CHILD DEVELOPMENT

Child Development, March/April 2013, Volume 84, Number 2, Pages 443–454

Did She Jump Because She Was the Big Sister or Because the Trampoline Was Safe? Causal Inference and the Development of Social Attribution

Elizabeth Seiver and Alison Gopnik
University of California at Berkeley

Noah D. Goodman
Stanford University

Children rely on both evidence and prior knowledge to make physical causal inferences; this study explores whether they make attributions about others' behavior in the same manner. A total of one hundred and fifty-nine 4- and 6-year-olds saw 2 dolls interacting with 2 activities, and explained the dolls' actions. In the person condition, each doll acted consistently across activities, but differently from each other. In the situation condition, the two dolls acted differently for each activity, but both performed the same actions. Both age groups provided more "person" explanations (citing features of the doll) in the person condition than in the situation condition. In addition, 6-year-olds showed an overall bias toward "person" explanations. As in physical causal inference, social causal inference combines covariational evidence and prior knowledge.

People explain human actions in different ways. They may attribute a person's actions to their internal, individual, enduring characteristics or to the effect of external situations. Social psychologists have found that these causal explanations and attributions have far-reaching consequences for other kinds of social cognition and behavior, such as motivation, achievement, assigning blame, mental health, and general emotional well-being in adults (e.g., Hong, Chiu, Dweck, Lin, & Wan, 1999; Levy & Dweck, 1998; Tamir, John, Srivastava, & Gross, 2007) and in children (Levy & Dweck, 1999; Patrick, Skinner, & Connell, 1993).

Especially in Western cultures, many adults tend to attribute the actions of others to individual, enduring traits of the person rather than to external situations (Jones & Harris, 1967; Na & Kitayama, 2011; Nisbett, 2004; Ross & Berkowitz, 1977). Some researchers have suggested that this is because these adults have developed an intuitive theory that explains action in terms of such traits (Molden, Plaks, & Dweck, 2006; Morris & Peng, 1994; Rosati et al., 2001). This existing theory would affect the observer's interpretation of new behavioral evidence. Just as a stubborn scientist will interpret and explain all evidence in terms of her pet theory, adults who have developed a strong prior belief that actions are the result of traits might show a bias toward trait explanations.

What kinds of evidence might lead to an attribution bias? Kelley originally suggested that reasoning from covariation evidence might play an important role in trait attributions (Kelley & Levine, 1967; Plaks, Grant, & Dweck, 2005). Empirical studies confirm that covariation—the degree to which two variables change together across contexts—plays a role in adult attribution (Cheng & Novick, 1990; Hewstone & Jaspars, 1987; Morris & Larrick, 1995; Orvis, Cunningham, & Kelley, 1975; Sutton & McClure, 2001); although see Malle, 2011, for a dissenting opinion on the person–situation dichotomy and role of covariation).

The developmental trajectory that leads to these adult attribution biases is still not known, however, and it is equally not known how the role of covariation evidence changes along that trajectory. Even very young children clearly can explain actions in terms of internal psychological causes; in fact, they preferentially explain action in terms of internal mental states (Flavell, Flavell, Green, & Moses, 1990). Very young children can also understand that these mental states may differ in different

This work was supported by funding from the McDonnell Causal Learning Collaborative Initiative and NSF Grant DLS0132487 to A.G. Portions of this paper were first presented as a poster at the biannual meeting of the Society for Research in Child Development in Boston, MA (2007). We would like to thank Chris Lucas, Alex Young, Chetan Amar, Adrienne Wente, and Rodolfo Cortes for help with data collection, and David Danks for helpful discussions and comments.

Correspondence concerning this article should be addressed to Elizabeth Seiver, Department of Psychology, University of California, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650. Electronic mail may be sent to seiver@berkeley.edu.

DOI: 10.1111/j.1467-8624.2012.01865.x

individuals. For example, 18-month-olds understand that someone else may have different desires than they do (see, e.g., Repacholi & Gopnik, 1997) and 2-year-olds can make these differences explicit in their explanations (Bartsch & Wellman, 1997). However, traits have a more complex causal structure than simple mental states. Beyond mental states themselves, traits also possess the qualities of (a) persistent differences across different individuals and (b) consistency within a particular individual over time and across different situations. When and why do children make causal attributions of this kind?

Previous research has shown that preschool children's explanations and predictions about behavior differ from those of adults in two important ways. First, preschoolers do not spontaneously use trait words to explain actions. Many researchers have demonstrated that children do not spontaneously explain actions in terms of traits until middle childhood and that these attributions increase over time (Alvarez, Ruble, & Bolger, 2001; Higgins & Bryant, 1982; Peevers & Secord, 1973; Rholes & Ruble, 1984; Ruble, Feldman, Higgins, & Karlovac, 1979; Shimizu, 2000). Second, preschool children, unlike adults, do not spontaneously predict that an individual actor will continue to display a particular type of behavior over time or across situations. For example, when they see someone behave in a nice or mean way once, they do not predict that that pattern will continue over time or in a new context (e.g., Rholes & Ruble, 1984).

These discrepancies between children and adults might lead to the conclusion that young children simply cannot make trait attributions at all. However, more recent research has shown that children think in more "trait-like" ways when they are given particular kinds of information. When preschoolers are shown an actor frequently exhibiting a particular behavior, they infer that the actor will continue to produce that behavior in the future (Boseovski & Lee, 2006). Similarly, if they are given a trait label (if they are told, e.g., that someone is nice or mean), they can infer the sort of behaviors the person will produce (Heyman & Gelman, 2000a, 2000b; Liu, Gelman, & Wellman, 2007). Conversely, if they see many instances of a trait-related behavior, they can infer the right trait label (Ferguson, Olthof, Luiten, & Rule, 1984; Heyman & Gelman, 1999; Matsunaga, 2002). In the earlier literature on consensus, preschool children were more likely to attribute a choice to the particular desire or preference of the actor when they saw many people making different choices. When different people made the same choice, they were more likely to attribute the choice to a feature of the object (Higgins & Bryant, 1982; Ruble et al., 1979).

