

“The Block Makes It Go”: Causal Language Helps Toddlers Integrate Prediction, Action, and Expectations about Contact Relations

Elizabeth Bonawitz¹, Alexandra Horowitz¹, Darlene Ferranti², & Laura Schulz¹

¹ (liz_b@mit.edu) Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139 USA

² Division of Hospital Medicine, Northwestern University; Chicago, IL 60611 USA

Abstract

Some researchers have suggested that correlation information and information about action are bound in a single representation: “causal knowledge”. If children have only observed correlation information, do they spontaneously try to generate the effect? Do they represent the relationship as potentially causal? We present three action and looking-time studies that suggest that even when toddlers (mean; 24 months) predict that one event will follow another, they neither initiate the first event to try to generate the second (as preschoolers, mean 47 months, do spontaneously), nor do they expect that the predictive relations will involve physical contact. Toddlers succeed at both of these inferences when the events are described using causal language. This suggests that causal language plays a role in helping children recognize the relationship between prediction, action, and contact causality.

Keywords: Causality; Language; Cognitive Development

No action from association

Classical conditioning and operant learning are two ways in which animals learn relationships in the environment. However, animals do not naturally bind these kinds of information together: Pavlov’s dog may learn to drool at the sound of a bell in anticipation of dinner, but will not spontaneously ring the bell to bring dinner (e.g., Gopnik & Schulz, 2004; Tomasello & Call, 1997; Blaisdell, 2008). In contrast, when adult humans observe correlated events, we can represent the events as potentially causally connected and can successfully intervene to determine whether or not one event is a direct cause of the other. This suggests that for adults, correlation information and information about action are bound in a single representation: “causal knowledge”. We do not know however, to what extent infants and young children can move from prediction to action, nor to what extent children infer that physical events that predict one another may involve spatial contact.

Researchers have long speculated about the relationship between correlation information and causal knowledge, and in particular, about the relationship between prediction and action. Philosophers have suggested that only a cognitively sophisticated being would recognize “that the very same relationship that he exploits in intervening also can be present both when other agents intervene and in nature even when no other agents are involved” (Woodward, 2004). Similarly, psychologists have suggested that causal knowledge requires understanding causal relations as stable relations among diverse events, not merely relations “that involve rewards or punishments (as in classical or operant conditioning), and not just events that immediately result from (one’s) own actions (as in operant conditioning or trial-and-error learning)” (Gopnik et al., 2004). The

implication is that human beings may be unique among animals in having the single representation (causal knowledge) that encodes the potential commonality between non-agentive covariation in the world and covariation between agent actions and outcomes.

In our first experiment, we look at whether older and younger children can use a predictive relationship between two events to initiate the target action and anticipate the outcome of their own action. In the second experiment, we test the hypothesis that causal language might help children recognize the relevance of predictive relations to their own interventions. In the third experiment, we eliminate the action measure and use a violation of expectation paradigm (manipulating the presence or absence of contact causality) to assess whether toddlers simply have difficulty bridging the gap between prediction and action or whether, in the absence of causal language, they also fail to expect that predictive relations respect principles of contact causality.

Experiment 1: From prediction to action

We introduce children to several trials of a novel event: all children see a block contact a base, after which a toy connected to the base lights up and spins. Following these observations, we ask whether children predicatively look to the toy when the block touches the base, and whether children (spontaneously or with prompting) touch the block to the base and then look to the toy. Our primary question is whether children use the evidence from the observation trials to infer that their action might generate the target outcome. Note that simply performing the action might not mean that the child expects the action to generate the outcome. To ensure that the child acts with the expectation that the outcome might result, we coded children’s predicative looks to the toy after performing the action.

Methods

Participants. Sixteen preschoolers (m=47mos; rng: 37-60 mos) and 14 toddlers (m=24 mos; rng: 19-30 mos) were recruited from a large metropolitan science museum.

Materials. A large stage blocked a confederate from view. A purple block was attached to a concealed lever which slid across a slit in the stage, creating a track for the block, and leading to a second block (base) which remained fixed to the left of the stage. An orange wire attached the base to a toy airplane in the stage’s upper left corner. The airplane was controlled by a button on the back of the airplane which could be surreptitiously activated by the confederate and which caused the toy to spin and light up.

