

## **UC Merced**

# **Proceedings of the Annual Meeting of the Cognitive Science Society**

### **Title**

Preschoolers' Trust is Sensitive to Variable Intentions

### **Permalink**

<https://escholarship.org/uc/item/5t25x1mn>

### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

### **ISSN**

1069-7977

### **Authors**

Hoicka, Elena  
Butcher, Jessica  
Malla, Felicity  
[et al.](#)

### **Publication Date**

2013

Peer reviewed

# Preschoolers' Trust is Sensitive to Variable Intentions

**Elena Hoicka (elena.hoicka@stir.ac.uk)**

Psychology Department, University of Sheffield, Western Bank, Sheffield, S10 2PT, UK

**Jessica Butcher, Felicity Malla**

Psychology, School of Natural Sciences, University of Stirling, Stirling, FK9 4LA, UK

**Paul L. Harris (paul\_harris@gse.harvard.edu)**

Harvard Graduate School, Harvard University, Appian Way, Cambridge, MA, 02138, USA

## Abstract

This research demonstrates that preschoolers flexibly trust and mistrust the same individuals, as preschoolers recognize that their intentions may vary. In Study 1 ( $N=101$ ) 3- and 4-year-olds trusted speakers based on their current, rather than previous, intentions to give in/correct information. Thus preschoolers infer the meanings behind different intentions and recognize that intentions change within individuals over time. In Study 2 ( $N=80$ ) 3- to 5-year-olds trusted speakers who were currently sincere, but previously intentionally inaccurate, rather than currently sincere, but previously ignorant, showing that preschoolers infer current knowledge from prior intentions. Preschoolers also trusted speakers who were currently knowledgeable, although previously ignorant, showing that they recognize knowledge is variable within individuals.

**Keywords:** Trust, Intention, Knowledge, Frame Problem, Preschoolers, Humor

## Introduction

A growing body of research suggests that children do not blindly trust just anyone; children consider *whom* to trust (e.g., Clément, Koenig & Harris, 2004; Corriveau, Meints & Harris, 2009; Koenig & Harris, 2005). However, in a world where people have variable knowledge and intentions, an important question is whether children are prepared to trust individuals on some occasions but not others (e.g., Nurmsoo & Robinson, 2009; Scofield & Behrend, 2008; Shafto, et al., 2012). Thus, the question becomes not just *whom* to trust but also *when* to trust a given person. This is important not only for how we acquire information, but also for how we dismiss uninformative or irrelevant information. This is the Frame Problem (e.g., Dennett, 1984), which is still proving difficult to solve in Artificial Intelligence (AI, e.g., Ekbia & Maguitman, 2001; Scherl & Levesque, 2003), but which may be a relatively easy problem for preschoolers to solve.

Preschoolers consider past behaviors when deciding whom to trust. For example, when learning new words, both 3- and 4-year-olds trust a speaker who previously labeled familiar objects correctly over a speaker who labeled them incorrectly. Thus, children trust accurate over inaccurate speakers (e.g., Clément, et al., 2004; Koenig & Harris, 2005). Children also trust knowledgeable over ignorant actors (Einav & Robinson, 2011).

Our first goal was to discover whether preschoolers trust an informant's claims depending on their current intentions and not just on their previous accuracy or apparent knowledge. More specifically, we asked: (1) whether preschoolers are flexible when they trust others; (2) whether mental states, specifically intentions, play a role in trust; and (3) whether preschoolers understand that people's intentions can change over time. Most of the research to date suggests that people who were previously accurate or knowledgeable can be trusted in future, whereas people who were previously inaccurate or ignorant cannot. However people are not statically trustworthy or untrustworthy (e.g., Nurmsoo & Robinson, 2009; Scofield & Behrend, 2008; Shafto et al., 2012). Rather, a person can be trustworthy at times, but not at others.

Joking is a clear example of intentionally saying or doing the wrong thing (e.g., Hoicka, Jutsum, & Gattis, 2008; Leekam, 1991), and so it is an occasion when the audience should not trust the information provided. Indeed, jokers want their audience to know about their falsehood, and they do not expect the audience to believe any part of it (e.g., Leekam, 1991). Thus, people provide cues when they are joking (e.g., Hoicka & Gattis, 2012; Mireault, et al., 2012). In particular, parents express greater disbelief when joking as compared to acting literally (Hoicka, et al., 2008). Additionally, everyone jokes. For example, all 3-year-olds in a survey were reported by their parents to have produced novel jokes (Hoicka & Akhtar, 2012). Therefore, joking is an ideal way to examine whether preschoolers use intent to say or do the wrong thing as a cue not to learn.