So evidence about frequency, variation or consensus, or the use of a trait label, can influence preschoolers' attributions. However, in all these cases, the attributions might be more like simple internal mental state attributions rather than having the distinctive features that characterize adult trait attributions. Thus, hearing a trait label, or witnessing that an action was frequently produced or varied across individuals, leads young children to infer that the mental state underlying this action is frequent or variable. However, these preschoolers still did not spontaneously construct "trait-like" explanations or use trait labels to make predictions about what different individuals would do across time or in new situations. They did not demonstrate that they interpreted trait labels as adults do, in terms of enduring and consistent features of individual people. Instead, they may have simply matched the frequency of behaviors in a particular individual, or the variation of behaviors among individuals, to relevant trait labels.

This might be because the data that children were given in these studies did not actually license the children to infer the full causal structure of adult attributions. To accurately infer internal or external causes for behavior, and to predict future behavior, it is important to track multiple people across multiple situations, not just to track the frequency of behavior in a single person, or the variance of behavior across people. This richer pattern could then be more confidently generalized to a novel person or novel situation. For example, if a scientific personality psychologist wanted to claim that an action was the result of a trait, she would have to show both that the action varied across individuals and was constant across situations—just one of these covariation patterns would be insufficient. This richer pattern of covariation, including both variation within and across individuals and variation within and across situations, would normatively support attributions with the causal structure of adult traits.

It might seem that tracking more complex covariation of this sort, and using it to infer causes, would be too difficult for very young children. Recently, however, a number of studies have shown that even very young children are surprisingly good at using covariation information to determine underlying physical causal structure, and that they do so in a rational way (Gopnik & Schulz, 2007; Gopnik, Sobel, Schulz, & Glymour, 2001; Gopnik et al., 2004). However, there is no systematic research on young children's use of this type of covariation, rather than simple frequency or variation, in the attribution and the explanation of action. Could covariation play a similar role in children's social inferences?

Such studies on covariation would extend our understanding of social cognition in young children, but they might also help us understand children's causal inference more generally. Recent developmental studies of causal inference from complex covariation have focused on somewhat narrow, specific causal attributions in the laboratory, for example, whether or not a block will cause a machine to activate. They also focus on deterministic (noiseless) causal relations—in fact, there is some evidence that preschoolers assume that physical causal relations are deterministic (Schulz & Sommerville, 2006).

Inferences about the causes of people's behavior play an important and general role in everyday life. They are also more likely to be probabilistic than purely physical causes—that is, we can never predict a person's behavior with complete certainty. Even if an observer makes trait attributions, she may not be surprised to see a timid person act bravely on occasion, but may expect that this behavior is more likely to occur in brave people. Person and situation explanations also require a fuller, more abstract causal schema rather than merely specific causal inferences. If children were to use covariation to infer traits, they would have to be capable of these more complex and general types of causal inference.

To explore these ideas, we presented 4- and 6-year-old children with different patterns of action covariation, including probabilistic covariation that would rationally support trait or situation inferences. We evaluated their causal explanations and predictions to determine if their attributions had a causal structure similar to the causal structure of adult trait attributions. In particular, we examined whether or not children attributed causes that both varied across individuals and were consistent and general within individuals. Both explanation and prediction are commonly used and valuable tools for insight into children's causal reasoning as well as social attributions, particularly when both measures are used together (see, e.g., Wellman & Liu, 2007).

The tasks were designed to be developmentally appropriate for these young children. We provided a simple scenario that did not rely on interplay between multiple people, containing data points easily tracked across trials. The relevant trait, risk taking, was less complex than characteristics previously studied in the attribution literature (such as intelligence or generosity). Risk taking is also less heavily valenced than traits such as "nice" or "mean," and thus less confounded by general value judgments (Alvarez et al., 2001; Miller, 1984; Morris & Peng, 1994). In this respect, it also more closely parallels the case of physical causation, where valence is not an issue.

Another question concerns the developmental course of such inference. Would older children, who do spontaneously describe people in terms of traits, reason differently than younger children? Earlier studies showed that middle school children were more likely to make trait attributions than younger children (e.g., Rholes & Ruble, 1984), in general, but there might be many reasons for this developmental pattern. Six-year-olds might simply make more accurate causal inferences from the evidence than 4-year-olds given their additional experience and better information-processing abilities, and so would be more likely to accurately infer traits from behavioral data. Or 4-year-olds might be biased against trait attributions and always prefer situation attributions, whereas 6-year-olds might simply have the opposite bias and prefer trait attributions, regardless of the evidence. A third possibility is that there would be a consistent interaction between the evidence and children's prior assumptions, of the sort described in Bayesian accounts of reasoning (Griffiths & Tenenbaum, 2009). Some recent studies suggest that this kind of interaction between evidence and prior knowledge can be found in children's physical causal reasoning (Kushnir & Gopnik, 2007; Schulz, Bonawitz, & Griffiths, 2007; Sobel, Tenenbaum, & Gopnik, 2004). For example, Kushnir and Gopnik (2007) pitted children's prior belief that contact is necessary for causal interaction against the evidence that children saw. Children initially believed that a block would have to be placed on a machine to make it activate, but they were able to gradually override that belief as they saw evidence that the block activated remotely. However, children continued to be biased toward the contact hypothesis—they were still more likely to say that the block would make the machine go when it made contact than not. In this case we might expect that 6-year-olds' responses would show an interaction between the evidence and an emerging trait bias, and that 4-year-olds would base their responses on the evidence. As children, at least North American children, get more evidence confirming a general "trait theory," they might develop a stronger "prior" for trait hypotheses, and require more evidence to overcome that prior. In that case, 6-year-olds might actually prove to be, rather surprisingly, less sensitive to behavioral evidence than 4-year-olds.

