlying preschoolers' selective word learning, in Experiment 2 we kept the familiarization and novel

labeling phases identical to Experiment 1, but in the test phase, we asked children to give the experimenter an object with a completely novel label. If response times in Experiment 1 were explained by the strength of the novel object-label pairings that may have resulted from registering and generalizing semantic errors, then by replacing the retention procedure with a disambiguation procedure, we would eliminate such an effect. Regardless of whether preschoolers have a clear recollection of the object-label pairing presented by the inaccurate labeler, they should remember the fact that the object was labeled (Sabbagh & Shafman, 2009). In this case, response times should not differ between conditions. Conversely, if children believed that the accurate labeler was more epistemically competent, we would expect to replicate our findings from Experiment 1.

Experiment 2 also addresses another concern with Experiment 1. Children's choices during test could have been biased towards the target objects because these objects were more familiar to children after being displayed in the video. The disambiguation procedure in Experiment 2 requires children to choose the unfamiliar object, not the familiar one. If children choose based on familiarity, their choices should not be different from Experiment 1.

Experiment 2

Participants

A new group of 64 preschoolers (32 female, 32 male, $M_{age} = 47.09$ months, range: 36.07–58.91 months,) was recruited and tested at the same local children's museum. All participants were native speakers of American English and had no known developmental disabilities. Four additional participants were tested but excluded from the analyses because of refusing to complete the test procedure. Demographic information was the same as in Experiment 1.

Apparatus & Materials

The apparatus and materials were the same as in Experiment 1.

Procedure

The procedure was identical to Experiment 1, except that during the test phase the experimenter asked children to choose an object corresponding to an unfamiliar novel label (see Figure 2). The unfamiliar labels were “dax” and “wug”. The unfamiliar novel object was considered the target in our analyses.

Data Coding and Analyses

Ten percent of videos were selected for reliability coding performed by a second coder. Inter-coder agreement about choice of object was 100%. Coders’ judgments of the response time varied on average by 6.67 ms (0.1% of the average response time). The average value of the coders’ judgments was taken to resolve the discrepancies.

Results and Discussion

In the accurate labeler condition, 14 children (out of 30) selected the target object on both trials, 10 children selected the target on one of the trials, and 6 children never selected the target object. In total, children chose the target object 63.33% of the time (see Figure 4). In the inaccurate labeler condition, these frequencies were 15 (out of 32), 14, and 3 respectively. In total, children chose the target object 69% of the time (Figure 4). Chi-square goodness of fit tests revealed that children chose the unfamiliar novel object (the target) on test trials significantly more often than chance in both conditions, $\chi^2(2, N = 30) = 7.60, p = .022$ and $\chi^2(2, N = 32) = 9.50, p = .008$, for the accurate labeler and the inaccurate labeler conditions respectively. As in Experiment 1, these findings were also significant when the full dataset was considered

(including the outliers described below). Preschoolers' above-chance performance suggests that their choice of object in Experiment 1 could not be explained by a bias to select familiar looking objects that were shown in the video.

As in Experiment 1, the response time data significantly deviated from normality, one-sample Kolmogorov-Smirnov test, $p < .001$. A logarithmic transformation was applied to normalize those data. After the transformation, Kolmogorov-Smirnov tests were not significant for both trials. Following the transformation, two additional children (both in the accurate labeler condition) were excluded from analyses as outliers because their response time was more than two standard deviations away from the mean on one or on both trials.

We ran preliminary analyses on choice of object and response time (GLMM with a logit link function for the choice of objects and GLMM with an identity link function for response time). These analyses revealed no significant effects of gender or the counterbalancing variables, all p -values $> .59$. These variables were therefore excluded from subsequent analyses of children's choice of object. There was a nonsignificant, but marginal effect of age, $z = 1.73$, $p = .084$ on children's choice of object and a significant effect of age on response time, $z = -2.56$, $p = .013$, with 4-year-olds selecting the target object more often than three-year-olds, 75% and 60% of the time respectively, and responding to the experimenter's questions faster than 3-year-olds, $M_{\text{response time}} = 11719$ ms and $M_{\text{response time}} = 8725$ ms for 3- and 4-year-olds respectively. Because of these effects, we included age in the main analyses.

