

Running head: Trust in previously inaccurate informants

Children's trust in previously inaccurate informants who were well- or poorly- informed:

When past errors can be excused.

Erika Nurmsoo and Elizabeth J. Robinson

University of Warwick

CHILD DEVELOPMENT IN PRESS SEPT 08

Author note. The research was supported by a grant from the Economic and Social Research Council, U.K. (RES-000-22-1847).

Correspondence concerning this article should be addressed to E.J. Robinson, Department of Psychology, Warwick University, Coventry, CV4 7AL, U.K. Electronic mail may be sent to e.j.robinson@warwick.ac.uk.

Abstract

Past research demonstrates that children learn from a previously accurate speaker rather than from a previously inaccurate one. This study shows that children do not necessarily treat a previously inaccurate speaker as unreliable. Rather, they appropriately excuse past inaccuracy arising from the speaker's limited information access. Children ($N = 67$) aged 3, 4 and 5 years aimed to identify a hidden toy in collaboration with a puppet as informant. When the puppet had previously been inaccurate despite having full information, children tended to ignore what they were told and guess for themselves: They treated the puppet as unreliable in the longer term. However children more frequently believed a currently well-informed puppet whose past inaccuracies arose legitimately from inadequate information access.

Children's trust in previously inaccurate informants who were well- or poorly- informed:

When past errors can be excused.

Much of our knowledge about the world is gained indirectly from what other people tell us, rather than from our own direct experience. This ability to learn from others confers great advantages over animals without language, but carries with it associated risks: Other people can deliberately deceive us, be mistaken, or be misunderstood. If we are to benefit overall from gaining knowledge from what other people say, we need to manage these risks. Ideally, we would believe what others tell us only when it is true. This suggests that children need a set of skills to assess the likely truth of what they are told. Without such skills they will be vulnerable to believing what is false or disbelieving what is true even if surrounded by people who intend to be cooperative and informative.

One useful predictor of a speaker's present reliability is his or her past accuracy. A neighbor who has been an accurate informant about garden pests in the past will be invited to give her diagnosis of my current problem. Another neighbor who previously offered inaccurate diagnoses is less likely to be consulted again. That is, accuracy of the speaker's past output is taken as a good predictor of current output, at least within a particular domain of knowledge. Four-year-olds, and 3-year-olds under some conditions, use this cue to speaker reliability. In a *naming game* developed independently by Koenig, Harris and colleagues (Koenig, Clément & Harris, 2004; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig & Harris, 2007), and by Birch and Bloom (Birch, Vauthier & Bloom, in press; see also Jaswal & Neely, 2006), children observe two speakers name

objects familiar to the child (e.g. a ball). On each of three trials, one speaker names an object accurately ('ball') and the other names it inaccurately ('shoe'). On subsequent test trials, the two speakers offer different novel names for an unfamiliar object (e.g. 'mido', 'loma'), or use the same novel name for two different unfamiliar objects. Children are then tested to find out which speaker's name they accept. Four-year-olds typically use the label given by the speaker who had named familiar objects accurately, while three-year-olds do so only under certain conditions (Koenig & Harris 2005). These authors conclude that children interpret the speaker's history of inaccuracy in terms of a lasting trait: The speaker is treated as an unreliable informant, at least within a particular domain.

Similarly, Birch et al. (in press) conclude that 3- and 4-year-olds keep track of speakers' histories of being accurate or inaccurate, and prefer to learn from someone with a good track record.

However, adults do not necessarily interpret inaccuracy as a sign that the speaker is unreliable in the longer term. Suppose I gave my neighbor only the briefest account of a pest infestation and encouraged her to produce a diagnosis. It would not only be unreasonable of me to dismiss her as unreliable if she turned out to be wrong, I might also miss out on accurate advice on a future occasion when she had full command of the facts. That is, when speaker inaccuracy can be explained in terms of the particular circumstances under which it occurred, it may be inappropriate to treat the speaker as an unreliable individual. Rather, the speaker is potentially reliable when circumstances change.

Hence while under some conditions it is appropriate to attribute inaccuracy to the speaker's lasting unreliability, and avoid learning from that speaker in the future, under

other conditions it is more appropriate to attribute inaccuracy to the speaker's particular circumstances, and to believe that speaker's future utterances when those circumstances have changed.

