Computational Modeling of Backwards-Blocking Reasoning in Human Adults

Deon T. Benton (deon_benton@brown.edu)

Department of Cognitive, Linguistics, and Psychological Sciences, Brown University, Providence, RI 02908

David H. Rakison (rakison@cmu.edu)

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

Causal reasoning is a fundamental cognitive ability that enables humans to learn about the complex interactions in the world around them. However, it remains unknown whether causal reasoning is underpinned by a Bayesian mechanism or an associative one. For example, some maintain that a Bayesian mechanism underpins human causal reasoning because it can better account for backward-blocking (BB) and indirect screening-off (IS) findings than certain associative models. However, the evidence is mixed about the extent to which learners engage in both kinds of reasoning. Here, we report an experiment and several computational models that examine to what extent adults engage in BB and IS reasoning using the blinket-detector design. The results revealed that adults' causal ratings in a backwards-blocking and indirect screening-off condition were consistent with associative rather than a Bayesian computational model. These results are interpreted to mean that adults use associative processes to reason about causal events.

Keywords: causal learning; computational modeling; backwards blocking; associative reasoning; Bayesian inference

Introduction

There is perhaps no ability that is more central for learning about how the world works than reasoning about cause-and-effect relations. Despite general consensus among researchers about the importance of causal reasoning, there is less agreement among theorists about the mechanism or mechanisms that underpin this ability. For example, it is unclear whether domain-general mechanisms such as associative learning underpin causal reasoning or whether causal reasoning is grounded in more rational processes such as Bayesian inference. One empirical finding about which domain-specific and domain-general theorists have disagreed considerably concerns whether an associative mechanism or a Bayesian mechanism can explain the finding that human learners engage in backwards-blocking (BB) and indirect screening-off (IS) reasoning in which learners block, discount, and screen-off causally redundant cues when other cues unambiguously produce effects (e.g., Shanks, 1985; Sobel, Tenenbaum, & Gopnik, 2004). The aim of the present investigation was to examine whether, to what extent, and under what conditions adults engage in BB and IS reasoning, and to investigate whether a Bayesian analytical model or one of two associative computational models best fits the data. We focus on BB and IS reasoning because they bear directly on the question of whether an associative or a Bayesian

mechanism underpins causal reasoning in adults and because the empirical evidence for BB and IS reasoning is mixed.

The emergence of BB reasoning

The ability to reason about causal events is thought generally to emerge between 18 months and 5 years of age (e.g., Gopnik et al., 2001; Sobel & Kirkham, 2006). Although researchers have used a variety of paradigms to examine causal reasoning in childhood (for an early review see Bullock, Gelman, & Baillargeon, 1982), here we focus on research that has used the blinket-detector design. We focus on this paradigm for two reasons. First, this paradigm and versions of it have been used to evaluate children's and adults' causal reasoning abilities, which may support cross-study and between-age comparisons. Second, the notion that human reasoners use Bayesian inference to reason about causal events and to engage in causal discounting was developed in the context of the blinket-detector paradigm and in concert with key advances in computer science, machine learning, and statistics (for a review, see Gopnik et al., 2004).

In studies that have examined the issue of BB and IS reasoning in children and adults (e.g., Griffiths, Sobel, Tenenbaum, & Gopnik, 2011), participants typically are introduced first to a machine called the "blinket detector" and are told that the machine lights up and plays music when certain objects—namely, "blickets"—are placed on it but not when other objects are placed on it.

Children are then shown together that two novel objects, A and B, make the detector activate and then that object A either fails to activate the detector (i.e., AB+; A-; IS condition) or activates the detector when placed alone on it (i.e., AB+, A+; BB condition). Children in both conditions are then asked which of the two objects are blickets and to make the machine go by placing the appropriate object on the detector. The results of these studies revealed that 4-year-olds and adults, but not 3-year-olds (although see Sobel & Munro, 2009), categorized correctly object B in the IS condition and intervened by placing it on the detector, whereas in the BB condition 4-year-olds and adults categorized only object A as the blinket and intervened by placing it on the detector.

