

The importance of decision making in causal learning from interventions

DAVID M. SOBEL

Brown University, Providence, Rhode Island

and

TAMAR KUSHNIR

University of Michigan, Ann Arbor, Michigan

Recent research has focused on how interventions benefit causal learning. This research suggests that the main benefit of interventions is in the temporal and conditional probability information that interventions provide a learner. But when one generates interventions, one must also decide what interventions to generate. In three experiments, we investigated the importance of these decision demands to causal learning. Experiment 1 demonstrated that learners were better at learning causal models when they observed intervention data that they had generated, as opposed to observing data generated by another learner. Experiment 2 demonstrated the same effect between self-generated interventions and interventions learners were forced to make. Experiment 3 demonstrated that when learners observed a sequence of interventions such that the decision-making process that generated those interventions was more readily available, learning was less impaired. These data suggest that decision making may be an important part of causal learning from interventions.

Causal knowledge plays an important role in everyday reasoning. Causal knowledge enables us to predict future outcomes, explain past events, control the present environment, and categorize novel information. A great deal of research suggests that both children and adults easily represent causal relations among events (Gopnik & Meltzoff, 1997; Murphy & Medin, 1985; Sperber, Premack, & Premack, 1995; Wellman & Gelman, 1992).

How are causal relations learned? Correlations among events can often be good indicators of the presence of some causal relations, but it is well known that observed associations are often insufficient to disambiguate causal structure (see, e.g., Bacon, 1620/1986). For this reason, much of causal learning in both formal (e.g., scientific) and informal settings takes place in the context of interventions. Interventions provide learners with conditional probability information that is often vital to causal learning (for formal accounts, see Pearl, 2000; Woodward, 2003). Interventions also provide learners with anticipatory information in the form of temporal priority: The variable intervened on should be considered a cause of any subsequent or concu-

rent event (see, e.g., Lagnado & Sloman, 2004). Several researchers have shown that people treat their own and others' intentional manipulative actions as interventions (in a formal sense) and learn causal relations better from these data than from observational data alone (Gopnik et al., 2004; Kushnir, Gopnik, Schulz, & Danks, 2003; Lagnado & Sloman, 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann & Hagmayer, 2005).

Research on causal learning from interventions has often focused on how interventions lead to causal structure learning in simple scenarios involving two or three related events. Causal learning in the real world often involves learning a complex network of relations among many events. In such situations, there may be an additional benefit of interventions—particularly, a learner's own interventions—in that learners can be active in their learning process. Interventions can provide learners with certain decision-making demands: To learn from interventions, one must first decide what intervention to make. This decision-making process allows learners to use interventions to disambiguate particular causal structures—namely, those that they have in mind as potential models of the causal system.

The idea that active examination of data facilitates learning more than passive observation of the same data does is not novel. In perception, active observers—those who can control what aspects of the environment they perceive—make more accurate judgments about depth and changes to the environment than do those who simply view a scene (see, e.g., Larish & Andersen, 1995; see also Gibson, 1979). In developmental psychology, both infants and young children are more sensitive to information garnered

A portion of this research was presented at the 2003 meeting of the Cognitive Science Society. The authors thank Alisa Breetz, Emily Blumenthal, Victoria Mazgalev, Kathryn Noe, and Kyle Schmitt for assistance with data collection and Tom Griffiths, Amy Hoff, David Lagnado, Laura Schulz, Steven Sloman, Mark Steyvers, Sean Stromsten, Joshua Tenenbaum, and the reviewers and editors for helpful discussions and/or comments on the manuscript. Correspondence concerning this article should be addressed to D. M. Sobel, Department of Cognitive and Linguistic Sciences, Brown University, Box 1978, Providence, RI 02912 (e-mail: david_sobel_1@brown.edu).

from their own interventions than to that garnered from their observations when making inferences about intentions, problem solving, and causality (Fireman, Kose, & Solomon, 2003; Kushnir & Gopnik, 2005; Sommerville, Woodward, & Needham, 2005). In cognitive psychology, many researchers have suggested that hypothesis testing is an important aspect of decision making and rule learning (e.g., Klayman, 1988; Klayman & Ha, 1987; Wason, 1968). We suggest that decision making may have similar benefits for causal learning.