We currently know rather little about how young children interpret inaccuracy, and whether they show sensitivity to an inaccurate speaker's circumstances. What evidence there is suggests children aged 3 to 7 years ignore the reasons for past inaccuracy when predicting future reliability. Nurmsoo and Robinson (in press) developed a variation of the typical procedure summarised above. In a game in which both speakers were required to name objects, one speaker named familiar objects inaccurately while wearing a blindfold, and a second speaker, who could see, either named the same objects accurately or named them inaccurately for no obvious reason. On test trials both speakers could see an unfamiliar object, and offered different labels. When one speaker had a history of accuracy, children endorsed the label offered by that speaker despite understanding why the previously blindfolded speaker had made errors. When both speakers were equally inaccurate, children showed no preference for either speaker's label, again despite understanding why the blindfolded speaker had erred. That is, contrary to the example above concerning an inaccurate neighbor, children in this research failed to take into account the fact that past inaccuracy occurred due to particular circumstances that no longer held: Errors made when temporarily wearing a blindfold were not excused.

In the research just described (Nurmsoo & Robinson, in press), the authors built on procedures developed to examine children's predictions based on a speaker's past history of accuracy or inaccuracy, inserting the variable of the speaker's access to

relevant information. In the research reported here, we did the opposite: We built on procedures in which young children are known to take into account a speaker's access to relevant information, and inserted the variable of speaker's history of accuracy or inaccuracy. (Robinson, Haigh & Nurmsoo, 2008; Robinson & Whitcombe, 2003; Whitcombe & Robinson, 2000) In Robinson and Whitcombe's (2003) procedure, children played a *tunnel game* with the Experimenter to identify which one of a pair of toys was hidden in a tunnel. Toys in a pair differed either in color (e.g. a red or a blue cat) or hardness (e.g. a hard or a soft caterpillar), and the hidden target toy could be felt through one end of the tunnel, or seen through a window in its side. When the target was identified by color, for example, a player who saw it could identify it reliably, while a player who felt it could only guess. Children who had only uninformative access (for example felt a toy identified by color) were more likely to believe what the Experimenter told them about the target's identity when the Experimenter had informative access (saw the toy), than when the Experimenter, like the child, had only uninformative access (felt it).

In the present study, we compared children's readiness to believe two currently well-informed speakers both with a history of three inaccuracies. One speaker was inaccurate while poorly informed, and the other erred despite being well-informed. There was no obvious reason for the inaccuracy in the latter case, as in Koenig & Harris (2005) and the many similar studies listed earlier, and no signs that the speaker was teasing or pretending, so children were expected to treat the speaker as unreliable in the longer term. Would children be more willing to believe the speaker whose inaccuracy could be excused on grounds of inadequate information access?

Method

Since our task involved a face to face game with the child, we decided to use a puppet rather than an adult as speaker in order to avoid the uncomfortable and perhaps unacceptable situation of an apparently fully informed adult giving obviously inaccurate information to the child participant. In previous research involving inaccurate but fully informed speakers, children have either watched adults on video, or puppets have been used.

Participants. Participants were 31 children (19 girls) from nursery classes, aged 3;9 to 4;7 mean 4;2, and 36 children (17 girls) in their first year of formal schooling (UK reception classes) aged 4;10 to 6;8, mean 5;1. The children attended schools serving predominantly white working and middle class areas of Warwickshire and the West Midlands regions of the U.K.

Materials. We used a tunnel with a curtain at the end through which the child could insert an arm, and a curtained window in one side through which the child could look. Seven pairs of toys were used, 2 pairs of worms for warm-up trials, one pair differing only in color and the other only in hardness; 3 pairs of animals for history trials, and 2 pairs for test trials. The animal pairs were drawn from a set including 3 pairs of bears (black and brown; blue and green; white and brown), and pairs of ducklings, ladybirds, and elephants (all differing only in hardness). A monkey hand puppet was used, which had arms and hands, so it could be shown to feel inside the tunnel.

Design and Procedure. Children were tested individually in a quiet corridor or corner of their school classroom. The procedure was heavily based on one developed previously (e.g. Robinson & Whitcombe, 2003; Robinson et al, 2008). On each of two

warm-up trials, children saw and felt the two toys in a pair and agreed on their properties. The Experimenter slipped one of the toys in secret into the tunnel. Children practiced looking at and feeling it. They were told explicitly that feeling was insufficient when the target toy was identified by color, and that looking was necessary. Similarly, they were told explicitly that looking was insufficient for a toy identified by hardness, and that feeling was necessary.

After the 2 warm-up trials, children had 3 history trials followed by 2 test trials. The purpose of the history trials was to give children experience of the puppet's inaccuracy. The toy pairs were ordered such that each child alternated feeling and seeing on sequential trials. Half the children started by feeling, and half started by seeing. Children were randomly assigned to the *Puppet Uninformed* and *Puppet Informed* conditions. These conditions differed in the history trials but had identical test trials.

In the *Puppet Uninformed* condition, participants witnessed the puppet make three sequential errors while poorly informed: On each of the 3 history trials, the puppet had access first: It took uninformative access (for example, it felt a toy identified by color), and said "I'll say which one I think it is. The (blue) one", identifying the target toy incorrectly. The child then had informative access (for example, saw the toy identified by color), and was asked by the Experimenter which toy was in the tunnel. As expected, children relied on their own informative access to identify the toy correctly.