Research has started to consider whether young children understand the various contexts in which others intend to say or do the wrong thing. From 25 months, toddlers copy wrong actions marked as jokes (laughter), but correct the same wrong actions marked as mistakes ("Whoops!"); Hoicka & Gattis, 2008). From 30 months, toddlers copy jokers who mislabel familiar objects, but they do not copy people who sincerely mislabel (Hoicka & Akhtar, 2011). In the case of trust, children as young as 3 years understand that pretending is not a reliable cue for acquiring correct information compared to, for example, having direct experience with the relevant information (Koenig, 2012).

A critical aspect of intention is that it is not a stable mental state. People's intentions change over time (e.g.,

Cohen & Levesque, 1990; Roy, 2009; van der Hoek, Jamroga, & Wooldridge, 2007). According to Cohen and Levesque (1990, p. 214), people “keep (or commit to) intentions, but not forever; [they] discharge those intentions believed to have been satisfied”. Thus, people can revise or complete their intentions, moving onto new intentions. When joking, the goal is to get a laugh at a moment in time, but not at everything forever.

Our first goal – examined in Experiment 1 – was to find out if children trust someone who currently intends to give correct information over someone who intends to give incorrect information, regardless of their past accuracy or intentions. A major goal of Experiment 2 was to discover whether children can infer a speaker’s knowledge from his or her intention. When people intend to do or say the wrong thing, through joking, lying, or pretending (e.g., Hoicka & Gattis, 2008), the speaker knows the correct information, but chooses not to say it (e.g., Leekam, 1991). Thus, if a joker previously said the wrong thing, he or she likely knew what the right answer was. By contrast, if an ignorant person said the wrong thing, this suggests that the person did not know the right answer. If children understand that joking is more likely to involve intentionally saying the wrong thing compared to being ignorant, then they should later trust a previous joker over a previously ignorant person when learning new information. Thus they may infer that the joker was more likely to know the information than the ignorant speaker, but withheld prior information.

## Experiment 1

Experiment 1 examined whether preschoolers base their trust on speakers’ current intentions rather than their past accuracy or intentions. In the final test trials, one actor named novel objects while giving sincere cues whereas the other named them while giving joking cues. Depending on the condition, children had previously seen the actors display the same intentions (i.e., sincere versus joking), a switch in intentions, or they had had no prior exposure to the actors’ intentions. The experimental question was whether children would be swayed by the actors’ current intentions – as expressed in the test trials – or by their past intentions.

## Method

**Participants** Fifty-three 3-year-olds (33 females,  $M = 3$  years, 5 months,  $range = 3$  years, 0 months – 3 years, 11 months) and 48 4-year-olds (25 females,  $M = 4$  years, 4 months,  $range = 4$  years, 0 months – 4 years, 11 months) were randomly assigned to one of three conditions: Consistent Intentions, Inconsistent Intentions, and No Prior Intentions. Children were of similar ages across conditions.

**Materials** The objects in the familiarization trials included a spoon, a bottle, a doll, and a brush. The objects in the action videos included a cookie, a cup, a scarf, and a hat. The objects in the test trials included a brown feathery cat toy, a

red, black and silver DIY object, a blue and white dog toy, and a red and white kitchen utensil. PowerPoint slideshows were made with each slide showing an object and/or actors, or a video (see procedure). Children’s responses were video-recorded directly onto the laptop computer.

**Design** This study was a between-subjects design in which there were three conditions. In the Consistent Intentions condition, one actor was consistently joking and one actor was consistently sincere throughout the familiarization trials, action videos, and test trials. In the Inconsistent Intentions condition, one actor joked in the familiarization trials but was sincere in the action videos and test trials. The other actor was sincere in the familiarization trials but joked in the action videos and test trials. By showing that actors had switched intentions in the action videos, we anticipated that children would be prepared to recognize their new intentions in the test trials. In the No Prior Intentions condition, children saw the objects but not the actors in the training trials, and then participated in the full test trials. The dependent variable was whether children trusted the joker or sincere actor at test trials when learning new labels.