The following experiments were based on a novel paradigm for causal learning in which we could manipulate whether learners (1) intervened themselves, (2) observed data from another's interventions (while also observing the temporal information inherent in those data), or (3) were forced to make a particular set of interventions. In Experiment 1, we examined the importance of decision-making abilities over conditional probability and temporal information alone by comparing Conditions 1 and 2. In Experiment 2, we examined the effect of decision making above and beyond kinesthetic or reinforcement effects by comparing Conditions 1 and 3. Both of these experiments employed a yoked design; the participants were matched on the actual data they observed. If decision making contributes to causal learning, we would expect that the participants in Condition 1 would be the most accurate at learning causal models.

Why might decision-making abilities improve learning? One possibility is that learners who can make decisions are able to focus on the particular interventions and outcomes that are important for their learning the causal structure. In Experiment 3, we gave people the minimal amount of intervention data necessary to specify a unique causal structure. We manipulated how accessible the observed interventions were to the learner by changing the order in which they were presented. In one condition, the learners observed data presented in a structured way; they first observed all of the dependence information and then all of the conditional independence information that would specify a unique causal model. In the other, the learners observed the same data, but in a random order. If learners benefit from focusing on interventions and outcomes relevant to learning a causal structure, having data presented in this systematic order should result in more accurate learning than would having the same data presented in a random order.

EXPERIMENT 1

Participants were asked to learn a set of causal systems in one of two conditions. In the *self-intervention* condition, the participants were able to generate interventions on the causal system. In the *observation-of-intervention* condition, the participants were asked to learn by observing another person make the same interventions as those made by the learners in the self-intervention condition. The interventions and outcomes were identical between the conditions. The participants in the observation-of-

intervention condition watched an animation of the intervention being performed; thus, the learners in both conditions had the same temporal information inherent to the interventions. What differed between these two conditions was the presence of decision-making abilities.

The learning environment contained four variables. The cover story indicated that any variable could be a cause or an effect of any other variable (following a procedure used by Sophian & Huber, 1984). The learning task was designed to ensure the most accurate learning possible. Each of the participants in the self-intervention condition controlled the number of interventions that he or she saw. The participants in both conditions were allowed to take notes. Also, the participants were allowed to perform two types of interventions: generative interventions and inhibitory interventions. This enabled them to learn causal structures that could not be fully determined by generative interventions alone, such as causal chains. Finally, in order to simplify the causal structure questions, the participants were asked to judge the causal relationship between every possible pair of events, as opposed to choosing a single representation of the causal structure.

Method

Participants

Thirty-four undergraduates were recruited from an urban area university. Approximately equal numbers of men and women participated in the experiment. They were compensated \$7 for their participation.

Materials

The participants were tested on a Dell Dimension 8100 desktop computer with a 19-in. monitor. The participants ran a Visual Basic program, which presented visual stimuli and asked the participants for responses (see below).

Procedure

The participants were seated at the computer and then given the following instructions:

In Dr. Science's laboratory, he has created a number of games. Each game has four lights, colored red, white, blue, and yellow. Each light also has zero, one, or many sensors. Some sensors are sensitive to red light, others to blue light, others to white light, and others to yellow light and sensors will activate the light that it is connected to. For example, if the red light is connected to a yellow sensor, then whenever the yellow light activates, the red light will also activate. But, because this happens at the speed of light, all you will see is the red and yellow lights activating together. It is also possible that a light has no sensors attached to it, and therefore is not activated by any other light. Sometimes Dr. Science is very careful about how he wires the lights together. At other times, he is not as careful, and the sensors do not always work perfectly.

