

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

The Development and Understanding of Selective Trust in Preschool Children /

Permalink

<https://escholarship.org/uc/item/2ch1w27c>

Author

Vanderbilt, Kimberly E.

Publication Date

2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Development and Understanding of Selective Trust in Preschool Children

A dissertation submitted in partial satisfaction of the requirements for the degree

Doctor of Philosophy

in

Psychology

by

Kimberly E. Vanderbilt

Committee in charge:

Professor Gail Heyman, Chair
Professor David Liu, Co-Chair
Professor Leslie Carver
Professor Sarah Creel
Professor Dana Nelkin

2013

Copyright

Kimberly E. Vanderbilt, 2013

All rights reserved.

The Dissertation of Kimberly E. Vanderbilt is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co- Chair

Chair

University of California, San Diego

2013

DEDICATION

This dissertation is dedicated to my mom, and to Julian, for their unwavering and unwarranted faith in me. This document would not exist without them.

TABLE OF CONTENTS

Signature Page.....	iii
Dedication.....	iv
Table of Contents.....	v
List of Figures.....	vi
List of Tables.....	vii
Acknowledgements.....	viii
Curriculum Vitae.....	x
Abstract.....	xi
Introduction.....	1
Chapter 1. The development of distrust.....	7
Chapter 2. In the absence of conflicting testimony young children trust inaccurate informants.....	36
Chapter 3. Young children’s judgment of informant reliability based on knowledge and intentions.....	65
General Discussion.....	87
References.....	94

LIST OF FIGURES

Figure 1.1: The percentage of trials on which children trusted helpers' and trickers' suggestions by age group.....	34
Figure 3.1: The percentage of children who distrusted the speaker's suggestion on each of six trials of the sticker search task.....	85

LIST OF TABLES

Table 1.1: The percentage of 3-, 4-, and 5-year-olds who gave the more positively valenced response to questions about helpers and trickers.....	33
Table 2.1: Summary of Experimental Design: Number and Type of Informant.....	63

ACKNOWLEDGEMENTS

First I would like to thank my advisors Gail Heyman and David Liu for their support and guidance. They provided me with amazing resources with which to investigate the work presented here and were instrumental in my development as student. I am grateful to them for taking a chance on me when I entered this program, and for supporting me throughout my time as a graduate student. I am especially appreciative of the time and energy they took to read endless drafts, listen to my ideas, and to help me become a better researcher. I am also grateful to Leslie Carver for lending helpful ideas and advice on many occasions, and to the other members of my dissertation committee, Sarah Creel and Dana Nelkin, for lending their invaluable time and guidance to this work and to me.

I am also very grateful to Heather Pelton and Amy Kaestner, the wonderful coordinators of the Liu Lab, for training research assistants, recruiting participants, and helping to keep this research organized and running. They made even the most complicated tasks seem simple. I am also greatly indebted to my awesome research assistants: Rowena Ng, Michael Rizzo, Carol Lee, Christina Callegari, Soo Lee, Jenny Wong, and so many more. I absolutely could not have completed these projects without them.

I would like to thank my amazing labmates: Catherine Hicks, Joe Dietzel and Lalida Sritanyaratana for being such good friends and for always providing me with the strength to move forward. I would also like to thank Caroline Gee and Annette Cluver for

providing me with guidance and reassurance that made succeeding in this program possible.

Finally I would like to thank my friends and family for their unwavering support and guidance. Thank you to my parents, Carolyn and Steve Accatino for their unconditional love, encouragement and understanding, and for cultivating in me a love of learning and teaching others. I could not have done any of this without you. And, I am especially grateful to Julian Parris for being such a selfless teacher, tireless supporter and invaluable friend.

Chapter 1, in full, is a reprint of the material as it appears in The Development of Distrust. *Child Development*, 82, 1372-1380. Vanderbilt, K.E., Liu, D., & Heyman, G.D. (2011). The dissertation author was the primary investigator and author of this paper. Permissions for use of this material have been obtained from John Wiley and Sons.

Chapter 2, in full, has been submitted for publication of the material as it may appear in In the absence of conflicting testimony, young children trust inaccurate informants in *Developmental Science*. Vanderbilt, K.E., Heyman, G. D., & Liu, D. (2013). The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Vanderbilt, K.E., Heyman, G. D., & Liu, D. The dissertation author was the primary investigator and author of this paper.

CIRRICULUM VITAE

- 2013 Doctor of Philosophy, Psychology
University of California, San Diego
- 2008 Master of Arts, Psychology
University of California, San Diego
- 2007 Bachelor of Science, Psychology
University of California, San Diego

PUBLICATIONS

- Vanderbilt, K. E., Liu, D., & Heyman, G. D. (2011). The development of distrust. *Child Development, 82*, 1372-1380.
- Heyman, G. D., Sritanyaratana, L., & Vanderbilt, K. E. (2013). Young Children's Trust in Overtly Misleading Advice. *Cognitive Science, 37*, 646–667.
- Liu, D., Vanderbilt, K. E., & Heyman, G. D. (2013). Selective trust: Children's use of intention and outcome of past testimony. *Developmental Psychology, 49*, 439-445.
- Liu, D. & Vanderbilt, K. E. (2013). Children learn from and about variability between people. In M. R. Banaji & S. A. Gelman (Eds.) *Navigating the Social World: A Developmental Perspective*. New York: Oxford University Press.
- Vanderbilt, K. E., Heyman, G. D., & Liu, D. (under review). In the absence of conflicting testimony young children trust inaccurate informants. *Developmental Science*.

ABSTRACT OF THE DISSERTATION

The Development and Understanding of Selective Trust in Preschool Children

by

Kimberly E. Vanderbilt

Doctor of Philosophy in Psychology

University of California, San Diego, 2013

Professor Gail Heyman, Chair
Professor David Liu, Co-Chair

The purpose of this dissertation was to investigate preschool children's selective trust and understanding of reliability. These experiments present new key findings in the area of selective trust.

Chapter 1 investigated children's reasoning about the reliability of sources based on honesty. Children received advice in a search task from speakers with either a history of helpfulness or deception. Three- and four-year-olds did not discriminate between helpful and deceptive sources, even though 4-year-olds identified and predicted the future behavior of deceptive sources. Five-year-olds systematically preferred advice from

helpers. These results suggest children distrust deceptive sources by 5 years of age, and that children may have knowledge of mental states they do not apply to their trust judgments. Children's ability to selectively trust honest sources was also related to their theory of mind understanding.

Chapter 2 investigated whether children make relative or absolute judgments about the reliability of inaccurate sources. Children heard testimony about novel object names from speakers with a history of correctly or incorrectly labeling familiar objects. Children tended to accept labels provided by inaccurate informants, unless another source (with a history of accurate labeling, or no known history) also provided conflicting testimony. These results suggest young children use prior accuracy to determine the *relative* reliability of informants, but still trust a single informant alone, regardless of that informant's past reliability.

Chapter 3 investigated whether children's selective trust is influenced by the reason a source is unreliable (bad knowledge or bad intentions), and the strength of cues about a source's unreliability. Children heard a puppet give inaccurate advice in a search task as a consequence of bad knowledge or bad intentions. With only minimal cues to speakers' mental states, children distrusted incompetent sources more than deceptive ones, but with stronger cues to unreliability, children distrusted both sources equally. This suggests children have more difficulty reasoning about unreliability resulting from bad intentions.

Collectively, this research highlights the importance of understanding children's trust judgments as well as the reasoning behind them, and suggests children's decisions to

trust sources may not always reflect knowledge they have about the mental states of others.

INTRODUCTION

The Development and Understanding of Selective Trust in Preschool Children

The ability to judge the reliability of sources is of the utmost importance in social life. Because social life requires that we absorb information at a rate far exceeding our ability to gain information firsthand, we often rely on other people to provide us with information about the world. For example, we commonly rely on others to teach us about the existence of germs and bacteria, world history and events, whether a particular animal is classified as a reptile or a fish, and even tomorrow's weather. Children especially, must rapidly acquire a great deal of information about the world, and often rely on other people to teach them such information. Although the transmission of information between individuals and across generation affords us many benefits (both as individuals and as a species) there are also inherent dangers to accepting information from others. As adults, we know to be skeptical of the information we receive from others, or risk being misinformed or misled. But when do children learn this critical social skill? When do children learn to be skeptical of the information others provide?

Background

Recently this question has attracted a great deal of interest in the domain of social cognitive development, and a growing body of research has begun to investigate how children learn to judge the reliability of sources. Despite their reputation as gullible and impressionable (Dawkins 1993), many studies suggest that young children do have some ability to trust more reliable over less reliable sources (Birch, Vauthier, & Bloom, 2008; Clement, Koenig, & Harris, 2004; Jaswal & Malone, 2007; Jaswal & Neely, 2006;

Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Lampinen & Smith, 1995; Lutz & Keil, 2002; Pasquini, Corriveau, Koenig, & Harris, 2007; Sabbagh & Baldwin, 2001). Much of this research has been conducted within the context of word learning. In a typical example of these studies, Koenig and Harris (2005) presented preschool-age children with two speakers who provided conflicting labels for objects. During familiarization trials, one speaker provided accurate labels for familiar objects and the other speaker provided inaccurate labels for the same objects. In subsequent test trials, a novel object was introduced and the speakers provided conflicting novel names for it. Results showed that 4-year-olds selectively endorsed labels from speakers who had accurately named the familiar objects and expected these speakers to continue to provide more accurate labels. Other studies have found similar results with respect to learning novel words (see Harris, 2007; Harris & Corriveau, 2011 for reviews) and novel object functions (Birch, et al., 2008; DiYanni & Keleman 2008; Jaswal, 2006; Koenig & Harris, 2005). Taken together, these studies suggest that preschool-age children are not only able to judge the relative accuracy of speakers, but are able to track accuracy over time and predict that speakers who were previously accurate will continue to be accurate in the future.

There is also evidence that by age 4 children are sensitive to the relative reliability of sources even when they don't have the clear contrast between a speaker who is always reliable and one who is never reliable. Pasquini, et al. (2007) showed that 4-year-olds (but not 3-year-olds) selectively relied on the testimony of a speaker who was correct 75% of the time over a speaker who was correct only 25% of the time. This suggests that

by 4 years of age, children are able to rely on the *more* reliable of two sources, and further, that preschool children track the relative number of errors made by speakers and subsequently use that information to judge those speakers' relative reliability.

There is also evidence that preschool children make assumptions about why some sources may be more knowledgeable than others (Lutz & Keil, 2002; Lampinen & Smith, 1995; Jaswal, 2006; Jaswal & Neely, 2006; Jaswal & Malone, 2007; Sabbagh & Baldwin, 2001; VanderBorghet & Jaswal, 2009). For example, preschool children are aware that a more confident speaker is likely to be a more reliable source: children as young as 3 years of age will selectively trust a more confident speaker over another speaker who expresses ignorance (Koenig & Harris 2005) or uncertainty (Birch, Akmal, & Frampton, 2009; Jaswal & Malone, 2007; Sabbagh & Baldwin, 2001) about the information they provide.

There is also some evidence that children take into account aspects of the mental states of sources. For example, children are aware of and keep track of speakers' knowledge-states. A study by Robinson, Champion, and Mitchell (1999) asked preschool-age children to guess which of two objects was hidden inside a container. Children then heard an adult speaker provide testimony that conflicted with their guess. Children were especially likely to alter their predictions to match those of a speaker whom they had previously seen look inside the container. This suggests that children are sensitive to a speaker's access to relevant knowledge when predicting reliability. Further, children will excuse inaccuracy with good cause, indicating that children are sensitive to the reason a source is incorrect. Work by Robinson and Nurmsoo (Robinson & Nurmsoo, 2009; Nurmsoo & Robinson, 2009a) showed that children are subsequently willing to

rely on a speaker who was previously blindfolded over a speaker who named the same objects wrong inexplicably. This shows that children understand at least some of the reasons underlying other people's inaccuracies. These studies also demonstrate evidence that children deeply process the reasoning behind a speaker's inaccuracies, rather than simply avoiding unreliable source as though it were "broken" (but see Nurmsoo & Robinson, 2009b for evidence that this is not always the case).

The Present Research

Although there is a great deal of research devoted to understanding of children's selective trust, there are still many unanswered questions about how children reason about the reliability of sources. For example, although we know a great deal about how children reason about the *competency* of sources, comparatively little research has addressed how children reason about the *honesty* of sources (but see Mascaro & Sperber, 2009; Heyman, Sritanyaratana, & Vanderbilt, 2013). Do children understand that a source may know the right answer, but convey the wrong answer on purpose? We also do not know a great deal about the nature of the inferences children draw about sources. We know that children are quite good at trusting more competent over less competent sources (Koenig & Harris, 2005; Pasquini et al., 2007), but we do not know what types of inferences children make about these sources individually; do children consider inaccurate sources to be wholly unreliable in these cases, or just slightly less reliable than the accurate source but still generally reliable? Current research in this area does not distinguish between these inferences (Liu & Vanderbilt, 2013). Additionally, although we know that children often take into account the reason for a speaker's inaccuracy

(Robinson & Nurmsoo, 2009; Nurmsoo & Robinson, 2009), we do not know how children use the mental states of speakers to make inferences about reliability. Are children sensitive to the fact that speakers can be unreliable for different reasons (incompetence, malevolence) and do they differentiate between sources based on these differences? Although such mental state information has varied between studies, a direct well-controlled comparison of how children use each of these factors in isolation has not yet been conducted. Understanding the answers to these questions would provide a more complete picture of children's reasoning about reliability, and how their understanding relates to other developing abilities, such as theory of mind.

The purpose of the research in this dissertation is to investigate children's selective trust and understanding of source reliability. We do this in three different ways: In Chapter One we investigate when children begin to distrust deceptive sources, the inferences they make about such sources, and how their selective trust relates to their understanding of mental states. In Chapter Two, we investigate the nature of the inferences children make about unreliable sources by testing whether children make relative or absolute judgments about source reliability. And in Chapter Three, we investigate how the reason for a source's unreliability (either bad knowledge or bad intentions), and the strength of cues to unreliability, influence children selective trust and inferences about unreliable sources. In addressing some of these questions we also include both verbal and nonverbal measures to better characterize children's developing abilities. Across all studies we take an approach of tightly controlling for factors other than the ones targeted for assessment. Taken together this research will allow us to better

identify specific difficulties young children have in reasoning about selective trust at different time points, pinpoint the specific factors that influence children's trust judgments, and gain insight into connections between reasoning about trust and conceptions of mental life. In these ways, the studies included in this dissertation expand and diversify our understanding of children's selective trust, giving a more complete and nuanced understanding of how children learn to judge the reliability of sources.

