Preschoolers continually adjust their epistemic trust based on an informant's ongoing accuracy

Samuel Ronfard

Harvard University

Jonathan D. Lane

Vanderbilt University

In press, *Child Development*

Correspondence concerning this article should be addressed to Samuel Ronfard.

Harvard Graduate School of Education
512 Larsen Hall, Appian Way
Cambridge, MA 02138.
Phone: 617 974 0215
Email: sar798@mail.harvard.edu

Acknowledgements. The study was conducted in the Living Laboratory at the Museum of Science, Boston and at the Lesley Ellis School, Arlington. We appreciate the museum and school staff and volunteers for their support, and we thank the parents and children for their participation. We also thank Tracy Elizabeth and Lynneth Solis for starring in our stimuli. Finally, we wish to thank Paul L. Harris and Patrick Shafto for comments on an earlier version of this manuscript, and three anonymous reviewers for their feedback on the submitted manuscript.

Abstract

Children aged 4 to 7 years (N = 120) played four rounds of a find-the-sticker game. For each round, an informant looked into two cups and made a claim about which cup held a sticker. At the end of each round, children guessed the sticker's location, and then the sticker's actual location was revealed. For three of the rounds, the informant accurately reported the sticker's location. But critically, for one round—either round 1, 2, or 3—she was inaccurate. Children continually adjusted their trust in the informant as they obtained more information about her accuracy. Relations between the informant's pattern of accuracy and children's trust were robust, neither mediated nor moderated by children inferences about her intent or traits.

Keywords: preschoolers, accuracy, epistemic trust

Preschoolers continually adjust their epistemic trust based on an informant's ongoing accuracy

Our ability to track the reliability of potential informants supports our ability to effectively learn from others. It allows us to seek information from people whose words and actions we can count on, and to appropriately trust (or mistrust) informants' claims (e.g., Harris, 2012). However, deciding whose information to trust and whose guidance to seek is challenging because individuals are not uniformly reliable or unreliable. In fact, their reliability is likely to fluctuate over time. As a result, we need to continually update our assessments of informants' reliability as we gain information about them.

Much recent work has been focused on identifying the criteria that children use to guide their epistemic trust, including informants' accuracy (e.g., Birch, Vauthier, & Bloom, 2008; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig, & Harris, 2007; Shafto, Eaves, Navarro, and Perfors, 2012), expertise (e.g., Landrum, Mills, & Johnston, 2013; Lane & Harris, 2015), in-group status (e.g., Kinzler, Corriveau, & Harris, 2011), and their personality traits and intent (e.g., Landrum et al., 2013; Lane, Wellman, & Gelman, 2013; Mascaro & Sperber, 2009). Relatively little work has examined children's ability to update their evaluation of a *single* informant's trustworthiness over time (but see: Jaswal, Croft, Setia, & Cole, 2010; Koenig & Woodward, 2010; Krogh-Jesperson & Echols, 2012; Nurmsoo & Robinson, 2009; Robinson & Nurmsoo, 2009; Sabbagh & Shafman, 2009; Vanderbilt, Liu, & Heyman, 2011).

In this study, we investigate the *trajectory* of children's trust in a single informant's verbal testimony over time. That is, we explore how children's trust in an informant *changes* as they gain knowledge about that informant's accuracy. This focus allows us to answer two questions: (1) Do children *continually adjust* their trust in an informant as they gain information about her accuracy? (2) Does the *timing* of an informant's inaccuracy– i.e., whether the informant is inaccurate the first time a child encounters her vs. if the informant is inaccurate in the child's most recent encounter with her—influence children's trust in that informant?

In the following sections, we review prior work on young children's ability to monitor and update their assessments of informants' reliability based on their accuracy. We then present alternative hypotheses on how the pattern of an informant's accuracy might influence children's trust.

The ability to use an informant's prior accuracy to guide learning is present in infancy and develops throughout the preschool years (Harris, 2012; Harris & Lane, 2014). Fourteen-month-olds prefer to track the gaze of a reliable rather than an unreliable looker (e.g., a looker who was excited when gazing into containers of toys rather than one who was excited when gazing into empty containers; Chow, Poulin-Dubois, & Lewis, 2008). A this age, infants also prefer to imitate actions from informants who are competent rather than incompetent (e.g., they imitate a model who knows how to put on shoes rather than a model who does not; Zmyj, Buttelmann, Carpenter, & Daum, 2010). Eighteen-month-olds can additionally use an informant's prior accuracy (when labelling familiar words) to learn *new* words and actions (Brooker & Poulin-Dubois, 2013). By 3-years of age, children can use informants' reliability in one domain (word labels) to decide whom to trust when learning about other domains, such as object functions (Koenig & Harris, 2005) and game rules (Rakoczy, Warneken, & Tomasello, 2009). By at least 4-years of age, children distinguish between informants who differ subtly in their accuracy. For example, they prefer the testimony of an informant who mislabeled 1 out of 4 objects rather than one who mislabeled 3 out of 4 objects (Pasquini et al., 2007). In sum, there is strong evidence that young children prefer to learn from more accurate and reliable informants over less accurate and reliable informants.

Most research that has revealed children's preference for more accurate informants has utilized a paradigm that requires children to choose between two informants who differ in their prior accuracy (e.g., Koenig & Harris, 2005; Pasquini et al., 2007). However, in their everyday interactions, children typically do not compare the conflicting testimony provided in succession by two or more informants; rather, they are often presented information by a single informant, often someone whom they have interacted with on many occasions (Lane & Harris, 2015; Vanderbilt, Heyman, & Liu, 2014). Thus, research examining the trajectory of children's trust in a *single* informant's verbal testimony is important because it reflects how children commonly confront new information. What little research has tracked children's trust in a single informant over time has been conducted with young preschoolers, revealing that 3- and 4-year-olds tend to the follow the testimony of an inaccurate informant, even after the informant has provided inaccurate information several trials in a row (e.g., Jaswal et al., 2010; Vanderbilt et al., 2011, 2014). Thus, it is largely unknown whether or how older preschoolers and school-age children update their trust in an informant as they interact with her and how the pattern of the informant's accuracies and inaccuracies influences children's trust. In the current study, we directly address these issues with children ranging from 4- to 7-years in age.

Pasquini and colleagues (2007) have argued that children older than 4-year-olds may distinguish between mostly-accurate and mostly-inaccurate speakers using one of two strategies: (1) a *raw frequency strategy* where children keep track of the *number* of inaccuracies, or (2) a *proportional strategy* where children keep track of the *proportion* of inaccuracies relative to the total amount of testimony. These two strategies lead to different predictions about the trajectory of children's trust in a single informant over time. If children track the number of inaccuracies rather than the proportion of inaccuracies then their trust in an informant should be the same when an informant is incorrect for 1-of-1 trials, 1-of-2 trials, or 1-of-3 trials. In contrast, if children track an informant's proportional inaccuracy then we would expect their trust to be lower at higher levels of proportional inaccuracy (i.e., least trusting when the informant is inaccurate for 1-of-1 trials compared to 1-of-2 trials, and most trusting when the informant is inaccurate for 1-of-3 trials).

Although these two strategies are plausible, they are not ideal from a social learning perspective because they do not account for all of the data that informants supply about their accuracy. Both the raw-frequency strategy and the proportional-inaccuracy strategy predict equivalent trust in an informant who has been *correct* for 0-of-0 trials, 1-of-1 trials, or 2-of-2 trials because these three informants do not differ in their number of inaccuracies or in their proportional inaccuracy. However, this prediction is unsatisfying. From an inferential statistical standpoint, someone who has provided correct information for 2-of-2 trials should be trusted more than someone who has been correct for 1-of-1 trials, and still more than someone who has been correct for 0-of-0 trials; if someone has been correct for 2-of-2 trials, we have more data from which to draw inferences and predict their behavior. This highlights the utility of a third strategy whereby children continually update their assessment of an informant's future testimony based on *all* of the information they have at their disposal (i.e., the number of incorrect *and* correct statements the informant has made). On this account, children's trust in informants should increase based on their proportional accuracy and based on their number of correct statements.