The participants were assigned to one of two conditions, using an alternation strategy to employ a yoked procedure. The 1st participant and subsequent odd-numbered participants were assigned to the self-intervention condition. They were told that Dr. Science wired up four buttons, each of which turned on one of the lights. In the learning phase, the participants were allowed to press the buttons at their own pace as many times as they wanted and to observe the results. When the participants pressed a button that turned on a light, it and any light caused by its activation appeared on the screen for 0.5 sec. The participants were also told that Dr. Science had given them a black bucket, which could be placed over any of the lights (and its sensors), thereby eliminating that light from the model. The

learners could also control the location of the bucket. The participants were instructed to turn lights on and manipulate the bucket until they believed that they knew the configuration of the lights and sensors. The participants were required to turn on lights at least 25 times and could not end the learning phase until they had done so.

The remaining participants were assigned to the observation-of-intervention condition. They were given similar instructions but were told that they would watch Dr. Science turn on the lights and move the bucket. The same buttons that made the interventions were on the screen, and the participants read what intervention was being made as they watched an icon representing Dr. Science physically move to each button and press it. The icon moved relatively slowly; it took approximately 3–5 sec for the icon to reach the particular button that was about to be pressed. This ensured that the learners were able to read what intervention was about to happen before they actually observed the intervention. The expectation that a certain light (i.e., the light intervened on) would activate on each trial was kept constant between the two conditions. Thus, the temporal information inherent to interventions was equivalent between the two conditions. For each participant in the observation-of-intervention condition, the interventions of Dr. Science and the subsequent data were yoked to those of a participant in the self-intervention condition. Thus, the conditional probability information that the participants observed between the two conditions was also identical.

Training phase. In both conditions, after the participants had read the instructions but before they were allowed to make or observe interventions, they received a short training session in using the bucket. They were shown three lights (green, purple, and gray) and three buttons, each of which activated one of the lights. They were told that in this circumstance, Dr. Science had wired these lights and sensors perfectly. They were shown, through a series of generative interventions, that purple caused green and that gray caused purple but were unable to tell whether gray was also a direct cause of green. They were then shown that by putting the bucket on purple and activating the gray light, they could distinguish whether gray was a direct cause of green or only an indirect cause of green (via the purple light). This intervention sequence revealed that gray did not cause green and, thus, that no direct link was present.

Test phase and questions. The participants were told that their goal was to learn the causal structure of the lights and sensors. Before learning each model, they were reminded that the configuration could have changed since the last learning phase. They were also reminded that sometimes the sensors did not work perfectly. This was done to ensure that the participants recognized that the causal relations were not deterministic.

The participants were asked to learn four different wirings of lights and sensors: a chain model (Figure 1A), a chain model with an A→D link (Figure 1B), a diamond model (Figure 1C), and a common effect and chain model (Figure 1D). These models were chosen because all of them possessed indirect causal links: They could not be learned only from dependence information; conditional probability information was needed to specify each of these models. Colors for the four lights were randomly assigned. Each sensor caused activation 80% of the time. When a sensor failed to fire, the light it was connected to did not activate. The models were presented in one of four quasirandom orders, counterbalanced across participants. To aid in memory, the participants were allowed to take notes during the test phase and to refer to those notes during the test questions.

After the participants had indicated that they had made a sufficient number of interventions to learn the causal structure of the model (or had observed the same data), they were asked a set of *causal structure* questions. The participants were asked whether each light had a sensor for each other color attached to it. The 12 causal structure questions were asked in a random order. These were asked as a set of *yes/no* questions: “Was there an x sensor on the y light?” To eliminate any potential errors in interpreting this question, a clarifying intervention statement accompanied each ques-




Figure 1. The causal models used in the two experiments. The four lights were randomly assigned for each model. Links were probabilistic with a weight of .8.

tion: “Put another way, if there were only an x light and a y light, if the x light activated, would the y light also activate?” The participants first answered this *yes/no* question, then rated their level of confidence in their answer on a scale of 0–100, with 0 being a *total guess* and 100 being *fully confident*.