This study differs from earlier studies of covariation and trait attribution in several ways. First, we give children the equivalent of a 2 × 2 covariation table: evidence that would support explanations with the causal structure of traits rather than simply giving evidence about marginal frequency or variation. Second, we see if children who receive these data will go beyond matching trait labels to patterns of frequency or variation and will generate spontaneous "trait-like" or "situation-like" explanations and make appropriately general predictions. Third, we include a probabilistic covariation condition to see whether and how children reason about noisy behavioral data. The current study thus integrates recent research on causal inference and on the development of social cognition.

Method

Participants

In the test conditions there were forty-eight 4-year-olds (M = 4.5 years, range = 4.0–5.2 years); 26 boys and 22 girls) and forty-eight 6-year-olds (M = 6.4 years, range = 6.0–6.9 years; 24 boys and 24 girls). The control conditions included thirty-one 4-year-olds (M = 4.4 years, range = 4.0–4.9 years) and thirty-two 6-year-olds (M = 6.6 years, range = 6.0–6.9 years). (For a full set of ages and gender by condition, see Table 1.) Recruitment and testing took place at a children's science museum and a local preschool. Although official demographic data were not collected, the participants were representative of the community in the surrounding area.

Table 1
Participants' Age and Gender by Condition

Study condition	Total participants	Average age (range) years	Boys	Girls
Deterministic				
Person				
4s	12	4.56 (4.00–5.24)	7	5
6s	12	6.53 (6.03–6.88)	6	6
Situation				
4s	12	4.59 (4.05–4.96)	7	5
6s	12	6.65 (6.33–6.99)	7	5
Control				
4s	16	4.48 (4.02–4.99)	9	7
6s	16	6.51 (5.96–6.98)	10	6
Probabilistic				
Person				
4s	12	4.47 (4.06–4.91)	6	6
6s	12	6.32 (5.73–6.93)	5	7
Situation				
4s	12	4.58 (4.12–4.95)	6	6
6s	12	6.39 (5.65–6.81)	7	5
Control				
4s	13	4.53 (4.11–4.89)	9	4
6s	16	6.60 (5.92–6.99)	7	9

Materials

Two small female dolls were used, as well as scaled three-dimensional colorful cardboard constructs of a diving board with swimming pool, a trampoline, and a bicycle.

Test Design

Person versus situation conditions. In each experimental condition, participants saw a total of eight engaging actions and eight backing away actions. However, the distribution of those actions either covaried with the situation or with the individual in a between-subjects design (see Table 2). We refer to the condition where behavior covaried with the dolls (but not with the situation) as the person condition and the condition where behavior covaried with the activities (but not with the dolls) as the situation condition. In the person condition, for example, Josie would consistently play on the bicycle and the trampoline, while Sally would consistently back away from both activities. In the parallel situation condition, both Sally and Josie would consistently play on the bicycle, but they would both consistently back away from the trampoline. Importantly, everything, including language, was held constant between these conditions except for the actual covariation pattern.

Deterministic versus probabilistic conditions. In the deterministic person and situation conditions, the dolls either engaged in or backed away from the activity consistently on all four trials. In the deterministic case, the experimenter referred to the appropriate mental state in narrating the events. On each trial, in both the situation and person conditions, she commented on the doll either playing (e.g., "Look, Josie's playing on the trampoline, she's not scared") or not playing (e.g., "Look Sally's not playing on the bicycle, she's scared"). In the probabilistic conditions, the dolls either engaged in the activity three of four times or backed away three of four times. The anomalous evidence occurred on

Table 2
Summary of Experimental Design by Condition

Study condition	Deterministic			Probabilistic		
		Doll			Doll	
	Activity	A	B	Activity	A	B
Proportion of playing trials to approaches						
Person	I	4/4	0/4	I	3/4	1/4
	II	4/4	0/4	II	3/4	1/4
Situation	I	4/4	4/4	I	3/4	3/4
	II	0/4	0/4	II	1/4	1/4
Control	I	4/4	—	I	6/8	—
	II	—	0/4	II	—	2/8

the third approach. The procedure was otherwise identical to the deterministic case, in both the person and situation condition, except that the experimenter's narration of the action was changed slightly to be more appropriate to the probabilistic context. When a scared person occasionally acts bravely, however, we tend to think that she is still scared but has overcome her fear on this occasion. So instead of saying "she's scared" the experimenter simply said, "Look! Sally's playing on the diving board" or "Look! Sally doesn't want to play on the diving board." This also meant that the mental state reference in the description of the event was much more indirectly related to the trait in this condition than in the deterministic condition.

Control condition. In addition to the experimental conditions, we ran a control measure to obtain a comparison baseline preference for explaining behavior and to assess the potential influence of prior knowledge. In the control conditions, we tested an additional group of children to see if they would prefer a person or situation explanation when frequencies differed, but in the absence of full 2×2 covariation information. One doll acted fearfully more frequently than the other in the person test condition, whereas each doll showed fear the same number of times in the situation test condition (see Table 2). This information alone might have caused children to make different attributions. Therefore, control condition vignettes matched this frequency difference—and were otherwise completely identical to the test conditions—but did not have the 2×2 covariation pattern (see Table 2). In these control conditions, the evidence by itself does not rationally support a particular inference about the cause of the characters' behavior, and children's explanations should be at chance if they are based solely on this evidence. However, prior knowledge might bias children toward preferring either person or situation explanations in these cases.

In the deterministic control, the procedure was identical to the deterministic test condition except that children observed one doll play at only one activity, and the other doll play on a second activity; for example children saw Sally approach and play on the trampoline four times, and then saw Josie approach and back away from the bicycle four times. Although the actions and language were identical to the test conditions, in this case the covariation evidence supports both causal hypotheses equally.