A reevaluation of the meaning of BB reasoning

Although researchers generally agree about the age at which BB and IS reasoning emerges, there is considerably less consensus among researchers about what is meant by, or how to measure, BB and IS reasoning. For example, Sobel et al. (2004) used the fact that a larger proportion of children chose

object B in the IS condition than in the BB condition to indicate the presence of BB and IS reasoning. Thus, according to Sobel et al. (2004), BB and IS reasoning represent the difference in B choices across the BB and IS conditions and, specifically, higher B choices or ratings in the IS condition than in the BB condition.

In contrast, others define BB and IS reasoning to mean differences in the causal ratings or treatment of object B between the BB and IS conditions and a similar control condition in which the causal efficacy of object A is not independently established (e.g. De Houwer, Beckers, & Glautier, 2002; Lovibond et al., 2003; Shanks, 1985). Despite the fact that these studies differ in their specifics, in each study participants were said to have engaged in BB reasoning if their causal judgements or choices of the redundant cue B in the BB condition were lower than those for the same object in a control condition that was identical to the BB condition in all respects except that one of the two cues was never shown alone (i.e., AB+; A+ vs. AB+). Likewise, participants were judged to have engaged in IS reasoning if their causal choices of object B were higher than those for the same object in the corresponding control condition (i.e., AB+; A- vs. AB+). Thus, according to these authors, lower causal ratings or choices of the redundant cue in the BB condition than of the same cue in the control condition constitutes evidence of BB reasoning whereas higher ratings of the same cue relative to the control condition constitutes evidence of IS reasoning.

Still others interpret BB and IS reasoning in terms of the predictions of simple Bayesian models (e.g., Griffiths et al., 2011). On this account, BB reasoning is conceptualized as a return to baseline in the causal rating of B between a midway rating phase and a final rating phase. In contrast, IS reasoning is conceptualized as a ceiling-level increase in the rating of object B but a floor-level decrease in the rating of object A between a midway rating phase and a post-rating phase (because A is shown unequivocally not to produce the effect).

One final group of theorists maintain that BB refers not to a decrease in the causal rating of the redundant cue between a midway and final rating phase but to an absolute decrease in the causal rating of the redundant cue between an initial rating phase and a final rating phase. For instance, in addressing the limitations of the traditional RW model, Van Hamme and Wasserman (1994) presented a modified version of the traditional RW model (i.e., the modified RW model) in which BB was conceptualized as a categorical drop in the rating of the redundant cause B between an initial rating phase (i.e., before the BB event was shown) and a final rating phase (i.e., after the entire BB sequence is modeled). Despite the fact that object B is treated identically across the BB and IS conditions, this model also predicts a linear increase to ceiling in the rating of A across the three rating phases.

The account that we endorse here is a combination of the latter frameworks. Specifically, we argue that BB reasoning implies the process of retrospectively reevaluating or discounting the final value of a redundant cue relative to its earlier value within the BB condition itself. In other words, BB reasoning implies a drop in the rating or choices of a

redundant cue within the BB condition itself between a final time point and earlier points. In contrast, IS reasoning implies the process of retrospectively reevaluating or screening off the treatment cue (A) in favor of the target cue (B) relative to their initial ratings; that is, IS reasoning implies a drop in the rating of A but an increase in the rating or treatment of B between an earlier time point and a final one. This is a better approach than using a between-condition comparison to evaluate BB and IS reasoning for a number of important reasons chief among those being that one need not make inferences about the existence of BB or IS reasoning based on comparing the causal ratings of A and B between the BB and IS conditions (for an argument, see McCormack, Butterfill, Hoerl, & Burns, 2009).