Results and Discussion

The participants received a score of 1 for each causal structure question they answered correctly and a score of 0 for each question they answered incorrectly. These scores were summed to reflect an overall accuracy score for each model (maximum = 12). These scores are shown in Table 1. We conducted preliminary analyses to exclude the influence of order; an ANOVA revealed that across the two source conditions, the order in which the models were presented did not influence performance.

A 4 (model) \times 2 (source condition) mixed ANOVA was performed, with model as a within-subjects factor and source condition as a between-subjects factor. A main effect of source condition was found: Across the four models, the participants were more accurate in the self-intervention condition than in the observation-of-intervention condition [$F(1,32) = 8.07, p < .01$]. No significant effect of model was found. No significant interactions were found. For all four models, the learners in the self-intervention condition were more accurate than those in the observation-of-intervention condition [$t(32) = 2.07, 2.51, 2.28, \text{ and } 2.92$ for the chain, diamond, chain with an A→D link, and common effect and chain models, respectively; all p values <

Table 1
Accuracy on the Causal Structure Questions and Percentages of the Learners Who Responded Correctly (%C) on All 12 Questions in Experiment 1

Model	Condition					
	Self-Intervention			Observation of Intervention		
	Accuracy		%C	Accuracy		%C
<i>M</i>	<i>SD</i>	<i>M</i>		<i>SD</i>		
Chain	11.18	1.29	59	9.88	2.23	35
Diamond	11.06	0.75	24	9.47	2.50	24
Chain with A→D link	11.06	1.20	53	9.35	2.85	24
Common effect and chain	11.65	0.86	82	9.18	3.38	35

Note—Accuracy responses are out of a maximum of 12.

.05]. A similar analysis was run on the participants' confidence scores. This analysis paralleled the accuracy analysis. As a result, confidence scores will not be discussed further.

We also examined these data by using a particular non-parametric analysis: whether the learners responded correctly on all 12 causal structure questions. This pattern of response is similar to choosing the correct causal model on a model select question, which was the test question used in previous research (Lagnado & Sloman, 2004; Steyvers et al., 2003). These data are also shown in Table 1. Looking across the four models, the participants responded with 100% accuracy 54% of the time in the self-intervention condition and 29% of the time in the observation-of-intervention condition. We set a criterion level of 100% accurate responses (i.e., scores equal to 12) on three out of the four models. More learners in the self-intervention condition met this criterion level than did learners in the observation-of-intervention condition [53% vs. 18%; $\chi^2(1, N = 34) = 4.64, p < .05$].

We next analyzed the number and type of interventions each participant in the self-intervention condition made (and hence were observed in the observation-of-intervention condition). These data were coded in two ways. First, we counted the total number of interventions for each learner and model. Second, we calculated the total number of interventions that offered the learner the particular conditional probability information necessary to discern that unique causal structure (deemed a critical intervention). For instance, in learning a chain model (Figure 1A), one must recognize that Light A activates independently of Lights C and D, given the absence of Light B. Thus, we counted the cases in which the participants activated Light A with the bucket on Light B as one example of a critical intervention. Table 2 shows the average number of interventions and the percentage of those interventions that were critical.

The number of interventions made by the participants did not correlate with their accuracy on the causal structure questions. The percentage of critical interventions also did not correlate with accuracy on the causal structure questions when the data from the two source conditions were combined. However, when the conditions were

examined separately, positive correlations were found. In the self-intervention condition, significant positive correlations were found between the percentages of critical interventions each participant generated and accuracy on the causal structure questions for the chain model [$r(17) = .470, p = .05$], the common effect and chain model [$r(17) = .782, p < .05$], and the chain with an A→D link model [$r(17) = .774, p < .05$]. These correlations were not found in the observation-of-intervention condition. These correlations do not appear to be the result of ceiling performance on the part of the participants in the self-generated intervention condition. On three out of the four models, the learners in this condition performed significantly lower than ceiling [all $t(16)$ values < -2.64 , all p values $< .05$]. The exception was the common effect and chain model, where this was only a trend in the data [$t(16) = -1.64, p = .11$].