A GLMM with a logit link function fitted to children's choice of object during test revealed no significant effects of condition or trial number on the likelihood of choosing the target object. The effect of age remained marginally significant, $z = 1.88$, $p = .06$. A GLMM with an identity link function fitted to the response time data revealed a significant effect of trial

number, $z = -18.07$, $p < .001$, with $M_{response\ time} = 14988$ ms and $M_{response\ time} = 4831$ ms on first and second trials respectively. The analyses also revealed a significant effect of age, $z = -2.81$, $p = .002$, with $M_{response\ time} = 11719$ ms and $M_{response\ time} = 8100$ ms for 3- and 4-year-olds respectively, and a marginally significant effect of condition, $z = 1.89$, $p = .063$, with $M_{response\ time} = 11053$ ms and $M_{response\ time} = 8766$ ms for the accurate labeler and the inaccurate labeler conditions respectively.

As in Experiment 1, we conducted additional analyses of children's behavior during familiarization to better understand if they could have treated one labeler as more epistemically competent than the other. As before, we did so by computing Spearman's correlation coefficients between labeler accuracy during familiarization and children's looks to caregivers and the experimenter, as well as their laughing, frowning, and responding to the labelers' questions (see Table 2). Neither of assessed correlations was statistically significant, $|r_s(122)\text{-values}| < .08$, $p\text{-values} > .36$ (see Table 1). We also evaluated the effects of children's social referencing behaviors on their choice of object and response time by adding the number of looks to caregivers or the experimenter to the GLMMs fitted to test performance measures. Consistent with the results of Experiment 1, there were no significant effects of children's social referencing behaviors on their choice of object or response time, all $p\text{-values} > .85$. These results provide additional evidence in support of our conjecture that children did not treat either labeler as less competent than the other.

Finally, we analyzed the effects of response time on children's choice of object. Exploratory analyses in Experiment 1 revealed that children in both conditions who took longer time to respond to the experimenter's question were less likely to select the target object. That result is consistent with our hypothesis that registering labeling errors weakens children's

subsequent encoding of the novel information, leading to longer response times and increasing the chance of selecting the distractor object on test trial. We hypothesized that in Experiment 2, the strength of encoding should not relate to children's test performance because they only need to remember that the familiar looking object had been labeled, but not to retrieve the specific object-label mapping. According to this hypothesis, in addition to not observing a significant effect of condition on children's response times, we should also not expect to observe an effect of response time on children's choice of object. We added log-transformed response time to the GLMM and fitted the new model to the choice of object data. Consistent with our prediction, there was no significant effect of response time on children's choice of object on test trials, $z = -0.02, p = .98$.

The results of Experiment 2 suggest that preschoolers considered both question-askers as reliable sources of word knowledge. The lack of significant difference in response times is consistent with the hypothesis that disambiguation inferences do not depend on the strength of information encoding. The lack of significant effects of response times on children's choice of object corroborates this conclusion. In sum, Experiment 2 suggests that children are influenced by both judgments of epistemic competence and by a different process that inhibits encoding of novel information upon registering semantic errors in labelers' questions.

General Discussion

The present studies examined the potential underlying mechanisms of children's selective word learning. We hypothesized that (1) by having speakers ask questions about the names of familiar objects with either their accurate or inaccurate label, we could minimize the contrast between the speakers' epistemic competence, while preserving the difference in error signal that

is registered by hearing inaccurate labels for familiar objects. Further, (2) if labeling errors influenced children's word learning, such errors would weaken the encoding of novel object-label mappings and manifest in a difference in response time as a function of whether the accurate or the inaccurate labeler introduced and labeled the novel objects.

Consistent with our first hypotheses, children learned equivalently from the two speakers in both experiments. One might object that in Experiment 2, children did not specifically map the novel labels provided by the speakers onto the novel objects, but the results of that experiment suggested that children realized that those objects already have labels, which they could only infer from the naming events. Consistent with our second hypothesis, response times in Experiment 1 revealed an effect of condition, with children taking longer to respond when novel labels were introduced by an inaccurate speaker. Moreover, there was a relation between response time and choice of object in Experiment 1: children who took longer to make their choice in both conditions were less likely to select the target object. The lack of a significant effect of condition on response time in Experiment 2 also supports our second hypothesis. In the disambiguation task, children did not have to retrieve the object-label mapping taught by one of the speakers in order to make this inference, and thus the strength of novel label encoding did not influence response times.