A second aim of the current study was to examine whether children's trust in an informant over multiple trials is subject to primacy or recency effects. That is, we asked whether children's trust in an informant would vary not only based on that informant's cumulative record of inaccuracy (and accuracy) but also as a function of *when* their inaccuracy occurred. For example, an informant's inaccuracy might be especially salient if it occurs the *first* time children receive information from that informant (a primacy effect) and might thus lead children to remain skeptical despite the informant's subsequent accuracy. Alternatively, trust in an informant might be lower if the informant has very *recently* provided incorrect information than if she had provided inaccurate information several trials ago (a recency effect). That is, because of its salience in memory, a recent inaccuracy might be especially powerful in reducing children's trust in an informant. Given the large literature demonstrating primacy and recency effects in adults' impression formation (e.g., Anderson & Barrios, 1961; Mayo & Crockett, 1964), we speculated that children's trust in an informant would vary based on the timing of that informant's inaccuracy.

To answer these two questions, the current experiment extends a paradigm used in prior research on children's epistemic trust (Jaswal et al., 2010; Vanderbilt et al., 2011), in which children are provided information by an informant and then decide whether or not to trust her. Note that this paradigm does not require children to trust one of two informants who differ in their prior accuracy (e.g., Koenig & Harris, 2005; Pasquini et al., 2007) but rather assesses how far they trust information provided by a single informant. Here, children played 4 rounds of a game in which an informant looks into two cups and states which cup contains a sticker. On each trial, children were asked to decide which cup contains the sticker, and then the contents of both cups were shown to children, thereby also revealing whether the informant had been accurate or inaccurate. Critically, the informant provided incorrect testimony (i.e., stated that the empty cup contained the sticker) on one of the four trials, and children were assigned to conditions such that the incorrect testimony was provided on either the first, second, or third trial.

Varying the timing of the informant's inaccuracy allowed us to more precisely test whether and how children continually update their assessment of an informant's accuracy. We hypothesized that children's acceptance of the informants' testimony would vary with both the overall pattern and the *amount* of data they had about the informant's accuracy. That is, children's trust in *inaccurate* informants should differ based on her proportional inaccuracy (i.e., inaccurate for 1-of-1 trials < 1-of-2 trials < 1-of-3 trials). In addition, when proportional inaccuracy is held constant, children's trust in informants should increase based on the number of times they have been correct (correct for 0-of-0 trials < 1-of-1 trials < 2-of-2 trials). This paradigm also allowed us to test for recency and primacy effects. If children's trust in informants is subject to a primacy effect then they should display lower trust on the very last trial (i.e., trial 4) if the informant's sole incorrect claim had been provided on the *very first* trial rather than on trial 2. On the other hand, if children's trust in informants is subject to a recency effect, they should display lower trust on the very last trial (i.e., trial 4) if the informant's sole incorrect statement had been provided on the *previous* trial (i.e., trial 3) rather than on trial 2.

To better understand the processes by which children's ongoing epistemic trust changes as a function of an informant's pattern of accuracy, we also asked children to make ongoing assessments of the informant's traits (i.e., niceness and smartness). Children rated the informant's traits at the very beginning of the study, immediately after she provided inaccurate information, and at the end of the study. This allowed us to test whether children's trait ratings mediate or moderate their trust in the informant. That is, children might interpret an informant's inaccuracies as a sign of meanness or as a lack of intelligence, and in turn this interpretation might reduce their trust in her. Indeed, preschoolers use behavioral information to make inferences about informants' traits (Boseovski & Lee, 2006; Boseovski, Chiu, & Marcovitch, 2013; Heyman, 2009), and informants' traits can influence children's epistemic trust (Lane et al., 2013). Five- and 6-year-old children require fewer exemplars than three- and four-year-old children to predict that actors who engage in negative behavior will do so again (e.g., Boseovski et al., 2013). Thus, we expected that after the informant had provided inaccurate information on a (single) trial, older (but not younger) children's evaluations of her niceness and intelligence would decrease, and that as a result their trust in the informant would also decrease to a greater extent than would younger children's trust.

Older preschoolers also adjust their trust in an informant based on the informant's intent (Mascaro & Sperber, 2009; Shafto et al., 2012; Vanderbilt et al., 2011). Thus, we also examined whether an informant's apparently *intentional* provision of inaccurate information would lead children to distrust her more than if she provided the inaccurate information by *mistake*. We experimentally manipulated the informant's apparent intent by having half of the children receive an apology from her after it became clear that she had provided inaccurate information ("I'm really sorry. I made a mistake"); the other half received no apology. We manipulated intent using an apology because children as young as 4-years understand the appeasement functions that explicit apologies serve (Smith, Chen, & Harris, 2010; Smith & Harris, 2012; Vaish, Carpenter, & Tomasello, 2011) and because children as young as 3-years can update their assessment of an informant's trustworthiness after they are provided information about the reason for the informant's inaccuracy (Nurmsoo & Robinson, 2009; Robinson & Nurmsoo, 2009).

Method

Participants

Children ranging from 4- to 7-years (*N* = 120) participated at a science museum in Boston (*n* = 97, 50 girls; M_{age} = 5 years, 10 months, *Range* = 4 years to 7 years and 11 months) and at a local preschool (*n* = 23, 13 girls; M_{age} = 5 years, *Range* = 4 years to 6 years). Two additional children were tested but excluded from analyses because they did not complete the study. Most families were of middle- to upper-middle socioeconomic status. Participants represented a variety of racial and ethnic backgrounds, but were predominately Caucasian.

Procedure

Participants were randomly assigned to watch an informant provide incorrect information on either the first trial (*n* = 39; 23 girls; M_{age} = 5 years, 9 months, *Range:* 4 years, 2 months – 7 years, 10 months), the second trial (*n* = 39; 19 girls; M_{age} = 5 years, 8 months, *Range:* 4 years – 7 years, 11 months), or the third trial (*n* = 42; 21 girls; M_{age} = 5 years, 7 months, *Range:* 4 years – 7 years, 10 months) of a find-the-sticker game. In addition, children were randomly assigned to either receive an apology after the provision of incorrect information (*n* = 63; 34 girls; M_{age} = 5 years, 7 months, *Range:* 4 years – 7 years, 10 months) or not receive an apology (*n* = 57; 29 girls; M_{age} = 5 years, 9 months, *Range:* 4 years – 7 years, 11 months). Children participated in four phases: (1) *meeting the players* (which included an initial trait assessment of the informant), (2) *learning the game*, (3) *playing the game*, and (4) *final trait assessment* of the informant. Phases were presented in a fixed order, and are described in greater detail below.

Meeting the players. Children were shown a picture of two women sitting on opposite sides of a table and were told that the woman facing the camera (and thus facing the participant) was named Tracy (the informant) and that the woman facing away from the camera was named

Lynn (the player). Lynn was included in the protocol in order to make the experiment consistent across children—Lynn always trusted Tracy, and thus it made sense for Tracy to offer an apology (for children in the Apology condition) regardless of whether or not a participant had trusted Tracy. As well, because children were shown pre-recorded videos, it made more sense for Tracy to interact with a play partner who was present in the same room. After being introduced to these two players, children were asked to identify Tracy and Lynn. If children misidentified the women, they were corrected; 95% (114/120) of children correctly identified both women on the first try. Then, children were shown a picture of Tracy and were asked to rate her niceness and smartness: “Do you think Tracy is nice/smart or mean/not smart?”, followed by,” Okay. Do you think Tracy is very ____ or just a little _____?”

Learning the game. Children were then shown the same picture of Tracy and Lynn sitting at a table, but now with two cups (one striped and one with dots) placed in front of Tracy. Children were told: “In this game, you and Lynn [*pointing to Lynn*] are going to try to find stickers hidden in cups. For each turn, there will be two cups on the table, and there will be *one* sticker hidden in *one* of the cups. Sometimes the sticker is hidden in the cup on this side [*point to left cup*] and sometimes it’s hidden in the cup on this side [*point to right cup*]. You and Lynn won’t look in the cups, but Tracy [*point*] *will* look in the cups. She’ll say something, about the sticker, and then you can choose where *you* think the sticker is. You can choose the same cup Tracy says or a different cup. It’s up to you. After you choose, Lynn will choose, and then everyone will get to see what’s inside the cups. Every time *you* pick the cup with the sticker, you get a point. We’ll use this sheet to show how many points you have.” This phrasing differs slightly from previous paradigms investigating children’s trust in verbal testimony. We chose to explicitly tell children that they could choose the same cup that the informant mentioned or a different cup in order to minimize the social pressure that children might feel to endorse the testimony of the informant. Children were then asked whether they were going to try to find as many stickers as they can. All children said “Yes.” They then saw a video of Lynn stating that she was going to try to find as many stickers as she can.