Chapter 1

The Development of Distrust

Kimberly E. Vanderbilt, David Liu, and Gail D. Heyman

University of California, San Diego

Acknowledgements. We would like to thank the children who participated and their parents. Special thanks to UCSD's Early Childhood Education Center, All Saint's Preschool, Little Steps Preschool, La Jolla Preschool Academy, Early Explorations Preschool, and Academy of Learning Preschool. This research was supported by a grant from NICHD (R01 HD048962).

Address correspondence to:
Kimberly Vanderbilt
Department of Psychology
University of California, San Diego
9500 Gilman Dr.
La Jolla, CA 92093-0109
kvanderbilt@ucsd.edu

Abstract

Preschool-age children's reasoning about the reliability of deceptive sources was investigated. Ninety 3- to 5-year-olds watched several trials in which an informant gave advice about the location of a hidden sticker. Informants were either *helpers* who were happy to give correct advice, or *trickers* who were happy to give incorrect advice. Three-year-olds tended to accept all advice from both helpers and trickers. Four-year-olds were more skeptical but showed no preference for advice from helpers over trickers, even though they differentiated between helpers and trickers on metacognitive measures. Five-year-olds systematically preferred advice from helpers. Selective trust was associated with children's ability to make mental state inferences.

The Development of Distrust

People often rely on others for information that would be difficult or impossible to acquire through direct experience. Typically, people who provide information are reliable and trusted by their audiences (Coady, 1992). However, people are not always reliable, and listeners must learn to differentiate good sources from bad ones. Recently, there has been substantial interest in the development of children's capacities in this domain (see Harris, 2007; Heyman, 2008). The present research focuses on development during the preschool years, when children must rely heavily upon information provided by others, but may lack the cognitive skills and social experience to exercise appropriate skepticism (Dawkins, 1993; Moses & Baldwin, 2005).

Prior research suggests that preschool-age children understand that certain individuals are more reliable than others (Birch, Vauthier, & Bloom, 2008; Clement, Koenig, & Harris, 2004; Koenig & Harris, 2005; Lampinen & Smith, 1995; Lutz & Keil, 2002; Pasquini, Corriveau, Koenig, & Harris, 2007; Jaswal & Malone, 2007; Jaswal & Neely, 2006; Sabbagh & Baldwin, 2001). Much of this research has been conducted within the context of word learning. For instance, Koenig and Harris (2005) presented preschool-age children with two speakers who provided conflicting labels for objects. During training trials, one speaker provided accurate labels for familiar objects and the other speaker provided inaccurate labels. In subsequent test trials, a novel object was introduced and the speakers provided conflicting novel names for it. Results showed that 4-year-olds selectively endorsed labels from speakers who had accurately named the familiar objects and expected these speakers to continue to provide more accurate labels.

Other work using related paradigms has demonstrated that preschool-age children show similar capacities when reasoning about the functions of novel objects (Birch et al., 2008; Clement et al., 2004).

Similar findings have been obtained in research investigating how young children use others as sources of information to help them identify hidden objects. Robinson, Champion, and Mitchell (1999) asked preschool-age children to guess which of two objects was hidden inside a container. Children then heard an adult provide testimony that conflicted with their guess. Children were more likely to alter their prediction to match the adult's when the adult had conspicuously looked inside the container, as compared to when the adult had not looked inside. This suggests that when attempting to identify hidden objects, young children are capable of understanding that not all sources are equally reliable. There is also evidence that young children make assumptions about why some individuals may be more knowledgeable sources than others (Jaswal, 2006; Jaswal & Malone, 2007; Jaswal & Neely, 2006; VanderBorghet & Jaswal, 2009; Lampinen & Smith, 1995; Lutz & Keil, 2002; Sabbagh & Baldwin, 2001). For instance, preschool-age children make distinctions between the knowledge possessed by others with respect to expertise (Jaswal, 2006; Lutz & Keil, 2002; Sabbagh & Baldwin, 2001) and age (Jaswal & Neely, 2006; Lampinen & Smith, 1995; Vanderborghet & Jaswal, 2009).

These studies suggest that young children consider the *knowledgability* of an informant when judging his or her reliability. However, assessment of reliability goes beyond mere consideration of a source's knowledge, consideration of his or her *motives*

and intentions is also critical. People are not always motivated to accurately communicate what they know – people lie. For example, speakers may deceive others to promote their own interests (Heyman & Legare, 2005) or the interests of others (Heyman, Sweet, & Lee, 2009). Although adults may sometimes have difficulty detecting deception, they recognize that deception is possible and that an individual's prior deceptive behavior is an important cue to whether that individual should be trusted in the future. Is this true of children as well?

Previous research suggests at least two distinct possibilities. One possibility is that even young children are likely to show substantial competence in this domain. Before age 2 years, children can make inferences about the intentions of others, and use those inferences to guide their behavior and expectations (Behne, Carpenter, Call, & Tomasello, 2005; Meltzoff, 1995; Woodward, 1998); and by age 3 years, children appreciate that their own deceptive behaviors can influence the behavior and beliefs of others (Chandler, Fritz, & Hala, 1989). Preschool-age children demonstrate earlier understanding of motives than knowledge states (Wellman & Liu, 2004) and show more sophisticated reasoning on false-belief tasks when deception is emphasized (Wellman, Cross, & Watson, 2001). There is also evidence that young children are highly sensitive to motive information when reasoning about traits (Heyman & Gelman, 1998; Yuill, 1992). This research suggests the intentions of others are salient to even young children, and that they serve as important cues for making generalizations about people and predicting their actions.

An alternative possibility is that young children have particular difficulty reasoning about deceptive communication and its implications. Early elementary school children often fail to critically evaluate the claims of others (Heyman, 2008; Heyman & Legare, 2005; Mills & Grant, 2008; Mills & Keil, 2005; Mills & Keil, 2008; Moses & Baldwin, 2005). Before second grade, young children often fail to appreciate self-serving biases that may underlie people's judgments (Mills & Keil, 2005), and do not consider that people's judgments may be biased due to their personal relationships, such as when the judge of a contest is a friend or adversary of a contestant (Mills & Grant, 2008; Mills & Keil, 2008). Furthermore, young children appear to have difficulty with reasoning about value-laden claims people make about themselves (Heyman, 2008). For example, Heyman and Legare (2005) found that 6- and 7-year-old children believe one can assess honesty simply by asking someone if he or she is honest, whereas 10- and 11-year-old children understand that such communication is subject to social desirability effects.

Young children's difficulty understanding potentially deceptive motives suggests they may have trouble understanding the implications of deception more broadly. Consistent with this possibility, Lee and Cameron (2000, Study 1) found that even when preschool-age children realize a source is lying, they nonetheless rely on the information the source presents. This finding suggests that even when young children recognize deception, they may not appreciate its implications.

Experiment 1

The present research examines preschool-age children's use of information about a source's history of deception to assess reliability. On each of several trials, children

watched a video in which an adult actor (the pointer) helped or tricked two other adult actors (finders), who were attempting to locate a sticker prize hidden in one of two boxes. On *helper* trials, participants observed the pointer helping finders locate the sticker, and on *tricker* trials the pointer was shown tricking finders into looking for the sticker in the wrong location. On each trial, the pointer either helped or tricked two other finders. Then, also via video, the pointer gave advice to the participant on the same sticker-finding task. Participants then decided whether to select the box recommended to them by the pointer, or the other box. Children were not given feedback after each trial, and were told at the end of the session they would be able to open all the boxes they selected and keep any stickers they found.

In addition to the primary experimental measure, two other sets of questions were included. The first was a set of metacognitive questions designed to provide information about children's understanding of the experimental manipulation and their attributions about the informants. For example, children were asked whether they believed the pointers were attempting to help or trick finders and to predict whether the pointers would suggest the right or wrong location to the next player.

The second set of questions was included to assess whether children's ability to selectively trust helpful versus deceptive sources is associated with their theory of mind development. Several researchers have suggested that young children's reasoning about sources of information hinges upon their understanding of mental life (Koenig & Harris, 2005; Moses & Baldwin, 2005). Indeed, many aspects of theory of mind understanding appear necessary for reasoning about whether another person will be helpful or harmful

when offering information, and we wanted to test this relationship with respect to reasoning about the intentions of others. DiYanni and Kelemen (2008) found a correlation between children's selective trust of knowledgeable speakers and their performance on false-belief tasks, but Pasquini et al. (2007) found no such correlation. Pasquini et al. suggested that to better explore this question, it would be useful to include a battery of tests, and suggested Wellman and Liu's (2004) Theory of Mind Scale as an example. The present study follows this suggestion and includes the Wellman and Liu scale.

Method

Participants

Ninety preschool-age children (39 boys, 51 girls) participated in the study: 30 3-year-olds ($M = 3$ years 7 months, range = 3 years 2 months to 3 years and 11 months), 30 4-year-olds ($M = 4$ years 6 months, range = 4 years to 4 years 11 months) and 30 5-year-olds ($M = 5$ years 3 months, range = 5 years to 5 years 11 months). There was a similar number of boys and girls in each age group. Children were recruited from preschools in a city in southern California. The sample was approximately 50% Caucasian, 25% Asian, 15% Latino, and 10% African-American.

Procedure and Materials

Sticker-finding task. Children completed six trials of the sticker-finding task. Children were instructed that they would be playing a game in which their goal (and the goal of other finders) would be to locate a sticker hidden in one of two boxes, that they had a chance to find a sticker on each trial and could keep all the stickers they found.

Trial type was manipulated within-subjects, such that each participant saw three *helper* and three *tricker* trials. Orders of presentation and trial type were counterbalanced between subjects, as was spatial position of the box indicated by each pointer. A different set of boxes was used for each trial.

To maintain consistency in the actions of pointers and finders across trials, their actions were presented in video vignettes. Each vignette depicted a different adult female pointer. All pointers were presented identically, verbally suggesting and pointing to one box. Helper and tricker trials differed only in whether the vignettes revealed a sticker or no sticker inside the opened box.

On helper trials, children observed the pointer helping two finders locate stickers hidden in one of the two boxes. On tricker trials, children observed the pointer tricking two finders into looking in the wrong box for stickers. In each trial, the pointer would point to one of the boxes and say to the first finder, “You should pick this one.” The first finder would then open the suggested box, revealing a sticker (on helper trials) or no sticker (on tricker trials). Finders would exclaim “yay!” upon finding stickers and, “aww!” when they failed to find stickers. Regardless of the outcome, the pointer would always respond, “Yes!” smiling to show she was pleased with the outcome. Thus, a helping pointer displayed positive affect after she succeeded in helping someone find a sticker, whereas a tricking pointer displayed positive affect after she succeeded in tricking someone into not finding a sticker. The pointer then repeated these same actions with a second finder, so that children observed either two instances of the pointer helping finders or two instances of the pointer tricking finders.

On each trial, after observing the actions of the pointer and two finders, it was the child's turn to be the finder. A new pair of differently colored boxes was placed in front of the child on each trial, and they were reminded, "Now, this pointer *knows* where the sticker is, and she'll tell you which box she wants you to pick. Pick the box *you* think the sticker is in." Children then observed the pointer providing them with a suggestion just as she had with the other finders: pointing to one of the boxes and saying, "You should pick this one." The experimenter then prompted children to select the box with the sticker inside, saying: "Which box do you think has the sticker?" Children's responses were scored for whether they followed the pointer's suggestion. During the test trial, boxes were depicted both in video vignettes so that pointers could clearly suggest one of the two boxes, and in real life so that children could physically choose the box they believed contained the sticker.

Metacognitive questions. Children answered four forced-choice metacognitive questions about each type of pointer (helper and tricker). Half answered metacognitive questions before the sticker-finding task, and half did so after the sticker-finding task. For metacognitive questions, children observed a new pointer helping two finders and a new pointer tricking two finders (as in the sticker-finding task). After observing each new pointer, children were asked two task-specific questions and two generalization questions. The two task-specific questions were an intention judgment question ("Was this pointer trying to help or trick the other person in the video?") and a same-context prediction question ("Will this pointer try to tell the next person the right box or the wrong box?"). The two generalization questions were a trait judgment question ("Is this

pointer nice or mean?") and a generalized-context prediction question ("Let's say this pointer just saw another person trip, fall and drop a lot of papers. Do you think this pointer would stop to help the person pick up the papers or just keep walking and not stop to help?").

Theory of Mind Scale. Following the sticker-finding task and metacognitive questions, children's theory of mind was assessed using the Theory of Mind Scale developed by Wellman and Liu (2004). This 5-item scale contains different tasks that capture the developmental progression of children's mental state understanding between 3 and 6 years of age. The tasks assess understanding of diverse desires, diverse beliefs, knowledge-ignorance, false beliefs, and false emotions. The diverse-desires task asks children to reason about a situation in which someone has a different preference than their own. The diverse-beliefs task asks children to reason about a situation in which another person has a different, explicitly stated, belief than their own. The knowledge-ignorance task asks children to reason about a situation in which they knew the contents of a non-descript container but another person did not. The false-belief task asks children to reason about a situation in which another person has a false belief about the location of an object. Lastly, the false-emotions task asks children to reason about a situation in which another person displays emotions counter to what they feel. Scores ranged from 0 (passing none) to 5 (passing all).

Results and Discussion

Preliminary analyses showed no significant effects of gender, task order (completing metacognitive questions or sticker-finding task first), or trial order

(presentation order of helpers and trickers); consequently, these variables were excluded from the remaining analyses.

In the sticker-finding task, as shown in Figure 1, 3-year-olds trusted helpers 94.4% of the time and trusted trickers 91.1% of the time. Four-year-olds trusted helpers 68.9% of the time and trusted trickers 60.0% of the time. Five-year-olds trusted helpers 71.1% of the time and trusted trickers 51.1% of the time. Tests against chance showed that 3-year-olds were above chance at trusting both helpers and trickers, $t(29) = 19.27, p < .001$, and $t(29) = 12.97, p < .001$, respectively; 4- and 5-year-olds were above chance at trusting helpers, $t(29) = 3.17, p < .01$, and $t(29) = 3.33, p < .01$ respectively, but were at chance at trusting trickers, $t(29) = 1.27, ns$, and $t(29) = .17, ns$ respectively.

A 2 (Intent: helper versus tricker) x 3 (Age: 3, 4, and 5 years) repeated-measures analysis of variance (ANOVA) was conducted to test whether children would selectively trust informants who had previously helped others over those who had previously tricked others. There was a significant main effect of Intent, $F(1, 87) = 8.149, p = .005$, showing that children trusted helpers more often than they trusted trickers. There was also a significant main effect of Age, $F(2, 87) = 13.451, p < .001$, indicating that younger children displayed more overall trust than did older children. The interaction was not significant.

