Retrospective Inferences in Selective Trust

Friederike Schütte^{1*}, Nivedita Mani^{2,3}, and Tanya Behne^{1,3}

¹Developmental Psychology, University of Göttingen, Waldweg 26, 37073 Göttingen,
Germany

²Psychology of Language, University of Göttingen, Goßlerstraße 14, 37073 Göttingen,
Germany

³Leibniz ScienceCampus 'Primate Cognition', Göttingen, Germany

*Author for correspondence (friederike.schuette@uni-goettingen.de)

Retrospective Inferences in Selective Trust

Abstract

Preschoolers learn selectively from others based on the speakers' prior accuracy. This indicates that they recognize the models' (in)competence and use it to predict who will provide the most accurate and useful information in the future. Here, we investigated whether 5-y.-o. children are also able to use speaker reliability retrospectively, once they have more information regarding their competence. They first experienced two previously unknown speakers who provided conflicting information about the referent of a novel label, with each speaker using the same novel label to refer exclusively to a different novel object. Following this, children learned about the speakers' differing labeling accuracy. Subsequently, children selectively endorsed the object-label-link initially provided by the speaker who turned out to be reliable significantly above chance. Crucially, more than half of these children justified their object selection with reference to speaker reliability, indicating the ability to explicitly reason about their selective trust in others based on the informants' individual competences. Findings further corroborate the notion that preschoolers are able to use advanced, metacognitive strategies (trait reasoning) to learn selectively. In contrast, since learning *preceded* reliability exposure and gaze data showed no preferential looking toward the more reliable speaker, findings cannot be accounted for by attentional bias accounts of selective social learning.

Keywords: Social learning, Selective trust, Trait ascription, Social cognition, Retrospective inferences

Retrospective Inferences in Selective Trust

Most of what young children acquire they learn from interaction with, testimony from, or observation of others (Harris, 2012). This reliance on social agents for information leaves young learners vulnerable to error because others might deceive them or unintentionally provide unconventional information (Clément, 2010). A plethora of recent findings has demonstrated that children are not as easily deceived as assumed but are quite selective in whom and what they trust. When confronted with two possible sources of information, children from four years on show robust capacities to accept new information selectively from, for example, the more knowledgeable, older, stronger, more familiar or nicer informant (for reviews, see e.g., Harris, Koenig, Corriveau, & Jaswal, 2018; Mills, 2013).

Preschool-aged children are also fairly sensitive to the informant's proven competence and trust more in information provided by previously reliable informants than by previously unreliable informants (Birch, Vauthier, & Bloom, 2008; Hermes, Behne, & Rakoczy, 2015; Jaswal & Neely, 2006; Koenig, Clément, & Harris, 2004; Koenig & Harris, 2005; Nurmsoo & Robinson, 2009; Pasquini, Corriveau, Koenig, & Harris, 2007). In the classic selective trust paradigm (Koenig & Harris, 2005), children first experience two informants label familiar objects. One of the informants is consistently accurate in labelling the objects while the other is consistently inaccurate. Subsequently, children preferentially accept new information from the previously accurate, and thus reliable informant. Preschool aged children thus predict future informativeness of sources on the basis of their previous reliability. But what is the underlying cognitive mechanism of such selective trust?

One proposal is that children make these inferences on the basis of conceptual knowledge (Sobel & Kushnir, 2013) including trait reasoning (Hermes, Behne, & Rakoczy, 2018). By the age of three, after listening to testimony from a consistently reliable and a consistently unreliable speaker, children can indeed identify the actor who was previously

reliable (Koenig et al., 2004). Further, only those children who successfully identify the reliability of the informants actually learn selectively from the more reliable one (Koenig & Harris, 2005). More importantly, preschool aged children selectively trust the model with the relevant competence for the tasks at hand (Hermes et al., 2015; Kushnir, Vredenburg, & Schneider, 2013). For example, after they experience two models differing in competence in two domains, e.g., labeling accuracy and physical strength, children select models in accordance with their attributes, preferring the strong model for strength-related tasks and the knowledgeable model for knowledge-related tasks (Hermes et al., 2015). Again, this pattern of selective model choice held only for those children who correctly identified the attributes of both models indicating that children's rational choice between two models competent in different domains depends on their knowledge about the underlying traits.