## Procedure

**Familiarization Trials:** At the start of the familiarization trials in the Consistent Intentions and Inconsistent Intentions conditions, children were shown a video of the two actors being asked to name an object (e.g., spoon). The joker laughed, named it incorrectly (e.g., duck) using a humorous intonation pattern (Hoicka & Gattis, 2012) and said, “I’m being silly, only joking.” The sincere actor labeled it correctly using a sincere intonation pattern (Hoicka & Gattis, 2012) while smiling. After watching the video children were shown a slide with pictures of the object and the two actors who had named it and were asked, “She called it a [e.g., duck] and she called it a [e.g., spoon]. Can you tell me what it’s called?” This repeated for the remaining three trials with different familiar objects. In the No Prior Intentions Condition, children were instead shown a slide with a picture of the familiar object, given two names for the object and then asked to name it. For example for the spoon, they were asked, “Is this a duck or a spoon?”

**Action Videos:** In the Consistent Intentions condition, the actor who had joked in the familiarization trials also joked in the action videos. Similarly, the sincere actor stayed sincere. In the Inconsistent Intentions condition, the actor who had joked in the familiarization trials became sincere during the action videos. Similarly, the actor who had been sincere in the familiarization trials became humorous.

For each action, each actor said the same line before performing an action such as, “I’m going to put this hat on”. The sincere actor then did the correct action (e.g., putting the hat on her head) and the joker did the action incorrectly, (e.g., putting the hat under her arm) and saying, “I’m being silly, I’m only joking” and laughing. This continued for the other three actions. The No Prior Intentions condition did not include action videos.

Test Trials: Children watched four videos where a third actor asked the two actors what a novel object was called. The sincere actor smiled and said, e.g., “That’s a mogo” with a sincere intonation pattern. The joker said, e.g., “That’s a sepa” in a humorous intonation pattern, then, “I’m being silly, only joking” and laughed. Following the video the child was shown a slide with a photo of the object and the two actors and told, “She called it a sepa and she called it a mogo. Can you tell me what it’s called?” This continued for the remaining three trials. In the Consistent Intentions condition, the actors played the same roles as they had for the familiarization trials and action videos. In the Inconsistent Intentions condition, the actor who had joked during the familiarization trials was now sincere (just as she had been during the action trials). By contrast, the actor who had been sincere during the familiarization trials was now joking (just as she had been during the action trials). In the No Prior Intentions condition, the test trials were the first time the children had seen the actors.

## Results

Data were analyzed with logit mixed effects models. Only significant effects and interactions are reported.

See Figure 1 for the percentage of trials on which children chose the sincere actor’s label over the joker’s, by Condition (Consistent Intentions, Inconsistent Intentions, No Prior Intentions) and Age. The base model was improved by Condition,  $X^2(2) = 13.91$ ,  $p = .0010$ , and Age,  $X^2(1) = 3.84$ ,  $p = .0501$ , as fixed effects. The resulting model ( $\log\text{-likelihood} = -149.17$ ,  $N = 340$ ) found children were significantly more likely to trust the sincere actor over the joker at test trials in the Consistent Intentions versus Inconsistent Intentions and No Prior Intentions conditions (both Odds-Ratio,  $OR > 2.43$ ,  $p < .0256$ ). Four-year-olds were marginally more likely to trust the sincere actor than 3-year-olds ( $OR = 1.78$ ,  $p = .0526$ ).

When each condition was tested individually, children in the Consistent Intentions and No Prior Intentions conditions (both  $\log\text{-likelihood} > -59.22$ ,  $N = 127/101$ ) were significantly more likely to trust the sincere actor than the joker at test trials (both  $OR > 2.66$ ,  $p < .0012$ ). Age improved the model for the Inconsistent Intentions condition,  $X^2(1) = 6.03$ ,  $p = .0141$ . The overall model ( $\log\text{-likelihood} = -49.84$ ,  $N = 112$ ) found 4-year-olds were more likely to trust the sincere actor than the joker at test trials than 3-year-olds ( $OR = 3.61$ ,  $p = .0180$ ). Follow-up tests found that both 3- and 4-year-olds (both  $\log\text{-likelihood} > -31.41$ ,  $N = 54/58$ ) were more likely to trust the sincere actor than the joker at test trials (both  $OR > 2.85$ ,  $p < .0160$ ).