Procedure. All children were tested individually in a quiet corner of the museum. There were three phases: an Observation Phase, an Action Phase, and (for those children who failed to intervene spontaneously) a Prompted Action Phase. (See Figure 1.) Throughout the experiment, the experimenter looked at the child so that the child would not follow the experimenter's eye gaze to the toy.

Observation Phase: At the beginning of the experiment, the experimenter elicited the child's attention by saying, "Watch my show!" The confederate, who was concealed behind the apparatus, slid the block towards the base, so that from the child's perspective the block appeared to move on its own. When the block contacted the base, the confederate immediately activated the airplane for 3 seconds. As soon as the block moved away from the base, the plane slowed to a stop. Pilot work confirmed that this provided a compelling causal illusion: adults believed that contact between the block and the base activated the plane.

This activation sequence was repeated four times. On the fifth trial (the Predictive look trial), the block contacted the base, but the confederate did not activate the plane. The experimenter observed the child to see if the child looked predicatively up to the plane. If the child failed, the experimenter added a sixth trial in which the plane activated, followed by a seventh trial in which it did not. If the child again failed to look predicatively towards the plane, they were excluded from the analyses. If the child looked predicatively towards the toy (on either trial five or seven), the experimenter concluded with a final trial in which the block contacted the base and the plane activated.

Action Phase: The experimenter slid the block towards the child, pointed to the plane and said, "Okay now it's your turn. Can you make it go?" Children were given 60 seconds to play freely. At no point was the plane activated for the child. If the child performed the target action during the 60 seconds of free play, the experiment ended; if the child failed to touch the block to the base during the 60 seconds, she or he moved onto the Prompted Action Phase.

Prompted Action Phase: The experimenter slid the block almost all the way into the base, stopping just short of the base and returned it to the child saying, "It's your turn." The child was given another 60 seconds to perform the target action. If the child failed to perform the complete action following the prompt, they were excluded from analyses.

Results & Discussion of Experiment 1

Two preschoolers and three toddlers were excluded and replaced for failing to make the initial predictive look. An additional three toddlers were excluded and replaced for failing to perform the action during the Prompted Action phase. The stringent inclusion criteria meant that we could be confident that all the children in the subsequent analyses had both learned the predictive association between the block and the toy and were able to perform the target action. Inter-coder agreement was 100%.

We coded success on the task generously: children who predicatively looked to the toy after performing the action,

regardless of whether they performed the action spontaneously or with prompting, were counted as passing the task. Children were coded as failing only if they never predicatively looked to the toy after performing the action.

While almost all preschoolers succeeded at the task (87.5%), none of the 14 toddlers did. That is, although the toddlers played freely with the block during the Action Phase, no toddler performed the target action spontaneously, and none predicatively looked to the toy after performing the prompted action. The preschoolers were significantly more likely to succeed at the task than the toddlers ($\chi^2 = 23.0, p < .0001$) and were also significantly more likely to generate the action spontaneously (and anticipate the outcome) ($\chi^2 = 13.1, p < .001$). (See Table 1.)

The striking discrepancy between the performance of the younger and older children suggests that only the older children believed the evidence of the Observation Phase indicated a possible causal relationship between the block and the toy¹. Though all toddlers both predicted the outcome of the observed action and performed the target action (under prompting), they did not show any indication of understanding that their own actions might activate the toy.

Experiment 2: Causal Language

In Experiment 2, we investigate the possibility that causal language might help toddlers represent the predictive relation as a potentially causal relation. There are two accounts by which causal language might support young children's causal reasoning. One possibility is that non-human animals (and arguably very young children) have at least two distinct systems for reasoning about correlated events: one for processing statistical associations among events in the world (as in classical conditioning), and another for processing associations between agent actions and outcomes (as in operant conditioning, trial and error learning, and imitative learning; Gopnik et al., 2004; Gopnik & Schulz, 2004; Woodward, 2004; Tomasello & Call, 1997). An intriguing proposal is that in any domain of knowledge where component abilities are ontogenetically early and phylogenetically broad, language may play a critical role in uniting otherwise separate inferential systems (e.g. Spelke, 2003). Because two core component abilities of causal inference—learning statistical associations between events and learning the relationship between one's own actions and their immediate outcomes—are present both in early infancy and in non-human animals, it is tempting to suggest that linguistic representations might support the integration of these component systems into adult-like causal reasoning.