Follow-up analysis showed there was a significant decrease in trust displayed between the ages of 3 and 4 years, $F(1, 58) = 18.647, p < .001$, but not between the ages of 4 and 5 years. Notably, 3-year-olds followed pointers' suggestions on almost all trials, even when pointers had previously tricked others. Because we were interested in the

specific ages at which children would selectively trust, we also analyzed each age group separately. Targeted analyses performed separately for each age group indicated that neither 3- nor 4-year-olds differentiated between helpers and trickers when deciding whom to trust. Only 5-year-olds were found to actively discriminate between helpers and trickers, trusting helpers significantly more often than trickers, $t(29) = 2.340, p = .026$. Thus, the manipulation of pointers' intent in past behaviors had an effect on 5-year-olds' judgments of whom to trust, but did not have an effect on 3- and 4-year-olds' judgments. One 3-year-old, four 4-year-olds, and eight 5-year-olds were highly consistent in their correct selective trust (answering 5 or more trials out of 6 correctly); this age trend is consistent with the above analyses showing that not until age 5 do children demonstrate significant selective trust of helpers versus trickers, but even their performance is not at ceiling.

Metacognitive Questions

A summary of responses to the metacognitive questions is presented in Table 1 (one 3-year-old did not answer these questions). We analyzed children's answers to the two task-specific and two generalization metacognitive questions using nonparametric tests because children answered each dichotomous question once for a helper and once for a tricker. For the task-specific questions, 3-year-olds did not distinguish between helpers and trickers on either item. However, 4- and 5-year-olds correctly differentiated between helpers and trickers on both task-specific questions: both 4- and 5-year-olds said that helpers were more likely than trickers to be trying to help, McNemar's $\chi^2(1) = 9.941, p = .002$, McNemar's $\chi^2(1) = 14.222, p < .001$, respectively; also both 4- and 5-

year-olds predicted that helpers would be more likely than trickers to tell the next person the correct location of the sticker, McNemar's $\chi^2(1) = 11.267, p < .001$, McNemar's $\chi^2(1) = 4.546, p = .033$, respectively. For the two generalization questions, 3- and 4-year-olds did not distinguish between helpers and trickers. However, 5-year-olds did differentiate between helpers and trickers on the trait judgment question, McNemar's $\chi^2(1) = 16.00, p < .001$, and the generalized-context prediction question, McNemar's $\chi^2(1) = 9.941, p = .002$.

In sum, the overall pattern of results on the metacognitive measures shows that 3-year-olds did not recognize the difference between helpers and trickers on any measure. Four-year-olds differentiated helpers and trickers on task-specific questions, but were unable to generalize these differentiations. Only 5-year-olds generalized the differentiation between helpers and trickers.

Theory of Mind Scale

Lastly, we examined whether children's performance on the sticker-finding task was correlated with theory of mind. On average, 3-year-olds scored 2.33 (out of 5) on the Theory of Mind Scale, 4-year-olds scored 3.07, and 5-year-olds scored 3.60. To perform this individual-differences analysis, we calculated a difference score for the sticker-finding task as a measure of how well each child differentiated helpers from trickers by subtracting the number of trials in which the child followed a tricker's suggestion (an incorrect response) from the number of trials the child followed a helper's suggestion (a correct response). Children's difference score on the sticker-finding task was positively correlated with their performance on the Theory of Mind Scale, $r(87) =$

.339, $p = .001$. This correlation remained significant even after controlling for age statistically, $r(87) = .290$, $p = .006$, which suggests that children's ability to selectively trust helpers versus trickers is associated with their developing understanding of mind. Just as we calculated a difference score for the sticker-finding task, we also calculated an overall difference score for children's performance on the metacognitive questions. Children's performance on the metacognitive questions was correlated with performance on the Theory of Mind scale, $r(86) = .336$, $p = .001$, and after controlling for age, the correlation was marginally significant, $r(86) = .197$, $p = .066$. Additionally, we did not find a correlation between children's performance on the sticker-finding task and their performance on metacognitive questions.

The pattern of results from Experiment 1 suggests that selective trust of deceptive versus helpful informants develops in preschool-age children. Whereas 5-year-olds significantly trusted helpers more than trickers, 3- and 4-year-olds were just as likely to trust helpers as they were to trust trickers. Interestingly, 4-year-olds trusted trickers, even though they expressed in their answers to the metacognitive questions that previously helpful informants would be more likely to provide accurate information than trickers.

Experiment 2

Why did 4-year-olds not incorporate their knowledge of helpers versus trickers into their decisions of whether to trust trickers? Before considering theoretical explanations, we wanted to rule out a possible explanation having to do with our procedures misleading children about the pointer's knowledge states. When children were asked to find the sticker, we reminded them that the pointer knew the sticker's

location, in order to emphasize that informants' past inaccuracies could not be attributed to lack of knowledge. However, it is possible that this mention of knowledge paradoxically cued children to focus on the question of the pointers' knowledge, rather than the pointers' motives. In order to rule out this possibility, we conducted a short control experiment with modified instructions.

Method

Participants

Fourteen 4-year-old children ($M = 4$ years 3 months, range = 4 years 1 months to 4 years and 10 months) participated in Experiment 2: seven children were boys and 7 were girls. Children were recruited from preschools in a city in southern California.

Procedure and Materials

The procedure and materials for the sticker-finding task were identical to those of Experiment 1, except that on each trial, after observing the pointers interact with the two finders, children were simply told "Pick the box you think has the sticker," and received no reminder about the informant's knowledge-state, as they did in Experiment 1.

Results and Discussion

Replicating Experiment 1, the 4-year-olds in Experiment 2 failed to differentiate, in their responses, helpers from trickers. Specifically, they trusted helpers on 83.3% of trials and trickers on 80.9% of trials $t(13) = .583$, *ns*. These results indicate that reminding participants that the pointer knew the sticker's location cannot explain the results of Study 1. The overall rates of trust between Experiment 2 and Experiment 1 were not different ($F(1, 42) = 2.56$, $p < .10$, *ns*, when we compared the rates across

experiments, and importantly, there was not a significant interaction between experiment and informant type.

General Discussion

The present work was designed to investigate preschool-age children's understanding that an informant's prior history of deception has implications for his or her reliability. We measured children's trust of speakers who had a prior history of helpful or deceptive behavior, and asked them specific metacognitive questions about each type of speaker. Although 5-year-olds significantly trusted helpers more than trickers, they demonstrated only chance performance on tricker trials. In contrast, 3-year-olds trusted both helpers and trickers almost all of the time. Four-year-olds were more skeptical than 3-year-olds overall, but still did not distinguish between helpers and trickers, despite the fact that they expressed the expectation that previously helpful speakers would be more likely to provide accurate information in the future.

Might it be that 3- and 4-year-olds failed to recognize that deception occurred altogether? This might explain the performance of the 3-year-olds, but 4-year-olds *did* recognize trickers' deceptive intent, as indicated by their metacognitive answers. Why then, did 4-year-olds not incorporate this knowledge into their decisions of whether to trust trickers?

We believe this mismatch between knowledge and behavior is best explained by children's failure to understand the implications of their knowledge and how to effectively apply it to their behavior. There are several ways this could manifest. Four-year-olds might not always apply what they know when reasoning about sources because

doing so conflicts with a general assumption that people will provide accurate information. Consequently, young children are likely to accept what they are told unless there are clear indications that the information may be unreliable, such as when adults indicate uncertainty (Jaswal & Malone, 2007; Sabbagh & Baldwin, 2001), or when new information conflicts with prior beliefs and the communicative intent is ambiguous (Jaswal, 2004). A related possibility is that young children hold a default assumption that adults are honest and have prosocial intentions, and this assumption can only be overridden by strong situational cues. Such an assumption may seem surprising in light of evidence that children already have some capacity to distinguish between prosocial and antisocial intentions before they reach their first birthday (Hamlin, Wynn, & Bloom, 2007), and are capable of identifying lies during the early preschool years (Lee & Cameron, 2000). However, preschool-age children often act based on what they desire to be true rather than what they believe to be true (Stipek, Roberts, & Sanborn, 1984), and both preschool- and early elementary-age children tend to hold highly positive views of others under circumstances in which older children and adults do not (Boseovski & Lee, 2008; Droege & Stipek, 1993; Heyman, 2009; Lockhart, Chang, & Story, 2002). Moreover, it is not uncommon for early elementary school children to argue that people can be trusted to tell the truth because they should tell the truth (Heyman & Legare, 2005). Thus, just as Lee and Cameron (2000) showed that young children can identify lies but nonetheless accept information the lies convey, our results suggest young children can identify deceptive intent but nonetheless accept information the deception provides.

Another possibility for why 4-year-olds recognized the pointer's deceptive intent, but were unable to incorporate this knowledge into their predictions of whom to trust is that children do not yet fully understand how to engage in the multi-step process that may be required for them to do so. A parallel finding has been observed in young children's difficulty with solving two-component inferences when making trait predictions (Liu, Gelman, & Wellman, 2005). In that study, preschool-age children were able to identify traits from others' past behaviors, and were able to predict future behaviors when provided with trait labels of others. Nevertheless, they were unable to put those two component inferences together to predict future behaviors from others' past behaviors. Similarly, in the current study, children were asked to predict future behaviors from others' past behaviors. Specifically, the 4-year-olds in our study were able to identify deceptive intent from others' past behaviors, but were unable to use this information to inform their behaviors. This parallels two of the three pieces of evidence found by Liu, Gelman, and Wellman (2005). What about the third piece of evidence? Would children be able to judge whether to trust someone, if they were provided with labels of others' deceptive intent (and they did not need to figure this out for themselves)? In fact, a recent study by Mascaro and Sperber (2009) showed that 4-year-olds knew not to trust those labeled as liars. The results from our study and Mascaro and Sperber (2009) suggests that young children might have difficulty solving two-component inference problems in selective trust, even when they are capable of solving each component inference by itself.

Overall, our findings show preschool-age children behaving more credulously than would be expected from previous studies of young children's selective trust of informants (e.g., Koenig & Harris, 2005). Even 5-year-olds were at chance when deciding whether to trust trickers. This is consistent with prior research demonstrating that children's critical thinking skills continue to develop after the preschool years (Heyman, 2008; Mills & Keil, 2005). We believe some of the issues discussed above contribute to the differences in findings among studies. Children's positive assumptions about people's motives may be harder to override than their assumptions about people's knowledge, and solving multi-step inference problems may be more difficult when reasoning about motives than when reasoning about knowledge.

There are also other factors that may contribute to discrepancies among studies. In the design of many selective trust studies, children observe two sources directly contradicting each other (e.g., Koenig & Harris, 2005). This contradiction may cue young children to more deeply process the possibility that specific individuals will give them inaccurate information. In the present study there were no such cues; children had to judge that a single informant's testimony would be incorrect, and there was nothing in their current environment to signal deceptive intent. Moreover, the fact that our instructions were accompanied by pointing may have made it easier for children to unquestioningly accept what they are told (Couillard & Woodward, 1999).

In future research it will be important to examine other possible influences on selective distrust among children of different ages and adults. One factor that is likely to be important is the nature of the evidence concerning whether a potential informant is

deceptive, including the number of times the individual was observed engaging in deception and how deceptive actions are construed. For example, young children may differ in their responses depending on whether actions construed as tricks, errors, or jokes, and these patterns of construal may be different than those of adults. Additionally, it will be important to examine what individuals may learn from others they observe being targeted for deception. In the present research participants consistently observed recipients trusting in these circumstances, which may have encouraged participants to do the same.

Four-year-olds differentiated between helpers and trickers in their predictions of future behavior, but not in their behavioral responses. This finding suggests that although behavioral evidence of a new understanding often precedes the child's ability to express it verbally (e.g., Alibali & Goldin-Meadow, 1993; Clements & Perner, 1994; Robinson, Champion & Mitchell, 1999; Ruffman, Garnham, Import, & Connolly, 2001) this is not always the case. Our results suggest that it is possible for children to express their knowledge of the social world verbally without using it to guide their behavior.

In addition to investigating preschool-age children's understanding of helpful and deceptive sources, we were interested in whether this understanding is related to their developing theory of mind. To test this relationship, we also presented children with the full Theory of Mind Scale (Wellman & Liu, 2004). In line with evidence from DiYanni and Keleman (2008), we found a correlation between children's selective trust of helpful over deceptive sources and their understanding of mental states, even after controlling for age. Because this evidence is correlational, we cannot conclusively show that having a

more advanced theory of mind leads children to have more sophisticated selective trust and distrust. Furthermore, since we did not control for other variables such as verbal ability or IQ, we cannot rule out the possibility that some other factors account for this relationship. However, this result provides some evidence for a link between children's theory of mind understanding and their ability to distrust deceptive sources.

In sum, our results suggest that young children are more credulous, in certain conditions, than the general impression given by the recent selective trust literature. They also suggest that children's reasoning about whom to trust is closely intertwined with their developing understanding of mental life. Further, our results demonstrate that children do not always show greater competence on behavioral measures than verbal measures, and suggest that children sometimes have difficulty understanding the implications of their knowledge or translating this knowledge into action.

References

- Alibali, M. W., & Goldin-Meadow, S. (1993). Gesture-speech mismatch and mechanisms of learning: What the hands reveal about a child's state of mind. *Cognitive Psychology, 25*, 468-523.
- Behne, T., Carpenter, M., Call, J., & Tomasello, M. (2005). Unwilling versus unable: Infants' understanding of intentional action. *Developmental Psychology, 41*, 328-337.
- 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.
- Boseovski, J. J., & Lee, K. (2008). Seeing the world through rose-colored glasses? Neglect of consensus information in young children's personality judgments. *Social Development, 17*, 399-415.
- Chandler, M., Fritz, A. S., & Hala, S. (1989). Small-scale deceit: Deception as a marker of two-, three-, and four-year-olds' early theories of mind. *Child Development, 60*, 1263-1277.
- Clement, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*, 360-379.
- Clements, W. A., & Perner, J. (1994). Implicit understanding of belief. *Cognitive Development, 9*, 377-395.
- Coady, C. A. J. (1992). *Testimony*. Oxford: Oxford University Press.
- Couillard, N. L., & Woodward, A. L. (1999). Children's comprehension of deceptive points. *British Journal of Developmental Psychology, 17*, 515-521.
- Dawkins, R. (1993). Viruses of the mind. In B. Dahlbom (Eds.), *Dennett and his critics: Demystifying mind* (pp. 13-27). Oxford, England: Blackwell.
- DiYanni, C., & Kelemen, D. (2008). Using a bad tool with good intention: Young children's imitation of adults' questionable choices. *Journal of Experimental Child Psychology, 101*, 241-261.
- Droege, K. L., & Stipek, D. J. (1993). Children's use of dispositions to predict classmates' behavior. *Developmental Psychology, 29*, 646-654.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature, 450*, 557-560.