## Discussion

Across all three conditions, both 3- and 4-year-olds were more likely to trust the actor who was sincere at test trials over the actor who was joking. This suggests children take into account the current intention of a speaker when deciding whether to learn from him or her. If the speaker’s intention is sincere, children will learn. If the speaker’s

intention is to joke, children will not learn. Thus, selective trust is not purely based on past experience with a speaker – it is also based on a speaker’s current intentions.

## Experiment 2

During the familiarization trials in both conditions in Experiment 2, one actor mislabeled familiar objects due to ignorance, while the other actor mislabeled familiar objects because she was joking. During test trials, the previously ignorant actor gave cues suggesting that she was now knowledgeable, whereas the joker continued to give joking cues (Knowledge Inconsistent condition). Alternatively, both actors gave cues that they were sincere (Knowledge Inferred condition).

The first experimental question concerned the Knowledge Inconsistent condition. Would children recognize that someone who was once ignorant could become knowledgeable, and would be better to trust than a previous joker who intended to continue saying the wrong thing? The second experimental question concerned the Knowledge Inferred condition. Would children distinguish the two types of inaccuracy during familiarization trials; more specifically, would they recognize that a previous joker was more likely to know the correct labels compared to an ignorant speaker, but chose not to say them. If so, during test trials children should trust a previous joker who becomes sincere, and intends to say the right thing, over a previously ignorant actor, who is also sincere, but is less likely to be knowledgeable.

## Method

**Participants** Thirty 3-year-olds (14 females,  $M = 3$  years, 5 months,  $range = 3$  years, 1 month – 3 years, 11 months), 28 4-year-olds (13 females,  $M = 4$  years, 5 months,  $range = 4$  years, 0 months – 4 years, 10 months), and 22 5-year olds (13 females,  $M = 5$  years, 6 months,  $range = 5$  years, 1 month – 5 years, 11 months) were randomly assigned to one of two conditions: the Knowledge Inconsistent condition and the Knowledge Inferred condition. Children were of similar ages across conditions.

**Materials** Same as Study 1, except that there were an additional two familiarization objects (car, pig).

**Design** This study was a between-subjects design in which there were two conditions. In the Knowledge Inconsistent condition, one actor was ignorant during training, but knowledgeable during testing, whereas the other actor consistently joked. In the Knowledge Inferred condition, again, one actor was ignorant during training whereas the other joked. During action and test trials, both actors were sincere. Action trials were included to show a change of intentions in the joker, as in Experiment 1.

There were six training trials in both conditions. The joker joked for four trials, and was knowledgeable for two trials. The ignorant actor was ignorant for four trials, and

knowledgeable for two trials. By showing that the ignorant actor could switch knowledge states during training, we anticipated that children would be prepared to recognize the ignorant actor's new knowledge in the test trials in the Knowledge Inconsistent condition. Jokers also had two knowledgeable trials in order to keep (in)accuracy consistent between actors. Training was the same in the Knowledge Inferred condition in order to ensure conditions were comparable. The dependent variable was who children trusted when learning new labels – the previous joker, or previously ignorant actor.

### Procedure

**Familiarization Trials:** The task proceeded in the same way as in Experiment 1. In both conditions, the joker gave incorrect labels paired with joking cues for four out of six familiarization trials and correct labels paired with knowledge cues for two familiarization trials. The ignorant actor gave incorrect labels paired with ignorance cues for four out of six familiarization trials and correct labels paired with knowledge cues for two familiarization trials. Humorous cues were the same as in Experiment 1. For ignorance cues, the actor shrugged her shoulders and labeled the object incorrectly saying, e.g., “I don't know, that's a train?” Knowledgeable cues involved displaying their knowledge and labeling an object correctly, e.g., “I know this one. That's a spoon.”

**Action Videos:** In the Knowledge Inferred condition only, children were shown action videos which were the same as those used in Experiment 1. The actor who had joked during the familiarization trials carried out the four sincere actions. The actor who had been ignorant during familiarization also carried out the four sincere actions.

**Test Trials:** The test trials were the same as in Experiment 1. In the Knowledge Inconsistent condition, the previous joker continued to joke during test trials saying, e.g., “That's a sepa, I'm being silly, only joking” and laughing, whilst the previously ignorant actor was now knowledgeable saying, e.g., “I know this one. It's a mogo”. In the Knowledge Inferred condition, both actors labeled the novel object giving sincere cues, where they would smile and say, “That's a mogo” or “That's a sepa”.