However, this meant there were fewer trials for each doll overall in the control than in the test conditions, and it was not clear whether this might make the task easier or harder for the children. Therefore, for the probabilistic control, one doll approached an activity two of eight times, and the other approached a second activity six of eight times, mirroring both the ordering and number of positive and negative trials in the probabilistic test condition for both persons and situations (see Table 2). Again, since language and frequency were the same across the test and control conditions, this ensured that responses were not simply an effect of the linguistic descriptions or the frequency of the actions.

Test Procedure

All children were tested in a quiet room by a female experimenter and randomly assigned to each experimental condition. Participants observed a vignette in which two dolls (named Sally and Josie) made a series of approaches to two activities (chosen from among the trampoline, bicycle, and diving board). The first doll would approach Activity A four times, followed by the second doll, which would also approach that activity four times. Then the first doll would approach Activity B four times, followed by the second doll. On each trial the doll would either engage in the activity (dive in the pool, jump on the trampoline, or ride the bike), or else would back away. Doll order and activity order were counterbalanced across participants.

For example, in the deterministic person condition, children might see Josie jump on the trampoline four times, then see Sally approach the trampoline and back away four times, then see Josie ride on the bicycle four times, and then see Sally approach the bicycle and back away four times. The experimenter narrated throughout on each trial as each doll either engaged in the activity or backed away (e.g. "Look! Josie's playing on the trampoline. She's not scared" in the deterministic condition or "Look! Josie's playing on the trampoline" in the probabilistic condition).

Explanation questions. At the end of the vignette, the experimenter asked two open-ended explanation questions, one for each doll's last action on the second activity (e.g., "Why did Josie jump on the trampoline?" and "Why didn't Sally jump on the trampoline?"). Children first answered for the doll most recently viewed, and then answered for the second doll. Children were free to explain the behavior as they wished. If the participant refused to answer, or gave an irrelevant answer, the experimenter would follow-up with a forced-choice question contrasting a person and situation attribution (e.g., "Why did Josie jump on the trampoline? Is it because she's the kind of person who does brave things, or because the trampoline is safe to play on?" or "Why didn't Sally jump on the trampoline? Is it because she's the kind of person who gets scared, or because the trampoline is dangerous to play on?"). Children sometimes responded to the explanation questions by simply saying, "Because she wanted to" or "because she didn't want to," especially in the probabilistic condition. The experimenter followed these responses with the question "Why did she want to?" The follow-up answer was used for coding the explanation type. A scoreable response was needed before moving on to the next question.

Prediction questions. In the person and situation conditions, the experimenter then asked two prediction questions about a new future event. In each of these conditions, the evidence facilitates a particular type of novel prediction. In the person condition, the evidence allows you to make a prediction about what each doll would do in a new situation, although it does not allow you to make predictions about what a new doll would do in the earlier situations. In the situation condition, this is reversed. We tested to see if children in each condition would make the appropriate generalization and so make correct predictions. In the person condition, the participants were asked to predict what each doll would do in a new situation ("Now let's pretend that Sally and Josie go over to the diving board. Do you think Josie will play on it? <child answers> Do you think Sally will play on it?"). In the situation condition, children were asked to predict what a new doll would do in each of the earlier situations ("Now let's pretend that Sally and Josie have a friend named Mary. Do you think she'll play on the trampoline? <child answers> Do you think she'll play on the bicycle?"). Children answered either "yes" or "no." For predictions to be scored as correct, children had to answer both questions correctly. This always entailed one "yes" and one "no" answer since playing and not playing were contrasted in both conditions. That is, children in the person condition had to say that each doll would act consistently (and differently from each other) in the new situation and children in the situation condition had to respond that a new doll would act consistently in (and differentiate between) each of the old situations. All other patterns of responses were scored as incorrect. Thus, if children were responding at chance, they would be scored as correct 25% of the time.

Explanation coding. As noted earlier, if children did not provide a relevant explanation spontaneously, they received the forced-choice question. Both forced-choice and open-ended explanation responses were coded by observers into two mutually exclusive response types (κ, interrater reliability = .732, $p < .001$). The observers were blind to study condition and only saw the explanations themselves. Therefore, differences across conditions would suggest that the coding scheme had some validity as well as reliability. (For examples of both types of explanations, see Table 3.) An explanation was coded as a "person" response if it attributed the doll's behavior to an internal cause specific to that doll and in contrast to others, similar to the Ruble et al. (1979) construct of "person attribution." This cause could involve the doll's mental states, such as consistent desires or beliefs, or refer to other stable characteristics of the person, such as personality, age, or size. This category included classic "trait" attributions, such as "she's brave," but also included a wider variety of person-specific attributions, such as "she's the big sister" or "she knows how to ride a bike" or "she likes to swim." Thus, trait attributions are subsumed into the broader category of "person attributions."

An explanation was coded as a "situation" response if it referred to an underlying cause that was outside the doll. This included both aspects of the physical situation (bounciness of the trampoline, pool temperature, etc.) and the social situation ("she did it because her friend did it"). All the spontaneous explanations could be coded into one of the two categories. Finally, forced-choice responses were coded "person" or "situation" based on which of the two options was chosen.

Some person and situation explanations might seem similar at first glance; for example, several children drew upon their own recent experiences about learning to ride a two-wheeler bike for both person explanations ("she doesn't know how to

Table 3
Examples of Person and Situation Explanations

Person
Mental states
She wanted to splash
She thinks there is a shark in the water
She thinks she might fall off
She learned how to ride her bike
She was in the mood for it
She liked it
She is afraid of heights
Physical attributes
She is younger
She is bigger
She does not have a helmet on
Situation
Physical object
It only has two wheels
It is too fast
It looks like fun
It looked scary
It might tip over
It is not over water and it is not high
Because it is red and blue
There is netting around it
Social situation
Josie (other doll) did not want to
Sally (other doll) played on it
Her friend did it
Her friend was not there

ride a bike") and situation explanations ("the bike only has two wheels"). However, interestingly, as we will see next, the prevalence of each explanation type consistently varied by study condition. Children chose to stress the person or situation differentially in their explanations, often as the subject of the sentence, even when the overall content of the explanations was similar. That would suggest that the coding captured genuine differences in attribution. When they faced different patterns of covariation, children produced different types of explanations—explanations that blind observers reliably classified as stressing persons or situations.