- Harris, P. (2007). Trust. *Developmental Science, 10*, 135-138.
- 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.
- Heyman, G. D. (2008). Children's critical thinking when learning from others. *Current Directions in Psychological Science, 17*, 344-347.
- Heyman, G. D., & Gelman, S. A. (1998). Young children use motive information to make trait inferences. *Developmental Psychology, 34*, 310-321.
- Heyman, G. D., & Legare, C. H. (2005). Children's evaluation of sources of information about traits. *Developmental Psychology, 41*, 636-647.
- Heyman, G. D., Sweet, M. A., & Lee, K. (2009). Children's reasoning about lie-telling and truth-telling in politeness context. *Social Development, 18*, 728-746.
- Jaswal, V. K. (2004). Don't believe everything you hear: Preschoolers' sensitivity to speaker intent in category induction. *Child Development, 75*, 1871-1885.
- Jaswal, V. K. (2006). Preschoolers favor the creator's label when reasoning about an artifact's function. *Cognition, 99*, B83-B92.
- Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development, 8*, 263-283.
- 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, 9*, 757-758.
- Koenig, M., Clement, F., & Harris, P. (2004). 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.
- Lampinen, J. M., & Smith, V. L. (1995). The incredible (and sometimes incredulous) child witness: Child eyewitnesses' sensitivity to source of credibility cues. *Journal of Applied Psychology, 80*, 621-627.
- Lee, K., & Cameron, C. A. (2000). Extracting truthful information from lies: Emergence of the expression-representation distinction. *Merrill-Palmer Quarterly, 46*, 1-20.

- Liu, D., Gelman, S. A. & Wellman, H. M. (2007). Components of Young Children's Trait Understanding: Behavior-to-Trait Inferences and Trait-to-Behavior Predictions *Child Development, 78*, 1543-1558.
- Lockhart, K. L., Chang, B., & Story, T. (2002). Young children's beliefs about the stability of traits: Protective optimism? *Child Development, 73*, 1408-1430.
- Lutz, D. J., & Keil, F. C. (2002). Early understandings of the division of cognitive labor. *Child Development, 73*, 1073-1084.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition, 112*, 367-380.
- Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology, 31*, 838-850.
- Mills, C. M., & Grant, G. (2008, February). *Understanding that judgments can be skewed: A developmental perspective*. Poster session presented at the annual meeting of the Society of Personality and Social Psychology, Albuquerque, NM.
- Mills, C. M., & Keil, F. C. (2008). Children's developing notions of (im)partiality. *Cognition, 107*, 528-551.
- Mills, C. M., & Keil, F. C. (2005). The development of cynicism. *Psychological Science, 16*, 385-390.
- Moses, L. J., & Baldwin, D. A. (2005). What can the study of cognitive development reveal about children's ability to appreciate and cope with advertising? *Journal of Public Policy & Marketing, 24*, 186-201.
- Pasquini, E. S., Corriveau, K. H., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*, 1216-1226.
- Robinson, E. J., Champion, H., & Mitchell, P. (1999). Children's ability to infer utterance veracity from speaker informedness. *Developmental Psychology, 35*, 535-546.
- Ruffman, T., Garnham, W., Import, A., & Connolly, D. (2001). Does eye gaze indicate implicit knowledge of false belief? Charting transitions in knowledge. *Journal of Experimental Child Psychology, 80*, 201-224.
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72*, 1054-1070.

- Stipek, D. J., Roberts, T. A., & Sanborn, M. E. (1984). Preschool-age children's performance expectations for themselves and another child as a function of the incentive value of success and the salience of past performance. *Child Development, 55*, 1983-1989.
- VanderBorgh, M., & Jaswal, V. K. (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development, 18*, 61-71.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development, 72*, 655-684.
- Wellman, H. M., & Liu, D. (2004). Scaling of theory of mind tasks. *Child Development, 75*, 523-541.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition, 69*, 1-34.
- Yuill, N. (1992). Children's conceptions of personality traits. *Human Development, 35*, 265-279.

Table 1.1: The percentage of 3-, 4-, and 5-year-olds who gave the more positively valenced response to questions about helpers and trickers.

<u>Metacognitive (response)</u>	<u>3-year-olds</u>		<u>4-year-olds</u>		<u>5-year-olds</u>	
	<u>helper</u>	<u>tricker</u>	<u>helper</u>	<u>tricker</u>	<u>helper</u>	<u>tricker</u>
Task Specific Questions						
intention judgment (<i>help</i>)	48.3	34.4	66.7	23.3	66.7	13.3
prediction (<i>right box</i>)	65.5	48.3	70	26.7	63.3	30
Generalization Questions						
trait judgment (<i>nice</i>)	86.7	80	90	70	83.3	30
prediction (<i>help</i>)	58.6	55.2	76.7	56.7	76.7	33.3




Figure 1.1: The percentage of trials on which children trusted helpers' and trickers' suggestions by age group.

Chapter 1, in full, is a reprint of the material as it appears in The Development of Distrust. *Child Development*, 82, 1372-1380. Vanderbilt, K.E., Liu, D., & Heyman, G.D. (2011). The dissertation author was the primary investigator and author of this paper.

Chapter 2

In the Absence of Conflicting Testimony Young Children Trust Inaccurate Informants

Kimberly E. Vanderbilt, Gail D. Heyman, and David Liu

University of California, San Diego

Abstract

The present research investigated the nature of the inferences and decisions young children make about informants with a prior history of inaccuracies. Across three experiments, 3- and 4-year-olds (total $N = 182$) reacted to previously-inaccurate informants who offered testimony in an object-labeling task. Of central interest was children's willingness to accept information provided by an inaccurate informant in different contexts of being alone, paired with an accurate informant, or paired with a novel (neutral) informant. Experiments 1 and 2 showed that when a previously-inaccurate informant was alone and provided testimony that was not in conflict with the testimony of another informant, children systematically accepted the testimony of that informant. Experiment 3 showed that children accepted testimony from a neutral informant over an inaccurate informant when both provided information, but accepted testimony from an inaccurate informant rather than seeking information from an available neutral informant who did not automatically offer information. These results suggest that even though young children use prior history of accuracy to determine the *relative* reliability of informants, they are quite willing to trust the testimony of a single informant alone, regardless of whether that informant had previously been reliable.

In the Absence of Conflicting Testimony Young Children Trust Inaccurate Informants

As humans, we rely heavily on others for information that would be too time-consuming or difficult to acquire firsthand. Children, especially, learn a great deal of natural and cultural knowledge from others (Gelman, 2009). Although social learning allows children to acquire knowledge quickly, there are obvious drawbacks. Others sometimes provide information that is inaccurate—informants might lack knowledge, have deceptive intentions, or simply make errors. To avoid being misinformed, it is critical for children to learn to critically evaluate both the trustworthiness of informants and the information they provide. Recent research suggests that even young children critically evaluate the trustworthiness of testimony (Koenig & Harris, 2005, Sabbagh & Baldwin, 2001); they preferentially trust individuals with a history of providing accurate testimony over those with a history of providing inaccurate testimony when given a choice (Birch, Vauthier, & Bloom, 2008; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig, & Harris, 2007; see Harris & Corriveau, 2011 for review). Nevertheless, the reasoning underlying children's decision of whom to trust is not well understood. The present research seeks to help fill this gap by examining what inferences children make when evaluating informants with a history of inaccuracy and the nature of children's decisions about whether to accept inaccurate informants' testimony.

Children's acceptance of an accurate informant's testimony over an inaccurate informant's testimony is consistent with two possible underlying inferences about the inaccurate informant. One possibility is that children infer that someone who has

previously provided inaccurate information is *absolutely* unreliable and should be strongly distrusted. Another possibility, however, is that children infer that someone who has previously provided inaccurate information is *relatively* less reliable than someone who has previously provided accurate information, but is still generally trustworthy.

In order to effectively evaluate inaccurate informants, children must first be able to identify inaccuracies and track them. Prior research suggests that these skills are in place early in life. Infants as young as 16 month of age look longer at, respond differently to, and correct speakers who labeled familiar objects incorrectly (Koenig & Echols, 2003; Koenig & Woodward, 2010; Pea, 1982). Beyond identifying inaccuracies, Corriveau, Meints, and Harris (2008) observed that children as young as 3 years of age track informants' inaccuracies. Nevertheless, prior research does not allow for a determination of whether children view inaccurate informants as unreliable only in a relative sense or also in an absolute sense.

Although both hypothesized inferences about the inaccurate informant are consistent with the observed results of children trusting accurate informants over the inaccurate ones (Birch, et al., 2008; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Pasquini, et al., 2007), they predict very different responses when children are deciding whether to trust a single inaccurate informant. If children infer that someone who got a few things wrong is absolutely unreliable, they will distrust testimony from the single inaccurate informant. Conversely, if children infer that someone who got a few things wrong is less reliable than others, but still generally reliable, they will trust testimony from the single inaccurate informant. In Experiment 1, we sought to investigate

which of the two hypothesized inferences children make when they are presented with an inaccurate informant by comparing children's responses when presented with a single inaccurate informant alone versus when presented with an inaccurate informant alongside an accurate informant.

Experiment 1

In a between-subjects design, we presented 3- and 4-year-olds with one of three conditions: 1) an accurate informant paired with an inaccurate informant, 2) a single inaccurate informant alone, or 3) a single accurate informant alone (as a control condition). Following previous selective trust studies (Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Sabbagh & Baldwin, 2001), children in the paired accurate-inaccurate informants condition observed an informant label familiar objects correctly and the other informant label familiar objects incorrectly. Children in the single inaccurate informant condition observed only an informant label familiar objects incorrectly, and children in the single accurate informant condition observed only an informant label familiar objects correctly. Then, importantly, children in all three conditions were tested with the same question with regard to their trust of the informants' labels for novel objects.

Method

Participants

Eighty-six 3- and 4-year-olds participated in Experiment 1; there were forty-three 3-year-olds (27 females and 16 males; M age = 3.6 years, age range: 3.0 to 3.9 years) and forty-three 4-year-olds (20 females and 23 males; M age = 4.4 years, age range: 4.0 to 5.0

years). Children were recruited from preschools and museums in southern California. The sample was approximately 68% Caucasian, 14% Asian, and 18% Hispanic.

Materials and Procedure

In the familiar objects history phase, children were introduced to two puppets or a single puppet, depending on the condition, and two pairs of familiar objects: a car with an apple and a ball with a crayon. Table 1 shows the types of speakers that were compared in each condition. In the *paired accurate-inaccurate informants* condition (N=29), children were introduced to a boy and a girl puppet; children observed one puppet demonstrate accuracy by correctly calling a car a “car” and a ball a “ball” and the other puppet demonstrated inaccuracy by incorrectly calling an apple a “car” and a crayon a “ball.” The gender of the accurate and inaccurate informants was counter-balanced between subjects. In the *single inaccurate informant* condition (N=26), children were introduced to either a boy or a girl puppet (counter-balanced between subjects); children observed the one puppet demonstrate inaccuracy by incorrectly calling an apple a “car” and a crayon a “ball.” To control for equal attention for the other object in each pair, the puppet also touched the car and the ball while saying, “There it is! Look at that thing.” In the *single accurate informant* condition (N=31), children were introduced to either a boy or a girl puppet (counter-balanced between subjects); children observed the one puppet demonstrate accuracy by calling a car a “car” and a ball a “ball.” To control for equal attention for the other object in each pair, the puppet also touched the apple and the crayon while saying, “There it is! Look at that thing.”

In the novel objects test phase of the task, children were presented with three test trials; on each trial, children were introduced to a pair of novel objects. For each test trial, in the paired condition, the accurate puppet referred to one of the two objects with a novel label (by touching it and saying, e.g., “That is a blicket! Look at the blicket.”), whereas the inaccurate puppet referred to the other object with the same novel label (by touching the other object and saying, e.g., “That is a blicket! Look at the blicket.”). In both the single inaccurate and the single accurate conditions, the puppet called one of the two objects a novel label (by touching it and saying, e.g., “That is a blicket! Look at the blicket.”) and, to control for equal attention for the other object in each pair, the puppet also touched the unlabeled object while saying, “There it is! Look at that thing.”

Next, for each test trial, and for all three conditions, children were asked to select the object associated with the novel label (e.g., “Can you put the blicket in this box?”). *Importantly*, children in all three conditions were asked the same question in a two-object situation. Across three test trials, their responses indicated whether they trusted the informants’ labels for novel objects in order to allow for clear comparisons between conditions. Regardless of what they chose, children received no feedback from the experimenter about their choice in any of the conditions.

At the end, we asked children a follow-up question about each informant: “If you wanted to know what this new thing was called, would he/she be a good person to ask?” Children in the paired condition were asked this question for both the accurate and the inaccurate informants separately. Children in the single inaccurate and the single accurate conditions were asked only for informant in their respective condition.

Results

All children completed three test trials, and each child's trust of a particular informant's labels was scored from zero to three. Preliminary analyses indicated no significant effects of participant age or gender; analyses also indicated no significant effect of puppet gender. Consequently, these variables were excluded from further analyses. Additionally we did not find any changes in children performance across trials. (Experiments 2 and 3 also did not show any effects of age, gender, puppet gender, or trial, and thus these variables were also excluded from the reported analyses of Experiments 2 and 3.)

In a replication of prior research (e.g. Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005), in the paired condition, children trusted the accurate informant ($M = 2.38$; 79% of the trials overall) over the inaccurate informant ($M = 0.62$; 21% of the trials overall). Of central theoretical interest is whether children distrust inaccurate informants in an absolute sense, such that they systematically reject the testimony of inaccurate informants. In order to address this issue, we determined that children's level of trust in the single inaccurate condition ($M = 2.46$; 82% of the trials overall) was significantly greater than would be expected by chance performance, $t(25) = 5.17, p < .001$. Additionally, we compared the two conditions (paired and single inaccurate) with a previously-inaccurate informant. We found that children's trust of the inaccurate informant was significantly greater in the single inaccurate condition than in the paired condition, $t(53) = 7.37, p < .001$. Finally, we compared children's level of trust in the single inaccurate condition with the single accurate condition ($M = 2.61$; 87% of the trials

overall) and found no difference, $t(55) = 0.67, ns$. These results showed that children trusted the inaccurate informant when he or she was not paired with an opposing (accurate) informant, and they did so at the same level as their trust of the accurate informant.

On the follow-up question, children overwhelmingly judged that both accurate and inaccurate speakers would be good people to ask about the label of a novel object. Specifically, 85% of the children in the single inaccurate condition and 84% of the children in the single accurate condition reported that the informant in question would be a good person to ask. This same general pattern was also seen in the paired condition where 86% of the children reported that the accurate informant would be a good person to ask compared to 79% of the children who reported that the inaccurate informant would be a good person to ask, McNemar's $\chi^2(1) = 0.73, ns$. In addition, children's explicit judgment of the inaccurate informant did not differ between the paired and single inaccurate conditions, McNemar's $\chi^2(1) = 0.26, ns$.

Overall, the results suggest that although children trusted the accurate informant over the inaccurate informant when asked to choose between them, they nevertheless considered the inaccurate informant to still be generally trustworthy. This was evident both in their willingness to accept information provided by the single inaccurate informant and by their explicit judgments that an inaccurate informant would be as good of a person to ask about a novel object label as an accurate informant.