### Results

We built logit mixed effects models as in Experiment 1. No gender or age (over, under 4.5 years) differences were found. See Figure 2 for the percentage of trials on which children chose the previous joker over the previously ignorant actor, by Condition (Knowledge Inconsistent, Knowledge Inferred). The base model was improved by Condition,  $X^2(1) = 50.05, p < .0001$  as a fixed effect. The resulting model (*log-likelihood* = -177.49,  $N = 307$ ) found an effect of Condition ( $OR = 6.21, p < .0001$ ). When each condition was tested individually, children in the Knowledge Inconsistent condition (*log-likelihood* = -73.42,  $N = 157$ ) were significantly more likely to trust the previously ignorant actor (now knowledgeable) versus the

previous joker (still joking) at test trials ( $OR = 4.77, p < .0001$ ). In the Knowledge Inferred condition (*log-likelihood* = -93.02,  $N = 137$ ) children were significantly more likely to trust the previous joker over the previously ignorant actor when both were sincere at test trials ( $OR = 1.40, p = .0505$ ).

### Discussion

Experiment 2 shows that children consider intentions in combination with knowledge when deciding whom to trust for information. When both actors were sincere during test trials in the Knowledge Inferred condition, children were more likely to trust the previous joker than the previously ignorant actor. Because both actors were equally inaccurate during the training trials, accuracy could not be used as a cue. Moreover, children in Experiment 1 did not trust the joker, even when no previous training was given, suggesting that children do not simply prefer jokers. Thus, children inferred that despite the joker previously being inaccurate, she likely actually knew the correct information, at least compared to the ignorant actor, and would thus express the correct information when being sincere.

Another possible way to explain the results is that children avoided learning from someone who was previously ignorant. However, this cannot be the case because children chose to learn from the previously ignorant actor rather than the previous joker when she showed signs of knowledge in the test trials in the Knowledge Inconsistent condition. This demonstrates that children are flexible in their trust, and understand that people's knowledge can vary. They acknowledge that sometimes people know words, and sometimes they do not.

### General Discussion

Experiment 1 showed that children trust speakers based on the speakers' current intentions, rather than their previous intentions or accuracy. Children trusted the actor who was currently being sincere versus joking when learning new labels, even when the actors' previous intentions were different or had not been made available to the preschoolers. Thus, preschoolers recognized that when speakers joke, they intend to say the wrong thing, and so should not be trusted to provide accurate information.

Experiment 2 showed that preschoolers combine intention and knowledge states to determine whether information is trustworthy. Specifically, children inferred that when both actors were sincere at test trials, the actor who previously said the wrong thing in the context of a joke was more likely to know the correct labels compared to the actor who previously said the wrong thing due to ignorance. Thus, children recognized that a joker is more likely to *know* the truth compared to an ignorant speaker, but chooses not to say it. However, when the previous joker continued to joke, and the previously ignorant actor showed that she was knowledgeable at test trials, children were flexible and preferred the previously ignorant actor, recognizing a change in the ignorant actor's knowledge.

## Stable Traits

Much of the research to date has portrayed children's trust as involving the attribution of a stable trait concerning previous accuracy or knowledge (e.g., Clément, et al., 2004; Corriveau, et al., 2009; Einav & Robinson, 2011; Koenig & Harris, 2005). This is the first empirical research to show that children consider not only previous accuracy or knowledge, but also the speaker's past and current intentions, when deciding whom to trust. This is an important skill to have because speakers shift rapidly in their intentions, joking at one moment and being sincere the next. This adds to a body of research showing that children are flexible in their trust (e.g., Nurmsoo & Robinson, 2009; Scofield & Behrend, 2008; Shafto, et al., 2012).

In the case of intentionally saying the wrong thing, such as joking, it is highly unlikely that someone would always joke, even if most people do joke at certain times (e.g., Hoicka & Akhtar, 2012) to trust them. The current findings converge with evidence from computational models, which suggest children consider intention and knowledge, and not just accuracy, when deciding whom to trust (Shafto, et al., 2012).