As noted earlier, children were asked two explanation questions, one about each doll. They were given a "person" score of 0, 1, or 2 depending on how many "person explanations" they provided.

Results

Explanation Results

Across study conditions, 4-year-olds provided 48 forced-choice responses and 106 explanations, and 6-year-olds provided 12 forced-choice responses and 148 explanations. Consistent with earlier studies, only a few children gave "classic" personality trait explanations, such as "she's brave." Importantly, however, all the children who gave two person explanations not only referred to internal states of the person but also differentiated between the two actors (e.g., they said Josie likes to swim, Sally does not like to swim). Moreover, both mental state and nonmental state person attributions often, although not always, implied some enduring feature, e.g., "she likes swimming," or "she's old."

We began by conducting a 2 (age: 4 vs. 6) $\times$ 2 (consistency: deterministic vs. probabilistic) $\times$ 3 (condition: person vs. situation vs. control) analysis of variance (ANOVA) on the person scores. A parallel ANOVA that excluded the forced-choice responses yielded the same pattern of results as the analysis that included them, suggesting that the results were not due to the particular question format. Therefore, we combined the two question types in our analyses.

As shown in Figure 1, there was a main effect of condition, $F(2, 153) = 19.242$, $p < .001$, $\eta^2 = .19$ overall; $F(2, 76) = 15.05$, $p < .001$, $\eta^2 = .28$ for 4-year-olds; and $F(2, 77) = 5.82$, $p < .01$, $\eta^2 = .13$ for 6-year-olds, with more person responses in the person condition ($M = 1.73$, $SD = .494$) than in the control condition ($M = 1.32$, $SD = .758$) and than in the situation condition ($M = .85$, $SD = .825$). There was also a main effect of age; 6-year-olds gave more person-based explanations than 4-year-olds across conditions, $F(1, 147) = 9.87$, $p < .01$, $\eta^2 = .05$. There was no effect of the deterministic versus probabilistic condition, (probabilistic: $M = 1.55$, $SD = .67$; deterministic: $M = 1.37$, $SD = .76$) and there were no interaction effects.

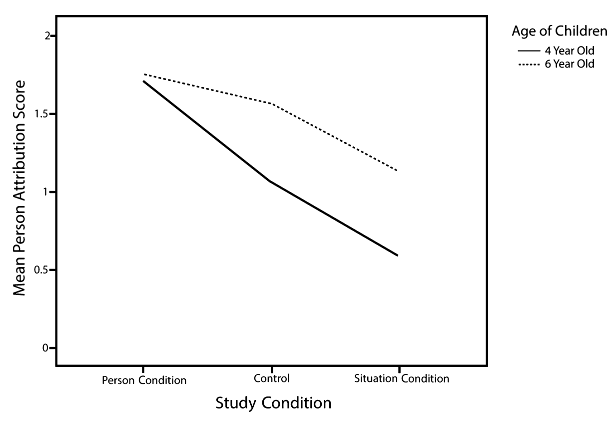

Figure 1. Four- and 6-year-olds' responses in each covariation condition.

Since there were neither significant main effects nor significant interactions involving the deterministic and probabilistic conditions, these two groups were collapsed for subsequent analyses. Note that there were also minor differences in the deterministic and probabilistic condition procedures as noted earlier—in particular, the probabilistic test condition contained fewer references to mental states, and the probabilistic control condition used eight rather than four examples of each behavior. These differences appeared to have no effect on the children's responses (nonparametric analyses yielded an identical pattern of results to those given by the ANOVAs).

Closer examination of the data showed that the 4-year-olds were actually more accurate than the 6-year-olds in the situation and control conditions. In the person condition, both age groups scored near ceiling, significantly different from chance, $t(23) = 6.3$, $p < .001$, $d = 2.63$ for 4-year-olds, and $t(23) = 8.3$, $p < .001$, $d = 3.46$ for 6-year-olds, and not significantly different from each other, $F(1, 46) = .084$, $p > .5$. In contrast, in the situation condition, 6-year-olds gave significantly more person explanations than 4-year-olds overall, $F(1, 70) = 5.7$, $p < .05$, $d = .11$. Four-year-olds gave significantly fewer person explanations in the situation condition than expected by chance, $t(23) = -2.8$, $p < .01$, $d = -1.17$, and 6-year-olds were at chance, $t(23) = .720$, $p > .4$. In addition, the distribution of their scores was at chance as revealed in a goodness-of-fit test, $\chi^2(2, n = 24) = 4.92$, $p > .08$, indicating that there were no consistent patterns within children; children were no more likely than chance to give either two person or two situation explanations.

The control conditions were identical in language and behavioral frequency to the respective test conditions, but the covariation information did not favor a person or situation attribution. Four-year-olds were at chance in the controls, $t(30) = .441$, $p > .5$. Six-year-olds, however, displayed a significant preference for person explanations, $t(31) = 5.14$, $p < .001$, $d = 1.85$. Both 4-year-olds, $t(53) = 2.29$, $p < .05$, $d = .63$, and 6-year-olds, $t(54) = 2.23$, $p < .05$, $d = .61$, produced significantly fewer person explanations in the situation condition than in the controls. For 4-year-olds, the person and control explanations also differed significantly, $t(53) = 3.33$, $p < .01$, $d = .91$, but the difference for 6-year-olds did not reach significance, $t(54) = 1.26$, $p > .2$.