We also observed effects of age in Experiment 2 – significant for response time and marginal for choice of object. Four-year-olds were more likely to select the target object and were faster in doing so than 3-year-olds. Unlike Experiment 1, where the effect of age was not observed, in Experiment 2 children had to engage in a disambiguation inference while keeping in mind two novel object-label mappings (or, at least, the fact that such mappings had been presented). Given the complexity of this task, it stands to reason that younger participants spent

more time making disambiguation inferences and produced more errors. Additionally, the pragmatic context of the situation may have prompted children to expect that the experimenter would request the objects from the video. As such, it might have been more difficult for younger learners to inhibit a preponderant response, causing more errors.

A concern with the present findings is that the speaker in the inaccurate label condition could have been viewed as less competent because she was making incorrect guesses about the labels of the familiar object, and thus indicated that she did not know the label of the familiar object nor the referent of the label she generated. In this case, the difference in response times shown in Experiment 1 would have been driven by a difference in children's confidence in the newly-learned labels because of an epistemic-level inference about the speakers, as opposed to a different mechanism for selective learning based on registering errors and using that error signal to gauge the strength of the word-object association.

Experiment 2, however, sheds some doubt on this interpretation. If performance in Experiment 1 was driven by children's inferences about labelers' epistemic knowledge, this effect would have remained in Experiment 2, in which children were asked to select objects corresponding to unfamiliar novel labels. If children were less sure about the reliability of the labels provided by the inaccurate labeler, they might have hesitated more when making their choice of object and potentially be less likely to select the target object in Experiment 2.

The results, however, show a different pattern. Across the two conditions, children were not different in their choices, and were faster to respond in the *inaccurate* condition as opposed to the accurate one (at a marginally significant level). If responses here were driven by only epistemic-level inferences, this would mean that children viewed the accurate labeler as a more puzzling case, because she acted uncertainly and questioned her knowledge despite mentioning

correct labels. The lack of differences in social referencing behaviors, as well as the general lack of reacting to the labelers' questions, however, suggest that questions did not result in a difference in children's evaluation of the speakers' epistemic competence. There remains a possibility that the difference in response times in Experiment 1 was driven by epistemic-level inferences and the complexity of disambiguation inference masked this effect in Experiment 2. This is an open empirical question that could be addressed by future investigations.

Luchkina et al. (2018) found that 18-month-olds in a similar condition to Experiment 1 did not treat either speaker as a reliable source of novel word knowledge. Here, we demonstrated a contrasting finding: Three- and 4-year-olds given a similar procedure treated both speakers as reliable, even though their response times differed. This contrast can be explained by 18-month-olds' perception of speakers' questions during familiarization. Because the speakers were instructed to emulate genuine requests for information, toddlers may have interpreted the situation as such. It is uncommon that adults request object labels from toddlers. Because toddlers possess some understanding that adults are typically more competent sources of information than children (Kachel, Moore, & Tomasello, 2018), it is possible that 18-month-olds interpreted labelers' questions as pragmatically incongruous or violating social conventions (e.g., Diesendruck & Shemer, 2006; see also Diesendruck, 2012, for a review on toddlers' understanding of conventionality). For this reason, toddlers may have either distrusted both labelers or were confused by the speakers' use of questions, which may have led to weaker encoding of novel information in both conditions.

Unlike toddlers, preschoolers have more experience hearing questions about object labels (at least, in the form of testing their knowledge). Having such experience may have led them to not register any pragmatic errors or violations of conventionality in labelers' questions, and

readily accept object names provided by those labelers. In addition, when children had to make a choice between the object from the video and an unfamiliar novel object, they may have had only a weak memory of the correct object-label mapping but remembered the labeling event. In this case, they could have chosen the correct object based on this recollection, whereas toddlers in Luchkina et al. (2018) study needed to remember the exact object-label mappings because both novel objects from the video were presented on the screen on test trials.