Present experiment

In light of the issues outlined above, the aims of the present experiment are threefold. The first aim—based on the mixed evidence for BB reasoning—was to determine whether adults engage in BB and IS reasoning using the Sobel et al.'s (2004) procedure. The rationale for testing adults was because neither adults' causal-reasoning abilities nor their ability to engage in BB reasoning have been tested in the original blicket-detector task. We also tested adults rather than children because, with few exceptions (e.g., Sobel et al., 2004), adults tend to be the target population when assessing BB and IS reasoning perhaps because adults possess improved information-processing abilities and generally have more real-world causal experience than children.

A second and related aim of this investigation was to implement pre-, mid-, and post-rating measures to determine whether participants' ratings of the candidate causes undergo changes between any of the three rating phases. This method of assessing both BB and IS reasoning represents a significant departure from previous methods because implementing multiple rating phases makes it possible to track how the ratings of various candidate causes change within the BB and IS conditions, without recourse to between-condition comparisons that should produce disparate evaluations based on cross-condition contamination.

The third and final aim of the investigation was to compare adults' causal ratings to the predictions of multiple competing computational models. This method of comparing adults' causal responses for competing causes to the predictions of different computational models is ideal because it is possible to determine the mechanism underlying adult causal reasoning abilities (cf., Griffiths et al., 2011). The three computational and analytical models that we will consider here are a simple Bayesian model, the traditional Rescorla-Wagner model, and the modified Rescorla-Wagner model.

We focus on the traditional RW model because it was among the earliest associative models of causal learning that failed to account for BB and IS reasoning, and this is true under all four interpretations of BB and IS reasoning outlined at the outset of this paper. On the basis of this failure, some have argued that a Bayesian inference mechanism must underlie human causal reasoning given that it can predict both

BB and IS reasoning especially under Sobel’s interpretation of BB and IS reasoning (e.g., Sobel et al., 2004; Griffiths et al., 2011). We focus on the modified RW model and a Bayesian model because unlike the traditional RW model, both models predict BB reasoning, albeit in slightly different ways. Indeed, although both models interpret BB as a drop of some kind between at least two different rating phases, the modified RW model interprets BB as an absolute drop in the rating of B between an initial and final rating phase, whereas a simple Bayesian model interprets BB as a return to baseline between a midway and final rating phase.

In the study, participants were introduced to a machine called the blicket detector and were told that their task was to determine which of two objects makes the machine activate by lighting up and playing music. Participants then received two trials—a backwards-blocking trial and an indirect screening-off trial—and were asked to provide three sets of causal ratings for objects A and B across a pre-, mid-, and post-rating phase.

Method

Subjects. Twenty-four college students were recruited for this experiment and were given course credit for their participation (11 males; 13 females).

Stimuli and Design. The device used in this experiment was similar to the blicket detector that was used in previous studies (e.g., Gopnik et al., 2001). The box was 12.7 cm x 17.78 cm x 7.62 cm, made of wood that was painted black, and had a white, transparent top. The machine operated via a remote control that was attached to the end of an electric wire that was attached to the side of the box. When the button on the remote control was pressed and the object that was predetermined to be a blicket was placed on the surface of the detector, the music began to play and the lights began to flash. If an object that was not predetermined to be a blicket was placed on the detector, the button was not pressed and the lights and music did not flash or play.

Adults were asked to provide three sets of causal ratings on the same 0-100 scale during the pre-, mid-, and post-rating phases for two objects (A and B). The stimuli for this experiment were eight cube and cylinder objects (four of each), each approximately 2.54 cm in diameter and of different colors. Within each trial, the objects were of different shapes and colors and the object that was designated as the blicket was counterbalanced across participants. Two unrelated objects—which were neither cubes nor cylinders—were used in the pretest phase of the experiment.

Procedure. Participants were tested in a quiet room on campus. At the beginning of the experiment, participants were introduced to the machine and were told it was called a blicket detector that activated only when objects that were predetermined to be blickets were placed on it. Participants were then informed that their goal was to determine which objects were blickets and which objects were not based on the pattern of activation observed during the experiment.