Overall, as Figure 1 shows, both 4- and 6-year-olds are sensitive to the pattern of covariation when deciding what kind of causal explanation to provide—they provide more person explanations in the person condition than in the control condition, and more in the control than in the situation condition. However, 6-year-olds, unlike 4-year-olds, have a consistent bias toward person explanations. This means that in the situation and control conditions, 4-year-olds were actually more accurate, in strictly covariational terms, than 6-year-olds.

Prediction Results

The prediction task was designed to test whether children had only inferred a single causal explanation from the data they had actually seen, or if they had inferred a more abstract causal schema. Trait attributions not only involve causes that are internal to the actor and contrast with other actors but also imply that these internal causes will lead the actor to behave in similar ways across situations and through time. Similarly, situation attributions not only imply that this actor behaved as she did because of the situation but that other actors in similar situations will behave in the same way.

If children in the person condition had inferred this more abstract "trait-like" causal scheme, then they ought to predict that each doll would behave consistently in a new situation. For example, if Josie failed to approach both the bicycle and the trampoline she should also be reluctant to dive off the board. In fact, in the person condition, more 4- and 6-year-olds predicted that both dolls' behavior would be consistent in a novel situation than would be predicted by chance (17 of 24 six-year-olds, $p < .001$, binomial test; 20 of 24 four-year-olds, $p < .001$, binomial test; 37 of 48 participants, $p < .0001$, binomial test; note that chance here was 25% since to be scored as correct, children had to respond to two forced-choice questions).

Similarly, in the situation condition, children who genuinely made a situation attribution should infer that a new person (Mary) would behave similarly to Josie and Sally in the two situations. If Josie and Sally both jumped on the trampoline, but avoided the diving board, then so should Mary. In fact, 4-year-olds tended to predict that the new doll would behave as the other dolls had done (12 of 24 participants, $p < .01$, binomial test) but 6-year-olds did not (9 of 24 participants, $p > .05$, Binomial test)—they tended to assume Mary would behave the same way across both situations. These results are consistent with the explanation results. They suggest that 4-year-olds may generalize from covariation information rationally to make new predictions about both persons and situations, but that 6-year-olds are more likely to make generalizations about persons than situations.

Discussion

In this study, 4- and 6-year-old children used covariation information to make corresponding inferences about the causes of human actions. Given the appropriate Covariation × Person evidence, even with probabilistic data, 4-year-olds explained actions in terms of internal, individual, and enduring causes. They also made appropriate predictions about an individual's behavior in a new situation. When covariation evidence supported a situation attribution, they would also make those attributions and predictions correctly. Six-year-olds also used covariation information to explain and predict, but in contrast to the younger children (and more like adults), they showed an overall bias for person explanations over situation explanations. This bias apparently led them to place less weight on the behavioral evidence that was presented in the vignettes.

In particular, the presence of appropriate covariation information leads even 4-year-old children to spontaneously infer causes with the full causal structure of traits, to use these causes to explain behavior, and to predict future behavior. This was true even though the only difference across conditions was the covariation pattern—language and other cues were held constant. Note also that children made these inferences significantly less frequently in the control conditions where frequency information was available, but the full covariation matrix was not.

In particular, 4-year-olds invoked a "trait-like" causal schema to generate consistent predictions about people in novel situations. There are two ways of thinking about this schema. Four-year-olds may already have a trait-like schema in place, but unlike adults, initially they may apply it only in restricted conditions; when it is explicitly described by a trait label, as in the Heyman and Gelman (1999) study, or when it is strongly supported by covariation information, as in the present study.

Alternatively, the covariation information, along with other evidence from everyday life, may actually lead the children to posit a "trait-like" schema. Even though the children did not attribute full-blown traits in the same way that adults do, they did explain actions in terms of enduring, consistent but individually variable internal causes, and their predictions revealed similar causal attributions. The fact that children seemed to spontaneously invent trait-like explanations (e.g., "she's bigger," "she knows how to ride a bicycle") may support this idea. Children may spontaneously invent such causal schemas to explain covariation patterns in particular cases and then generalize those schemas. This kind of inference fits the pattern of general schema inference described by Kemp, Goodman, and Tenenbaum (2008).

In contrast, for the 6-year-olds, this schema may have been confirmed by covariation data many times and across many contexts and become an entrenched intuitive theoretical framework. Explanatory hypotheses that fit this framework would receive a higher probability at the outset, and require more data to defeat them. This sort of default theoretical framework could lead to a person bias.

What kinds of evidence could lead to this developmental change? One interesting hypothesis is that the developments at age 6 years are related to the increase in peer group interaction in middle childhood. In peer interaction, individual traits, rather than social roles or situations, will account for much of the variance in behavior. In a classroom of 20 otherwise similar children placed in a similar situation on the playground, some will consistently take risks and others will not. Children will see more trait-based covariation as they pay increasing attention to their peers, and acquire rich datasets across individuals and situations to draw upon.

Similarly, cross-cultural differences in covariation evidence may influence the development of attribution. Miller (1984) suggested that children across cultures began with similar attribution patterns and then diverged toward the more extreme adult patterns as they grew older, a claim that has been supported by further studies with children (Gonzalez, Zosuls, & Ruble, 2010; Kalish, 2002; Lockhart, Nakashima, Inagaki, & Keil, 2009). Again, these results suggest a mechanism by which cultural differences may influence the course of attribution. This may either be because members of different cultures actually do behave differently or, more probably, because culture and experience influence the information children receive from adults about traits, such as adult trait language. This evidence is especially relevant to the development of causal schemata. If people within a culture tend to describe behavior in terms of traits, then this will lead to covariation between certain behaviors and trait labels, which might itself provide evidence for a trait schema (see Kemp et al., 2008). If children are using covariation information about people's behavior *and* adult trait language to infer both specific causes and more general causal schemas, such differences in the data could affect their adult social cognition. It would be very interesting to see if children in a less trait-based culture (such as mainland China) would show a similar pattern of results. One might predict that in such a culture 4-year-olds would show a similar pattern, but 6-year-olds would not manifest the same trait bias. We are currently conducting such studies.