If so, it is possible that even during the course of a short task, hearing the same words used to describe both events might help children bind together information about

¹ Toddlers might have been confused or frightened by the spontaneously moving block (i.e., because it violates core object principles; Spelke, 1990), however, their willingness to play freely with the block during the Action Phase and their success in the subsequent experiments argues against this construal.


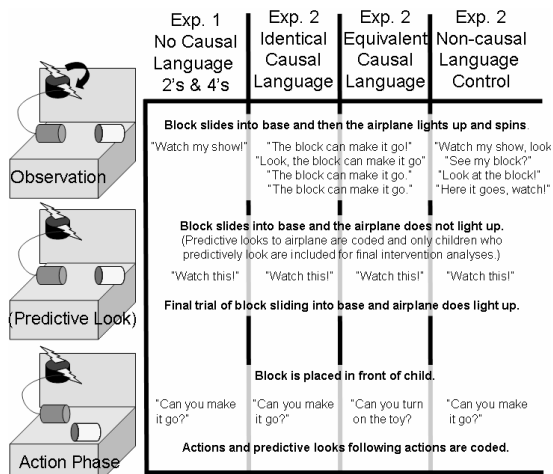


Figure 1: Method and Results for Experiments 1 and 2.

predictive relations and information about action. Describing the observed correlation (“The block makes the toy go”) with the same verb as the invitation to act (“Can you make the toy go?”) might help children recognize the relevance of observational evidence to their own interventions. That is, children may infer that the same words are used because they refer to the same underlying concept. Eventually children might thus come to recognize the possibility of causal relations among a broad range of non-agentive physical events (including cases where no verbal description of events is provided).

A weaker version of how language could affect children’s causal representations might suggest that children do form common representations of agentive and non-agentive correlations but fail to use predictive information as a basis for action if they have no additional reason (beyond the covariation evidence itself) to assume events are causally related. Many researchers have suggested that neither adults nor children draw causal inferences merely from covariation information (Ahn, Kalish, Medin, & Gelman, 1995; Koslowski, 1996; Shultz, 1982). In the absence of any knowledge about the causal mechanism that might underlie the observed relationship, children might realize that one event predicts another but fail to conclude that the events are causally related. If so, children would have no basis for action. Under this interpretation, the tendency to represent correlated events as viable candidates for intervention depends upon prior knowledge of a plausible mechanism

| Action: | Success (Predictive Look) | | Failure (No Predictive Look) | |
|-----------------|---------------------------|----------|------------------------------|----------|
| | Spontaneous | Prompted | Spontaneous | Prompted |
| Exp. 1: 4’s | 10 | 4 | 0 | 2 |
| Exp. 1: 2’s | 0 | 0 | 0 | 14 |
| Exp. 2: ICL | 7 | 1 | 5 | 3 |
| Exp. 2: ECL | 9 | 1 | 0 | 6 |
| Exp. 2: Control | 1 | 0 | 5 | 10 |

Table 1: Results of Experiment 1 & 2 by action type and success.

linking the events. In this case, causal language (by which we mean here language familiar to toddlers: “make go”, “turn on”) might facilitate causal learning by testifying that an observed relation is indeed causal. That is, children might treat events as causal (i.e., as supporting manipulation) simply because they are told that they are causal (that the block, rather than for instance, some unobserved common cause, does indeed make the toy go).

Although these accounts are theoretically distinct, they make a common prediction with respect to this experiment: if language facilitates young children’s ability to move from prediction to action – either by helping toddlers bind together correlational information with representations of their own actions, and/or because 24-month-olds can rapidly learn novel causal relations from testimony – then describing the events of the Observation Phase using causal language might improve children’s performance.

If so, we can then ask whether the facilitating effect of causal language is fragile and depends on using *precisely* the same words (“The block makes it go”; “Can you make it go?”) or whether language acts as a fairly robust cue to children’s causal learning and non-identical but semantically equivalent words suffice (“The block makes it go”; “Can you turn it on?”). However, if language merely improves children’s performance by increasing their attention to the events, then, relative to Experiment 1, toddlers’ performance should also improve in a non-causal language control condition where language is used merely to attract children’s attention (“Look at my block! Let’s watch my show! Here it goes!”).