Playing the game. Children played four trials of the game. The informant provided correct information about the location of the sticker for 3 of the 4 trials. Children were randomly assigned to receive incorrect information about the location of the sticker on the 1st, 2nd, or 3rd trial. This allowed us to manipulate the frequency with which the informant had proven to be accurate prior to her provision of incorrect information. In addition, children were randomly assigned to either receive or not receive an apology from the informant immediately after she provided incorrect information. The trials were presented on videos, shot from the same perspective as the photos previously described, over the shoulder of Lynn (the player) who was sitting across a table from the informant (Tracy). The video began with an empty table. A third woman, whose face was not visible, placed in the middle of the table two different solid-colored cups upright (their contents were not visible to Lynn or to participants). The video was paused and children were asked about the color of each cup; children were corrected when necessary. The informant then stood up and slowly looked over into each cup. Whether the informant looked inside the right cup or the left cup first was counterbalanced across trials. The informant then sat down and made a statement about the location of the cup: “The sticker is in the [color] cup”. At that point the video was paused and children were asked: “Where do *you* think the sticker is?” The experimenter recorded children’s answer and said: “Let’s see what Lynn chooses”. Lynn then proceeded to choose the cup suggested by the informant. Lynn always followed the suggestion of the informant in order to create a context where it made sense for the informant to apologize on the trial when she provided inaccurate information. The woman who initially set up the cups came back and tilted each cup towards the camera, starting with the empty cup, saying "It's not in this one", followed by the cup with the sticker, saying "It's in this one." The experimenter started the next trial by saying: "Someone is going to put new cups on the table. Let's see what happens." On the trial when the informant provided inaccurate information, the video then differed depending on whether children were assigned to the apology or no apology condition. In the *no apology* condition, the informant looked up without saying a word. In the *apology* condition, the informant looked up and said: "I'm really sorry. I made a mistake." Immediately after the incorrect trial, children were asked to complete the same *niceness* and *smartness* ratings about the informant that they completed before playing the game. Children were also asked about the informant's *intent*: "Do you think Tracy named that cup *on purpose* or by *mistake*?" The three trials when the informant provided a correct recommendation were identical except that children did not answer the trait or intent questions. The actual location of the sticker on each trial (left or right cup) was counterbalanced across participants. To do this, we generated two orders for the location of the sticker that did not follow a predictable patter (left, right, right, left, and right, left, left, right) and randomly assigned children to these patterns. Each trial used different colored cups, and thus no two colors were used on more than one trial.

After children had completed all four trials of the game, they were asked to rate the niceness and smartness of the informant a third, and final time using the same questions as in the initial assessment.

Results

Our focal analyses are centered on children's trust in the informant's testimony, but we begin by examining whether children's trait attributions and their attribution of intent varied based on the informant's pattern of inaccuracy and the informant's provision of an apology following her inaccuracy. Trait and intent attributions could conceivably mediate effects of the informant's prior accuracy and her provision of an apology on children's trust. Thus, we first explore whether trait and intent attributions are directly related to the informant's prior accuracy or to the informant's provision of an apology. Following these analyses, we turn to our primary analysis of how the informant's pattern of accuracy—as well as children's trait inferences, and their intent attributions—influence children's trust in the informant. Throughout our analyses, we model the informant's history of accuracy as a categorical predictor rather than as a continuous predictor. By modeling patterns of accuracy as a categorical predictor we make no assumptions as to the shape of the relations between patterns of accuracy and our outcomes of interest. We report results for logistic regression models using odds-ratios. We report the same results in log-odds in the Supporting Online Materials 2, 4, 6a, 6b, and 8. All logistic analyses were conducted using the –logit– command in Stata 14.

Trait attribution

Children were asked to rate the informant's niceness and smartness three times: (1) at the start of the experiment, (2) immediately after they had received incorrect information, and (3) at the end of the experiment. To examine whether the timing of the informant's inaccuracy and the provision of an apology influenced children's ratings of the informant's traits, we converted children's trait ratings at each time point into numerical scores and treated them as a continuous variable: "very mean / not smart" was scored as 0; "just a little mean / not smart" was scored as .33; "just a little nice / smart" was scored as .67; and "very nice / smart" was scored as 1.

Receiving an apology and the timing of the inacuracy on children's trait ratings. To investigate the effect of the informant's apology and of the timing of her inacuracy on children's trait ratings, a multilevel linear regression model (using the –xtmixed– command in Stata 14) regressed children's trait ratings on the predictors of interest. A multilevel model was used to account for repeated measures: children were asked about two traits (i.e., niceness and smartness) at three time points (i.e., baseline, post-inaccuracy, final). Analyses revealed that there were no interactions between trait type and other predictors and no main effect of trait type. Nonetheless, trait type was retained as a control variable in all models. The results of these analyses are displayed in Table 1. Coefficients represent mean differences in trait ratings.

Two models were tested. The first model tested whether differences in children's trait ratings at different time points (i.e., baseline, after receiving inaccurate information, and the end of the experiment) differed as a function of children's assignment to the Apology and Prior-Accuracy conditions. In the second model, we tested whether differences in children's trait ratings at these different time points varied with age. Previous research has found that older children need less behavioral evidence to make trait inferences then younger children (Boseovski et al., 2013). Thus, we predicted that the provision of incorrect information might have a greater impact on older than younger children's trait attributions.

The first model included trait type, age, whether children received an apology (or not), whether children were assigned to receive inaccurate information on the 1st, 2nd, or 3rd trial (with trial 1 as the reference category), and the interaction between the timing of the inaccuracy and apology. Finally, to test whether differences in children's trait ratings at the three time points differed as a function of their assignment to the six conditions, the timing of children's trait assessments (baseline, post-inaccuracy, final) was included as a categorical predictor (baseline assessments were the reference category) and the model also included interactions between the timing of children's trait assessments and their assignment to the six conditions (i.e., the two apology conditions and 3 inaccuracy timing conditions). This set of interactions failed to predict a significant amount of additional variance as assessed by a General Linear Test, $\chi^2(10) = 13.76$, $p = .18$. Thus, the impact of the timing of children's trait assessments (whether it came at the start of the experiment, immediately after the receipt of incorrect information, or at the end of the experiment) on their trait ratings did not differ based on whether children received an apology or experienced the inaccuracy on their first, second, or third trial. These interactions were thus removed from the model. The resulting model is presented in Table 1 (Model 1).

This model reveals a significant interaction between the timing of the inaccuracy and receiving an apology. This interaction is displayed in Supporting Online Information 1. Children who received incorrect information on the third trial *and* then received an apology provided trait ratings that were, on average, significantly lower than those provided by children assigned to the other conditions. Because the previous modeling demonstrated that differences in children's trait ratings at the three time points *did not* interact with children's condition assignment, this difference between conditions likely reflects baseline differences in children's trait ratings.

More importantly, this model reveals that children's ratings of the informant's traits varied significantly based on the timing of those ratings (i.e., whether ratings were made at baseline, after the informant's inaccuracy, or after the very last trial). Children provided significantly lower trait ratings immediately after they had received incorrect information (see Table 1, Model1). However, children's ratings of the informant's traits recovered by the final assessment, and were not statistically different from their baseline ratings.

The second model tested for an interaction between age and the timing of children's trait ratings (see Table 1, Model 2). Immediately after receiving incorrect information older children lowered their ratings of the informant more than younger children. However, by the final assessment, there were no differences between younger and older children's trait ratings. This interaction is displayed in Figure 1. Inspection of Figure 1 illustrates this pattern and shows that even though older children's ratings of the informant's traits were lower after she provided inaccurate information, such ratings remained relatively positive. On average, children's ratings never dropped below "just a little nice / smart".

In summary, regardless of whether children received an apology, or whether they received incorrect information on their first, second, or third trial, children's ratings of the informant's smartness and niceness decreased significantly after they received incorrect information relative to their baseline ratings. This decrease was most pronounced among the oldest children. However, this decrease was short lived. By the end of the experiment, children's ratings of the informant's traits had increased significantly relative to their ratings immediately after receiving incorrect information and did not differ from their baseline ratings.