Following this introduction, participants were given two test trials in counterbalanced order. The trials were the backwards-blocking (BB) and indirect screening-off trials (e.g., Gopnik et al. 2001; Sobel et al., 2004).

In the BB trial, objects A and B were placed together on the machine twice, which caused the machine to activate both times. Object A was then placed alone on the machine, which once more caused the machine to activate (i.e., AB++; A+). The IS trial was similar to the BB trial except that object A failed to cause the machine to activate when it was placed alone on it (i.e., AB++; A-). In both test trials, participants were instructed to rate on a scale that ranged from 0 (definitely not) to 100 (definitely is) that each object in the pair was a blicket both before, midway, and after a trial for a total of three causal ratings. The exact points at which participants were asked to provide the three sets of ratings for objects A and B is shown below in Table 1.

Schematic of the Experiment					
	Pre		Mid		Post
BB	*	AB++	*	A+	*
IS	*	AB++	*	A-	*

Table 1. The three points at which participants were asked to provide pre-, mid-, and post-ratings in the across the four conditions. The asterisks correspond to where in a particular sequence participant were asked to provide causal ratings of objects A and B. The number of +/- signs corresponds to the number of times a particular sequence was demonstrated and whether the effect happened (+) or did not happen (-).

Results

All analyses were conducted in R. Figure 1 shows the mean pre- and post-causal ratings of objects A and B across the four conditions. Given evidence of non-normality and unequal variance in the causal-rating data, we used non-parametric analyses with 4,000 replications each for hypothesis testing and to estimate confidence intervals. In particular, the Shapiro-Wilks test indicated that the data were not statistically normally distributed, all p ’s < .05. Likewise, the Levene’s test indicated heterogeneity of variance, $F(3, 572) = 15.92, p < .001$.

Figure 1

Indirect screening-off. The first analysis examined whether participants engaged in IS reasoning. A mixed-effects model was fit to participants’ causal ratings in the IS condition. Participants were included as a random-effect factor and phase (pre vs. mid vs. post) and object type (A vs. B) were included as within-participants fixed-effects factors. This analysis yielded main effects for Phase, $F(2,115) = 9.32, p < .001$ and Object Type, $F(1,115) = 260.49, p < .0001$,

which were qualified by a Phase x Object Type interaction, $F(2,115) = 189.27, p < .0001$.

Follow-up planned comparisons using permutation tests with 4,000 simulations for object A revealed that although participants' pre-ratings of object A ($M = 47.58$, 95% CI [42.09, 53.07]) did not differ from their mid-ratings of object A ($M = 56.13$, 95% CI [49.42, 62.83]), $p = .82$, 95% CI [-17.26, 0.18], participants provided higher pre-ratings of A than post-ratings of A ($M = 1.88$, 95% CI [-0.38, 4.13]), $p < .0001$, 95% CI [39.75, 51.66] and higher mid-ratings of A than post-ratings of it, $p < .0001$, 95% CI [47.11, 61.39].

Planned comparisons for object B revealed that although participants provided lower pre-ratings of B ($M = 51.33$, 95% CI [44.05, 58.62]) than post-ratings of B ($M = 97.08$, 95% CI [94.14, 100.02]), $p < .0001$, 95% CI [-53.58, -37.92] and lower mid-ratings of B ($M = 62.79$, 95% CI [55.48, 70.1]) than post-ratings of it, $p < .0001$, 95% CI [-42.12, -26.47], participants' pre- and mid-ratings of B did not differ.

Backwards-blocking. The final analysis examined whether adults engaged in BB reasoning. Similar to the preceding analysis, a mixed-effects model was fit to participants' causal ratings in the BB condition. Participants were included as a random-effect factor and phase (pre vs. mid vs. post) and object type (A vs. B) were included as within-participants fixed-effects factors. This analysis yielded main effects for Phase, $F(2,115) = 17.21, p < .0001$ and Object Type, $F(1,115) = 21.61, p < .0001$, which were qualified by a Phase x Object Type interaction, $F(2,115) = 29.29, p < .0001$.