In contrast, attentional bias accounts (such as the 'associative account' by Heyes, 2017) argue that children's selective trust is not necessarily based on higher socio-cognitive processes, but can be based on low-level, attentional differences alone. When children learn about their informants' reliability, it is argued, they develop an attentional bias toward the more reliable speaker since her output matches children's visual input whereas the unreliable speaker's utterances constantly cause prediction errors. This bias then leads children to be more attentive to the information provided by the reliable speaker. In the classical selective trust paradigm, children thus develop an attentional bias toward the speaker labeling objects accurately and are thus more likely to acquire the novel label-object link she provides.

To test if young children are indeed able to learn selectively on the basis of trait reasoning and not (only) due to attentional differences during learning, we investigate children's selective trust in a retrospective task in which the novel information provided by two unknown informants *precedes* accuracy exposure. Thus, at the time when speakers provide novel information, children are still ignorant about the speakers' (un)reliability and a potential attentional bias developed during accuracy exposure cannot impact the selective

endorsement of novel information. If children selectively endorse the information that had been provided by the speaker who subsequently turned out to be reliable, children's trust cannot be easily explained by such low-level attentional biases. Instead, it would indicate that children are able to make (metacognitive) inferences on the basis of informants' competence (trait reasoning) and thus, to strategically learn from one source over another.

Previous research indicates that by 4 to 5 years of age children are able to make retrospective inferences in general (Sobel, Tenenbaum, & Gopnik, 2004) as well as engage in belief revision (Kimura & Gopnik, 2018; Macris & Sobel, 2017). In fact, one previous study by Scofield & Behrend, (2008) suggests that some children aged 4 may revise their trust when informants proved unreliable after they had taught a new label to children. In this study, an unknown speaker used a novel label to refer to one of two novel objects, and at test, all 4-year-olds selected this object as the referent of the novel label. Subsequently, this first speaker proved to be unreliable (labelling familiar objects incorrectly) and a second speaker appeared who used the same novel label to refer to the other novel object (as well as labelling familiar object correctly). When children were now again asked for the referent of the novel label, half the 4-year-olds changed their original label-object link and chose the object that the second speaker had referred too, while the other half stuck with the initial label-object link provided by the unreliable source. In other words, some 4-year-olds may have revised their judgement retrospectively in accordance with the source's respective reliability, or children's performance as a group may have dropped to chance level (once both objects had been treated as potential referents). It is, thus, an open question whether children aged 4 to 5 years are indeed able to revise their trust retrospectively.

When children encounter conflicting novel information provided by two unfamiliar speakers, will they use subsequent information about the speakers' respective reliability to evaluate the speakers' earlier claims and selectively endorse the novel information by the speaker who turned out to be the more reliable source? In the present study, 5-year-old

children first experienced two unknown speakers who provided conflicting information about the referent of a novel label, with each speaker using the same novel label to refer exclusively to a different novel object. Only then did children learn about the speakers' labeling accuracy (i.e. that one of them labels familiar objects correctly, the other incorrectly).

Following the assumption that children selectively trust on the basis of trait reasoning, we predicted that they would be able to retrospectively infer that the speaker who turned out to be reliable initially provided the correct object-label link. Thus, when subsequently asked to select the correct referent for the novel label, children should be more likely to select the object initially labelled by the speaker who later turned out to be reliable than by the one who turned out to be unreliable.

We expect this effect both in a test in which objects are displayed together with the speakers, indicating selective endorsement of the label-object-link provided by the reliable speaker and in a subsequent test when the objects are displayed without the speakers being present, indicating source-independent selective endorsement of the label-object-link provided by the reliable speaker.