Methods

Participants Forty-eight toddlers were assigned to each of three conditions; the *Identical Causal Language Condition (ICL)* (m=24.5mos; rng: 19-29mos), the *Equivalent Causal Language Condition (ECL)* (m=23.6mos; rng: 18-30 mos), and the *Non-Causal Language Control Condition (Control)* (m=23.6mos; rng: 18-30mos).

Materials & Procedure. The same materials used in Experiments 1 and 2 were used. The procedure was identical to the procedure in Experiment 1 except for the previously noted language changes, (See Figure 1).

Results & Discussion of Experiment 2

Coding and inclusion criteria were as with Experiment 1. Intercoader agreement was high across all conditions (*ICL*:

92%; *ECL*: 97%; *Control*: 98%). Three toddlers were excluded and replaced: 1 for failing to make the initial predictive look and 2 for failing to perform the target action.

Children in the *Control* condition replicated the failure of toddlers in Experiment 1; only 1 of the 16 toddlers (6%) succeeded at the task, intervening spontaneously and then looking predictively. *Control* children were no more likely to succeed on the task and no more likely to generate the action spontaneously (and anticipate the outcome) than the toddlers in Experiment 1 ($\chi^2 = .905, p = NS$).

However, the pattern of results reversed for children in the Identical Causal Language (*ICL*) condition and the Equivalent Causal Language (*ECL*) condition. Eight of the 16 toddlers (50%) in the *ICL* condition succeeded at the task and 10 of the 16 toddlers (62%) in *ECL* condition succeeded at the task. In both *ICL* and *ECL* conditions, the toddlers were significantly more likely to succeed on the task than the toddlers in Experiment 1 (*ICL*: $\chi^2 = 9.55, p < .01$; *ECL*: $\chi^2 = 13.1, p < .001$) and more likely than toddlers in the *Control* condition (*ICL*: $\chi^2 = 7.58, p < .01$; *ECL*: $\chi^2 = 11.2, p < .001$). The toddlers were also significantly more likely to generate the action spontaneously (and anticipate the outcome) than were the toddlers in Experiment 1 (*ICL*: $\chi^2 = 8.00, p < .01$; *ECL*: $\chi^2 = 11.2, p < .001$), and also more likely than the toddlers in the *Control* condition (*ICL*: $\chi^2 = 6.0, p = .01$; *ECL*: $\chi^2 = 9.3, p < .01$). There were no differences between the *ICL* and *ECL* conditions on either of these measures (Overall Success: $\chi^2 = .508, p = NS$; Spontaneous action success: $\chi^2 = .50, p = NS$), and no differences between the *ECL* condition and the preschoolers from Experiment 1 (Overall Success: $\chi^2 = 2.67, p = NS$; Spontaneous action success: $\chi^2 = 0.13, p = NS$). Additionally, toddlers in the *ICL* condition were just as likely to generate the action spontaneously (and anticipate the outcome) as preschoolers in Experiment 1 ($\chi^2 = 1.13, p = NS$), but were less likely than preschoolers to succeed overall ($\chi^2 = 2.67, p = .02$). (Table 1).

These results suggest that describing observed events with causal language supports children's ability to recognize that non-agentive events support manipulation. Toddlers who were given causal language cues were just as successful as the four-year-olds in Experiment 1. The effect of language was relatively robust, surviving minor changes in wording as long as the meaning was preserved. Critically, the effect of language is not entirely general; merely calling children's attention to events did not improve their performance.