Discussion

In Experiment 1 children accepted the novel labels of the inaccurate informant more often when he or she was the only informant than when he or she was paired with an accurate informant. Furthermore, children accepted the novel labels of the single inaccurate informant at a level greater than chance. These findings suggest that children's specific inference about the inaccurate informant is that someone who got a few things wrong is less reliable than an accurate informant, but still generally reliable. That is, it appears children lower their level of trust in inaccurate informants only slightly. As such, our results suggest that children will tend to trust testimony from a single inaccurate informant alone long after they have acquired the ability to encode and track an informant's previous inaccuracies (Corriveau et al., 2008; Koenig & Echols, 2003; Koenig & Woodward, 2010; Pea, 1982). Indeed, other studies have also found that preschool-age children will continue to trust single inaccurate informants who offer advice (Couillard & Woodward, 1999; Jaswal, Croft, Setia, & Cole, 2010; Vanderbilt, Liu & Heyman, 2011), and Krogh-Jespersen and Echols (2012) found that 24-month-olds also accepted labels from single inaccurate and ignorant speakers for novel objects. This decision to accept information from a single inaccurate informant stands in contrast to children's rejection of information from the inaccurate informant in our paired condition, as well as their behavior in previous selective trust studies that have provided participants with contrasting informants on a within-subjects basis (Birch, et al., 2008; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Pasquini, et al., 2007).

One interpretation of our results is that children are generally quite willing to accept information from inaccurate speakers, unless their testimony conflicts with that of

accurate (and perhaps neutral) speakers. Specifically, this account explains the difference between reasoning about inaccurate speakers in the single versus paired conditions during the test phase of the task, when two speakers offer conflicting advice to the child.

However, an alternative possibility is that children need to observe the contrast between accurate and inaccurate speakers during the history phase (which occurred in the paired but not the single inaccurate condition) to highlight the inaccurate speaker's deficits.

Perhaps children simply need a working comparison of a speaker's past accuracy relative to others to inform trust decisions, rather than conflicting offers of testimony during the test phase. If this second possibility is correct, we would expect to see children reject testimony from a single speaker with a history of inaccuracy, as long as they had also seen another speaker label the same objects correctly during the history phase.

Experiment 2 tested this possibility by presenting children with the contrast of accurate and inaccurate informants during the history phase, but then testing children on only one of the informants.

Experiment 2

In a between-subjects design, we presented 3- and 4-year-olds with paired accurate and inaccurate informants during the history phase, but with only one of those informants during the test phase. In one condition, the informant during the test phase was the one who had been accurate, and in the other condition, the informant was the one who had been inaccurate.

Method

Participants

Forty-four 3- and 4-year-olds participated in Experiment 2; there were twenty-two 3-year-olds (7 females and 15 males; M age = 3.5 years, age range: 3.0 to 3.9 years) and twenty-two 4-year-olds (10 females and 12 males; M age = 4.5 years, age range: 4.0 to 4.9 years). Children were recruited from preschools and museums in southern California. The sample was approximately 77% Caucasian, 9% Asian, and 14% African-American.

Materials and Procedure

The materials for Experiment 2 were identical to Experiment 1. The familiar objects history phase of the paired condition in Experiment 1 was combined with the novel objects test phase of either the single inaccurate condition in Experiment 1 (for the Experiment 2 *paired-single inaccurate* condition) or the single accurate condition in Experiment 1 (for the Experiment 2 *paired-single accurate* condition); see Table 1. Lastly, as in Experiment 1, children were asked an explicit follow-up question about the informant.

Results

Results showed that children systematically accepted the testimony of the inaccurate informant. Children trusted the inaccurate speaker at a rate ($M = 2.76$; 92% of the trials overall) that was significantly greater than chance performance, $t(20) = 10.73$, $p < .001$. Additionally, their trust of the informant in the paired-single inaccurate condition was not significantly different than in the paired-single accurate condition ($M = 2.35$; 78% of the trials overall), $t(42) = 1.55$, ns . As in Experiment 1, on the follow-up question, children in both conditions reported that the informant would be a good person to ask for labels of novel objects; 76% of children in the paired-single inaccurate

condition and 78% of children in the paired-single accurate condition judged that the informant would be a good person to ask.

A goal of Experiment 2 was to test the possibility that having the contrast between an accurate and an inaccurate informant highlights the inaccurate informant's inaccuracies compared to only observing the inaccurate informant alone during the history phase. Thus, we compared children's trust of the single inaccurate informant in the paired-single inaccurate condition of Experiment 2 and the single inaccurate condition of Experiment 1. Children's trust of a single inaccurate informant did not differ between the conditions, $t(45) = 1.29$, *ns*. It appears that having the contrast of an accurate informant during the history phase did not impact children's inference about the reliability of the inaccurate informant. Thus, offered conflicting testimony during the test phase may carry more weight in trust decisions.

Discussion

The results of Experiment 2 showed that the 3- and 4-year-olds accepted the novel labels of the single inaccurate informant as often as the novel labels of the single accurate informant. This was the case even though they observed the contrast between the accurate and the inaccurate informants during the history phase. Also, children accepted the novel labels of the single inaccurate informant as often in the paired-single inaccurate condition of Experiment 2 as in the single inaccurate condition of Experiment 1. That is, children trusted the single inaccurate informant regardless of whether there was an opposing accurate informant during the history phase to highlight the inaccuracies. Lending further support for these findings are children's explicit judgments on the follow-up questions.

Across Experiments 1 and 2, children judged the inaccurate informant as someone who would be good to ask for information regardless of whether there was an opposing accurate informant during the history or the test phase.

Experiment 3a

In Experiments 1 and 2, we demonstrated that young children readily accepted information from an informant with a history of inaccuracy unless they were presented with conflicting information from an accurate informant during test trials. Experiment 2 showed that having the contrast of an accurate informant during the history phase did not impact children's inference about the reliability of the inaccurate informant. Thus, the presence of another informant during the *test phase* appears to play a critical role in trust decisions. In Experiment 3, we examined children's trust decisions about an inaccurate informant's testimony when the alternative informant during the test phase is a novel, neutral informant without any performance history. We examined two ways in which the neutral informant is paired with the inaccurate informant (see Table 1). In Experiment 3a, the neutral informant offers testimony that opposes the inaccurate informant during the test phase. In Experiment 3b, the neutral informant was available for children to seek information from during the test phase, but did not proactively offer testimony.

Of interest in Experiment 3a was whether children would distrust the testimony of an inaccurate informant when conflicting testimony is from a neutral informant with unknown accuracy history. Corriveau et al. (2008) provide suggestive evidence that this is the case. This study found that children trusted testimony from a neutral source over testimony from an inaccurate one in cases where the neutral source provided only

uninformative statements (e.g., “let me look at that”) during their history phase. Of interest in the present experiment is whether their finding would replicate under the procedures of the present research when no information from the neutral source was provided during the history phase. Thus, the only information children received about the alternative source in our experiment was that they provided alternative testimony to the inaccurate informant.

Method

Participants

Twenty 3- and 4-year-olds participated in Experiment 3a; there were nine 3-year-olds (3 females and 6 males; M age = 3.63 years, age range: 3.13 to 3.93 years) and eleven 4-year-olds (4 females and 7 males; M age = 4.52 years, age range: 4.02 to 5.05 years). Children were recruited from preschools and museums in southern California. The sample was approximately 65% Caucasian, 20% Asian, and 15% Hispanic.

Materials and Procedure

The materials for Experiment 3a were identical to Experiment 1. The familiar objects history phase was identical to that phase of the single inaccurate condition in Experiment 1. During the novel objects test phase, a second (neutral) puppet was presented. In each of three test trials, as in Experiment 1, the inaccurate puppet called one of two novel objects a novel label, the novel puppet then gave the same label for the other object, and children were asked to endorse one of the labels provided, by selecting the correct object. Children’s responses were coded for whether they accepted the novel

labels of the inaccurate informant. Critically, children answered the same test question as in the previous 2 experiments.

Results and Discussion

Consistent with the results of Corriveau et al., (2008), children trusted the novel informant at above chance levels $t(19) = 2.15, p = 0.045$ ($M = 2.05$; 68% of the trials overall). These results indicate that when another source offers conflicting testimony alongside an inaccurate source, children will rely on the alternate source, even when that source has no prior history of accuracy or inaccuracy. In line with the results from Experiment 1, these results also suggest that young children are adept at relative trust judgments, and will disregard the testimony of an inaccurate source when alternative testimony is provided.

On the follow-up question, children reported that both the inaccurate and the neutral informant would be good people to ask for labels of novel objects; 80% of children reported that the inaccurate source would be a good person to ask, and 90% of children reported that the neutral informant would be a good person to ask. This suggests that children considered both informants to be generally reliable, despite the lack of evidence to confirm that assumption, or (in the case of the inaccurate informant) direct evidence to the contrary.

Experiment 3b

Experiment 3a showed that as long as children were given alternative testimony to that of the inaccurate informant, they showed appropriate selective trust. Of interest in Experiment 3b was whether the neutral informant in question had to proactively provide

alternative testimony, in order to cue children to distrust the inaccurate informants' testimony, or whether the mere availability of an alternative source was sufficient. We addressed this question by providing children with only the testimony of an inaccurate informant, but in this case we made it clear that there was another source available for additional testimony if children sought that out.

Method

Participants

Thirty-two 3- and 4-year-olds participated in Experiment 3b; there were eighteen 3-year-olds (6 females and 12 males; *M* age = 3.60 years, age range: 3.08 to 3.98 years) and fourteen 4-year-olds (7 females and 7 males; *M* age = 4.58 years, age range: 4.07 to 5.02 years). Children were recruited from preschools and museums in southern California. The sample was approximately 71.5% Caucasian, 12.5% Asian, and 15% Hispanic.

Materials and Procedure

The materials for Experiment 3b were identical to Experiment 1, and the procedure was the same as in Experiment 3a, except that the novel source remained silent during test trials and did not offer information unless asked by the child. All children were informed that the new puppet was there to answer any questions they had during the game, and all children confirmed verbally that they understood this. Children's responses were coded for whether they accepted the novel labels of the inaccurate informant, and for whether or not they chose to solicit advice from the novel source. Again, the test question was identical to that from Experiments 1, 2 and 3a.

Results and Discussion

Children in this condition trusted the inaccurate informant the majority of the time ($M = 2.03$; 68% of the trials overall) at a level above that expected by chance $t(31) = 2.83, p = .008$. Because we wanted to understand the effect of a potential alternate neutral source during the test phase versus an alternate neutral source who proactively offers conflicting testimony, we compared children's trust of the inaccurate informant in Experiment 3b with their trust in the inaccurate source in Experiment 3a. Children trusted the inaccurate informant in Experiment 3b (68% of the trials overall) more than the inaccurate informant in Experiment 3a (32% of the trials overall), $t(50) = 3.47, p = .001$. These results suggest that proactively offered conflicting testimony cues children's selective trust, whereas simply having an alternative source present may not. Further, *no* children in this experiment took the opportunity to ask the neutral source for advice about which object to choose. Instead, *all* children relied solely on the testimony provided by the inaccurate informant to make their decision, by either accepting or going against the inaccurate informant's testimony. These results indicate that children did not feel the need to ask another source for advice after receiving the inaccurate informant's testimony¹.

The results of Experiments 3a and 3b suggest that children do not accept testimony from an inaccurate informant simply because that informant is the only source available. Instead, it appears that children are happy to accept an inaccurate informant's testimony unless that source is actively opposed by conflicting testimony. This provides

further evidence that children do not infer an inaccurate informant is untrustworthy in an absolute sense.

General Discussion

The present research examines young children's inferences about informants who provide testimony, with a focus on informants who have shown a history of inaccuracies. In a series of experiments, the present research is the first to systematically examine the nature of the inferences and decisions young children make about informants with a prior history of inaccuracies. We observed whether children trusted the testimony of speakers who had been accurate or inaccurate in the past. In Experiment 1, children selectively distrusted inaccurate informants when accurate informants provided conflicting information, but they trusted inaccurate informants who were presented alone. In Experiment 2, children who had previously observed the contrast between inaccurate and accurate informants, again systematically trusted inaccurate informants who offered testimony alone. Finally, in Experiment 3, children relied on testimony from a neutral speaker over testimony from an inaccurate speaker when both provided a label, but children accepted testimony from an inaccurate speaker rather than seeking information from an available source that did not automatically offer a label. Taken together, the present results demonstrate that young children's inferences about inaccurate informants are of their relative reliability, not their *absolute* reliability. Further, these results suggest that preschool children may exercise selective trust only in situations where they must choose between at least two instances of offered testimony. This possibly is also consistent with findings of Krogh-Jespersen and Echols (2012), showing that although

24-month-olds were able to reject labels from single inaccurate and ignorant speakers for familiar objects, they were not able to do so for novel objects. Thus, children appear to exercise selective trust when testimony from an inaccurate informant conflicts with either another informant's offered testimony or their own knowledge.

Explicit judgment data also supports our claim that young children generally trust inaccurate informants. In all Experiments in which children were asked more than three quarters of children responded that the inaccurate source would be a good person to ask for novel object labels (79% and 85% in Experiment 1, 76% in Experiment 2, and 80% in Experiment 3a), and these responses were not significantly different than children's rates of response to accurate informants. This was the case even in the paired condition of Experiment 1, when children were asked about both an accurate and an inaccurate speaker. In this case, 79% of the children reported that the inaccurate informant would be a good person to ask, which was not significantly different than the 86% of the children who reported that the accurate informant would be a good person to ask. These explicit judgments likely differ from what have been observed in prior research on selective trust of paired informants (e.g., Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005) because we asked children to rate the absolute reliability of each individual informant, rather than to make a relative comparison between the two informants.

Our results suggest that young children may not assign a large penalty in response to past inaccuracies and instead lower their level of trust only slightly. Therefore, although children encode and track inaccuracies (Corriveau et al., 2008) and trust inaccurate informants relatively less than accurate informants (Birch, Vauthier, & Bloom,

2008; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Pasquini, et al., 2007), they are generally quite willing to accept testimony from inaccurate informants when that is the only testimony offered to them.

Why might children fail to appropriately question what they learn from previously-inaccurate individuals when conflicting testimony is not explicitly available to them? One possibility is that children identify the source as inaccurate, but trust their testimony nonetheless because that source's testimony is all that is available. The results of Experiment 3, and children's responses to explicit judgment questions indicate that this is not the case. Children in Experiment 3 accepted testimony from an inaccurate speaker even when another source was present and waiting to answer their questions. This suggests that children consider the inaccurate source to be a completely acceptable source of information, and do not feel the need to verify information they provide. Further, children's responses to explicit judgment questions indicate that they do not consider sources with a history of inaccuracy to be unreliable sources. Thus, it is not the case that children are accepting testimony from these sources against their better judgment. They simply believe these sources are generally reliable, despite their past inaccuracies.