An important feature of the procedure was that the game required child and Experimenter to take turns to say which toy they thought was in the tunnel whether they had informative or uninformative access. There was therefore no reason to construe the

puppet in the puppet uninformed condition as untrustworthy for offering a suggestion despite being poorly informed.

The *Puppet Informed* condition was the same except that on all three history trials, the puppet had informative access (seeing a target identified by color or feeling a target identified by hardness), before giving the wrong judgment. The child then had the same informative access and had the opportunity to give the correct judgment.

Two *test trials* followed immediately. As on the history trials, children agreed on the properties of the two toys in a pair, and one was hidden inside the tunnel. On test trials, however, children had access first: They had uninformative access to the target toy, said which one they thought it was (for example “The hard one” having only seen it), and then the puppet had informative access and contradicted the child (“The soft one”, having felt it). Finally, the Experimenter asked the child which the target was (“the hard one or the soft one?”). Children could either repeat their original guess, or switch to agree with the now well-informed puppet. Regardless of whether the child guessed correctly on test trials, the puppet gave a contradicting judgment. Children were not given the opportunity to check the accuracy of the puppet’s judgment on test trials. The purpose of the test trials was to find out whether or not children believed the puppet’s suggestion when it was better informed than they were.

Results

There was no difference between the two age groups, $t(65) = 1.29, p = 0.20$, and so data were combined for analysis. Table 1 shows the frequencies of children in each condition who believed the well-informed puppet on 0, 1 and 2 test trials. Children in the *Puppet Informed* condition performed no differently from chance, $\chi^2(2, N = 32) = 4.5, p =$

0.11; they behaved as would be expected if they ignored what the puppet said and simply guessed for themselves. In contrast, 71% of the children in the *Puppet Uninformed* condition believed the puppet's suggestion on both test trials. Children were significantly more likely to believe the puppet's suggestion when its inaccuracy was excusable on the grounds of inadequate access to the target, than when it was not: $t(65) = 3.19$, $p < 0.003$, $d = 0.78$. These results suggest that children excused past inaccuracy when the speaker had been poorly informed, but not when there was no such explanation.

Discussion and Conclusions

We have confirmed the published findings that young children do not uncritically accept that whatever they are told is true. They are not passive recipients of testimony from others, but actively evaluate it when deciding whether or not to believe it. One variable they attend to is speakers' history of accuracy or inaccuracy: As in the previous research, children in the Puppet Informed condition predicted the reliability of a speaker's current output on the basis of his or her previous output.

Our tasks differed in several potentially important respects from those used previously to assess children's attention to speaker's past accuracy (Birch et al., in press; Jaswal & Neely, 2006; Koenig et al, 2004; Koenig & Harris, 2005; Nurmsoo & Robinson, in press). First, in those studies, children chose which of two speakers to believe, one of whom was previously accurate and the other inaccurate for no obvious reason. Children consistently preferred the previously accurate speaker. It could have been that the contrast between speakers was important for directing children to the relevance of past accuracy. This appears not to be the case. Here we show an effect of

past accuracy in the Puppet Informed condition when children were faced with only a single well-informed speaker with a history of unexplained inaccuracy.

Second, in previous studies children learned about conventional, generalizable knowledge such as the names or functions of unfamiliar objects (but see Clement, Koenig & Harris, 2004, who included items on objects' color). Csibra & Gergely (2005) argue that infants and young children are sensitive to cues that an adult is about to pass on generalizable knowledge, whilst adults are particularly ready to pass on their expertise in ways accessible to children. It might have been that sensitivity to speakers' history of accuracy or inaccuracy was confined to such knowledge, for which children are heavily dependent on adult informants, and for which the costs of learning something false may be high. However, our findings show that children are also sensitive to a speaker's history of accuracy when informed about the identity of a hidden target, non-generalizable knowledge about a particular event that the child could easily have found out herself.

Third and most importantly, our results go beyond those of previous research in showing that children do not simply attend to the speaker's history of accuracy or inaccuracy, but also pay attention to the reasons for it. When they understood that the inaccuracy occurred due to the particular circumstances at the time, they appropriately did not predict that the speaker would be unreliable when circumstances changed. Children took into account speakers' input (their information access) to interpret the significance of inaccurate output (what they said). This implies that they engaged in mentalistic reasoning, and did not simply treat potential informants as they might treat an accurate or an inaccurate clock. Interestingly, both Harris (2007) and Birch et al. (in press) argue that selective trust in previously accurate speakers is based on mentalistic

reasoning, although hitherto there has been no evidence to support that strong interpretation.