Attribution of intent to the informant

Recall that immediately after learning that the informant had provided incorrect information children were asked whether she had provided that information on purpose or by mistake. We examined whether children's attributions of intent [i.e., their judgments that she provided incorrect information unintentionally ("by mistake," coded as 1) rather than intentionally ("on purpose," coded as 0)] were related to the informant's proportional accuracy (entered as a categorical variable; 0-for-1, 1-for-2, 2-for-3; with 0-for-1 as the reference category), whether she apologized for her inaccuracy, children's attribution of smartness and niceness (after her inaccuracy), and children's age and gender. Table 2 displays parameter estimates using odds ratios with corresponding 95% confidence intervals (parameters estimates in log-odds with corresponding standard errors are provided in Supporting Online Information 2). We began by entering all 7 predictors into a logistic regression and found that the only two significant predictors were children's age and their receipt of an apology. Children were more likely to say that the informant made a mistake if they were younger and received an apology. In subsequent models, not pictured in Table 2, we found no significant interaction between gender and receiving an apology or between age and receiving an apology. Thus, the informant's apology conveyed to children that her inaccuracy was unintentional, regardless of children's age or gender.

Children's trust in the informant's testimony

Our focal aim was to examine how children's trust in the informant—as indexed by their tendency to select the cup that she indicated—varied depending on the informant's prior accuracy. We also considered whether children's trust in the informant was affected by whether they received an apology as well as their attributions of traits and intent to the informant.

To answer these questions, we conducted three sets of analyses. We first examined children's trust in the informant on the trial *immediately after* the trial when she provided inaccurate information (i.e., when the informant had, up to that point, been accurate for 0, 1 or 2 trials). Next, we examined whether the timing of the informant's inaccuracy influenced children's trust, by examining their trust in the informant on the *last* trial, at which point the informant had the same proportional accuracy across all conditions (2 out of 3 trials). This allowed us to test whether children's trust in the informant was subject to recency or primacy effects or whether children simply took into account the overall accuracy of the informant.

Finally, we pooled data across conditions to more precisely examine how children's trust in the informant varied as children obtained more information about her. We describe results from each of these analyses in turn. Note that preliminary analyses (detailed in Supporting Online Information 3) revealed no significant differences in children's trust in the first cup as a function of children's assignment to the Prior Accuracy and Apology conditions.

Trust immediately after the provision of incorrect information. Table 3 displays parameter estimates (odds-ratios) with corresponding 95% confidence intervals for logistic regression models predicting children's trust in the informant's recommendation immediately after she had provided incorrect information (parameters estimates in log-odds with corresponding standard errors are provided in Supporting Online Information 4). The informant's proportional accuracy was entered as a categorical variable (0-for-1, 1-for-2, 2-for-3), with 1-for-2 as the reference category. Therefore, the coefficient for 0-for-1 is the ratio of the odds that a child trusted the informant's recommendation when she was 0-for-1 rather than 1-for-2, and the coefficient for 2-for-3 is the ratio of the odds that a child trusted the informant's recommendation when she was 2-for-3 rather than 1-for-2. A significant coefficient below 1.00 would indicate lower trust in the informant's claim at that level of accuracy relative to when the informant was 1-for-2. A significant coefficient above 1.00 would indicate greater trust in the informant's claim at that level of accuracy relative to when the informant was 1-for-2.

Our analytic strategy was to first introduce our main predictor of interest (i.e., the informant's proportional accuracy) before testing for potential mediators or interaction effects of interest. Model 1 included the informant's proportional accuracy, which significantly predicted children's trust in the informant (see Table 3, Model 1). In Model 2, we controlled for children's age, gender, whether they received an apology, the intent children attributed to the informant, and the traits (smartness and niceness) that children attributed to the informant (following the informant's provision of incorrect information). Accounting for these factors, the informant's proportional accuracy remained a significant predictor of children's trust in the informant, and none of these additional factors significantly predicted children's trust (see Table 3, Model 2). Children trusted the informant's recommendation significantly more when she was 1-for-2 rather than 0-for-1, and when she was 2-for-3 rather than 1-for-2. In subsequent models, we tested a selected number of interactions (not presented in Table 3, but detailed in Supporting Online Materials 5), none of which were statistically significant. Thus, as displayed in Figure 2, children's trust in the informant following an inaccuracy was greatest if she had made that inaccurate claim in later trials, and thus had greater proportional accuracy. This effect was quite robust; it was not mediated by or moderated by children's age, trait attributions, attributions of intent, or the informant's apology.

Trust in the informant's recommendation on the last trial. In Table 4, we display parameter estimates (odds-ratios) with corresponding 95% confidence intervals for fitted logistic regression models predicting children's trust in the informant's recommendation on trial 4 (parameters estimates in log-odds with corresponding standard errors are provided in Supporting Online Information 6a). The timing of the informant's provision of inaccurate information was entered as a categorical variable (1^{st}, 2^{nd}, or 3^{rd} trial), with trial 2 as the reference category. This allowed us to test for primacy and recency effects. In Table 4, the coefficient for the 1^{st} trial represents the ratio of the odds that children trusted the informant's recommendation on the final trial when the informant provided incorrect information on the 1^{st} trial rather than on the 2^{nd} trial. A significant coefficient below 1.00 would indicate a primacy effect; that is, lower trust when the inaccurate information is provided in the first encounter with the informant rather than in the second encounter. The coefficient for the 3rd trial represents the ratio of the odds that children trusted the informant's recommendation on the final trial when the informant provided incorrect information on the 3rd trial rather than on the 2nd trial. A significant coefficient below 1.00 would indicate a recency effect; that is, lower trust when the inaccurate information was provided in the most recent encounter (i.e., trial 3) rather than in an earlier encounter (i.e., trial 2).

Our analytic strategy was to first introduce our main predictor of interest (i.e., the timing of the informant's inaccurate statement) before testing for potential mediators or interaction effects of interest. In our first model, we examined whether there was evidence for primacy or recency effects, by entering the timing of the informant's inaccurate claim (see Table 4, Model 1). There were no significant differences in children's trust in the informant on the fourth trial as a function of the timing of the inaccuracy--regardless of whether children had received incorrect information on trial 1, 2, or 3, they displayed a similar level trust on the fourth trial, by which time the informant had been accurate on 2 of the 3 previous trials. Thus, children's trust in the informant showed no signs of primacy or recency effects. In Model 2, we added children's age, gender, whether they received an apology (coded as 1) or not (coded as 0), the intent children attributed to the informant, and the traits (smartness and niceness) that children attributed to the informant (following the informant's provision of incorrect information). With the addition of these variables, the coefficient on the 3rd trial became statistically significant (see Table 4, model 2). Children were more likely to trust the informant on the last cup if she had been incorrect on the 3rd rather than 2nd trial. This rules out the possibility that there is a recency effect of having been provided with inaccurate information. Such an effect would predict that children would trust the informant *less* if she had provided incorrect information on the previous trial rather than two trials prior. However, it raises the possibility that even though children had the same statistical evidence across conditions on the 4th cup (an informant who was 2-for-3) children were more trusting of an informant who provided incorrect information later rather than earlier (a primacy effect). Such an effect seems unlikely given that we found that children were equally trusting on the 4th cup whether the informant provided incorrect information on the first or second cup. Nonetheless, we tested whether children were more trusting on the 4th cup if the informant provided incorrect information on the third cup rather than the first cup. We found no significant differences (the complete regression table for this analysis is provided in Supporting Online Information 6b). Thus, the only significant difference on the 4th trial is between children who received incorrect information on the 3rd cup (greater trust) rather than on the 2nd cup (less trust) when we control for age, gender, receiving an apology, the intent children attributed to the informant, and the traits (smartness and niceness) that children attributed. Note that this difference is not statistically significant without these controls. Thus, we find no evidence for robust recency or primacy effects on children's trust in the informant on the last trial.

Of the additional variables we added in Model 2, only children's age and their final attribution of smartness were significant predictors of trust on the last trial. Trust on the last trial was higher among younger children, and among children who attributed greater smartness to the informant. In subsequent models, we tested a selected number of interactions (not included in Table 4, but detailed in Supporting Online Information 7), none of which were statistically significant. In sum, on the last trial, trust was higher among younger children and among children who interpreted the informant as smarter. However, there was little evidence that the timing of the informant's inaccuracy affected children's final level of trust—there was neither a primacy effect nor a recency effect.