Crucially, we assume that children can only use the reliability of a speaker to make (retrospective) inferences if they also ascribe (un)reliability to each speaker in the first place. Therefore, only those children who were able to identify the speakers' reliability should selectively accept the object labelled by the reliable speaker as the referent of the novel label in the retrospective object selection tasks (both when speakers are present and absent). We do not expect such retrospective endorsement in children who were not able to correctly identify the speaker's reliability. This would also be in line with previous findings by Koenig and Harris (2005) as well as Hermes, Behne, and Rakoczy (2015).

After the retrospective object selection task (with both speakers present), children were asked to justify their object decision. Following the trait-based explanation of selective learning, we expect those children who refer to speaker reliability to be more likely to endorse

the information by the reliable than by the unreliable speaker in this retrospective task. In addition, we appended the “classic selective trust paradigm” (prospective task), in which children experienced the speakers’ labelling accuracy *before* encountering the conflicting novel information the two speakers provided. Replicating the classical selective trust finding (e.g., Koenig & Harris, 2005), children in this prospective task should also be more likely to endorse the information provided by the previously reliable than the previously unreliable speaker, but again, only if they are able to identify the speakers’ reliability.

Finally, we hypothesize that making retrospective inferences (in the context of selective trust) is more demanding than predicting future (in)accurate behavior on the basis of past reliability (see Scofield & Behrend, 2008). Thus, we predict that more children will pass the prospective than the retrospective object selection task. Further, children who pass the retrospective task should also pass the prospective task (i.e. classical selective trust paradigm) but children who do not pass the prospective task should tend to not pass the more demanding retrospective task. All hypotheses, procedural details, statistical tests, sample size and exclusion criteria were preregistered on OSF (https://osf.io/nxbu2/?view_only=b5cb3950268f4500bad9733018483265).

Method

Participants

Forty-two¹ children between 60 and 71 months ($M_{age} = 65.21$ months; 18 girls) were recruited from a database of parents who had volunteered to participate in studies on child development and came from mixed socio-economic backgrounds. Parents gave their written consent for the participation of their children. Ten additional children were tested but excluded from the analysis due to experimenter error ($n = 1$), technical error ($n = 2$), or failure to correctly answer the control questions ($n = 7$). These exclusion criteria had previously been

specified in our preregistered analysis plan. All children were tested at their daycare center in a quiet room.

Materials and design

Two novel two-syllable words, 'Mido' and 'Toma' were used for label training. Four novel objects were retrieved from the *Novel Object & Unusual Name (NOUN) Database* (Horst & Hout, 2014) based on their low familiarity and name-ability scores (see Fig. 1). During the accuracy exposure, children saw prototypical photographs of six familiar objects.

For the introduction, label trainings and accuracy exposure trials, video clips were prepared displaying two female actors on the left and right side of the screen (see Fig. 1). The lengths of each of the videos from the two speakers in each of the trials was matched. Further, the order of familiar objects named by the actors in the accuracy exposure trials was held constant across participants (baby, car, ball, bike, apple, dog).

Reliability of the speaker was counterbalanced across children regarding the following three aspects: The order in which the speakers appeared in label training and accuracy exposure (reliable first vs unreliable first), the actor assigned to a speaker role (speaker A vs speaker B), and the object assigned to a speaker role (orange and colorful vs blue and green).

Procedure

Object familiarization. Children were first familiarized with both objects through a "game" in which they had to catch each novel object on a screen. The experimenter then explained that she builds toys and, according to her manual, she now needed a part of a toy that is labelled 'Mido' and one of the objects they had just seen was such a 'Mido', but that the experimenter didn't know which one. She asked the child to help her find out which object was the Mido.