## Intention

Although much research has considered toddlers' understanding that people intend to do the right thing (e.g., Carpenter, Akhtar, & Tomasello, 1998), understanding complex intentions, such as intentions to do the wrong thing, may be a more refined test of intention understanding as it involves considering *why*, and not just *whether*, someone would do something intentionally. Our experiments show that preschoolers respond appropriately to complex intentions from 3 years. Specifically, they recognize that people can intend to do different things for different reasons. They can intend to say the right thing to teach others, or they can intend to say the wrong thing to joke.

A growing body of research suggests that preschoolers understand that people can intend to do the wrong thing (Hoicka & Akhtar, 2011; Hoicka & Gattis, 2008; Rakoczy, Tomasello, & Striano, 2004). The current experiments extend this prior research by showing that preschoolers can make use of this insight when learning new information. Children are thus flexible learners, accepting new information only when appropriate. Preschoolers are therefore able to solve the Frame Problem (e.g., Dennett, 1984) to a relatively sophisticated degree, tracking speakers' prior and current accuracy, knowledge, and intentions to decide when to accept versus reject information. An important question that follows is *how* they solve the Frame Problem. In the current studies, social cues clearly helped. Indeed, in the case of humor and humorous intentions, parents scaffold infants' and toddlers' understanding through cues and explicit expressions of disbelief (Hoicka & Gattis, 2012; Hoicka, et al., 2008; Mireault, et al., 2012). Thus, the Frame Problem may be solved to some extent through social cues and parental scaffolding.

## Knowledge

Analysis based in philosophy and AI shows that intention is not a stand-alone mental state. Rather, to have an intention, one must also have other mental states such as beliefs and knowledge (Cohen & Levesque, 1990; van der Hoek, et al., 2007). Thus, for children to truly understand others' intentions, they must also understand others' beliefs or knowledge. Experiment 2 provides the first experimental evidence that children as young as 3 years infer knowledge from intentions, and use inferred knowledge to learn from a previously inaccurate person later on.

Although past research demonstrates that children can infer intentions to do the wrong thing when joking (e.g., Hoicka & Akhtar, 2011; Hoicka & Gattis, 2008), it was not clear from this research whether children understood that the actor actually knew the correct information. The current research shows that they consider intention alongside other mental states, specifically knowledge.

Experiment 2 also suggests that preschoolers can attribute knowledge to people who were previously ignorant if they later demonstrate cues showing knowledge. This is consistent with previous findings on perceptual access. Children did not trust an informant who could not perceive the information that they needed, but later trusted the same informant when he or she could perceive that information (Nurmsoo & Robinson, 2009; Robinson, et al., 2011). This flexibility makes sense. For example, sometimes people forget information, but not always. Sometimes people have some knowledge in a domain, but not all knowledge, for example vocabulary, for which there is variation amongst parents (e.g., Huttenlocher, 1998). Thus, a speaker may know some labels, but not all labels. Being sensitive to cues which suggest when someone has knowledge and when they do not, even within the same domain, would thus be a useful tool in selectively trusting, and acquiring information, from others.

## Figures

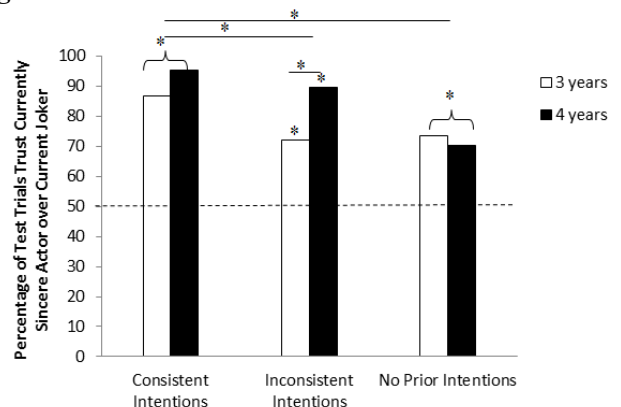


Figure 1: Percentage of test trials for which children trusted the currently sincere actor over the current joker in Study 1. \* $p < .05$ . Lines indicate where differences were examined. Parentheses indicate results summed across groups.