Overall influence of prior accuracy on trust in the informant's recommendation. Because we found no effect of the informant's apology on children's trust, we combined data across the two Apology conditions to more precisely examine how children's trust in the informant varied as children obtained more information about her. We binned children's decisions to trust the informant across conditions based on the informant's prior number of accuracies and inaccuracies. This allowed us to examine children's trust in the informant when she was 0-for-0, 0-for-1, 1-for-1, 1-for-2, 2-for-2, or 2-for-3. All 120 children provided data on the first trial at which point the informant had previously been accurate on zero of zero trials (0-for-0) and on the last trial at which point she had been accurate on two of three trials (2-for-3). Children for whom the informant was incorrect on the first trial (n = 39) provided data on their trust when her prior accuracy had been 0-for-1 (on their 2nd trial) and 1-for-2 (on their 3rd trial). Likewise, children for whom the informant was incorrect on the second trial (n = 39) provided data when her prior accuracy had been 1-for-1 (on their 2nd trial) and 1-for-2 (on their 3rd trial). Finally, children for whom the informant was incorrect on the third trial (n = 42) provided data for when the informant's prior accuracy had been 1-for-1 (on their 2nd trial) and 2-for-2 (on their 3rd trial). In Figure 3, we display children's trust in the informant across six different accuracy ratios, notably when the informant's prior accuracy had been 0-for-0 (n = 120), 0-for-1 (n = 39), 1-for-1 (n = 81), 1-for-2 (n = 78), 2-for-2 (n = 42), or 2-for-3 (n = 120).

Inspection of Figure 3 reveals that (1) children were inclined to trust the informant when they did not yet have any information about her accuracy – when she was 0-for-0, and (2) children's subsequent decisions to heed the informant's advice were strongly related to her pattern of accuracy. When children did not have any information about the informant's accuracy, i.e., when she was 0-for-0, children trusted her significantly above chance (50%; Binomial Test,

$p < .001$). In contrast, when her prior accuracy was 0-for-1, children *distrusted* her significantly below chance (Binomial Test, $p < .01$), and when her prior accuracy was 1-for-2, children's trust was no different than chance, Binomial Test, $p = .99$. When her prior accuracy was 2-for-3, 1-for-1, or 2-for-2, children trusted her well above chance levels, Binomial Tests, $p < .001$.

In order to compare children's initial trust in the informant (0-for-0) to different levels of prior accuracy, we regressed children's trust in the informant's claim on the six different prior accuracy ratios (as categorical predictors) using a logistic regression model with 0-for-0 as the reference category. In Table 5, we display parameter estimates (odds-ratios) with corresponding 95% confidence intervals for this logistic regression (parameters estimates in log-odds with corresponding standard errors are provided in Supporting Online Information 8). This allowed us to compare children's initial trust in the informant (when she was 0-for-0) and their subsequent trust in her at different accuracy ratios. For example, the coefficient on 0-for-1 represents the ratio of the odds that children trusted the informant's recommendation when the informant was 0-for-1 rather than 0-for-0. Similarly, the coefficient on 1-for-2 represents the ratio of the odds that children trusted the informant's recommendation when the informant was 1-for-2 rather than 0-for-0. Thus, significant coefficients below 1.00 represent lower trust on that trial relative to baseline (when the informant was 0-for-0) and coefficients above 1.00 represent greater trust on that trial relative to baseline.

As depicted in Table 5, children's initial trust in the informant was significantly greater than their trust when she was 0-for-1, and 1-for-2. Children's trust in the informant when she was 2-for-3 was no different from their initial trust in her. Children's trust in the informant when she was 1-for-1 was greater than their initial trust in her (though that difference was not statistically significant at an $\alpha = .05$), and their trust in the informant when she was 2-for-2 was significantly greater than their initial trust in her. Thus, children's trust in the informant's recommendation varied depending on the informant's pattern of accuracy. Children began with considerable trust in the informant but reduced their trust if she proved inaccurate, especially if she was inaccurate on one of the first two trials. As the informant's accuracy increased, so did children's trust.

Discussion

In order to learn effectively from others' testimony, children need to identify whether or not their informants are trustworthy (Harris, 2012). This already difficult task is complicated by the fact that an informant's accuracy is typically not fixed; rather, an informant's accuracy changes over time. The current study was designed to answer two main questions: (1) Do children continually adjust their trust in an informant as they gain more information about her accuracy? (2) Does the timing of an informant's inaccuracy—e.g., whether the informant is inaccurate the first time the child encounters her vs. if she is inaccurate in children's most recent encounter with her—influence children's trust in that informant? We found that children continually updated their trust in the informant as they gained more information about her accuracy (and inaccuracy) across multiple episodes. Moreover, children's decisions to trust the informant did not reveal recency or primacy effects. Instead, children's trust in the informant's testimony was equivalent when she had been correct 3 of 4 times regardless of *when* her single inaccuracy had occurred. Children used the informant's behavior (i.e., her accuracy and whether or not she apologized for her inaccuracy) to infer her intent, niceness, and smartness. However, these attributions did not mediate or moderate the observed differences in trust associated with the informant's history of accuracy. In the following sections, we discuss these findings and their implications in detail.

The informant's prior accuracy had a strong and robust influence on children's epistemic trust. Our results are therefore consistent with earlier research on the relation between an informant's accuracy and children's trust in that informant; children are more trusting of claims offered by more accurate informants (e.g., Birch et al., 2008; Koenig & Harris, 2005; Pasquini et al., 2007). However, our findings also extend this research in three ways. First, we find that children's trust in the claims provided by a single informant changes as children learn more about the informant's accuracy. This was true for children across the entire age range of 4- to 7-years. These results go beyond findings from studies, discussed earlier, that have focused on children's trust in conflicting claims made by multiple informants—learning situations that children will rarely encounter outside of the laboratory. Indeed, children typically learn from one informant at a time; often someone whom they interact with over a prolonged period (e.g., a parent, teacher, or sibling). Drawing on the current findings, we now have evidence suggesting that preschoolers and young school-age children continually update their impressions of such informants' trustworthiness over time, as a function of their accuracy.

Second, our results also provide insight into *how* children use accuracy information when learning from other people. Children's trust in the informant's claims varied with her proportional accuracy, suggesting that children were *not* simply basing their decisions on the informant's *number* of inaccuracies (i.e., a raw-frequency strategy). In addition, children's greater trust in the informant's recommendation as a function of the number of times she was *accurate* (i.e., 2-for-2 > 1-for-1 > 0-for-0) indicates that children could not have relied *solely* on a proportional-inaccuracy strategy because in these three cases the informant's proportional inaccuracy did not differ. Thus, children are either using two different strategies (i.e., a proportional inaccuracy strategy for informants who have been inaccurate, and an accuracy-counting strategy when the informant has never been inaccurate) or they are using a statistical strategy that takes into account all of the information at their disposal (i.e., the total number of accurate and inaccurate statements made by the informant). We currently favor the latter hypothesis. A single strategy seems more efficient and more capable of dealing with real-world learning situations where no one informant is always accurate. A challenge for the two-strategies hypothesis is to explain when and how children switch from using one strategy to another. For example, if the informant is accurate two times in a row but then makes a mistake, children would need to switch over from the accuracy-counting strategy to the proportional-inaccuracy strategy and transfer information about the number of accurate trials to calculate the proportional inaccuracy. Research on children's trust in verbal testimony has benefitted from attempts to formalize the processes underlying children's trust using Bayesian approaches (e.g., Shafto, Eaves, Navarro, & Perfors, 2012). In addition to further experimental work, statistical modelling of such data might be helpful in determining whether and when children are using different strategies when tracking the accuracy of an informant over time.

A third important finding is that children *regained* trust in a previously inaccurate informant who went on to offer accurate information, regardless of when she had provided inaccurate information. That is, we found no evidence of recency or primacy effects. This means that children not only update their inferences about the reliability of a previously inaccurate informant when they are provided a legitimate reason for why the informant had been inaccurate (e.g., if they learn that she held a false belief; Nurmsoo & Robinson, 2009; Robinson & Nurmsoo, 2009) but also when they subsequently have positive epistemic evidence about her.