There is another reason for the contrast between preschoolers' performance in Experiment 1 and toddlers' chance-level looking behavior in Luchkina et al. (as well as earlier research that demonstrated preschoolers' selectivity in learning novel labels from accurate and inaccurate informants, see e.g., Brooker & Poulin-Dubois, 2013; Koenig & Woodward, 2010). One of the purposes of Experiment 2 was to see whether performance in Experiment 1 was driven by children's tendency to give the experimenter the object demonstrated by a labeler. The results of Experiment 2 suggest that it is unlikely to be the case because children succeeded on the disambiguation task. However, given the difference in children's performance on the two experiments, it is possible that at least *some* children in Experiment 1 were selecting the objects based on their visual familiarity rather than their labels. In this case, the proportion of the target object choices in Experiment 2 (66% across the two conditions) is a more conservative estimate of children's word learning outcomes.

A potential limitation of the present study includes the relatively homogeneous sample of participants, the majority of whom were Caucasian and from middle- to upper-middle-SES families. Existing literature on children's selective trust suggests that SES can affect the trajectory of mechanism development in children. For example, Corriveau, Kurkul, and Arunachalam (2016) showed that higher- but not lower-SES preschoolers employed syntactic

complexity of speakers' utterances as a cue to reliability. That effect was likely driven by the differences in the complexity of language preschoolers from different SES group were routinely exposed to. Given that the current study extensively relied on children's word knowledge, there is chance that the current results would have been different with a more diverse sample of participants. We believe, however, that such a chance is rather small because only syntactically simple utterances and highly familiar object-label pairs, typically known even by 12-18-month-olds, were used as stimuli. Another potential limitation is that all participants spoke English as their dominant language and were a part of a WEIRD (Western, Educated, Industrialized, Rich and Democratic; see Henrich, Heine, & Norenzayan, 2010) culture. Lucas, Lewis, Pala, Wong, and Berridge (2013) demonstrated that children who speak languages that highlight the relevance of the source of information, such as Turkish, have an advantage on selective trust tasks when compared to children who speak other languages. This work suggests that testing children from different cultural backgrounds could yield different findings in selective trust and selective learning experiments. While it is difficult to predict how our result would have been modified by testing preschoolers from different linguistic communities, we acknowledge that our findings might not generalize to all languages and cultures.

Finally, it is possible that testing multilingual instead of monolingual preschoolers could have affected our findings. While it is unlikely that a different set of mechanisms governs word learning in multilingual preschoolers, the relative contributions of such mechanisms might follow a different trajectory. Ardal, Donald, Meuter, Muldrew, and Luce (1990) reported reduced frontal mismatch negativity in response to the stimuli presented in the second language of bilingual adults. It is possible that for children acquiring several languages at once, the mismatch negativity (N400) response is reduced for the non-dominant language. If tested in the non-

dominant language on the task reported in the present study, multilingual preschoolers might show a reduced effect of label accuracy. Such a result would be consistent with our interpretation of the present findings and informative for our understanding of language processing and word learning in bilinguals. This possibility invites potential future investigations of mechanisms underlying selective word learning in multilingual preschoolers.

To conclude, our findings are consistent with the idea that preschoolers' accuracy in selective learning reflects their ability to make inferences about others' epistemic competence. The lack of selectivity in these experiments, as measured by the proportion of the target object choices, suggests that children are highly sensitive to pragmatic aspects of speech and do not attribute different levels of epistemic competence when they hear speakers mention the labels of familiar objects in the context of questions. Children's response times, however, reveal that they do register semantic errors in speakers' labels, which affect their encoding of that information—a potential mechanism underlying selective learning in addition to judgments about speaker's epistemic competence. In sum, the present work demonstrates that epistemic-level inferences influence selective word learning in parallel with more implicit mechanisms that keep track of labeling errors and potentially inhibit the encoding of novel labels provided by those who produced such errors. This evidence is in line with the dual-process account of selective trust (Hermes, Behne, & Rakoczy, 2018), according to which children's selectivity is driven by two processes: a fast and implicit heuristic process and a systematic slow one. The differential encoding mechanisms could be viewed as a product of the heuristic-level process and epistemic-level inferences could be viewed as a product of the effortful slow inference process. As such, our findings provide additional support for a dual-process account and shed light on the nature of children's selectivity in social learning.