Planned comparisons for object A revealed that although participants provided lower pre-ratings of A ($M = 49.38$, 95% CI [42.33, 56.42]) than post-ratings of A ($M = 98.54$, 95% CI [96.35, 100.73]) and lower mid-ratings of A ($M = 57.5$, 95% CI [47.83, 67.17]) than post-ratings of it, $p < .0001$, 95% CI [-51.02, -31.06], participants' pre- and mid-ratings of A did not differ. A final set of planned comparisons for object B revealed that ratings of object B did not differ across the three rating phases, all p 's $> .2$. The goal of the next section is to report the predictions of each model to facilitate model comparison and to determine with which model participants' causal ratings of objects A and object B in the BB and IS conditions were consistent.

The Bayesian model. A fundamental assumption of the Bayesian inference approach is that causal induction is a process that involves representing the entire space of candidate causal hypotheses—which are expressed as parameterized causal graphical models with nodes that are connected by edges that encode the Markov condition—and then choosing the hypothesis that is most consistent with the data by application of Bayes' rule.

Formally, it is assumed that at the beginning of a task an ideal Bayesian learner represents all possible candidate hypotheses, H , whereby each hypothesis, $h \in H$, is assigned some prior probability, $p(h)$. This prior probability represents the learners' confidence that the observed data were generated by a given causal hypothesis. Following

observations of data, d , the learner then uses Bayes' rule to compute and assign posterior probabilities to each hypothesis, $p(h|d)$,

$$p(h|d) = \frac{p(d|h)p(h)}{\sum_{h'} p(d|h')p(h')},$$

where $p(d|h)$ represents the likelihood or the probability of the data d under a given hypothesis h normalized by members of the set. Because the hypotheses in this experiment are assumed to be deterministic (i.e., objects either produce or do not produce detector activation), the likelihoods are set to 1 whenever a link (i.e., causal relation) exists in the hypothesis (Figure 2) and is consistent with the observed data; otherwise, they are set to 0.

The first step in defining a model of the task presented here was to specify the hypothesis space H and the hypotheses h that comprise that space. This step is necessary before Bayes' rule can be used to determine the hypothesis with the largest posterior probability. Because this experiment uses two objects, the hypothesis space for this experiment consists of

Figure 2

four hypotheses (see Figure 2 below).

The specific parameterization of each hypothesis in the space is specified by the activation law, which states that the blinket detector will activate if, and only if, a blinket object contacts it. The second step in defining this model is to specify the prior probabilities of the four hypotheses in this experiment. If we assume that the probability that a particular object is a blinket is independent of the probability that other objects are blinkets, then prior probabilities for this experiment for both the BB and IS conditions can be found below in Table 2. These prior probabilities are then used to compute posterior probabilities for each hypothesis when new data is observed according to Bayes' rule.

Lastly, because the data that are presented here are deterministic, whenever a causal link exists the likelihood is set to 1 and 0 otherwise. Once we have determined whether such a link exists for a particular object, we can compute the likelihood that object A or B is a blinket in the present experiment by taking the product of the likelihood that a particular object activated the detector under each hypothesis (according to the aforementioned procedure) and the prior probability of each hypothesis and then summing this product. To determine the probability that object B is a blinket, for example, we can compute the following equation:

$$p(B \rightarrow E|d) = \sum_h p(B \rightarrow E|h)p(h),$$

where $p(B \rightarrow E|d)$ is 1 if a causal link between B and E exists for a specific hypothesis h ; otherwise, $p(B \rightarrow E|d)$ is 0.

Below are the predictions of this model for adults' treatment of objects A and B across the three rating phases.