Consistent with prior research, children used the informant's behavior (i.e., her accuracy and whether or not she apologized) to infer her intent, niceness, and smartness (e.g., Heyman, 2009; Mills, 2013). We found that after the informant provided inaccurate information on a single trial, older (but not younger) children's evaluations of her niceness and intelligence decreased. This is consistent with prior research reporting that older children need fewer behavioral examples than younger children to make trait attributions (Boseovski et al., 2013). However, these attributions did not moderate or mediate the observed differences in children's trust associated with the informant's history of accuracy. It might be the case that an informant's intent and traits simply have smaller effects than accuracy on children's trust, particularly when children have a small sample of data to draw from. Indeed, an effect of smartness on trust did emerge on the last trial, at which point all children had identical information about her proportional inaccuracy: children who judged the informant to be smarter trusted her more than children who judged her to be less smart. Thus, when accuracy was held constant, children relied on other aspects of the informant's profile to make epistemic inferences. However, it is possible that had we manipulated the informant's traits and intent more strongly (for example, by having the informant provide inaccurate information multiple times or by emphasizing the informant's meanness or niceness) those factors would have had stronger influences on children's trust in the informant. For example, it is possible that the failure of the informant's apology to influence children's trust decisions indicates that the informant's apology did not adequately account for her inaccuracy. It is possible that providing a reason for her mistake, rather than simply acknowledging it, might influence children's trust. This could be tested by manipulating whether informants offer apologies with an explanation for their mistake: "I'm sorry I made a mistake" vs. "I'm sorry but I couldn't see the sticker".

Our results raise several additional questions for future research. For one, do children continually monitor informants' accuracy across domains to generate a holistic accuracy index for their informants or do they monitor accuracy within domains to generate domain-specific accuracy indices? Because preschoolers take into account informants' expertise when deciding whether or not to trust them (Landrum et al., 2013; Lane & Harris, 2015), we suspect that children might track the accuracy of individuals within specific domains. Such a strategy would be more benefical than simply tracking accuracy across domains because an informant might be highly accurate in one domain but relatively innacurate in others.

Another question is whether children will regain trust in informants who have provided inaccurate information on multiple occasions or whose inaccuracies result in more substantial consequences for the child. If children do regain trust, how quickly do they do so? Would we have seen the same recovery of trust if the informant had been wrong multiple times in a row?

A third question pertains to the development of children's ability to track the proportional accuracy of informants. Pasquini et al. (2007) found that 4-year-olds but not 3-year-olds prefer to learn novel name from someone who has been correct 3 out of 4 times rather than from someone who has been correct 1 out of 4 times. Is this age-related difference a function of younger children's inability to track informant's proportional accuracy or simply their inability to track the proportional accuracy of *two* informants simultaneously?

We chose to present children with videos of our informant rather than have children interact with "live" informants. We made this decision because our protocol requires that the actors' behaviors were standardized across participants, and this was best achieved using pre-recorded video. More importantly, our protocol required that the informant interact with the confederate (Lynn) so that it would always make sense for the informant to apologize following the provision of inaccurate information in the Apology condition (as Lynn always followed the informant's testimony and was thus always incorrect on that trial). Nevertheless, we do not have reason to believe that children would have behaved very differently in a live setting. In fact, a live game might have introduced greater noise because children would have felt more social pressure to follow the informant's testimony (even if they did not actually trust the informant).

In sum, our data suggest that young children track informants' accuracies over time and use that history and other behavioral information to build and update their cognitive profile of the informant – an aggregate of their prior accuracy, traits, and intent. This profile is most subject to change earlier in children's encounters with someone new, when they have less data to draw upon. However, as children gain more information, they adjust that profile accordingly. Thus, although there are circumstances in which young preschoolers are prone to follow the claims of inaccurate informants (e.g., Jaswal et al., 2010; Vanderbilt et al., 2011), by the end of the preschool years children's trust in informants is remarkably flexible and attuned to their informants' histories of accuracy. This ability supports children's social learning in an ever-changing social world.

References

Anderson, N. H., & Barrios, A. A. (1961). Primacy effects in personality impression formation. *The Journal of Abnormal and Social Psychology*, *63*, 346–350. doi: 10.1037/h0046719

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*, 1018–1034. doi: 10.1016/j.cognition.2007.12.008

Boseovski, J. J., Chiu, K., & Marcovitch, S. (2013). Integration of behavioral frequency and intention information in young children's trait attributions. *Social Development*, *22*, 38–57. doi: 10.1111/sode.12008

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioral prediction. *Developmental Psychology*, *42*, 500-513. doi: 0.1037/0012-1649.42.3.500

Brosseau-Liard, P. E., & Birch, S. A. (2010). 'I bet you know more and are nicer too!': what children infer from others' accuracy. *Developmental Science*, *13*, 772-778. doi: 10.1111/j.1467-7687.2009.00932.x

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. doi: 10.1111/infa.12027

Chow, V., Poulin-Dubois, D., & Lewis, J. (2008). To see or not to see: infants prefer to follow the gaze of a reliable looker. *Developmental Science*, *11*, 761–770. doi: 10.1111/j.1467-7687.2008.00726.x

Danovitch, J. H., & Keil, F. C. (2004). Should you ask a fisherman or a biologist?: Developmental shifts in ways of clustering knowledge. *Child Development*, *75*, 918–931. doi: 10.1111/j.1467-8624.2004.00714.x

Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge, MA: Belknap Press/Harvard University Press.

Harris, P. L., & Lane, J. D. (2014). Infants understand how testimony works. *Topoi: An International Review of Philosophy, 33*, 443-458. doi: 10.1007/s11245-013-9180-0

Heyman, G.D. (2009). Children's reasoning about traits. In P.J. Bauer (Ed.), *Advances in child development and behavior* (Vol. 37, pp. 105-143). New York: Academic Press.

Jaswal, V. K., Croft, A. C., Setia, A. R., & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science*, *21*, 1541–1547. doi: 10.1177/0956797610383438

Jaswal, V. K., Pérez-Edgar, K., Kondrad, R. L., Palmquist, C. M., Cole, C. A., & Cole, C. E. (2014). Can't stop believing: Inhibitory control and resistance to misleading testimony. *Developmental Science, 17*, 965-976. doi: 10.1111/desc.12187.

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers: Selective trust in native-accented speakers. *Developmental Science*, *14*, 106–111. doi: 10.1111/j.1467-7687.2010.00965.x

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*, 1261–1277. doi: 10.1111/j.1467-8624.2005.00849.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, 815. doi: 0.1037/a0019664

Krogh-Jespersen, S., & Echols, C. H. (2012). The influence of speaker reliability on first versus second label learning. *Child development*, *83*, 581-590. doi: 10.1111/j.1467-8624.2011.01713.x

Landrum, A. R., Mills, C. M., & Johnston, A. M. (2013). When do children trust the expert? Benevolence information influences children's trust more than expertise. *Developmental Science, 16*, 622-638. doi: 10.1111/desc.12059

Lane, J. D., Wellman, H. M., & Gelman, S. A. (2013). Informants' traits weigh heavily in young children's trust in testimony and in their epistemic inferences. *Child Development*, *84*, 1253-1268. doi: 10.1111/cdev.12029

Lane, J. D., & Harris, P. L. (2015). The roles of intuition and informants' expertise in children's epistemic trust. *Child Development*, *86*, 919-926. doi: 10.1111/cdev.12324

Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, *73*, 1073–1084. doi: 10.1111/1467-8624.00458

Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition*, *112*, 367–380. doi: 10.1016/j.cognition.2009.05.012

Mayo, C. W., & Crockett, W. H. (1964). Cognitive complexity and primacy-recency effects in impression formation. *The Journal of Abnormal and Social Psychology*, *68*, 335–338. doi: 10.1037/h0041716

Mills, C. M. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology*, *49*, 404-418. 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*, 23–27. doi: 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*, 1216–1226. doi: 10.1037/0012-1649.43.5.1216

Rakoczy, H., Warneken, F., & Tomasello, M. (2009). Young children's selective learning of rule games from reliable and unreliable models. *Cognitive Development*, *24*, 61–69. doi: 10.1016/j.cogdev.2008.07.004

Shafto, P., Eaves, B., Navarro, D.J. & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15, 436-447. doi: 0.1111/j.1467-7687.2012.01135.x

Tagar, M. R., Federico, C. M., Lyons, K. E., Ludeke, S., & Koenig, M. A. (2014). Heralding the authoritarian? Orientation toward authority in early childhood. *Psychological Science*, *25*, 883–892. doi: 10.1177/0956797613516470

Robinson, E. J., & Nurmsoo, E. (2009). When do children learn from unreliable speakers? *Cognitive Development*, *24*, 16–22. doi: 10.1016/j.cogdev.2008.08.001

Sabbagh, M. A., & Shafman, D. (2009). How children block learning from ignorant speakers. *Cognition*, *112*, 415-422. doi: 10.1016/j.cognition.2009.06.005