Language in Representation and Processing

We proposed that causal language could improve children's performance either by helping children providing a common representation for predictive looking events and children's own actions, or by testifying that an observed relation is genuinely causal. However, the results of Experiments 1 and 2 are also consistent with a more deflationary account: the toddlers' failure might be one not of competence but of performance. Myriad developmental studies suggest that children's apparent understanding of a concept depends on whether the dependent measure involves looking or acting

(e.g. Hood, Cole-Davies, & Dias, 2003; Ahmed & Ruffman, 1998). Such results have led some researchers to suggest that intentional action might, *in general*, lag behind predictive looking: either because the demands of planning and executing motor responses interfere with children's ability to access task-relevant information (Baillargeon et al, 1990; Diamond & Goldman-Rakic, 1989; Thelen & Smith, 1994), or because stronger representations are necessary for acting than for looking (see Munakata, 2001 for review). Although there are important theoretical distinctions between these two claims, they are united in suggesting that a gap between children's ability to make successful predictions and their ability to perform effective actions might reflect changes in children's ability to manifest their knowledge under complex task demands. These accounts assume that difficulties in moving from prediction to action are due primarily to competing information processing demands. If so, any information (such as causal language) that strengthens the representation of a causal relationship might make it more likely to withstand the demands of planning and executing a motor act.

One way to look at whether causal language merely facilitates children's ability to move from prediction to action, or whether causal language genuinely changes children's underlying representations is to use a dependent measure of children's causal understanding that does not involve action. It has long been suggested that infants and children expect physical causal events to involve contact between the agent and patient. While this has primarily been demonstrated with respect to motion events (Leslie & Keeble, 1987; Oakes & Cohen, 1990), recent research suggests that children also assume that agent-initiated events resulting in an object's change of state (e.g., a hand moving towards a box and the box breaking, or making music) also require contact (Muentener, in review; Kushnir & Gopnik, 2007)². If toddlers require causal language to represent the predictive relationship between the block and the plane as causal, then in the absence of causal language, they should make no predictions about whether or not the block contacts the base; in the presence of causal language, they should assume the block contacted the base when the toy is on and did not contact the base when the toy is off.

Experiment 3

In Experiment 3, we use a violation of expectation paradigm to test whether children infer contact causality from predictive relations. Using the same stage as in Experiments 1 and 2, we slide the block behind an occluder that partially covers the base. On a final trial, we remove the occluder to reveal the block in contact or not in contact with the base.

If language helps children form a 'causal representation' of the events, then in the absence of causal language, children might fail to form a causal representation of the

² Critically, infants only seem to expect contact when the change of state is initiated by an agent (e.g., a human hand); consistent with our findings here, if the event is initiated by another object, the infants have no expectations about contact causality.

events (consistent with the results of Experiments 1 and 2) and should thus show no looking-time differences. By contrast, when causal language is used, children should look longer at the no contact case than the contact case when the plane activates but show the reverse pattern when the plane fails to activate. However if the processing account is correct and toddlers' failures in Experiment 1 and the *Control* condition of Experiment 2 are due only to the difficulty involved in initiating motor responses, then children should expect the block to contact the base when the toy activates even in the absence of causal language.

Methods

Participants Seventy-two toddlers were randomly assigned to one of six conditions ($m = 23\text{mos}$; $\text{rng: } 18\text{-}30\text{mos}$).

Materials & Procedure. Materials were identical to Experiments 1 and 2 with two exceptions: a blue box was used to occlude the track during the familiarization trials; a larger board occluded the whole stage between trials.

Children were assigned to the *Language Control* condition, the *Causal Language Toy On* condition or the *Causal Language Toy Off* condition. Within each of these conditions, half the children were assigned to the Contact outcome and half to the No-contact outcome. The experimenter first pointed out the elements of the stage to the child (the occluder, block, base, and airplane), then placed the occluder in front of the track so that the base was visible, but so the trajectory of the block was hidden.

During the Observation phase, the block slid toward the base, behind the occluder, and appeared to move on its own. On half the trials, following the block's movement, the plane activated, and on half it did not activate. Children's looks to the plane were coded on the non-active trials, and only children who predictively looked on at least one trial were included in analyses. After each trial, the whole stage was occluded by the board and the stage reset (the plane stopped, the block slid back). At the start of each trial, children in the *Language Control* were told "Look at this! Watch the block!" and children in the *Causal Language* conditions were told, "Look! The block can make it go!"

For the test trial, the procedure was repeated. Once the block moved behind the occluder, the plane activated in the *Language Control* and *Causal Language Toy On* conditions; it did not activate in the *Causal Language Toy Off* condition. The experimenter then removed the occluder revealing either the Contact or No Contact outcome and waved her hand non-specifically toward the stage saying, "Look at this!" to elicit the child's attention. The child's looking time was monitored through the camera to avoid influencing the child's attention. The trial ended when the child looked away from the stage for 2 consecutive seconds.