Indeed, there is evidence that children do not engage in mentalistic reasoning with procedures very similar to those used by Harris and colleagues (e.g. Koenig & Harris, 2005) and by Birch et al. (in press): As mentioned in the introduction, Nurmsoo and Robinson (in press) found that children did not excuse a speaker whose inaccurate naming of familiar objects was due to the temporary wearing of a blindfold. Faced with two speakers who offered contrasting names for an unfamiliar object, one of whom had previously labelled three familiar objects inaccurately despite being able to see them, and the other of whom had done the same whilst blindfolded, children showed no preference for the label offered by the previously blindfolded speaker (who could now see). That is, children attended only to the accuracy of a speaker's previous output when deciding whether or not to believe her current suggestion about the name of an unfamiliar object; they did not take into account the reason for the prior inaccuracy. Importantly, as in the procedure used in the Experiment reported here, the demands of the game meant that the naming objects despite being blindfolded could not be construed as foolish or irresponsible.

Nurmsoo and Robinson (in press) discuss possible reasons for children's failure to take into account the reasons for speaker inaccuracy in the typical naming task. Perhaps communicative and social cues necessary to engage mentalistic reasoning are absent from the typical procedure. This may be important to check in further research, since procedures of this kind are now widely used in investigations of children's trust in speakers, and it is possible that the findings have limited relevance to children's learning

from others in real life. A second possibility is that children are particularly intolerant of speakers' inaccuracy when they learn generalizable knowledge such as the names or functions of objects. That is, they may be overly cautious and unwilling to excuse inaccuracy for any reason, rather than risk acquiring false generalizable information. In contrast, when learning specific information such as the identity of a hidden toy, as in the task used here, they may weight the benefit of gaining new information more heavily and be less averse to the risk of learning something false.

Whatever the reasons for the children's failure to take into account the reasons for speaker inaccuracy in the typical naming game task, we have the clear demonstration in this research that children aged 3- to 4 years *can* show appropriate balance between risking believing an unreliable speaker, and missing out on learning from a speaker whose short-term circumstances led to inaccuracy. This is confirmed in a further study (Nurmsoo & Robinson, 2007) in which children were even prepared to abandon their well-founded expectation about the contents of a box, and to believe instead what they were told by a previously unreliable puppet, but only when the previous unreliability could be excused in terms of limited information access. Further research is needed to examine the limits of this well-balanced behavior.

References

- Birch, S. A. J., Vauthier, S. A., & Bloom, P. (in press). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*.
- Clement, F., Koenig, M., & Harris (2004) The ontogenesis of trust. *Mind and Language*, 19, 360-379.
- Csibra, G., & Gergey, G. (2005). Social learning and social cognition: The case for pedagogy. In M.H. Johnson & Y. Munakatas (Eds.) *Processes of Change in Brain and Cognitive Development. Attention and Performance XXI*. (pp. 249-274). Oxford: Oxford University Press.
- Harris, P. L. (2007). Trust. *Developmental Science*, 10, 135-138.
- Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, 17, 757-758.
- Koenig, M. A., Clement, F., and Harris, P. L. (2004) Trust in testimony: Children's use of true and false statements, *Psychological Science*, 15, 694–698.
- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76, 1261-1277.
- Nurmsoo, E., & Robinson, E.J. (2007, October). Children trust speakers with a history of excusable inaccuracy. Poster presented at meetings of the Cognitive Development Society, Santa Fe, N.M.
- Nurmsoo, E. & Robinson, E.J. (in press) Identifying unreliable informants: Do children excuse past inaccuracy? *Developmental Science*.

- Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, 43, 1216-1226.
- Robinson, E. J., Haigh, S. N., & Nurmsoo, E. (2008). Children's working understanding of knowledge sources: Confidence in knowledge gained from testimony. *Cognitive Development*, 23, 105-118.
- Robinson, E. J., & Whitcombe, E. L. (2003). Children's suggestibility in relation to their understanding about sources of knowledge. *Child Development*, 74, 48-62.
- Whitcombe, E. L., and Robinson, E. J. (2000) Children's decisions about what to believe and their ability to report the source of their belief, *Cognitive Development*, 15, 329-346.

Table 1

Incidence of Believing Puppet's Suggestion on 0, 1 and 2 Test Trials in each Condition.

Condition	Age group	Frequency of accepting Puppet's suggestion			Mean (sd)
		0	1	2	
Puppet Uninformed	Nursery, 3 – 4 years	1	3	13	
	Reception, 4 – 5 years	3	3	12	
	Total	4	6	25	1.60 (0.69)
Puppet Informed	Nursery, 3 – 4 years	4	4	6	
	Reception, 4 – 5 years	7	6	5	
	Total	11	10	11	1.00 (0.84)