These results are also interesting because they may point to broader mechanisms for learning about traits and situations. Recent computational work outlines how attributional learning might take place. In particular, causal Bayes net learning mechanisms (Pearl, 2000; Spirtes, Glymour, & Scheines, 1993) can be used to model causal reasoning and learning in adults (e.g., chapters in Gopnik & Schulz, 2007; Rehder & Hastie, 2001; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann, Martignon, Gernsbacher, & Derry, 1998), preschool children (e.g., Gopnik et al., 2001; Gopnik et al., 2004), and even infants (Sobel & Kirkham, 2006). These models predict which causal inferences should rationally be made from different patterns of covariation and prior knowledge.

Bayesian models of causal learning (e.g., Griffiths & Tenenbaum, 2009), in particular, suggest that children make new inferences by systematically combining prior knowledge and current covariation evidence to arrive at the right causal hypothesis. Learners can select hypotheses rationally in the light of data by using Bayes's rule to combine the prior probability of different causal hypotheses and the probability of the current evidence given each hypothesis. Several recent studies (Kushnir & Gopnik, 2007; Schulz et al., 2007; Sobel et al., 2004) suggest that preschoolers can combine prior knowledge with covariation evidence in this Bayesian way. Moreover, recent work shows that this kind of inference can be used not only to develop specific causal hypotheses but also to construct more abstract causal schemas or "framework theories" (Griffiths & Tenenbaum, 2009; Kemp et al., 2008; Schulz et al., 2007).

This suggests a potential mechanism for the development of attribution. Children may begin by forming theories based on both people's behavior and how adults explain such behavior. They continue to overweight the evidence that confirms a culturally conferred hypothesis or abstract causal schema, particularly the hypothesis that internal traits cause actions, although they underweight the evidence that contradicts this hypothesis. Once that schema has been highly confirmed, it will be more difficult to overturn in future, although it might still be overturned with sufficient evidence. Eventually, in adulthood, this may result in a consistent "trait bias" that is difficult to overcome.

Whether or not this account of how children learn this bias is correct, the current study shows that some of the prerequisites for such an account are in place. Children as young as 4 years of age can use covariation evidence to make behavioral attributions, and 6-year-olds combine that evidence with prior biases to arrive at similar (but slightly skewed) conclusions. This mirrors children's ability to infer causes in the physical domain using both prior knowledge and evidence. Additional research is needed to explore a potential broader underlying framework of causal inference connecting the social and physical domains. Nonetheless, we can see the origins of Kelley and Levine's (1967) social schemata even in preschoolers.

References

Alvarez, J. M., Ruble, D. N., & Bolger, N. (2001). Trait understanding or evaluative reasoning? An analysis of children's behavioral predictions. *Child Development*, *72*, 1409–1425. doi:10.1111/1467-8624.00356

Bartsch, K., & Wellman, H. (1997). *Children talk about the mind*. New York: Oxford University Press.

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioral prediction. *Developmental Psychology*, *42*, 500–513. doi:10.1037/0012-1649.42.3.500

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545–561. doi:10.1037/0022-3514.58.4.545

Ferguson, T. J., Olthof, T., Luiten, A., & Rule, B. G. (1984). Children's use of observed behavioral frequency versus behavioral covariation in ascribing dispositions to others. *Child Development*, *55*, 2094–2105. doi:10.2307/1129782

Flavell, J. H., Flavell, E. R., Green, F. L., & Moses, L. J. (1990). Young children's understanding of fact beliefs versus value beliefs. *Child Development*, *61*, 915–928. doi:10.2307/1130865

Gonzalez, C. M., Zosuls, K. M., & Ruble, D. N. (2010). Developmental change in the understanding of trait terms. Traits as dimensions or categories? *Developmental Psychology*, *46*, 1078–1088. doi:10.1037/a0020207

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32. doi:10.1037/0033-295X.111.1.3

Gopnik, A., & Schulz, L. (Eds.). (2007). *Causal learning: Psychology, philosophy, and computation*. New York: Oxford

University Press. doi:10.1093/acprof:oso/9780195176803.001.0001

Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology*, *37*, 620–629. doi:10.1037/0012-1649.37.5.620

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*, 661–716. doi:10.1037/a0017201

Hewstone, M., & Jaspars, J. (1987). Covariation and causal attribution: A logical model of the intuitive analysis of variance. *Journal of Personality and Social Psychology*, *53*, 663–672. doi:10.1037/0022-3514.53.4.663

Heyman, G. D., & Gelman, S. A. (1999). The use of trait labels in making psychological inferences. *Child Development*, *70*, 604–619. doi:10.1111/1467-8624.00044

Heyman, G. D., & Gelman, S. A. (2000a). Preschool children's use of novel predicates to make inductive inferences about people. *Cognitive Development*, *15*, 263–280. doi:10.1016/S0885-2014(00)00028-9

Heyman, G. D., & Gelman, S. A. (2000b). Preschool children's use of trait labels to make inductive inferences. *Journal of Experimental Child Psychology*, *77*, 1–19. doi:10.1006/jecp.1999.2555

Higgins, T. E., & Bryant, S. L. (1982). Consensus information and fundamental attribution error: The role of development and in-group versus out-group knowledge. *Journal of Personality and Social Psychology*, *43*, 889–900. doi:10.1037/0022-3514.43.5.889

Hong, Y., Chiu, C., Dweck, C. S., Lin, D. M. S., & Wan, W. (1999). Implicit theories, attributions, and coping: A meaning system approach. *Journal of Personality and Social Psychology*, *77*, 588–599. doi:10.1037/0022-3514.77.3.588

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, *3*, 1–24. doi:10.1016/0022-1031(67)90034-0

Kalish, C. W. (2002). Children's predictions of consistency in people's actions. *Cognition*, *84*, 237–265. doi:10.1016/S0010-0277(02)00052-5

Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 15, pp. 192–238). Lincoln: University of Nebraska Press.

Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2008). Theory acquisition and the language of thought. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the Cognitive Science Society* (pp. 1606–1611). Austin, TX: Cognitive Science Society.

Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, *43*, 186–196. doi:10.1037/0012-1649.43.1.186

Levy, S. R., & Dweck, C. S. (1998). Trait- versus process-focused social judgment. *Social Cognition*, *16*, 151–172. doi:10.1521/soco.1998.16.1.151

Levy, S. R., & Dweck, C. S. (1999). The impact of children's static versus dynamic conceptions of people on stereotype formation. *Child Development*, *70*, 1163–1180. doi:10.1111/1467-8624.00085

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. doi:10.1111/j.1467-8624.2007.01082.x

Lockhart, K. L., Nakashima, N., Inagaki, K., & Keil, F. C. (2009). From ugly duckling to swan? Japanese and American beliefs about the stability and origins of traits. *Cognitive Development*, *23*, 155–179. doi:10.1016/j.cogdev.2007.08.001

Malle, B. F. (2011). Time to give up the dogmas of attribution: An alternative theory of behavior explanation. *Advances in Experimental Social Psychology*, *44*, 297–311. doi:10.1016/B978-0-12-385522-0.00006-8

Matsunaga, A. (2002). Preschool children's inferences about personality traits. *Japanese Journal of Developmental Psychology*, *13*, 168–177.

Miller, J. G. (1984). Culture and the development of everyday social explanation. *Journal of Personality and Social Psychology*, *46*, 961–978. doi:10.1037/0022-3514.46.5.961

Molden, D. C., Plaks, J. E., & Dweck, C. S. (2006). "Meaningful" social inferences: Effects of implicit theories on inferential processes. *Journal of Experimental Social Psychology*, *42*, 738–752. doi:10.1016/j.jesp.2005.11.005

Morris, M. W., & Larrick, R. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review*, *102*, 331–355. doi:10.1037/0033-295X.102.2.331

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology*, *67*, 949–971. doi:10.1037/0022-3514.67.6.949

Na, J., & Kitayama, S. (2011). Spontaneous trait inference is culture-specific: Behavioral and neural evidence. *Psychological Science*, *22*, 1025–1032. doi:10.1177/0956797611414727

Nisbett, R. E. (2004). *The geography of thought: How Asians and Westerners think differently and why*. New York: Free Press.

Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology*, *32*, 605–616. doi:10.1037/0022-3514.32.4.605

Patrick, B. C., Skinner, E. A., & Connell, J. P. (1993). What motivates children's behavior and emotion? Joint effects of perceived control and autonomy in the academic domain. *Journal of Personality and Social Psychology*, *65*, 781–791. doi:10.1037/0022-3514.65.4.781

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. New York: Cambridge University Press.

Peevers, B. H., & Secord, P. F. (1973). Developmental changes in attribution of descriptive concepts to persons. *Journal of Personality and Social Psychology*, *27*, 120–128. doi:10.1037/h0034469

Plaks, J. E., Grant, H., & Dweck, C. (2005). Violations of implicit theories and the sense of prediction and control: Implications for motivated person perception. *Journal of Personality and Social Psychology*, *88*, 245–262. doi:10.1037/0022-3514.88.2.245

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, *3*, 323–360. doi:10.1037/0096-3445.130.3.323

Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: Evidence from 14- and 18-month-olds. *Developmental Psychology*, *33*, 12–21. doi:10.1037/0012-1649.33.1.12

Rholes, W. S., & Ruble, D. N. (1984). Children's understanding of dispositional characteristics of others. *Child Development*, *55*, 550–560. doi:10.2307/1129966

Rosati, A. D., Knowles, E. D., Kalish, C. W., Gopnik, A., Ames, D. R., & Morris, M. W. (2001). The rocky road from acts to dispositions: Insights for attribution theory from developmental research on theories of mind. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition* (pp. 287–303). Cambridge, MA: MIT Press.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173–220). New York: Academic Press. doi:10.1016/S0065-2601(08)60357-3

Ruble, D. N., Feldman, N. S., Higgins, E. T., & Karlovac, M. (1979). Locus of causality and the use of information in the development of causal attributions. *Journal of Personality*, *47*, 595–614. doi:10.1111/j.1467-6494.1979.tb00211.x

Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology*, *43*, 1124–1139. doi:10.1037/0012-1649.43.5.1124

Schulz, L., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development*, *77*, 427–442. doi:10.1111/j.1467-8624.2006.00880.x

Shimizu, Y. (2000). Development of trait inference: Do young children understand the causal relation of trait, motive, and behavior? *Japanese Journal of Educational Psychology*, *48*, 255–266.

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology*, *42*, 1103–1115. doi:10.1037/0012-1649.42.6.1103

Sobel, D. M., Tenenbaum, J., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, *28*, 303–333 doi:10.1207/s15516709cog2803_1

Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. New York: Springer-Verlag.

Steyvers, M., Tenenbaum, J., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489. doi:10.1016/S0364-0213(03)00010-7

Sutton, R. M., & McClure, J. (2001). Covariational influences on goal-based explanation: An integrative model. *Journal of Personality and Social Psychology*, *80*, 222–236. doi:10.1037/0022-3514.80.2.222

Tamir, M., John, O. P., Srivastava, S., & Gross, J. J. (2007). Implicit theories of emotion: Affective and social outcomes across a major life transition. *Journal of Personality and Social Psychology*, *92*, 731–744. doi:10.1037/0022-3514.92.4.731

Waldmann, M. R., & Martignon, L. (1998). A Bayesian network model of causal learning. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the twentieth annual conference of the Cognitive Science Society* (pp. 1102–1107). Mahwah, NJ: Erlbaum.

Wellman, H. M., & Liu, D. (2007). Causal reasoning as informed by the early development of explanations. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation*. New York: Oxford University Press.