	Pre	Mid	Post
Model 0	p^2	$p/(2-p)$	p
Model 1	$p(1-p)$	$(1-p)/(2-p)$	$1-p$
Model 2	$p(1-p)$	$(1-p)/(2-p)$	0
Model 3	$(1-p)^2$	0	0
Probabilities of objects A and B			
	Pre	Mid	Post
Object A	p	$1/(2-p)$	1
Object B	p	$1/(2-p)$	p

Table 2. Bayesian model predictions for objects A and B as well as the four models in Figure 1.

Rescorla-Wagner model. In contrast to the Bayesian model, the traditional Rescorla-Wagner (RW) model (Rescorla & Wagner, 1972) uses error correction to produce causal ratings of the test events. On this model, error correction is computed according to the following equations:

$$\Delta V_i = \alpha_i \beta_1 (\lambda - \Sigma V_k)$$

and

$$V^{n+1} = V_i + \Delta V_i,$$

where ΔV_i is the change in the associative weight between the candidate cause i and effect e , α is the salience of the candidate cause i , β is the salience of the effect e , V is the current associative strength between the candidate cause i and effect e , and ΣV_k is the summed associative strength between all candidate causes i to k and the effect e . The predictions of this model and the modified Rescorla-Wagner model (below) were obtained via simulation.

Note that because the RW model does not inherently espouse a parameter that accounts for participants' baseline beliefs about causes (prior to the experiment) we set the weight for A and B to equal .5 to reflect that fact participants' pre-ratings of both objects did not differ reliably from 50%. Below are the predictions of this model.

	Pre	Mid	Post
Object A	p	p	1
Object B	p	p	p

Table 3. RW model predictions for objects A and B across the three rating phases.

Modified Rescorla-Wagner model. Despite similar learning rules, the Rescorla-Wagner and modified Rescorla-Wagner models differ in their treatment of the salience parameter, α . Whereas α is set to 0 for absent cues in the traditional RW model because such cues are assumed to have no salience, α is set to a nonzero (negative) value because absent cues are assumed to be negatively correlated with the outcome in the modified RW model. The training procedure

for this model was identical to that for the traditional RW model. Below are the predictions of this model.

	Pre	Mid	Post
Object A	p	p	1
Object B	p	p	0

Table 4. The modified RW model predictions for objects A and B across the three rating phases.

It should be clear from these predictions that the sole model with which adults' causal ratings were consistent was the traditional RW model: adults' causal ratings were not consistent with the predictions of the Bayesian model. This is evidenced by the fact that, in the BB condition, adults' causal ratings of objects A and B were consistent with the predictions of the traditional RW model across all three rating phases. As such, the traditional RW model, but neither the simple Bayesian model nor the modified RW model, provides the best fit to the present BB causal-rating data.

Evidence of associative reasoning for the IS condition.

Given that adults' causal ratings of object B did not change across the three rating phases of the BB condition, the results presented thus far suggest that adults do not engage in BB reasoning in a manner that is consistent with the predictions of a simple Bayesian model. However, before it can be concluded definitively that adults do not rely on a Bayesian mechanism to reason about causal events in the context of two objects, it is necessary to show that their causal ratings of objects A and B also do not conform with the predictions of the Bayesian model for the IS condition across the three rating phases. Thus, similar to above, we present the predictions of the three models, but only for objects A and B.

Table 5. The Bayesian model predictions for IS

	Pre	Mid	Post
Object A	p	$1/(2-p)$	0
Object B	p	$1/(2-p)$	1

Table 6. The traditional RW model predictions for IS.

	Pre	Mid	Post
Object A	p	p	0
Object B	p	p	p

Table 7. The modified RW model predictions for IS.