References

- Ardal, S., Donald, M. W., Meuter, R., Muldrew, S., & Luce, M. (1990). Brain responses to semantic incongruity in bilinguals. *Brain and Language*, 39(2), 187-205.
[https://doi.org/10.1016/0093-934X\(90\)90011-5](https://doi.org/10.1016/0093-934X(90)90011-5)
- 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, E46-E68. <https://doi.org/10.1111/inf.12027>
- 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>
- Corriveau, K. H., Kurkul, K., & Arunachalam, S. (2016). Preschoolers' preference for syntactic complexity varies by socioeconomic status. *Child Development*, 87(5), 1529-1537.
<https://doi.org/10.1111/cdev.12553>
- de Saussure, F. (1916/1966). *Course in General Linguistics*. McGraw-Hill Book Company.
- Diesendruck, G. (2012). The scope and origins of children's assumption of conventionality. In M. Siegal & L. Surian (Eds.) *Access to language and cognitive development*, 116-134. Oxford University Press
- Diesendruck, G., & Shemer, G. (2006). Young children's expectation of competence in word learning. *Journal of Child Language*, 33(2), 321-338.
<https://doi.org/10.1017/S0305000906007343>
- Feinman, S. (1982). Social referencing in infancy. *Merrill-Palmer Quarterly* 28(4), 445-470.
- Friedrich, M., & Friederici, A. D. (2004). N400-like semantic incongruity effect in 19-month-olds: Processing known words in picture contexts. *Journal of cognitive neuroscience*, 16(8), 1465-1477. <https://doi.org/10.1162/0898929042304705>

- Golinkoff, R. M., Hirsh-Pasek, K., Bailey, L. M., & Wenger, N. R. (1992). Young children and adults use lexical principles to learn new nouns. *Developmental Psychology*, 28(1), 99.
<https://psycnet.apa.org/doi/10.1037/0012-1649.28.1.99>
- Golinkoff, R. M., Hirsh-Pasek, K., Cauley, K. M., & Gordon, L. (1987). The eyes have it: Lexical and syntactic comprehension in a new paradigm. *Journal of Child Language*, 14(1), 23-45. <https://doi.org/10.1017/S030500090001271X>
- Halberda, J. (2003). The development of a word-learning strategy. *Cognition*, 87(1), B23-B34.
[https://doi.org/10.1016/S0010-0277\(02\)00186-5](https://doi.org/10.1016/S0010-0277(02)00186-5)
- 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>
- Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive foundations of learning from testimony. *Annual Review of Psychology*, 69, 251-273.
<https://doi.org/10.1146/annurev-psych-122216-011710>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, 466(7302), 29. <https://doi.org/10.1017/S0140525X0999152X>
- 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.
<http://dx.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., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, 17(9), 757-758.
<http://dx.doi.org/10.1111/j.1467-9280.2006.01778.x>
- Kachel, G., Moore, R., & Tomasello, M. (2018). Two-year-olds use adults' but not peers' points. *Developmental Science*, 21(5), e12660. <https://doi.org/10.1111/desc.12660>
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Koenig, M. A. (2012). Beyond semantic accuracy: Preschoolers evaluate a speaker's reasons. *Child Development*, 83(3), 1051-1063. <https://doi.org/10.1111/j.1467-8624.2012.01742.x>
- Koenig, M. A., & Echols, C. H. (2003). Infants' understanding of false labeling events: The referential roles of words and the speakers who use them. *Cognition*, 87(3), 179-208.
[https://doi.org/10.1016/S0010-0277\(03\)00002-7](https://doi.org/10.1016/S0010-0277(03)00002-7)
- 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., & 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>
- Koenig, M. A., & Woodward, A. L. (2010). Sensitivity of 24-month-olds to the prior inaccuracy of the source: possible mechanisms. *Developmental Psychology*, 46(4), 815.
<https://psycnet.apa.org/doi/10.1037/a0019664>
- Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, 15(10), 694-698.
<https://doi.org/10.1111/j.0956-7976.2004.00742.x>