Label training retrospective task. Children saw both speakers introducing themselves one after the other in either the left or right upper corner of the screen, by saying "hello" and waving their hands. Then, the two novel objects appeared on the left and right side of the

screen underneath each of the speakers. Speaker 1 claimed "There is *one* 'Mido' here" (stressing her point by holding up her thumb to indicate 'one'), to emphasize that the objects could not both be 'Midos'. She gazed at both the objects and then pointed towards the one underneath her saying "*This* is the 'Mido'" (see Fig. 1). Speaker 2 repeated the same procedure using the same words and gestures, but claimed that the object underneath *her* was the 'Mido'. Afterwards, each speaker, in turn, looked at both objects again and claimed again "*This* is the 'Mido'", pointing to the object presented underneath each of them, respectively. The alternating naming was aimed to avoid order confounds and to establish a "balanced conflict" in which neither speaker had an immediate advantage. In order for the children to potentially select the object labelled by the speaker who was later revealed as more reliable, it was crucial that children remembered which speaker had labelled which object.

In a subsequent first episodic check, children were tested on their learned speaker-object association and prompted to point towards the respective object when asked "Which one did she [Speaker 1/2] say is the 'Mido'?". Only if children were able to answer this question correctly for both speakers, did the experimenter proceed with the accuracy exposure. If children failed at least in one of the two questions, the second part of the label training was repeated in which both speakers labelled each of their objects as 'Mido', after which children were given the same episodic check again. Seven children failed at the first episodic check and had to re-watch the label training. Three children did not pass this control question even after repetition of the label training and were thus excluded from the study.

Accuracy exposure. The experimenter left the child alone at the PC, pretending to need to organize something. Thus, the child was on her quest by herself and could not ask the experimenter for help and was considered the "expert" later on, being the only person to have watched the video. Each speaker labelled the same four familiar items, but one speaker consistently labelled each item inaccurately (e.g., calling the baby "a car") and the other consistently labelled each item accurately. In order to reinforce traits of specific speakers and

the contrast between them, they labelled these objects in the following order: a block of 3 items labelled by speaker 1, then the same block by speaker 2, then the last object first labelled by speaker 1 then speaker 2.

Test phase retrospective task. The experimenter returned and then asked the child the test question: "Which one is the 'Mido'?" referring to the screen showing the two speakers and objects (Endorsement test). If children hesitated to give an answer, the experimenter explained that she needed to know which of the two objects (pointing at both, always starting with the one on the left), the child thought was more likely to be the 'Mido'.

Children who made a choice were then asked for a reason for their decision (Justification question). Subsequently, to ensure that children still remembered which speaker had provided which kind of information, children went through another, second episodic check on their object-speaker association ("Which speaker labelled this object as a 'Mido' earlier on?"). Two children did not pass these control questions and were, thus, excluded from the study.

Two following Trait questions were aimed to test whether children were aware of the speakers' differing accuracy ("Which of the women is better at labelling things? And how did the other woman label things?"). Lastly, children were asked the test question "Which one is the 'Mido'?" again, but this time while being presented with the two objects on a sheet of paper in a different spatial arrangement (on top of each other instead of side by side, see Fig. 1) and without any of the speakers presented, neither on the sheet nor onscreen (Novel context test). This additionally tested if children had selectively acquired and retained an object-label link that they could use in a somewhat different context without any supporting cues to the original source of knowledge.

Prospective task. Next children participated in a classic selective trust task (e.g., Koenig & Harris, 2004), in which the information about the speakers' respective reliability is followed (not preceded) by the two speakers providing novel information. To set up this task,

the experimenter shortly explained to the child that she just realized that she also needed a second part, which was called 'Toma' but that she didn't know what a 'Toma' was. Children first experienced a rerun of the accuracy exposure. Each of the two speakers labelled the same two objects (car, dog), with the reliable speaker consistently labelling the two objects accurately and the unreliable speaker consistently doing so inaccurately (in alternating order).

Label training prospective task. Immediately after, two novel objects appeared underneath where the speakers were located onscreen. In succession, each speaker looked at both objects and claimed, "*This* is the 'Toma'", pointing to the object presented underneath each of them, respectively.

Test phase prospective task. Another, third, episodic check tested children's object-speaker association ("Which speaker labelled this part as a 'Toma'?"). Two children did not pass these control questions and were, thus, excluded from the study. The children who were able to pass these questions could proceed to the test question: "Which one is the 'Toma'?". If children gave an answer, they were also asked for a reason for their decision.