Shafto, P., Eaves, B., Navarro, D.J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15, 436-447. doi: 1111/j.1467-7687.2012.01135.x

Smith, C. E., Chen, D., & Harris, P. L. (2010). When the happy victimizer says sorry: Children's understanding of apology and emotion. *British Journal of Developmental Psychology*, *28*, 727–746. doi: 10.1348/026151009X475343

Smith, C. E., & Harris, P. L. (2012). He didn't want me to feel sad : Children's reactions to disappointment and apology. *Social Development*, *21*, 215–228. doi: 10.1111/j.1467-9507.2011.00606.x

Vaish, A., Carpenter, M., & Tomasello, M. (2011). Young children's responses to guilt displays. *Developmental Psychology*, *47*, 1248–1262. doi: 10.1037/a0024462

Vanderbilt, K. E., Heyman, G. D., & Liu, D. (2014). In the absence of conflicting testimony young children trust inaccurate informants. *Developmental Science*, *17*, 443–451. doi: 10.1111/desc.12134

Vanderbilt, K. E., Liu, D., & Heyman, G. D. (2011). The development of distrust. *Child Development*, *82*, 1372–1380. doi: 10.1111/j.1467-8624.2011.01629.x

Zmyj, N., Buttelmann, D., Carpenter, M., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. *Journal of Experimental Child Psychology*, *106*, 208–220. doi: 10.1016/j.jecp.2010.03.002

Table 1

Multi-level linear regression models predicting children's trait attributions to the informant.

	Model 1			Model 2		
	b	z scores	95% CI	*b*	z scores	95% CI
Trait Type	.03	1.71	.00, .06	.03	1.72	.00, .06
Age (Years)	-.01	-.81	-.04, .01	.00	.09	-.03, .03
Apology	-.02	-.50	-.12, .07	-.02	-.50	-.12, .07
Inaccurate 2nd trial	.08	1.62	-.02, .17	.08	1.62	-.02, .17
Inaccurate 3rd trial	.08	1.56	-.02, .17	.08	1.56	-.02, .17
Apology X Inaccurate 2nd trial	.05	.69	-.09, .18	.05	.69	-.09, .18
Apology X Inaccurate 3rd trial	-.18**	-2.67	-.31, -.05	-.18**	-2.67	-.31, -.05
Rating after inaccuracy	-10***	-5.61	-.14, -.07	.12	1.17	-.08, .31
Final rating	-.04	-1.95	-.07, .00	-.05	-.54	-.25, .14
Rating after inaccuracy X Age				-.04*	-2.25	-.07, -.01
Final rating X Age				.00	.18	-.03, .04
Constant	.91***	10.72	.74, 1.07	.84***	8.25	.64, 1.04
Random-Effects Variance		.015***			.015***	
Interclass Correlation		.27			.28	
X^2		71.07***			78.86***	
Model df		9			11	

* $p < .05$, ** $p < .01$, *** $p < .001$.

Note. $n_{children} = 120$; $n_{ratings} = 720$. Trait type is coded as 1 for niceness and 0 for smartness. Receiving an apology is coded as 1 and not receiving an apology is coded as 0. The timing of the inaccuracy is coded as a categorical variable (1st trial, 2nd trial, 3rd trial) with 1st trial as the reference category. Whether children made attributions about the informant's traits at baseline, after the inaccuracy, or at the end of the experiment is coded as a categorical variable with baseline ratings as the reference category.

Table 2

Logistic regression model predicting whether children believed the informant's provision of incorrect information was unintentional.

	Odds-Ratios	**z scores**	**95% CI**
Accuracy (1-for-2)	1.44	.62	.45, 4.67
Accuracy (2-for-3)	2.05	1.26	.67, 6.29
Age (years)	.32***	-4.34	.19, .53
Gender	.99	-.03	.39, 2.48
Apology	6.92 ***	3.66	2.46, 19.49
Niceness Rating Post Inaccuracy	2.91	1.10	.44, 19.31
Smartness Rating Post Inaccuracy	5.42	1.80	.86, 34.21
Constant	19.73	1.73	.67, 583.02
X^2		50.48***	
Model df		7	
-2 Log Likelihood		115.8	
Pseudo R^2		.30	

*** $p < .001$.

Note. n = 120. Children's interpretation of the informant's intent in communicating her inaccuracy was coded such that 1 = unintentional and 0 = intentionally misleading. The proportional inaccuracy of the informant is coded as a categorical variable (0-for-1, 1-for-2, 2-for-3) with 0-for-1 as the reference category. Receiving an apology is coded as 1 and not receiving an apology is coded as 0.

Table 3

Logistic regression models predicting children's trust in the informant's claim after the provision of incorrect information.

	Model 1			Model 2		
	Odds-Ratios	**z scores**	**95% CI**	**Odds-Ratios**	**z scores**	**95% CI**
Accuracy (0-for-1)	.33*	-2.29	.13, .85	.35*	-2.02	.13, .97
Accuracy (2-for-3)	5.7***	3.19	1.96, 16.59	6.36***	3.25	2.09, 19.37
Age (years)				1.26	1.02	.81, 1.98
Gender				.83	-.41	.35, 1.98
Apology				.73	-.66	.29, 1.84
Negative Intent				1.35	.56	.47, 3.85
Niceness Rating Post Inaccuracy				.82	-.22	.14, 4.85
Smartness Rating Post Inaccuracy				1.70	.56	.27, 10.66
Constant	1.05	.16	.56, 1.97	.27	-.75	.01, 8.29
X^2		32.26***			34.32***	
Model df		2			8	
-2 Log Likelihood		132.89			130.83	
Pseudo R^2		.20			.21	

* $p < .05$, *** $p < .001$.

Note. $n = 120$. The proportional inaccuracy of the informant is coded as a categorical variable (0-for-1, 1-for-2, 2-for-3) with 1-for-2 as the reference category. Receiving an apology is coded as 1 and not receiving an apology is coded as 0. Children's attribution of intent is coded as 1 if children reported that the informant provided incorrect information unintentionally and 0 if they believed she was intentionally misleading.

Table 4

Logistic regression models predicting children's trust in the informant's claim on the last cup (i.e., 4th cup).

	Model 1			Model 2		
	Odds-Ratios	**z scores**	**95% CI**	**Odds-Ratios**	**z scores**	**95% CI**
Inaccurate information (1st trial)	.89	-.24	.34, 2.30	1.76	.95	.55, 5.59
Inaccurate information (3rd trial)	2.67	1.75	.89, 8.00	3.94*	2.15	1.13, 13.76
Age (years)				.54*	-2.52	.33, .87
Gender				.47	-1.54	.18, 1.23
Apology				.89	-.23	.32, 2.43
Intent				1.75	.97	.56, 5.45
Niceness Rating Post Inaccuracy				1.02	.02	.16, 6.51
Smartness Rating Post Inaccuracy				8.71*	2.25	1.32, 57.51
Constant	2.25*	2.34	1.14, 4.44	33.80	1.88	.86, 1321.89
X^2		4.87			25.79**	
Model df		2			8	
-2 Log Likelihood		132.2			111.3	
Pseudo R^2		.04			.19	

* $p < .05$, ** $p < .01$.

Note. $n = 120$. The proportional inaccuracy of the informant is coded as a categorical variable (0-for-1, 1-for-2, 2-for-3) with 1-for-2 as the reference category. Receiving an apology is coded as 1 and not receiving an apology is coded as 0. Children's attribution of intent is coded as 1 if children reported that the informant provided incorrect information unintentionally and 0 if they believed she was intentionally misleading.

Table 5

Logistic regression model comparing children's trust in information provided by the informant at different ratios of accuracy, with 0-for-0 as the reference category.