	Pre	Mid	Post
Object A	p	p	0
Object B	p	p	1

As can be seen from Tables 5-7 above, participants' ratings of both objects did confirm with the predictions of the modified RW model but not with the predictions of the simple Bayesian model or the traditional RW model. There are two reasons for this. First, participants' pre- and mid-ratings of objects A and B did not differ from each other in both the BB and IS conditions. Second, participants' ratings

of object B increased between the mid- and post-rating phase in the IS condition but remained unchanged across all three rating phases in the BB condition. The distribution of responses to object A also accorded with the predictions of the modified RW model, but not with the predictions of the other two models, across all three rating phases in both the IS and BB conditions. These results suggest that, together, the traditional RW model and modified RW model provided a better fit to the data than did the Bayesian model, although neither model alone accounted for adults' causal-rating performance in both conditions. Nonetheless, the results from Experiment 1 suggest that an associative mechanism, rather than a Bayesian mechanism, may underlie adult causal reasoning for two objects.

General Discussion

The present investigation examined to what extent adults engaged in BB and IS reasoning and tested the predictions of the three computational models for a causal task that used two objects. The results showed that adults' responses in the BB and IS conditions were consistent with the predictions of the modified and traditional RW models but not with the Bayesian model. Taken together, these results indicate that adults did not engage in BB reasoning; that is, they do not block object B such that B undergoes a drop between any two of the three rating phases as would be predicted either by the modified RW model or a simple Bayesian model. That adults' ratings did not conform with the predictions of the simple Bayesian model further suggests that a Bayesian-inference mechanism may not underpin adult causal reasoning in a blinket-detector context that consists of two objects. It will be recalled that if such a mechanism did underpin adult causal reasoning, then their rating of object B should have increased between the pre- and mid-rating phases but then return to baseline (i.e., the pre-rating value) between the mid- and post-rating phases. Likewise, if this mechanism underpinned adults' processing of the IS event, then their ratings of objects A and B should have increased between the pre- and mid-rating phase. However, their ratings of object A should have dropped between the mid- and post-rating phases, whereas their ratings of object B should have increased between these two phases. One plausible objection to this conclusion is that the reason the adults did not engage in BB reasoning was simply because the experiment was underpowered. However, this criticism is weakened by the fact that an a priori power analysis for a *t*-test indicated that 16 participants would be sufficient to detect a medium-size effect assuming that such an effect existed. This conclusion is also weakened by the fact that a Bayes' factor analysis indicated that the observed differences between participants' pre-, mid-, and post-ratings of objects A and B represented substantial evidence in support for the null than under the alternative hypothesis.

This experiment represents a key theoretical contribution to studies of BB reasoning. Our argument is that when BB is taken to represent a drop in the rating of causally redundant cues between any two phases, BB is a phenomenon that can

be accounted for by associative models but not by a Bayesian model. This extends previous theorizing on this topic that has maintained that associative models cannot account for BB reasoning presumably because models such as the traditional RW model fail to account for the differential treatment of object B across the BB and IS conditions (e.g., Sobel et al., 2004) or the observed drop in rating of object B between a mid- and post-rating phase (e.g., Griffiths et al., 2011), among other things. The present results show that, in fact, adults do not treat objects A and B differently when a different metric is used to assess BB and IS reasoning. Indeed, when formulated as a drop in the rating of a redundant cause between any two rating phases, it is clear that adults do not engage in BB reasoning and that their causal inferences may be subserved by an associative mechanism.

References

- 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(5), 620.
- 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(1), 3.
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and blinkets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*, 35(8), 1407-1455.
- Lovibond, P. F., Been, S., Mitchell, S.J., Bouton, M.E., Fronhardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory and Cognition*, 29(1), 97.
- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *The Quarterly Journal of Experimental Psychology Section B*, 37(1b), 1-21.
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive science*, 28(3), 303-333.
- Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and motivation*, 25(2), 127-151.
- McCormack, T., Butterfill, S., Hoerl, C., & Burns, P. (2009). Cue competition effects and young children's causal and counterfactual inferences. *Developmental psychology*, 45(6), 1563.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2, 64-99.
- Sobel, D. M., & Kirkham, N. Z. (2006). Blinkets and babies: the development of causal reasoning in toddlers and infants. *Developmental psychology*, 42(6), 1103.