Results & Discussion of Experiment 3

Two toddlers were excluded and replaced for failing to make a predictive look during the Observation phase. We coded how long children looked at the stage following the

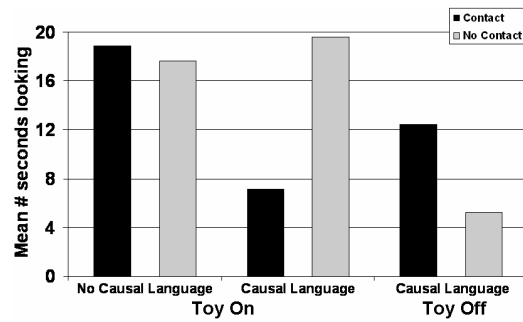


Figure 2: Looking results for Experiment 3 revealing no difference in looking between Contact and No Contact outcomes when children do not receive causal language, but showing children in the *Causal Language Toy On* condition looking significantly less following the Contact trial (Mean = 7.10) than children in the No-contact trial (Mean = 19.16) ($t(22) = -4.15, p < .001$), and a reverse of this pattern of looking for children in the *Causal Language Toy Off* condition (Contact Mean = 12.44; No Contact Mean = 5.25; $t(22) = 2.58, p < .01$).

removal of the occluder on the test trial. Inter-coder agreement was high (96%).

In the *Language Control* condition children were just as likely to look following the Contact test trial (Mean=18.98s) as children who observed the No-contact test trial (Mean = 17.61s) ($t(22)=0.24, p = NS$). We ran a two-way analysis of variance for independent samples on children in the *Causal Language* conditions, with activation (toy on or off) as the first between subjects variable and block-to-base relationship (Contact, No Contact) as the second. Comparisons between conditions revealed an effect of toy activation (averaging across Contact/No Contact outcomes, children looked less overall when the toy did not activate ($F = 4.31, p = .04$)), but there was no main effect of block-to-base outcome (averaging across the two conditions by toy's activation, children who saw the Contact outcome looked as long as children who saw the No Contact outcome). Importantly, comparisons revealed a significant interaction: children in *Causal Language* conditions spent less time looking at the stage when the evidence was consistent with contact causality than when it was inconsistent ($F(1, 47) = 21.79, p < .0001$). (See Figure 2).

When the toy activated but no causal language was used, children looked equally long at the stage whether the block was in contact with the base or not. This suggests that even when toddlers are freed from the necessity of making a motor response, they fail to form a causal representation of these events. By contrast, children in the *Causal Language* condition showed the predicted pattern of differential looking; toddlers seemed to expect the block to make contact with the base when the toy activated and to fail to make contact when the toy did not. This suggests that in the presence of causal language, children assume not only that predictive relations support intervention, but also that the events respect the principles of contact causality. These results are consistent with the Representation account, suggesting that language helps children form a genuinely causal representation of the events.

Conclusions

Despite having no difficulty learning the predictive relationship between events and performing the relevant action, toddlers succeeded at moving directly from prediction to action only when the observed events were described with causal language. Similarly, toddlers succeeded at moving from prediction to expectations consistent with contact causality only when the predictive relations were described with causal language.

As noted, our results are consistent with both a stronger and a weaker form of the representational change account. The stronger version suggests a genuine discontinuity between linguistically unsophisticated toddlers and older children (and potentially a similar discontinuity between adult human and non-human animals). Initially, infants and toddlers might recognize predictive relations, the ability to support intervention, and physical contact relations as independent features of events, but not have access to an adult-like concept of causality, (which entails all three). Providing a common term for relations with these features might help children develop an integrated representation.

The weaker account suggests that children have the same concept of causality as adults but expand their repertoire of what 'counts' as evidence for a causal relation. Infants, for instance, might treat only agent-initiated action and Michottean events as 'causal' and only gradually (aided by adult testimony) include a broader range of events.