Another possibility for why children fail to reject testimony from speakers with a history of inaccuracy is that their judgment is part of a broader tendency to view people in a positive light (Boseovski, 2010; Boseovski & Lee, 2008; Heyman & Giles, 2004; Lockhart, Chang, and Story, 2002). For example, children require less behavioral evidence to make positive attributions than negative attributions (Boseovski & Lee,

2006). It may be that children are reluctant to view people in negative terms without very strong evidence. A related possibility is that children have learned that people's testimony is usually accurate, and so their default assumption is to accept what they hear in the absence of explicit alternatives.

We believe it is especially important to understand how children evaluate testimony from individual informants, because this is arguably a more common situation than having to resolve simultaneously conflicting testimony between two or more informants. Although young children may have to resolve discrepancies in testimony provided at different times, they probably hear most inaccurate testimony without the concurrent benefit of explicit, conflicting testimony. Of course, as in Experiment 3, children are likely to have opportunities to seek out additional information when an informant has a history of inaccuracy, but our results suggest that they often fail to do so, even in cases where they are specifically reminded that such an option is available. This suggests that in everyday interactions children receiving testimony from inaccurate informants would be even less likely to seek out additional information, because they would have to generate the option of seeking out additional information. Indeed, our results are consistent with other research indicating that children often rely on single unreliable sources when information is presented in a familiar and conventional way (Couillard & Woodward, 1999; Jaswal et al., 2010; Krogh-Jespersen & Echols, 2012).

Furthermore, in everyday interactions, the actual process of seeking out information can involve multiple cognitive demands that are not required in the present research, such as the ability to formulate appropriate questions (Mills, Legare, Bills, &

Mejias, 2010). This suggests that young children may be even less likely to seek out alternatives to inaccurate informants in everyday situations.

Future research is needed to examine whether circumstances that can help explain how an informant's prior inaccuracies, might affect the inferences young children make about inaccurate informants. One crucial question is how many instances of inaccurate information informants need to provide for children to judge them as being inaccurate. Research by Cameron and Lee (2000) suggests that children require several instances of consistent unreliability before they will reject information provided by a speaker, but further work is needed to address this question directly. Additionally, further work is needed to understand how children's knowledge about the reason for inaccuracy might affect the inferences they draw (Nurmsoo & Robinson, 2009). In the critical condition of Nurmsoo and Robinson (2009), 3- to 5-year-olds discounted the past inaccuracies of an informant because the informant had the wrong type of perceptual access to hidden objects (e.g., providing color information after touching, but not seeing the objects). Interestingly, in the other condition of this study, children were asked to reason about an inaccurate informant (who provided inaccuracies but had the correct type of perceptual access) alone, and children did not systematically trust the single inaccurate informant. This finding appears to be inconsistent with the findings of the present research. However, we believe this apparent inconsistency can be explained by the fact that in the test trials presented in Nurmsoo and Robinson (2009), children were asked to guess the feature (e.g., the color) of each object before the inaccurate informant provided testimony that always contradicted the children's guess. That is, the children were instructed to

overtly generate, ahead of time, an alternative to the testimony of the inaccurate informant. This interpretation is consistent with our suggestion that children require explicit alternatives to dissuade them from trusting the inaccurate informant. It is also consistent with Krogh-Jespersen and Echols (2012), who found that children accepted labels from inaccurate and ignorant speakers for novel objects but not for familiar objects they already had labels for (and thus had alternatives in their own knowledge). Additional studies are needed to examine what other forms such explicit alternatives can take.

In the present research, we asked children to reason about informants with reference to learning about object labels. We investigated testimony about object labels because a substantial amount is known about children's selective trust in this domain (Harris & Corriveau, 2011; Koenig & Harris, 2005), and our use of object labels facilitates comparison with this body of work. However, further research is needed to extend this work to other domains, such as learning how things work. In addition, more work is needed to understand the precise circumstances that affect children's willingness to seek out information from others.

In conclusion, our results suggest that young children do assign some consequence to informant's inaccuracies, lowering their relative level of trust in these individuals. However, at least when there is only a single informant present to provide information, children do not appear to consider previously inaccurate sources to be unreliable in an absolute sense. It seems that children are generally trusting of a single informant alone, regardless of the informant's prior history of reliability.

References

- Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*, 1018-1034.
- Boseovski, J. J. (2010). Evidence of "rose-colored glasses": An examination of the positivity bias in young children's personality judgments. *Child Development Perspectives*, *4*, 212-218.
- Boseovski, J. J., & Lee, K. (2008). Seeing the World Through Rose-Colored Glasses? Neglect of Consensus Information in Young Children's Personality Judgments. *Social Development*, *17*, 399-415.
- Boseovski, J. J., & Lee, K. (2006). Preschoolers' use of frequency information for trait categorization and behavioral prediction. *Developmental Psychology*, *42*, 500-513.
- Lee, K., & Cameron, C.A. (2000). Extracting truthful information from lies: Emergence of the expression-representation distinction. *Merrill-Palmer Quarterly*, *46*, 1-20.
- Corriveau, K.H., Meints, K., & Harris, P.L. (2009). Early tracking of informant accuracy and inaccuracy by young children. *British Journal of Developmental Psychology*, *27*, 331-342.
- Couillard, N.L., & Woodward, A.L. (1999). Children's comprehension of deceptive points. *British Journal of Developmental Psychology*, *17*, 515-521.
- Gelman, S.A. (2009). Learning from others: Children's construction of concepts. *Annual Review of Psychology*, *60*, 115-140.
- Harris, P.L. & Corriveau, K.H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*, 1179-1187.
- Heyman, G.D., & Giles, J.W. (2004). Valence effects in reasoning about evaluative traits. *Merrill-Palmer Quarterly*, *50*, 86-109.
- Jaswal, V. K., Croft, A. C., Seftia, A. R., & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science*, *21*, 1541-1547.
- Koenig, M., Clement, F., & Harris, P. (2004). Children's Use of True and False Statements. *Psychological Science*, *15*, 694-698.

- Koenig, M. A., & Echols, C. H. (2003). Infants' understanding of false labeling events: The referential roles of words and the speakers who use them. *Cognition*, *87*, 179–208.
- Koenig, M., & Harris, P. (2005). Preschoolers Mistrust Ignorant and Inaccurate Speakers. *Child Development*, *76*, 1261-1277.
- Koenig, M. A., & Woodward, A. L., (2010). 24-month-olds' sensitivity to the prior inaccuracy of the source: Possible mechanisms. *Developmental Psychology*, *46*, 815–826.
- Krogh-Jespersen, S., & Echols, C. E. (2012). The influence of speaker reliability on first versus second label learning. *Child Development*, *83*, 581-590.
- Lockhart, K. L., Chang, B., & Story, T. (2002). Young children's beliefs about the stability of traits: Protective optimism? *Child Development*, *73*, 1408-1430.
- Mills, C. M., Legare, C. H., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development*, *11*, 533-560.
- 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.
- Pea, R. (1982). Origins of Verbal Logic: Spontaneous Denials by Two-and Three-Year-Olds. *Journal of Child Language*, *9*, 597-626.
- Pasquini, E., Corriveau, K., Koenig, M., & Harris, P. (2007). Preschoolers Monitor the Relative Accuracy of Informants. *Developmental Psychology*, *43*, 1216-1226.
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development*, *72*, 1054–1070.
- Vanderbilt, K. E., Liu, D., & Heyman, G. D. (2011). The development of distrust. *Child Development*, *82*, 1372–1380.

Footnote

¹We collected data for a control condition ($N = 72$) in which the test question was changed to explicitly remind children of the option to ask the neutral informant. In this version, 45% of children chose the object labeled by the inaccurate informant, 50% chose to ask the neutral informant, and 5% chose the opposite object of the one labeled by the inaccurate informant. These results show that even when explicitly reminded of the option to ask another informant, children still didn't systematically reject the testimony of the inaccurate informant.

Table 2.1: Summary of Experimental Design: Number and Type of Informant

Group	History Phase Informants		Test Phase Informants	
	Number	Type	Number	Type
Experiment 1				
Paired accurate-inaccurate	2	A and I	2	A and I
Single inaccurate	1	I	1	I
Single accurate	1	A	1	A
Experiment 2				
Paired single inaccurate	2	A and I	1	I
Paired single accurate	2	A and I	1	A
Experiment 3a	1	I	2	I and N
Experiment 3b	1	I	2	I and N ^a

Note. A = accurate informant; I = inaccurate informant; N = neutral informant. All informants actively provided testimony, except as specifically noted for the neutral informant in Experiment 3b.

^aNeutral informant did not actively provide testimony; children were informed at the start of the test phase that the neutral informant is available for them to ask about the test objects.

Chapter 2, in full, has been submitted for publication of the material as it may appear in
In the absence of conflicting testimony, young children trust inaccurate informants in
Developmental Science. Vanderbilt, K.E., Heyman, G. D., & Liu, D.

Chapter 3

Young Children's Judgment of Informant Reliability Based on Knowledge and Intentions

Kimberly E. Vanderbilt, Gail D. Heyman, and David Liu

University of California, San Diego

Abstract

The present research investigated whether children's inferences about unreliable informants are influenced by the mental state causing unreliability (either bad knowledge or bad intentions) and by the strength of the cues children receive about those mental states. Children ($N = 170$) played a searching game with a puppet who gave them inaccurate advice as a consequence of having bad knowledge or bad intentions. Results showed that when presented with minimal mental-state cues, children had greater initial distrust toward an unreliable source with bad knowledge than an unreliable source with bad intentions. When presented with stronger cues to unreliability, children showed equal initial distrust toward sources with bad knowledge and bad intentions. These findings suggest that children may have more difficulty reasoning about unreliability resulting from bad intentions. Although children showed effects of mental-state type and cue strength on initial trust, the effects went away after receiving feedback on their performance, indicating that once children understand a source *is* unreliable they may no longer pay attention the *reason* a source is unreliable.

Young Children's Judgment of Informant Reliability Based on Knowledge and Intentions

Children often rely on other people to teach them about the world around them, acquiring a great deal of their knowledge through the testimony and advice of others. A consequence of this, however, is that children need to develop ways to judge whether someone is trustworthy as an informant. As such, it is important to understand the development of how children evaluate the reliability of sources they encounter. Sources might provide inaccurate information due to having poor knowledge or bad intentions, and thus, children need to consider these two critical mental states when evaluating the reliability of sources. Of primary interest in the present study is to compare the inferences children draw about individuals who provide incorrect information due to each of these two mental states: poor knowledge or bad intentions.

Previous research has documented that preschool-age children selectively trust speakers with a history of accurate performance over those with a history of inaccurate performance (see Harris & Corriveau, 2011 for review). This pattern has been observed with reference to a range of content domains, including learning novel object names (Clement, Koenig, & Harris, 2004; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig, & Harris, 2007; Sabbagh & Baldwin, 2001), novel object functions (Birch, et al., 2008; DiYanni & Keleman, 2008; Jaswal, 2006; Koenig & Harris, 2005), the location of hidden objects (Robinson & Nurmsoo, 2009; Nurmsoo & Robinson, 2009a), and the contents of opaque containers (Robinson, Champion, and Mitchell, 1999). An important question

raised by these findings of young children's selective trust is how this type of reasoning is associated with their developing understanding of mental states.

It is possible that young children's selective trust has nothing to do with reasoning about informants' mental states. When inaccurate performance is observed of potential sources, there is often no clear specific reasons for why they provided inaccurate information. Moreover, it is not even clear that children would always consider mental states when they selectively trust one source over another; just as one might distrust a computer that provides inaccurate information without inferring or caring why it is unreliable, it is possible to distrust a person who provides inaccurate information without inferring or caring why he or she is unreliable (Heyman, 2008). Indeed, Liu, Vanderbilt, and Heyman (2013) found that young children strongly weigh the external outcomes of informants' past performance, regardless of the internal mental states (e.g., children trusted a malevolent informant with past accurate performance more than a benevolent informant with past inaccurate performance). Nevertheless, Liu, Vanderbilt, and Heyman (2013) also found that when the external past performance was equated young children based their trust on informants' mental states (e.g., children trusted a benevolent informant more than a malevolent informant when both showed past inaccurate performance). This suggests that young children might not always or solely consider mental states in their selective trust, but they do so at least sometimes or in part.

A small number of studies have examined the link between selective trust and mental-state understanding by correlating measures of selective trust with measures of theory of mind. Some studies taking this approach have documented significant positive

correlations (DiYanni & Keleman, 2008, DiYanni, Nini, Rheel, & Livelli, 2012), but this pattern has not been observed in all studies (Pasquini, Corriveau, Koenig, & Harris, 2007). Other studies have addressed this question by examining whether children's trust judgments are influenced by experimental manipulation of mental-state cues. Nurmsoo and Robinson (2009) found that 3- to 5-year-olds showed some tendency to excuse bad information from sources when the bad information was provided in a context that could explain the inaccuracy, such as when an individual who was not able to name an object's color when blindfolded. This suggests that children discounted inaccurate performance when it could be attributed to a lack of knowledge access. Einav and Robinson (2011) found that 4- and 5-year-olds placed greater trust in sources who had previously given correct answers without getting help than they did in sources who had previously given correct answers who had consistently relied on getting help. These studies and others (Jaswal & Malone, 2007; Lutz & Keil, 2002; Mascaro & Sperber, 2009; Sabbagh & Baldwin, 2001; Vanderbilt, Liu, & Heyman, 2011) suggest that children make some use of cues about informants' knowledge or informants' intentions when judging whether to trust them. However, because no study has directly compared the two mental states, it is unknown whether children make use of them to evaluate sources on the same developmental time-table.

Beliefs and desires are the two fundamental mental-state categories in people's folk understanding of mental states and in children's theory of mind development (Wellman, 1990). Research in theory of mind suggests that children develop an explicit understanding of desires/intentions before an explicit understanding of beliefs/knowledge

(Wellman & Liu, 2004). Wellman and Liu (2004) found that young children develop an understanding of diverse desires (that different people can want different things) before developing an understanding of diverse beliefs (that different people can believe different things); they also develop an understanding of diverse desires before an understanding of knowledge access (that one requires perceptual access to have knowledge about an object). Even within the hotly debated topic of infant theory of mind, studies using looking-time measures suggest that infants demonstrate implicit understanding of intentions and goal-directed actions (Meltzoff, 2006; Woodward, 2009) before they demonstrate implicit understanding of knowledge and beliefs (Baillargeon, Scott, & He, 2010). There is also evidence that reasoning about deceptive intentions enhances children's understanding of beliefs (Chandler, Fritz, & Hala, 1991; Wellman, Cross, & Watson, 2001). Also consistent with the possibility that children may develop to trust selectively based on a speaker's bad intentions before developing to trust selectively based on a speaker's bad knowledge, is evidence that when children are provided with information about both a source's intentions and knowledge, information about bad intentions often overrides information about knowledge (Landrum, Mills & Johnston, 2013). This suggests that cues about intentions may be more salient to young children in their reasoning about informant reliability.