To summarize, when learning a set of causal models, learners were more accurate when they themselves generated interventions than when they observed another generate those interventions, even though the conditional probability information and temporal information that the learners observed in these two conditions were identical. However, there might have been at least three reasons for the accuracy differences between the conditions. First, there might have been an additional benefit in the self-intervention condition from kinesthetic learning effects. Educational research suggests that some learners are kinesthetic or *active* learners: They benefit from the ability to act on the environment, as opposed to simply observing it. This individual difference is critical in kinesthetic domains (e.g., learning to dance), and several researchers have suggested similar benefits in nonkinesthetic domains (e.g., Furuhashi, 1999; Helm, 1991; Reis & Wolf, 1987).

Second, it is also possible that the learners who generated their own interventions found those data more salient than did the learners who observed another generate interventions. This difference in salience could also account for the difference between the learning conditions. Third, research on operant conditioning suggests that when effects are seen as reinforcement, the number of times an effect occurs overall affects both human and animal judgments of causal relationships (see, e.g., Dickinson & Shanks, 1995). In Experiment 1, the learners in the observation-of-intervention condition were not specifically reinforced for their actions, whereas one could imagine that the learn-

Table 2
Mean Numbers of Interventions (with Standard Deviations, Minimums, and Maximums) and Percentages of Critical Interventions Generated or Observed (%Critical) by the Learners in Experiment 1

Model	Interventions				
	<i>M</i>	<i>SD</i>	Min	Max	%Critical
Chain	44.29	16.33	25	79	15.38
Common effect and chain	40.53	18.09	25	72	21.27
Chain with A→D link	42.82	22.28	25	101	22.27
Diamond	45.53	18.68	25	96	25.39

ers in the self-intervention condition were reinforced for the interventions that they generated, because they themselves activated the lights.

In Experiment 2, we addressed these problems by introducing a forced intervention condition in which a group of learners actively made interventions but had no control over which interventions were made. We could then see whether the improved accuracy in Experiment 1 was due to any non-decision-related aspects of interventions, such as differences due to kinesthetic learning, salience of the data, or reinforcement effects.

EXPERIMENT 2

Experiment 2 utilized a learning environment similar to that in Experiment 1. The participants either were allowed to intervene on the data freely or were instructed to make specific interventions. As in Experiment 1, the interventions and the data were yoked together. If the differences between the self-intervention and the observation-of-intervention conditions observed in Experiment 1 were motivated by kinesthetic learning, salience, or reinforcement effects, accuracy should not differ between these two conditions. However, if decision-making abilities benefit causal learning, the learners who were forced to make intervention should show a reduction in accuracy.

Method

Participants

Forty-eight undergraduates were recruited from an urban area university. Approximately equal numbers of men and women participated in the experiment. The participants were paid \$7 for their participation.

Materials

The participants were tested on the same computer and monitor as those used in Experiment 1 and interacted with a Visual Basic script similar to that used in Experiment 1.

Procedure

The participants were seated at the computer. Those in the *self-intervention* condition were given instructions identical to those in Experiment 1. The participants in the *forced intervention* condition were also given these instructions; however, they were told that Dr. Science would instruct them to press particular buttons. These participants were not allowed to intervene on their own; instead, instructions to press each button appeared on the screen one at a time. When the learners pressed the appropriate button, the instruction disappeared, so the screen and data were identical to those in the self-intervention condition. If the learners pressed an incorrect button, they were given an error message, and the instruction was presented again. The sequence of interventions in the forced intervention condition was yoked to a participant in the self-intervention condition. All other aspects of the experiment were identical to those in Experiment 1.