- Lucas, A. J., & Lewis, C. (2010). Should we trust experiments on trust? *Human Development*, 53(4), 167-172. <https://doi.org/10.1159/000320044>
- Lucas, A. J., Lewis, C., Pala, F. C., Wong, K., & Berridge, D. (2013). Social-cognitive processes in preschoolers' selective trust: Three cultures compared. *Developmental Psychology*, 49(3), 579. <https://doi.org/10.1037/a0029864>
- 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>
- Mangardich, H., & Sabbagh, M. A. (2018). Children remember words from ignorant speakers but do not attach meaning: evidence from event-related potentials. *Developmental Science*, 21(2), e12544. <https://doi.org/10.1111/desc.12544>
- Merriman, W. E., & Schuster, J. M. (1991). Young children's disambiguation of object name reference. *Child Development*, 62(6), 1288-1301. <https://doi.org/10.1111/j.1467-8624.1991.tb01606.x>
- Mervis, C. B., & Bertrand, J. (1994). Acquisition of the novel name–nameless category (N3C) principle. *Child Development*, 65(6), 1646-1662. <https://doi.org/10.1111/j.1467-8624.1994.tb00840.x>
- Mills, C. M. (2013). Knowing when to doubt: developing a critical stance when learning from others. *Developmental Psychology*, 49(3), 404. <http://psycnet.apa.org/doi/10.1037/a0029500>
- 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>

- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72*(4), 1054-1070. <https://doi.org/10.1111/1467-8624.00334>
- 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>
- Sheehan, E. A., Namy, L. L., & Mills, D. L. (2007). Developmental changes in neural activity to familiar words and gestures. *Brain and Language, 101*(3), 246-259.
<https://doi.org/10.1016/j.bandl.2006.11.008>
- Smirnov, N. (1948). Table for estimating the goodness of fit of empirical distributions. *The Annals of Mathematical Statistics, 19*(2), 279-281. <http://www.jstor.org/stable/2236278>.
- Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development, 81*(2), 669-679. <https://doi.org/10.1111/j.1467-8624.2009.01422.x>
- 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), 779-797.
<https://doi.org/10.1037/a0034191>
- Soderstrom, M., Ko, E. S., & Nevzorova, U. (2011). It's a question? Infants attend differently to yes/no questions and declaratives. *Infant Behavior and Development, 34*(1), 107-110.
<https://doi.org/10.1016/j.infbeh.2010.10.003>
- Sternberg, S. (1969). Memory-scanning: Mental processes revealed by reaction-time experiments. *American Scientist, 57*(4), 421-457. <https://www.jstor.org/stable/27828738>

- Tummeltshammer, K. S., Wu, R., Sobel, D. M., & Kirkham, N. Z. (2014). Infants track the reliability of potential informants. *Psychological Science*, 25(9), 1730-1738.
<https://doi.org/10.1177/0956797614540178>
- VanderBorght, M., & Jaswal, V. K. (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development: An International Journal of Research and Practice*, 18(1), 61-71. <https://doi.org/10.1002/icd.591>
- Warreyn, P., Roeyers, H., & De Groote, I. (2005). Early social communicative behaviours of preschoolers with autism spectrum disorder during interaction with their mothers. *Autism*, 9(4), 342-361. <https://doi.org/10.1177/1362361305056076>
- Werker, J. F., Cohen, L. B., Lloyd, V. L., Casasola, M., & Stager, C. L. (1998). Acquisition of word-object associations by 14-month-old infants. *Developmental Psychology*, 34(6), 1289-1309. <https://doi.org/10.1037/0012-1649.34.6.1289>

Figure 1. Novel objects introduced during the novel label training (top row) and the unfamiliar novel distractor objects (bottom row)

Figure 2. Experimental procedure in Experiments 1 and 2

Figure 3. Example of intonation contour of speakers' utterances during familiarization

Figure 4. Children's choice of object on test trials

Figure 5. Response time of children's choice on test trials

Table 1. Average frequency of social referencing and engagement with the video behaviors across 4 familiarization trials of each type (accurate vs. inaccurate labeler presents objects)

	Social referencing	Laughing	Frowning	Responding to labelers' questions
Experiment 1 – accurate labeler trials	0.31	0.18	0.00	0.85
Experiment 1 – inaccurate labeler trials	0.32	0.15	0.02	0.91
Experiment 2 – accurate labeler trials	0.78	0.29	0.00	0.94
Experiment 2 – inaccurate labeler trials	0.63	0.22	0.00	0.91

Table 2. Spearman's rank correlation between labeler accuracy during familiarization trials and the average frequency of social referencing and engagement with the video behaviors

Behavior	Experiment 1	Experiment 2
	r_s (p -value)	r_s (p -value)
Social referencing	0.05 (0.58)	0.08 (0.36)
Laughing	0.0001 (0.99)	-0.02 (0.85)
Frowning	-0.09 (0.32)	n/a
Responding to labelers' questions	0.01 (0.95)	0.01 (0.95)