INSERT FIGURE 1 ABOUT HERE

Results

Retrospective task performance. Children selected the object initially labelled by the reliable speaker significantly above chance in both retrospective tests (Figure 2). In the Endorsement test, in which objects were shown in the presence of speakers, 78% of children selected the object initially labelled by the speaker who subsequently proved more reliable (29 out of the 37 children who made a choice, $p < .001$ one-sided binomial²). A similar pattern emerged in the subsequent Novel context test, in which the speakers were absent: 78% of children selected the object initially labelled by the reliable speaker (31 out of the 40 children who made a choice, $p < .001$, one-sided binomial). Hence, children selectively

endorsed the referent-label link provided by the speaker who later proved to be reliable rather than unreliable (see Figure 2). Children showed high consistency in their choice of object across the two retrospective tests (Linear weighted kappa = .93). In fact, those children who selected an object in both tests always selected the same object.

INSERT FIGURE 2 ABOUT HERE

Retrospective task and object choice justification. The majority of children who passed the retrospective tests justified their object selection with reference to the speakers' reliability as demonstrated in the accuracy exposure. Specifically, 16 out of the 29 children who passed the Endorsement test responded to the open question by explicitly referring to the speakers' reliability (see Table 1). Those children who explicitly referred to speaker performance in labelling familiar objects were significantly more likely to have endorsed the label provided by the reliable rather than the unreliable speaker ($p < .001$, one-sided McNemar), with the same pattern holding for the Novel context test ($p < .001$, one-sided McNemar).

INSERT TABLE 1 ABOUT HERE

Gaze duration analyses. Since label-object learning preceded the accuracy exposure, children's selective endorsement cannot be explained by an attentional bias developed during accuracy exposure. However, in the retrospective endorsement test directly following the accuracy exposure, pictures of the speakers were presented accompanied by their respective objects on the same side as in the preceding accuracy exposure (see Fig. 1). If children developed an attentional bias toward the reliable speaker in the accuracy exposure, this bias could potentially still transfer to the following test and lead children to select the object on the

side of the reliable speaker, without considering who initially labelled which object. Thus, we analyzed the relationship between children's gaze pattern during the accuracy exposure and their performance in the retrospective test. The first analysis, which we had planned, operationalized children's gaze duration during the accuracy exposure as a binary variable (looking longer at the reliable vs. the unreliable speaker). Children's object selection in the Endorsement test was not significantly related to their relative gaze duration at the speakers during accuracy exposure ($p = .405$, two-sided McNemar). Further, exploratory analyses that looked at children's gaze duration as a continuous variable showed the same pattern of results (see Figure 3). Thus, we could not detect any relation between children's visual attention during the accuracy exposure and their performance in the retrospective endorsement test.

INSERT FIGURE 3 ABOUT HERE

Prospective test performance and relation between prospective and retrospective tests.

In the prospective test, children selected the object labelled by the reliable speaker significantly above chance (30 of those 39 children who made a choice, $p < .001$, one-sided binomial³). 12 of these children explicitly referred to speaker performance in labelling familiar objects when asked to justify their object selection and those children were significantly more likely to have endorsed the label provided by the reliable rather than the unreliable speaker ($p < .001$, exploratory two-sided McNemar). Their pattern of performance in this test was similar to their performance in the retrospective tests (see Table 2). Contrary to our predictions, children were not more likely to fail a retrospective test but pass the prospective task than to show the reverse pattern of performance (Endorsement test: $p = .5$, one-sided McNemar; Novel context test: $p = .73$, one-sided McNemar).

INSERT TABLE 2 ABOUT HERE

Trait ascription and performance on the selective trust tasks. The trait ascription task was passed by 95% of children. In other words, all but two children answered both trait questions correctly. Hence, for each retrospective test, significantly more children were able to answer both trait questions, but still failed the retrospective test, than showed the reverse pattern of performance (Retrospective endorsement test, $p = .002$; Retrospective novel context test, $p = .006$; One-sided McNemar's). The same pattern holds when looking at the relation of the trait ascription task to the prospective task ($p = .003$, one-sided McNemar).