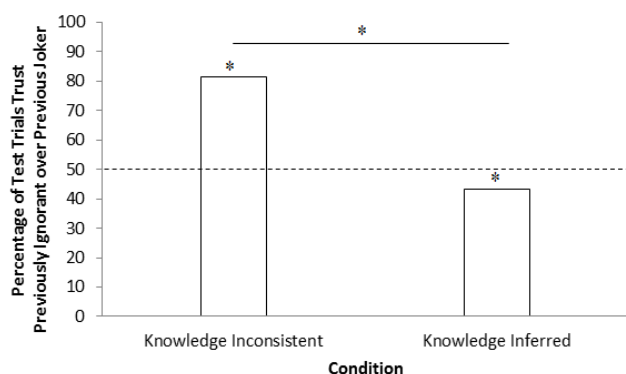


Figure 2. Total percentage of test trials during which children trusted the previously ignorant actor over the previous joker, by condition. \* $p \leq .05$

### Acknowledgments

We thank parents and toddlers for participating in our research. This research was supported by a British Academy Small Research Grant SG110479 awarded to Elena Hoicka.

### References

Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior & Development, 21*, 315-330.

Clément, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*, 360-379.

Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence, 42*, 213-261.

Corriveau, K. H., Meints, K., & Harris, P. L. (2009). Early tracking of informant accuracy and inaccuracy. *British Journal of Developmental Psychology, 27*, 331-342.

Dennett, D. C. (1984). Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, machines and evolution*. Cambridge, United Kingdom: Cambridge University Press.

Einav, S., & Robinson, E. J. (2011). When being right is not enough: Four-year-olds distinguish knowledgeable informants from merely accurate informants. *Psychological Science, 22*, 1250-1253.

Ekbja, H. R., & Maguitman, A. M. (2001). Context and relevance: A pragmatic approach. In V. Akman (Ed.), *Modeling and using context*. Berlin: Springer Verlag.

Hoicka, E., & Akhtar, N. (2012). Early humor production. *British Journal of Developmental Psychology, 30*, 586-603.

Hoicka, E., & Akhtar, N. (2011). Preschoolers joke with jokers, but correct foreigners. *Developmental Science, 14*, 848-858.

Hoicka, E., & Gattis, M. (2012). Acoustic differences between humorous and sincere communicative intentions. *British Journal of Developmental Psychology, 30*, 531-549.

Hoicka, E., & Gattis, M. (2008). Do the wrong thing: How toddlers tell a joke from a mistake. *Cognitive Development, 23*, 180-190.

Hoicka, E., Jutsum, S., & Gattis, M. (2008). Humor, abstraction, and disbelief. *Cognitive Science: A Multidisciplinary Journal, 32*, 985-1002.

Huttenlocher, J. (1998). Language input and language growth. *Preventive Medicine: An International Journal Devoted to Practice and Theory, 27*, 195-199.

Koenig, M. A. (2012). Beyond semantic accuracy: Preschoolers evaluate a speaker's reasons. *Child Development, 83*, 1051-1063.

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development, 76*, 1261-1277.

Leekam, S. R. (1991). Jokes and lies: Children's understanding of intentional falsehood. In A. Whiten, & A. Whiten (Eds.), *Natural theories of mind: Evolution, development and simulation of everyday mindreading*. Cambridge, MA US: Basil Blackwell.

Mireault, G., Poutre, M., Sargent-Hier, M., Dias, C., Perdue, B., & Myrick, A. (2012). Humour perception and creation between parents and 3- to 6-month-old infants. *Infant and Child Development, 21*, 338-347.

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.

Rakoczy, H., Tomasello, M., & Striano, T. (2004). Young children know that trying is not pretending: A test of the 'behaving-as-if' construal of children's early concept of pretense. *Developmental Psychology, 40*, 388-399.

Robinson, E. J., Butterfill, S. A., & Nurmsoo, E. (2011). Gaining knowledge via other minds: Children's flexible trust in others as sources of information. *British Journal of Developmental Psychology, 29*, 961-980.

Roy, O. (2009). A dynamic-epistemic hybrid logic for intentions and information changes in strategic games. *Synthese, 171*, 291-320.

Scherl, R. B., & Levesque, H. J. (2003). Knowledge, action, and the frame problem. *Artificial Intelligence, 144*, 1-39.

Scofield, J., & Behrend, D. A. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development, 23*, 278-290.

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.

van der Hoek, W., Jamroga, W., & Wooldridge, M. (2007). Towards a theory of intention revision. *Synthese, 155*, 265-290.