Although these accounts are theoretically distinct, the distinction may diminish to the extent that young children's initial understanding of causality is constrained. We find it striking, for instance, that although the setup provided a compelling causal illusion to both adults and preschoolers, and even toddlers readily learned the predictive relationship, toddlers in Experiment 1 never once attempted to bang the block into the base, despite being prompted to "make the toy go" and having just seen the block repeatedly strike the base and the toy activate. If children initially only recognize a very limited set of relations as potentially causal, then learning more about particular causal relations may affect their representation of causality in general. Future work might investigate the extent to which very young children engage in such joint inferences (see e.g., Kemp & Tenenbaum, 2008; Schulz, Goodman, Tenenbaum & Jenkins, 2008 for related work). The present study, however, suggests that even during the course of our short task, causal language helps two-year-olds perform better than Pavlov's dog, by helping them integrate information about prediction, intervention, and contact causality.

Acknowledgments

Thanks to Paul Muentener, members of the Early Childhood Cognition Lab, the Museum of Science, Boston, and the Boston Children's Museum. Research supported from the Elizabeth Munsterberg Koppitz Fellowship from the American Psychological Foundation to E.B., and the McDonnell Foundation and James H. Ferry Fund grant to L.S.

References

- Ahmed, A., & Ruffman, T. (1998). Why do infants make A not B errors in a search task, yet show memory for the location of hidden objects in a non-search task? *Developmental Psychology*, *34*, 441-445.
- Ahn, W., Kalish, C.W., Medin, D.L., & Gelman, S.A. (1995) The role of covariation ersus mechanism information in causal attribution. *Cognition*, *54*(3), p. 299-352.
- Baillargeon, R., Graber, M., DeVos, J., & Black, J. (1990). Why do young infants fail to search for hidden objects? *Cognition*, *36*, 225-284.
- Blaisdell, A. P. (2008). Cognitive dimension of operant learning. In Roediger (Ed.), *Cognitive Psychology of Memory. Vol. 1 of Learning and Memory*. Oxford: Elsevier.
- Diamond, A., & Goldman-Rakic, P. (1989). Comparison of human infants and rhesus monkeys on Piaget's A-not-B task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research*, *74*, 24-40.
- Hood, B., Cole-Davies, V., & Dias, M. (2003) Looking and Search Measures of Object Knowledge in Preschool Children. *Developmental Psychology*, *39*(1), p. 61-70.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 1-31.
- Gopnik, A. & Schulz, L. (2004). Mechanisms of theory-formation in young children. *Trends in Cognitive Science*, *8*(8), 371-377.
- Kemp, C. and Tenenbaum, J. B. (2008). The discovery of structural form. *PNAS* *105*(31), 10687-10692.
- Koslowski, B. (1996). *Theory and evidence: The development of scientific reasoning*. Cambridge, MA: MIT Press.
- 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*, *44*, 186-196.
- Kushnir, T, Wellman, H., & Gelman, S., (in press). A self-agency bias in children's causal inferences. *Developmental Psychology*.
- Leslie, A.M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*, 265-288.
- Muentener, P. & Carey, S. (in review). Infants' causal representations of state change events.
- Munakata, Y. (2001) Graded representations in behavioral dissociations. *Trends in Cognitive Sciences*, *5* (7), 309-315
- Oakes, L.M., & Cohen, L.B., (1990) Infant perception of a causal event. *Cognitive Development*, *5*, 193-207.
- Shultz, T.R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, *47* (Serial No.194).
- Schulz, L.E., Goodman, N.D., Tenenbaum, J.B., & Jenkins, C.A. (in press). Going beyond the evidence: Abstract laws and preschoolers' responses to anomalous data. *Cognition*.
- Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29-56.
- Spelke, E. S. (2003). What makes us smart? Core knowledge and natural language. In D. Gentner and S. Goldin-Meadow (Eds.), *Language in Mind: Advances in the Investigation of Language and Thought*. Cambridge, MA: MIT Press.
- Thelen, E. & Smith, L. B. (1994) *A dynamics systems approach to the development of perception and action*. MIT Press.
- Tomasello, M. & Call, J. (1997) *Primate Cognition*, Oxford University Press, New York
- Woodward, J. (2004) Interventionist Theories of Causation in Psychological Perspective. In Gopnik & Schulz (eds), *Causal Learning*. Oxford University Press.