Discussion

In order to learn efficiently about the world around them, it is crucial for children to develop the ability to critically evaluate their sources and in case of conflict or uncertainty, select some information over others. Although a plethora of previous research suggests that preschool aged children selectively accept information from previously reliable over unreliable informants (see e.g., Harris, 2012), the cognitive processes underlying this selective behavior are still unclear. To further investigate which mechanism drives children's learning, we here used a reversed version of the classic selective learning paradigm and demonstrate that preschoolers use speaker reliability to selectively accept the more reliable source *in retrospect*. Children selectively endorsed the object-label-link initially provided by the speaker who later turned out to be reliable. This even transferred to a somewhat novel context that did not include the original sources.

These main findings using the reversed paradigm suggest that children's selective trust in the more reliable source cannot be easily explained by an attentional bias toward the reliable speaker biasing children to attend to new information s/he provides since label introduction was *prior* to accuracy exposure.

To pass the first endorsement test in which the two novel objects appeared underneath the speakers, one might still argue that it was possible for the children to merely select the

object on the side of the speaker they paid more attention to during accuracy exposure. However, we obtained the same reliability effect in the subsequent, Novel context test, in which objects were spatially rearranged and speakers' pictures were absent. Furthermore, additional analyses of gaze data did not find any relation between children's visual attention during the accuracy exposure and their performance in the retrospective endorsement test, further supporting the notion that their selective behavior might not be a result of attentional processes alone.

In contrast, more than half of the children justified their object selection with reference to speaker reliability, indicating advanced metacognitive awareness and the ability to explicitly reason about their selective trust in others based on the informants' individual competences. Thus, our findings are in line with conceptual accounts of selective trust (Sobel & Kushnir, 2013) including trait reasoning (Hermes et al., 2015). Children's retrospective weighting of information based on informants' reliability here are in accordance with related findings on children's learning from ignorant vs knowledgeable speakers (Mangardich & Sabbagh, 2018). Extending previous research, our study suggests that even younger children use trait reasoning to learn selectively from others and do this even in situations in which reliability has to be inferred (from previous accuracy) rather than explicitly stated (proclaimed knowledge / ignorance).

A very recent study supports this claim: Luchkina, Corriveau and Sobel (2018) presented preschoolers with one speaker who labelled two novel objects and then continued to label familiar objects. In a between-subject design, the speaker either turned out to be accurate (i.e., she labelled the familiar objects correctly) or inaccurate (she labelled the familiar objects incorrectly). Children relied on the novel information the speaker had provided significantly more in the accurate than the inaccurate speaker condition. Thus, using somewhat different methodologies, the research findings obtained all suggest that preschool-aged children

discount information from inaccurate speakers retrospectively (Luchkina et al. 2018, Scofield & Behrend, 2008).

How does children's performance in the retrospective task relate to their performance on standard selective trust tasks? In our prospective task, in which children received novel information after experiencing the speakers' accuracy, they also selectively endorsed the information provided by the reliable speaker. Thus, our study additionally replicates previous findings showing that 5-year-old children selectively predict future accuracy on basis of speakers' prior reliability (e.g., Koenig et al., 2004; Koenig & Sabbagh, 2013). We assumed that this prediction is less cognitively demanding than making retrospective inferences in selective trust. The prospective task could potentially be solved by associative learning on the basis of biased attention alone. However, the number of children passing both tasks was similar and passing the retrospective task was not predicted by children's performance on the prospective task. Although the prospective task might be cognitively less demanding, this advantage could be counterbalanced by the fact that there is a stronger pressure to resolve the (initially presented) conflict in the retrospective task, leading more children to use reliability as a first useful cue to resolve this conflict between competing referents for the same label. This might enable even younger children to deploy informant reliability to select useful information over others, which is for future research to follow up on. Similarly, further studies might disentangle if children use speaker reliability only if available (as in the present study) or actively search for informative clues in such situations of conflict.