Results and Discussion

The causal structure questions were scored in the same manner as in Experiment 1. Table 3 shows the mean scores for both learning conditions. Preliminary analyses similar to those done in Experiment 1 revealed no order effects.

A 4 (model) \times 2 (source condition) mixed ANOVA was performed on accuracy scores, with model as a within-subjects factor and source condition as a between-subjects factor. This analysis revealed a main effect of source condition: The learners were more accurate in the self-intervention condition than in the forced intervention condition [$F(1,46) = 8.74, p < .01$]. No main effect of model was found. No significant interactions were found. The difference between the source conditions held for each model individually [$t(46) = 2.07, 2.01, 3.24, \text{ and } 2.20$ for the chain, chain with an A \rightarrow D link, common effect and chain, and diamond models, respectively; all p values $\leq .05$].

We next examined the frequency with which learners responded correctly to all 12 causal structure questions. These data are shown in Table 3. Looking across the four models, the participants responded in this manner 49% of the time in the self-intervention condition and 20% of the time in the forced intervention condition. A criterion of 100% accurate responses on at least three out of the four models was set. More learners in the self-intervention condition met this criterion level than did those in the forced intervention condition [46% vs. 13%; $\chi^2(1) = 6.45, p < .05$].

The interventions made by each participant were analyzed in the same manner as in Experiment 1. These data are shown in Table 4. For each model, the total number of interventions made or the percentage of those interventions that were designated as critical did not correlate with accuracy on the causal structure questions. However, when the two learning conditions were examined individually, a significant positive correlation was found in the self-intervention condition between the percentage of interventions designated as critical and accuracy on the causal structure questions for the chain with an A \rightarrow D link model [$r(24) = .633, p < .01$]. This correlation was a nonsignificant trend for each of the other three models [$r(24) = .337, .335, \text{ and } .358$ for the chain, common effect and chain, and diamond models, respectively; all p values $\leq .11$]. No significant correlations were found in the forced intervention condition. These correlations were not due to the learners' responding at ceiling in the self-intervention condition. The learners responded signifi-

Table 3
Accuracy on the Causal Structure Questions and Percentages of the Learners Who Responded Correctly (%C) on All 12 Questions in Experiment 2

Model	Condition					
	Self-Intervention			Forced Intervention		
	Accuracy			Accuracy		
	<i>M</i>	<i>SD</i>	%C	<i>M</i>	<i>SD</i>	%C
Chain	10.88	1.45	46	9.88	1.87	21
Diamond	10.96	1.08	25	9.88	2.15	17
Chain with A \rightarrow D link	11.08	1.10	50	9.96	2.51	29
Common effect and chain	11.33	1.47	75	9.79	1.82	13

Note—Accuracy responses are out of a maximum of 12.

Table 4
Mean Numbers of Interventions (with Standard Deviations, Minimums, and Maximums) and Percentages of Critical Interventions Generated or Observed (%Critical) by the Learners in Experiment 2

Model	Interventions				%Critical
	<i>M</i>	<i>SD</i>	Min	Max	
Chain	49.04	19.40	25	89	15.92
Common effect and chain	44.79	30.51	25	166	25.85
Chain with A→D link	48.71	25.96	25	120	23.94
Diamond	52.08	42.90	25	234	25.15

cantly below ceiling performance on all four models [all $t(16)$ values < -2.23 , all p values $< .05$].

The participants in the forced intervention condition received the same conditional probability and temporal information as did the learners in the self-intervention condition but lacked the ability to decide what interventions to make. The learners were more accurate when they could decide what interventions to make. This also suggests that the differences observed in Experiment 1 and in previous research were not the result of kinesthetic learning, salience, or differences in reinforcement. Rather, these data suggest that causal learning is benefited by decision making.