Selective trust research examining children's reasoning about bad intentions or about bad knowledge in isolation, however, points to the opposite prediction—that children may be more sophisticated at reasoning about inaccuracies based on knowledge rather than intentions. This prediction is based on findings that children as young as age 3

or 4 years show substantial sophistication in making distinctions based on a source's knowledge (Jaswal & Malone, 2007; Lutz & Keil, 2002; Sabbagh & Baldwin, 2001), but that children of the same age (Mascaro & Sperber, 2009; Vanderbilt, Liu & Heyman, 2011) or even older (Heyman & Legare, 2005) show evidence of having difficulty reasoning about the implications of having bad intentions.

The apparent differences across studies in selective trust, however, might be due to methodological differences between studies. Previous studies have varied along a range of dimensions that could affect children's inferences: the number of sources who provide testimony to children, the type of testimony sources provide, the way in which testimony is provided by sources, the way in which children are asked to respond, and the way in which children receive information about sources. The possibility that differences in reasoning about knowledge versus intentions is an artifact of differing methodologies across studies is also supported by evidence documenting that study design can sometimes dramatically influence results. For example, Vanderbilt, Heyman, and Liu (2013) demonstrated that children selectively relied on testimony from an accurate speaker over an inaccurate speaker, but children overwhelmingly relied on testimony from an inaccurate speaker when that speaker was the only one to offer testimony. This points to a need for a direct comparison between how children reason about knowledge versus intentions on the same selective trust task. In the present research we made such direct comparisons.

A secondary goal of the present research was to examine whether the strength of reliability cues provided to young children would influence their level of trust. Research

has suggested that children are hesitant to draw negative conclusions about others in the absence of very salient cues (Boseovski, 2010) and that providing children with stronger evidence of the unreliability of sources may help them to overcome this tendency (Heyman, Sritanyaratana, & Vanderbilt, 2013). Consistent with this possibility is evidence that young children showed substantially lower rates of distrust of deceptive sources when cues were weak (Vanderbilt et al., 2011) than in a study when cues were strong (Heyman et al., 2013). Specifically, in Vanderbilt et al. (2011), the only reliability cue was the source's deceptive behavior, but in Heyman et al. (2013), multiple cues were provided: the source was a well known bad character (i.e., the big bad wolf) and participants were explicitly told and repeatedly reminded of his bad motives. However, as is the case with attempting to draw conclusions by comparing intentions and knowledge across studies, it is difficult to draw conclusions about the effect of the strength of reliability cues without a direct comparison within the same study. As such, the present research also systematically varied the amount of information provided to children about sources. In low cue conditions, children were provided with only minimal cues to the unreliability of the speakers, whereas in high cue conditions, children were provided with stronger cues to the reliability of the speaker as well as its situational consequences.

Method

Participants

One hundred seventy 3- and 4-year-olds participated in this experiment; there were eighty-three 3-year-olds (46 females and 37 males; *M* age = 3.5 years, age range: 2.8 to 3.9 years) and eighty-seven 4-year-olds (44 females and 43 males; *M* age = 4.5

years, age range: 4.0 to 5.2 years). The four conditions (*low cue-bad intention*, *low cue-bad knowledge*, *high cue-bad intention*, and *high cue-bad knowledge*) were assigned randomly between subjects; mean age did not differ significantly between conditions (M age = 4.1 years, 4.0 years, 4.1 years, and 4.0 years, respectively). Children were recruited from preschools and museums in southern California. The sample was approximately 76% Caucasian, 16% Asian, 3% African American, and 5% Hispanic.

Materials and Procedure

The procedure was adapted from Heyman, Sritanyaratana, and Vanderbilt (2013). Children were seated at a table across from the experimenter and familiarized with a sticker search task. Two boxes, one yellow and one green, were placed in front of the child. Children were told that a sticker would be hidden in one of the two boxes and if they picked the correct box, they could keep the sticker. Next, children were introduced to a monkey puppet (operated by a second experimenter), who would be playing the game with them. The puppet was introduced with descriptions (detailed below) in accordance with one of four conditions (*low cue-bad intention*, *low cue-bad knowledge*, *high cue-bad intention*, or *high cue-bad knowledge*). Children then received advice from the puppet on each of six trials of the sticker search task. On each trial, the puppet touched one of the boxes and said “The sticker’s in the [green] box.” The color and spatial location of the suggested box were counterbalanced between trials. After the puppet’s testimony, children were asked to select the box they believed contained the sticker. Children’s responses were coded for whether or not they followed the advice of the puppet.

When the puppet was introduced, children heard one of two mental state descriptions—both descriptions indicated that the puppet was a bad source of information, but for different reasons. In the *bad intention* condition, the puppet was described as mean but knowledgeable, and in *bad knowledge* condition, the puppet was described as nice, but lacking knowledge.

The strength of the cues to the puppet's unreliability (high or low) was varied orthogonally to the mental state (bad intention or bad knowledge). In the *low cue-bad intention* condition, children were told that the puppet “knew where stickers were hiding, but did not want people to find stickers.” In the *low cue-bad knowledge* condition, children were told that the puppet “wanted people to find stickers but did not know where stickers were hiding.” For both low cue conditions, children received the following low cue feedback about their choice after each trial: the box selected was immediately opened (revealing a sticker or no sticker). If there was a sticker, it was placed aside for the child to take home. If there was no sticker inside, the other box was opened to reveal the sticker, and the child was not allowed to keep the sticker.

In the *high cue-bad intention* condition, children were told that the puppet “wants to be mean and does not want people to find stickers” and that he “tricks people so they cannot find stickers.” In the *high cue-bad knowledge* condition, children were told that the puppet “wants to help people find sticker but never knows where stickers are hiding” and that he “always thinks the stickers are in one when really they are in the other.” Importantly, children in both high cue conditions also received information about why the mental-state descriptions of the puppet were relevant to them: all children were told

that the puppet's unreliable mental state meant that he "always says the wrong answer." For both high cue conditions, children received the following high cue feedback about their choice after each trial: the box selected was immediately opened (revealing a sticker or no sticker). If there was a sticker, it was placed aside for the child to take home, and the experimenter said, "he told you the wrong answer, but you found the sticker anyway." If there was no sticker inside, the other box was opened to reveal the sticker, and the child was not allowed to keep the sticker, and the experimenter said, "the sticker's not in this box, it's in this other one. He told you the wrong answer." After the feedback from the experimenter, the puppet also proclaimed, "I always say the wrong answer!"

Results

Figure 1 shows the proportion of children who trusted the informant across trials (collapsed across condition and age). As can be seen from Figure 1, informant trust decreased over trials. The decrease was especially pronounced from trial 1 to trial 2, which was likely a result of the feedback provided after each trial. Informant trust was significantly greater on trial 1 (62%) than on each of the later trials: trial 2 (45%), McNemar's $\chi^2(1) = 16.00, p < .001$, trial 3 (41%), McNemar's $\chi^2(1) = 21.25, p < .001$, trial 4 (42%), McNemar's $\chi^2(1) = 18.15, p < .001$, trial 5 (37%), McNemar's $\chi^2(1) = 25.57, p < .001$, trial 6 (37%), McNemar's $\chi^2(1) = 28.02, p < .001$. We did not expect such significant trial effects. However, because our central theoretical question was whether children reason about unreliability differently for bad intentions and bad knowledge, we analyzed initial response (trial 1) separately from the post-feedback

responses (trials 2 to 6) to isolate the effects of informant mental state from the effects of feedback.

Initial trust. A logistic regression, with age, mental state (bad intention versus bad knowledge), cue strength (high versus low), and their interactions as independent variables, was conducted on children's first trial response. Results showed significant main effects of age, Wald $\chi^2(1) = 8.00, p = .005$, and cue strength, Wald $\chi^2(1) = 9.12, p = .003$. Distrust of the unreliable informant was greater for 4-year-olds than 3-year-olds, and distrust of the unreliable informant was greater with stronger cues. There was a marginally significant main effect of mental state, Wald $\chi^2(1) = 3.69, p = .055$, with children distrusting the bad knowledge informant more than the bad intention informant. There was a significant interaction between mental state and cue strength, Wald $\chi^2(1) = 4.17, p = .041$. Follow-up analyses indicated that the effect of the mental state manipulation was significant in the low cue condition, $\chi^2(1) = 7.52, p = .006$, but not in the high cue condition, $\chi^2(1) = 0.05, ns$. Children distrusted the bad knowledge informant more than the bad intention informant in the low cue condition, but they did not differentiate between informants with the two mental states in the high cue condition. These results suggest that when there is minimal information about the mental states of sources, young children are better at making trust judgments based on informant knowledge than on informant intentions; however, when there are stronger information about sources, young children are equally able to make trust judgments about informant knowledge and intentions. None of the other interactions in the logistic regression were significant.

Trust after feedback. For the post-feedback responses, we conducted a 2 (Age: 3 or 4 years) x 2 (Cue Strength: High versus low) x 2 (Mental State: bad intention versus bad knowledge) analysis of variance (ANOVA) on total distrust score for trials 2 through 6. There was a significant main effect of age in which 4-year-olds distrusted the unreliable informant (75.9% of the time) more than 3-year-olds (42.7% of the time), $F(1, 162) = 33.96, p < .001$. None of the other main effects were significant, and there were no interactions. These results suggest that once children receive feedback, the distinctions between mental state conditions and between cue strength conditions no longer influenced children's reasoning.

Discussion

In the present study we examined how readily young children come to the conclusion that an individual providing inaccurate information is untrustworthy. We did so with a focus on examining the effects of what children are told about the mental state of the source. We found that mental state information affected initial trust judgments, with children showing more sophisticated reasoning about mistakes based on knowledge rather than about mistakes based on intentions. This finding is consistent with previous work collectively showing that children as young as age 3 or 4 show substantial sophistication in making distinctions based on a source's knowledge (Sabbagh and Baldwin, 2001; Jaswal & Malone, 2007; Lutz & Keil, 2002), but that children of the same age or older often fail to systematically use information about an individual's intentions to make inferences about trust (Heyman & Legare, 2005; Mascaro & Sperber, 2009; Vanderbilt, Liu & Heyman, 2011). However, isolating mental states as the key factor

distinguishing these results was previously difficult due to differences between studies, and recent research suggests that differences across studies can make a large difference in the pattern of results observed (Vanderbilt, Heyman, and Liu, 2013). The present study controlled for these factors and indeed found that it is more difficult for children to reason about bad intentions than bad knowledge prior to receiving feedback. This indicates that children may indeed begin to be skeptical of sources who lack knowledge before becoming skeptical of sources who intend to deceive. This may be the case because understanding the implications of trusting a source described as deceptive is more difficult for young children, or requires more advanced theory of mind.

The fact that children showed more sophisticated reasoning about bad knowledge than bad intentions raises the question of why the results of some studies suggest that we might have observed the reverse pattern. Specifically, there is clear evidence that young children show a strong tendency to successfully reason about desires before they successfully reason about beliefs (Wellman & Liu, 2004). This may be because there is a critical distinction to be made between understanding bad desires and understanding other kinds of desires. For example, children may readily accept that others can have different food preferences that cause them to make particular food choices before they understand that people can have bad intentions that make them inflict negative consequences on others. Such a relative difficulty reasoning about bad intentions could reflect general tendencies to avoid viewing people negatively in a moral sense (Boseovski, 2010) or to a difficulty in distinguishing between what children know about

how someone *should* act versus how they do act (Heyman, Fu, & Lee, 2013; Heyman & Legare, 2005).

In addition to our primary findings we found direct evidence that children showed more sophisticated reasoning when given strong unreliability cues than weak reliability cues. Prior research had suggested that children might make such a distinction (Heyman et al, 2013; Vanderbilt et al, 2011), and the direct comparison done in the present research allowed us to document it. Cue strength may have influenced children's assessments for a few reasons. First, as previously suggested, stronger cues to the negative qualities of speakers may have helped children overcome a reluctance to make negative assessments of others who demonstrate unreliability. Second, stronger cues may have helped to prompt skepticism in children with only a weak understanding of the implications of unreliability.

Interestingly, neither mental state information, nor cue strength, made a difference in children's judgments after children began receiving feedback on the sticker task. This is particularly notable because the manipulation of cue strength included information about the characteristics of the speaker, as well as stronger feedback cues. We interpret this to mean that the feedback children receive tends to override other information they receive, possibly because direct feedback is easier for children to reason about than the internal qualities of others. This suggests that young children may only consider the mental states of others when they don't have other cues to use as shortcuts to inform their trust judgments.

The explanation that when young children have available shortcuts they no longer focus on theory of mind is also consistent with results showing an effect of mental state on initial trust judgments in the low cue but not the high cue condition. This is because the high cue condition provided children with information about the source that outlined the consequences of the speakers' mental states (i.e., that the speaker would say the wrong answer) that could have lessened their need to attend to mental states. This possibility suggests that once children understand the consequences of a speaker's unreliability, they no longer attend to the cause of unreliability. This finding is consistent with results from Liu, Vanderbilt and Heyman (2012) indicating that children tend to attend to the outcome of their search in the sticker task, rather than the intentions of their partners.

It is also possible that children did not discriminate between types of sources in the high-cue condition because hearing the extra cues to speaker unreliability, especially that the speaker would "say the wrong answer," really drove home for children that the deceptive sources meant to deceive them. In line with this idea, hearing about the consequences of speakers' mental states might have been especially helpful for young children reasoning about bad intentions but not about bad knowledge because children may already be familiar with the consequences of bad knowledge, but may not yet have a full understanding of the implications of having bad intentions.

These findings also have important implications for debates about the relation between theory of mind capacities and trust (Heyman, 2008): It may be that these

capacities predict trust judgments primarily in situations where children rely on mental states to distrust rather than on other cues.

The results presented here are the first to directly compare children's reasoning about speakers who are unreliable due to bad knowledge and children's reasoning about speakers who are unreliable due to bad intentions. This direct comparison revealed that children have more difficulty distrusting deceptive sources. This was also the first study to directly investigate how cue strength influences children's trust judgments. We found that cue strength can facilitate reasoning especially with reference to individuals intentionally providing inaccurate information. Further, this study highlights the importance of directly comparing different aspects of tasks used to measure selective trust in order to most clearly assess children reasoning in this domain. Future research will be needed to continue to equate different aspects of selective trust tasks and to directly investigate the role feedback plays in children's trust judgments.