Young learners' ability to be selective is crucial for social and, more generally, cultural learning (Harris, 2012; Mills, 2013). Being able to evaluate sources and selectively endorse information not only pro- but also retrospectively, allows for greater flexibility and ecological validity of selective trust. The present study thus further suggests that preschool aged children are already well equipped to use the characteristics of social agents to efficiently learn about the world around them.

Acknowledgements

We thank the families and their children for participating in the study. Further, we would like to thank Marlen Kaufmann, Joana Lonquich, and Konstanze Schirmer for their help with recruitment and coding.

Ethical Statement

Ethics approval was granted by the ethics committee of the Georg-Elias-Müller-Institute for Psychology, University of Göttingen (Project no. 213). Parents signed a written informed consent form for their child.

Funding statement

This research was supported by the German Research Foundation (DFG) as part of the ‘RTG 2070 Understanding Social Relationships’ and the Leibniz-WissenschaftsCampus ‘Primate Cognition’.

Data Accessibility

Preregistration, data, code, and stimuli are available via the Open Science Framework under https://osf.io/bqhrx/?view_only=e1e114eab28747ef87b518aeff3e2234

Competing Interests

We have no competing interests.

Authors' contributions

N.M. and T.B. had the original research idea and acquired funding. F.S., N.M. and T.B. designed the study. F.S. collected and analysed the data with T.B.'s support. F.S. and T.B. wrote the manuscript. F.S., N.M. and T.B. discussed the results and commented on the final version of the manuscript.

References

- Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*(3), 1018–1034. <https://doi.org/10.1016/j.cognition.2007.12.008>
- Clément, F. (2010). To Trust or not to Trust? Children's Social Epistemology. *Review of Philosophy and Psychology*, *1*(4), 531–549. <https://doi.org/10.1007/s13164-010-0022-3>
- Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge, MA: Harvard University Press.
- 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>
- Hermes, J., Behne, T., & Rakoczy, H. (2015). The role of trait reasoning in young children's selective trust. *Developmental Psychology*, *51*(11), 1574–1587. <https://doi.org/10.1037/dev0000042>
- 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. (2017). When does social learning become cultural learning? *Developmental Science*, *20*(2), e12350. <https://doi.org/10.1111/desc.12350>
- Horst, J. S., & Hout, M. C. (2014). The Novel Object and Unusual Name (NOUN) Database: a collection of novel images for use in experimental research. *Unpublished Manuscript*.
- Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best. *Psychological Science*, *17*(9), 757–758. <https://doi.org/10.1111/j.1467-9280.2006.01778.x>

- Kimura, K., & Gopnik, A. (2018). Rational Higher-Order Belief Revision in Young Children. *Child Development*, 00(0), 1–7. <https://doi.org/10.1111/cdev.13143>
- 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>
- 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., & Sabbagh, M. A. (2013). Selective social learning: new perspectives on learning from others. *Developmental Psychology*, 49(3), 399–403. <https://doi.org/10.1037/a0031619>
- Kushnir, T., Vredenburg, C., & Schneider, L. A. (2013). “Who can help me fix this toy?” The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology*, 49(3), 446–453. <https://doi.org/10.1037/a0031649>
- Luchkina, E., Corrieau, K. H., & Sobel, D. M. (2018, November 9). I don't believe what you said before: Preschoolers retrospectively discount information from inaccurate speakers. <https://doi.org/10.31219/osf.io/x2uaj>
- Macris, D. M., & Sobel, D. M. (2017). The Role of Evidence Diversity and Explanation in 4- and 5-Year-Olds' Resolution of Counterevidence. *Journal of Cognition and Development*, 18(3), 358–374. <https://doi.org/10.1080/15248372.2017.1323755>
- 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), 1–14. <https://doi.org/10.1111/desc.12544>
- Mills, C. M. (2013). Knowing when to doubt: developing a critical stance when learning from others. *Developmental Psychology*, 49(3), 404–418. <https://doi.org/10.1037/a0029500>
- Mundry, R. (2014). Statistical Issues and Assumptions of Phylogenetic Generalized Least Squares. In L. Z. Garamszegi (Ed.), *Modern Phylogenetic Comparative Methods and Their Application in Evolutionary Biology* (pp. 131–153). Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-662-43550-2_6
- 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>