Experiment 2 also replicated a manipulation performed by Lagnado and Sloman (2004), but with different results: Lagnado and Sloman found no difference between learning from a self-generated intervention and that in a similar forced intervention condition. However, this discrepancy was possibly due to a difference in the learning environment: Their participants learned in a fully probabilistic environment, in which events occurred spontaneously. Learning in this environment was difficult: Accuracy was less than 20% in both conditions. This might have prevented differences from emerging.

One concern about the present findings is that a yoked procedure might have concealed individual differences in learning abilities (see Church, 1964). It is possible that the learners in the observation-of-intervention and forced intervention conditions were simply weaker learners than their yoked "partners." Thus, it seems important to control the amount of data that all learners can observe.

Experiments 1 and 2 suggested that deciding what interventions to make benefited causal learning, but they do not suggest how. One possibility is that the learners in the self-generated condition benefited because they were able to decide how much total evidence to observe: The learners could collect evidence until they were satisfied that they understood the causal structure. There are two pieces of anecdotal evidence that suggest that this was not the case. First, in both Experiments 1 and 2, there was no correlation between accuracy and the total number of interventions the learners could make (either overall or in each condition independently). Second, Lagnado and Sloman (2004) also performed an experiment contrasting learners in self-intervention and observation-of-intervention conditions (with results similar to those in the present Experi-

ment 1), in which they fixed the number of trials both groups observed.

A second possibility is that when the learners chose their own interventions, they might have been inclined to consider what data they would observe, given each intervention; they recognized that some interventions were more informative than others. In doing so, the learners might have tried to choose interventions that were critical to learning the causal structure (see Steyvers et al., 2003, for a model that relies on this process). The attention required to choose one intervention among many possibilities should have helped the learners focus on the relation between particular interventions and outcomes. Since learning causal structure requires learning these relations, this focus, present in the self-intervention conditions, should have benefited causal learning. Conversely, in the observation-of-intervention and forced intervention conditions, the learners had no reason to think about the range of interventions they might perform. This might have impaired their accuracy in recognizing the relation between any particular intervention and its particular outcome. This interpretation is supported by a finding, in both Experiments 1 and 2, that the number of critical interventions made/observed was related to accuracy only in the self-intervention condition. It would appear that without decision-making abilities, the learners were not able to recognize and, thus, benefit from interventions critical to distinguishing between causal models.

This hypothesis suggests that if learners in an observation-of-intervention condition were presented with data that enabled them to identify interventions that were critical to learning the causal structure and then observed those data, they might show an improvement over learners who were given the same data, but without access to that decision process. To test this, in Experiment 3, we asked the participants to learn the same causal structures twice. One time, the data were structured in such a way that all the dependencies among the lights could first be observed, and then the learners observed the critical interventions that disambiguated a single causal structure from that dependence information. The other time, the learners observed the same data, but in a random order, so that this process was unavailable.

This procedure also allowed us to address potential problems with the yoked design of Experiments 1 and 2. In Experiment 3, the data presented to all the participants were the same and were predetermined by the experimenter, thus ensuring that individual learning abilities among the participants could not affect the type and amount of data presented. In this experiment, we also eliminated the training phase with the bucket. In real-life learning environments, it is rare that learners are given explicit training in how various kinds of interventions can elicit conditional probability information. Indeed, we doubt that this training is necessary: When faced with a similar causal learning problem, preschoolers generate such interventions spontaneously, without any training (Schulz, 2003). This leads us to hypothesize that learning will be relatively accurate even without this training.

EXPERIMENT 3

The participants were asked to learn causal structures from observing interventions in either a structured or a random manner. When the data were structured, the participants could first identify what interventions were necessary to specify a unique causal model, and then they observed exactly those interventions. The learners should have been more easily able to decide what conditional independence data it was necessary to observe in order to learn the causal structures. In the random case, although this decision process was possible, it was much more difficult. We expected that learning in the structured environment would be superior, even though the temporal information and conditional probability information were identical.