References

- Baillargeon, R., Scott, R. M., & He, Z. (2010). False-belief understanding in infants. *Trends in Cognitive Sciences, 14*, 110.
- 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.
- Boseovski, J. J. (2010). Evidence of "rose-colored glasses": An examination of the positivity bias in young children's personality judgments. *Child Development Perspectives, 4*, 212-218.
- Hala, S., Chandler, M., & Fritz, A. S. (1991). Fledgling Theories of Mind: Deception as a Marker of Three-Year-Olds' Understanding of False Belief. *Child Development, 62*, 83-97.
- Chandler, M., Fritz, A. S., & Hala, S. (1989). Small-scale deceit: Deception as a marker of two-, three-, and four-year-olds' early theories of mind. *Child Development, 60*, 1263-1277.
- Clement, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*, 360-379.
- DiYanni, C., & Kelemen, D. (2008). Using a bad tool with good intention: Young children's imitation of adults' questionable choices. *Journal of Experimental Child Psychology, 101*, 241-261.
- DiYanni, C., Nini, D., Rheel, W., & Livelli, A. (2012). 'I Won't Trust You if I Think You're Trying to Deceive Me': Relations Between Selective Trust, Theory of Mind, and Imitation in Early Childhood. *Journal of Cognition and Development, 13*, 354-371.
- 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.
- Harris, P.L. & Corriveau, K.H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences, 366*, 1179-1187.
- Heyman, G. D. (2008). Children's critical thinking when learning from others. *Current Directions in Psychological Science, 17*, 344-347.

- Heyman, G. D., Fu, G., Lee, K. (2013). Selective skepticism: American and Chinese children's reasoning about evaluative academic feedback. *Developmental Psychology*, *49*, 543-553.
- Heyman, G. D., & Legare, C. H. (2005). Children's evaluation of sources of information about traits. *Developmental Psychology*, *41*, 636-647.
- Heyman, G. D., Sritanyaratana, L., & Vanderbilt, K. E. (2013). Young Children's Trust in Overtly Misleading Advice. *Cognitive Science*, *37*, 646-667.
- Pasquini, E., Corriveau, K., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, *43*, 1216-1226.
- Jaswal, V. K. (2006). Preschoolers favor the creator's label when reasoning about an artifact's function. *Cognition*, *99*, B83-B92.
- Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development*, *8*, 263-283.
- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*, 1261-1277.
- 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*, 1-17.
- Liu, D., Vanderbilt, K. E., & Heyman, G. D. (2013). Selective trust: Children's use of intention and outcome of past testimony. *Developmental Psychology*, *49*, 439-445.
- Lutz, D. J., & Keil, F. C. (2002). Early understandings of the division of cognitive labor. *Child Development*, *73*, 1073-1084.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition*, *112*, 367-380.
- Meltzoff, A. N. (2007). 'Like me': a foundation for social cognition. *Developmental Science*, *10*, 126-134.
- 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.

- Pasquini, E. S., Corriveau, K. H., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*, 1216-1226.
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72*, 1054-1070.
- Vanderbilt, K. E., Liu, D., & Heyman, G. D. (2011). The development of distrust. *Child Development, 82*, 1372-1380.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development, 72*, 655-684.
- Wellman, H. M., & Liu, D. (2004). Scaling of theory of mind tasks. *Child Development, 75*, 523-541.
- Woodward, A. L. (2009). Infants' grasp of others' intentions. *Current Directions in Psychological Science, 18*, 53-57.




Figure 3.1: The percentage of children who distrusted the speaker's suggestion on each of six trials of the sticker search task

Chapter 3, in full, is currently being prepared for submission for publication of the material. Vanderbilt, K.E., Heyman, G. D., & Liu, D. The dissertation author was the primary investigator and author of this paper.

GENERAL DISCUSSION

In this set of experiments, our goal was to investigate preschool children's understanding of unreliability and their ability to selectively trust more reliable sources. Our research suggests several new findings in the domain of selective trust.

In Chapter 1, we were primarily interested in whether or not children demonstrate selective trust of helpful over deceptive speakers, and how children's selective trust relates to their understanding of mental states. We presented children with speakers who had either a past history of helpfulness or deception, and measured whether children followed the advice of these speakers when looking for a hidden sticker. We found that 5-year-olds, but not 4- and 3-year-olds selectively relied on advice from more helpful speakers. Although there was a significant increase in the amount of distrust displayed in deceptive sources between 3- and 4-year-olds, it was not until 5 years of age that children discriminated between different types of sources. Interestingly, we also found that both 4- and 5-year-olds were able to verbally identify more deceptive sources and accurately predict the future behavior of these sources in the same context. 5-year-olds were also able to make generalized assessments of these speakers, by pointing out who was generally nice or mean, and predicting the behavior of sources in alternative contexts. Notably, 4-year-olds were able to verbally predict that previously deceptive sources would try to deceive them in the future, but still did not discriminate between helpful and deceptive sources on behavioral measures. This suggests that children may possess knowledge about a speaker's unreliability before they are able to use that information to inform their trust judgments. Finally, we found a relationship between children's

selective trust of more helpful sources and their understanding of mental states, indicating that children's understanding of reliability was related to their theory of mind understanding.

In Chapter 2, we were primarily interested in investigating whether children make relative or absolute judgments about a speaker's unreliability. Children heard speakers accurately or inaccurately label familiar objects. Then, the same speakers labeled novel objects; either an individual source provided labels or two sources provided conflicting labels, and children's endorsements of these labels were measured. We found that children systematically rejected the advice of unreliable speakers when they provided testimony that conflicted with testimony from another speaker, but that children systematically *relied* on advice from unreliable speakers when those sources were the only ones to proactively offer advice. Further, children in all cases believed that unreliable speakers would be good people to ask for information in the future. These results indicate that children make relative judgments about the reliability of sources they encounter, and do not consider unreliable sources to be wholly unreliable, but only slightly less reliable than any other source. This suggests that children may have a weak or incomplete understanding of the reliability of sources.

In Chapter 3, we were primarily interested in whether children's inferences about unreliable speakers would be influenced by the mental state that caused the speaker to be unreliable (either bad knowledge or bad intentions) and the strength of the cues children received about those mental states. Children were presented with a speaker who gave bad advice either because he lacked knowledge or because he had bad intentions, and children

heard descriptions that involved strong or weak cues to the mental state of the speaker. We found that children showed more initial distrust toward sources who were unreliable due to bad knowledge, and that stronger cues to unreliability were particularly helpful for children when reasoning about deceptive sources. These results indicate that children may have more difficulty reasoning about unreliability resulting from bad intentions, as children showed great improvement in their reasoning about deceptive sources when they had more information. We also found that after receiving feedback on their trust judgments, children appeared to use external feedback cues to make their decisions in lieu of focusing on mental states. This indicates that children may use feedback cues as a shortcut to reasoning about reliability, rather than reasoning about the mental states of others, and that children may not consider the reason for unreliability once they know a source is incorrect.

Taken together, our research suggests that children may not always use information they possess about the mental states of speakers to make trust judgments. We saw this in some form in each of the three chapters discussed. In Chapter 1, 4-year-olds were able to identify sources who wanted to deceive them and predict their future behavior, but still trusted these sources as much as they trusted helpful sources. This suggests that 4-year-olds may be able to make inferences about speakers, but not yet always be able to use that knowledge of mental states to inform their actions. This finding supports the idea that young children's ability to make inferences about the mental states of others may arise separately from their ability to apply that information to future predictions and choices (Flavell, 1974). Children may require a stronger understanding of

reliability before they are able to use their knowledge of mental states to inform their trust judgments. In Chapter 2, we observed that children often believed inaccurate sources to be reliable, and did so even in situations in which they rejected information provided by those sources. This suggests that children may not have a full understanding of the implications of unreliability displayed by others. Here as well, children's ability to make inferences about the reliability of a source does not always seem to align with their ability to make trust judgments. Woolley (2006) proposes that such mismatches in children's verbal and behavioral assessments shed light on the quality of their understanding in a given domain, suggesting that when children show strong reasoning on one type of measure (either verbal or behavioral), but weakness on the other, it indicates children may not yet have a full understanding of the domain in question. This corroborates our suggestion that children may have a weak or developing understanding of reliability and its implications. In Chapter 3, we observed that children did not rely on their knowledge of a source's mental state when other cues to reliability were available. Children in this study focused on external feedback cues as a shortcut to assessing reliability, rather than thinking about the mental state of the speaker. This suggests that children do not always rely on knowledge they have about a source when making trust judgments and is consistent with findings from Liu, Vanderbilt and Heyman (2013) indicating that children often pay attention to outcome information rather than the internal characteristics of others. Further, this result indicates that children may have trouble understanding why it is critical to attend to a source's mental states when making trust judgments. Collectively the findings from these three chapters suggest that children

may not have a full understanding of unreliability and its consequences (even when they demonstrate correct reasoning on selective trust tasks), and do not always take into account the mental states of others when making trust judgments.

Our results also address methodological aspects of experimentation in selective trust in two ways. First, these findings highlight the importance of including multiple measures when investigating children's understanding of selective trust. In Chapters 1 and 2 we used both verbal and behavioral measures in our tasks. This allowed us to assess children's selective trust, as well as the reasoning behind their choices, and in turn revealed many nuances of their understanding which would not have been apparent from a single measure of selective trust, and had not previously been addressed by the literature. For example, in Chapter 1, through the use of multiple measures we were able to determine that children had made inferences about the deceptive intentions of speakers even though they did not display distrust toward these sources. Importantly, as in this example, we often found differences in our measures that led us to a better characterization of how children were reasoning about sources during our tasks. Second, we found evidence that even small methodological changes sometimes had large effects on children's trust judgments. In Chapter 2 we observed that conflict in testimony, rather than testimony itself, played a large role in prompting selective trust in children. In Chapter 3 we found a large shift in children's trust judgments as a result of the strength of the information available about the mental state of the speaker. These results provide evidence that children are indeed attending to specific details in order to assess reliability

in others accurately. However, it also means that small changes can have a large influence on children's decisions while their understanding is developing.

Finally, our results highlight the idea that the development of children's understanding of reliability is a gradual and nuanced process. Rather than developing all at once, children's understanding in this area seems to develop slowly throughout the preschool years and may develop in separate steps or components that build toward a full understanding of reliability. As we have observed, 5-year-olds, but not 3- or 4-year-olds selectively rely on more helpful sources, as well as make inferences and predictions about the behavior of such speakers in both similar and generalized contexts. Four-year-olds can identify and predict the future behavior of deceptive sources, and can selectively trust more competent sources, but not more honest sources. And 3-year-olds can selectively rely on more competent sources when both a competent and an incompetent source both provide conflicting testimony. However, 3- and 4-year-olds still consider sources with a history of unreliability to be mostly trustworthy, and do not always use their knowledge of speaker's mental states to inform their trust judgments, possibly because they do not fully understand the implications of unreliability or the consequences of trusting an unreliable source. Together the results presented here suggest that children's understanding may develop as components of a larger ability to assess and deploy selective trust, and to understand the implications of doing so.

In closing, from the research presented here we now have a clearer and more descriptive picture of how children understand the reliability of others. We now have evidence that children begin to distrust deceptive sources around 4 or 5 years of age and

that selective trust is related to mental state understanding. We have evidence that children's selective trust of more knowledgeable sources likely arises before their selective trust of more prosocial sources, and that children make relative trust judgments prior to absolute trust judgments. And we have evidence that children may not fully understand the implications of unreliability even when they can identify, or even distrust, a bad source. It is apparent from these results that children's understanding in this area is complex and takes several years to develop, and that all aspects of this understanding may not develop at the same time. Though much more research will be needed to fully understand children's reasoning in this domain, the research presented in this dissertation has made a significant contribution to our understanding of children's reasoning about the reliability of sources and the factors that influence children's judgments of such sources.

REFERENCES

- Birch, S. A. J., Akmal, N., & Frampton, K. L. (2009). Two-year-olds are vigilant of others' nonverbal cues to credibility. *Developmental Science, 13*, 363 – 369.
- 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.
- Clement, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19*, 360-379.
- Dawkins, R. (1993). Viruses of the mind. In B. Dahlbom (Eds.), *Dennett and his critics: Demystifying mind* (pp. 13-27). Oxford, England: Blackwell.
- DiYanni, C., & Kelemen, D. (2008). Using a bad tool with good intention: Young children's imitation of adults' questionable choices. *Journal of Experimental Child Psychology, 101*, 241-261.
- Flavell, J.H. (1974), "The Development of Inferences About Others," In Theodore Mischel, (Ed.) *Understanding Other Persons*. Oxford, UK: Basil Blackwell, 66-116.
- Harris, P. (2007). Trust. *Developmental Science, 10*, 135-138.
- Harris, P.L. & Corriveau, K.H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences, 366*, 1179-1187.
- Heyman, G. D., Sritanyaratana, L., & Vanderbilt, K. E. (2013). Young Children's Trust in Overtly Misleading Advice. *Cognitive Science, 37*, 646–667.
- Jaswal, V. K. (2006). Preschoolers favor the creator's label when reasoning about an artifact's function. *Cognition, 99*, B83–B92.
- Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development, 8*, 263-283.
- 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, 9*, 757-758.
- Koenig, M., Clement, F., & Harris, P. (2004). 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.
- Lampinen, J. M., & Smith, V. L. (1995). The incredible (and sometimes incredulous) child witness: Child eyewitnesses' sensitivity to source of credibility cues. *Journal of Applied Psychology, 80*, 621-627.
- Liu, D. & Vanderbilt, K. E. (2013). Children learn from and about variability between people. In M. R. Banaji & S. A. Gelman (Eds.) *Navigating the Social World: A Developmental Perspective*. New York: Oxford University Press.
- Liu, D., Vanderbilt, K. E., & Heyman, G. D. (2013). Selective trust: Children's use of intention and outcome of past testimony. *Developmental Psychology, 49*, 439-445.
- Lutz, D. J., & Keil, F. C. (2002). Early understandings of the division of cognitive labor. *Child Development, 73*, 1073-1084.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition, 112*, 367-380.
- Nurmsoo, E., & Robinson, E. J. (2009a) Children's trust in previously inaccurate informants who were well or poorly-informed: When past errors can be excused. *Child Development 80*, 23-27.
- Nurmsoo, E., & Robinson, E. J. (2009b) Identifying unreliable informants: do children excuse past inaccuracy? *Development Science 12*, 41-47.
- Pasquini, E. S., Corriveau, K. H., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*, 1216-1226.
- Robinson, E. J., Champion, H., & Mitchell, P. (1999). Children's ability to infer utterance veracity from speaker informedness. *Developmental Psychology, 35*, 535-546.
- Robinson E.J., & Nurmsoo, E. (2009) When do children learn from unreliable speakers? *Cognitive Development 24*,16-22
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72*, 1054-1070.
- VanderBorgh, M., & Jaswal, V. K. (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development, 18*, 61-71.

Woolley, J. D. (2006). Verbal-behavioral dissociations in development. *Child Development, 77*, 1539-1553.