	Odds-Ratios	z scores	95% CI
0-for-1	.17***	-4.32	.07, .37
1-for-2	.46 **	-2.62	.25, .82
2-for-3	1.38	1.13	.79, 2.42
1-for-1	1.81	1.77	.94, 3.50
2-for-2	4.57 **	2.71	1.52, 13.72
Constant	2.08***	3.75	1.42, 3.04
X^2		60.47***	
Model df		5	
-2 Log Likelihood		550.59	
Pseudo R^2		.10	

Note. $n = 480$

** $p < .01$, *** $p < .001$.

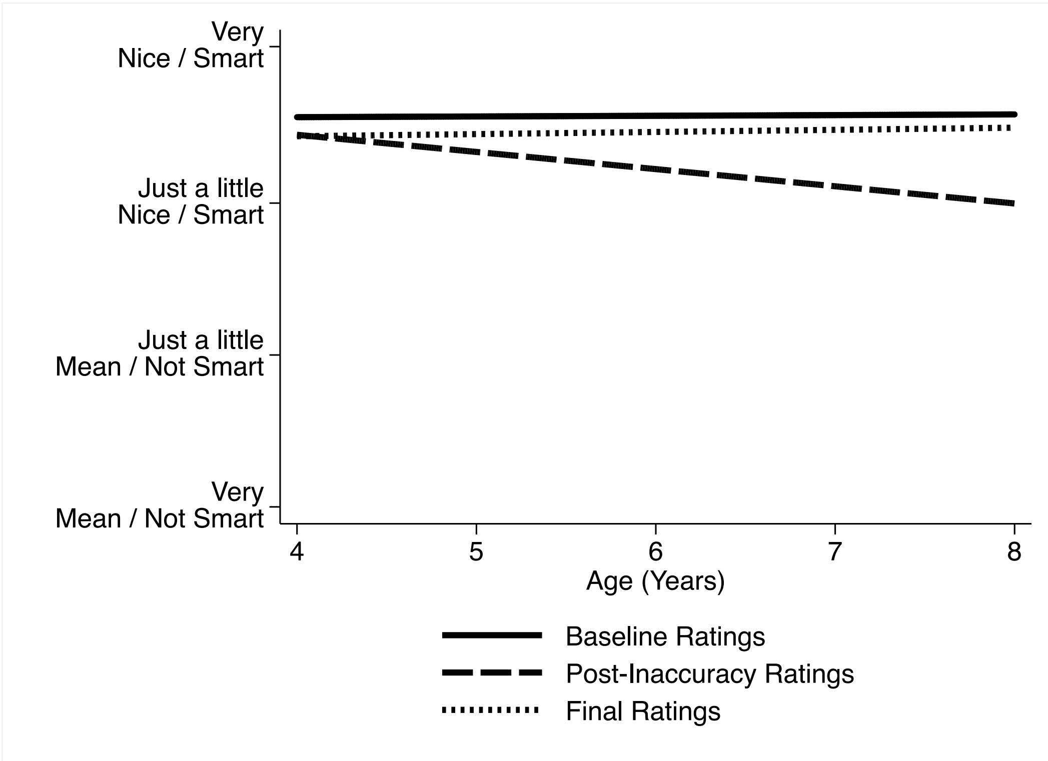

Figure 1. Relation between children's trait attributions and when children were asked to make these attributions (baseline, after receiving inacurate information, at the end of the experiment) as a function of children's age (Table 1, Model 2).

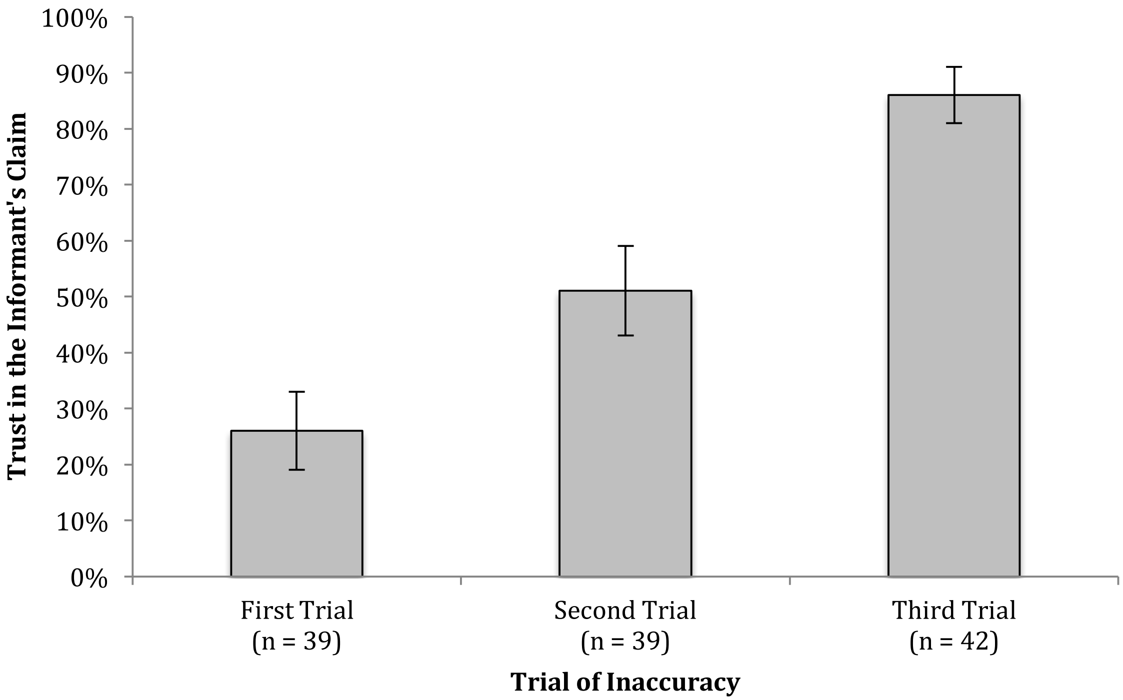

Figure 2. Fitted probability (Table 2, Model 1) of a child trusting the informant immediately after the informant's provision of incorrect information as a function of whether the informant provided incorrect information on the 1^{st} trial, 2^{nd} trial, and 3^{rd} trial. Error bars represent +/- 2 standard errors of the mean.

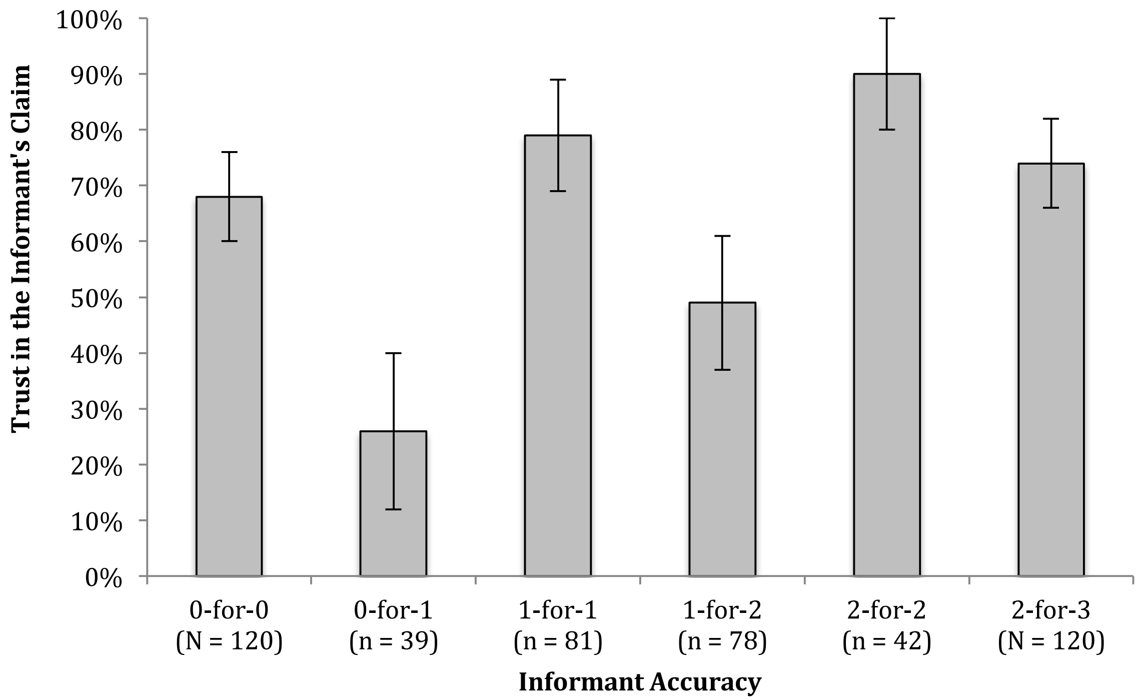

Figure 3. Probability of a child trusting the informant as a function of the informant's prior accuracy (data pooled across all conditions). For each condition (e.g., 1-for-2) the first number reflects the informant's number of accurate claims, and the second number reflects the informant's total number of claims. Error bars represent +/- 2 standard errors of the mean.