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers Monitor the Relative Accuracy of Informants. *Developmental Psychology*, 43(5), 1216–1226. <https://doi.org/10.1037/0012-1649.43.5.1216>

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), 779–797. <https://doi.org/10.1037/a0034191>

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/10.1016/j.cogsci.2003.11.001>

Footnotes

¹ Required sample size based on a power analysis (with G*Power). Our goal was to obtain .95 power to detect a small effect size of .25 at the standard .05 alpha error probability with an exact binomial one-sided test.

² Note that we also preregistered a multiple logistic regression analysis in which object choice (blue/orange) would be predicted by the object labelled by reliable speaker (blue/orange) and performance on trait questions (pass/fail). As explained in the preregistration, this analysis should only be performed if the less frequent response (object selection: blue vs orange) occurred at least 20 times (No of predictors x 10; Mundry, 2014). Since the less frequent outcome (blue object selection) occurred only 18 times in both retrospective tests, these analyses were not performed.

³ Note that we also preregistered a multiple logistic regression analysis in which object choice (green/colorful) would be predicted by the object labelled by reliable speaker (green/colorful) and performance on trait questions (pass/fail). As explained in the preregistration, this analysis should only be performed if the less frequent response (object selection: green vs colorful) occurred at least 20 times (No of predictors x 10; Mundry, 2014). Since the less frequent outcome (colorful object selection) occurred only 16 times, these analyses were not performed.

Tables

Table 1.

		Explicit reference to speaker reliability	
		Yes	No
Retrospective endorsement test	Pass	16	13
	Fail	1	12
Retrospective novel context test	Pass	16	15
	Fail	1	10

Table 2.

		Prospective test	
		Pass	Fail
Retrospective endorsement test	Pass	22	7
	Fail	8	5
Retrospective novel context test	Pass	25	6
	Fail	5	6

Figures

Figure 1.

Figure 2.

Figure 3.

Figure and table captions

Figure 1. Procedural steps (left) with corresponding original screens (right). See main text for details of speakers' and experimenter's statements or question(s) in each step.

Figure 2. Object selection in the retrospective tests. In each retrospective test, significantly more children chose the object initially labelled by the reliable rather than the unreliable speaker. (In the Endorsement test, 37 children selected one of the two objects, 2 children selected both objects and 3 selected none; in the Novel context test 40 children made an object choice, 1 child selected both, and 1 child selected none.).

Figure 3. Standard boxplot displaying children's looking times at the reliable and unreliable speaker during the main accuracy exposure as a function of task performance in the retrospective endorsement test (i.e., with the speakers present). (Passing task = choosing the object initially labelled by the reliable speaker, failing = choosing the object initially labelled by the unreliable speaker, both, or none of the objects). Analyses showed that children who passed the Endorsement test did not look significantly longer to the reliable speaker ($M = 27.96s$; $SD = 5.92s$) than the unreliable speaker during the accuracy exposure ($M = 28s$; $SD = 6.63$; $p = .971$, paired t-test). Furthermore, children who endorsed the object-label-link provided by the reliable speaker did not look significantly longer to that speaker during accuracy exposure than children who did not endorse this link ($M = 27.84s$; $SD = 5.92$; $p = .946$, independent t-test).

Table 1. *Overview of children's object choice justifications (with vs without reference to the speakers' reliability) divided into test performances in both retrospective tests. Note that the one child that failed the retrospective test but still justified her decision with reference to the speakers' reliability is the same child for both endorsement and learning test.*

Table 2.. *Children's performance in the two retrospective tests as a function of their performance in the prospective task. Note that "fail" in all tests here includes all children who did not exclusively chose the object initially labelled by the reliable speaker, and thus, all children who chose the object initially labelled by the unreliable speaker as well as those who did not choose any object or chose both objects.*