Method

Participants

Twenty-four undergraduates were recruited from an urban area university. Approximately equal numbers of men and women participated in the experiment. The participants were paid \$7 or received course credit for their participation.

Materials

The same computer and monitor as those in Experiments 1 and 2 were used. The participants interacted with a similar Visual Basic script.

Procedure

The participants were given the same procedure as those in the observation-of-intervention condition in Experiment 1, with several exceptions. They were asked to learn four causal structures: two chain models (Figure 1A) and two common effect and chain models (Figure 1D). Each model was learned under a *structured* condition and a *scrambled* condition (presented in one of four quasirandom orders).

For each model, Lights A, B, C, and D were randomly assigned to the four colors, with the provision that the same assignments for the two models were not made for each learner. In the structured condition, interventions were presented in a particular order. First, Dr. Science intervened on each light individually eight times. This way, the learners could observe the dependencies among the lights. Then he generated four examples of each critical intervention (as defined in Experiment 1) necessary to specify that particular causal structure. For instance, in the chain model, Dr. Science covered Light B with the bucket and activated Light A. Likewise, Dr. Science covered Light C and activated Light B. All together, the learners in this condition observed exactly 40 interventions on the causal system (slightly less than the mean for the learners in the self-generated conditions in Experiments 1 and 2). In the scrambled condition, the participants observed the same 40 interventions, but presented in a random order. All other aspects of the learning environment were the same as those in Experiment 1 (including the learners' ability to take notes), except that no training in how to use the bucket was provided. The learners were told about the bucket only in the instructions.

Results and Discussion

Causal structure scores were generated in the same manner as in Experiments 1 and 2. These data are shown in Table 5. A preliminary ANOVA revealed that the scores were unaffected by the order in which the models were learned. A 2 (model: chain vs. common effect and chain) \times 2 (intervention order: structured vs. scrambled) repeated measures ANOVA was performed. This revealed

Table 5
Accuracy on the Causal Structure Questions and Percentages of the Learners Who Responded Correctly (%C) on All 12 Questions in Experiment 3

Model	Condition					
	Structured Interventions			Scrambled Interventions		
	Accuracy			Accuracy		
	<i>M</i>	<i>SD</i>	%C	<i>M</i>	<i>SD</i>	%C
Chain	11.54	1.22	83	11.25	1.26	67
Common effect and chain	11.79	0.59	88	11.13	1.23	54

Note—Accuracy responses are out of a maximum of 12.

a significant main effect of intervention order [$F(1,23) = 8.90, p < .01$]. No other significant main effects or interactions were found.

Looking at performance on the individual models, the learners were more accurate in learning the structured common effect and chain model than in learning the same scrambled model [$t(23) = 2.80, p = .01$] and showed a trend in this direction in learning the chain model [$t(23) = 1.66, p = .11$]. When only 100% correct performance was considered, the learners answered all 12 questions correct on 85% of the structured models, as compared with only 60% of the scrambled models (Wilcoxon signed rank test, $z = -2.76, p < .01$).

The results showed higher accuracy when the learners were given a structured, rather than a random, presentation of data. Since this experiment showed differences in the absence of a yoked procedure, it suggests that the use of such a procedure does not account for the findings in Experiments 1 and 2.

One concern with these data is that learning in the structured condition might have required less memory demands than did learning in the scrambled condition, since the data were presented to the learners in batches. This reduced memory load might have facilitated learning. However, in this experiment, the learners could take notes and proceeded at their own pace, so converting all the data to table format was possible (although an inspection of the notes revealed this to be a rare strategy).

In general, the participants in this experiment were more accurate than the learners in the observation-of-intervention condition in Experiment 1. We suspect that this improved performance resulted from the participants' observing all of the critical interventions necessary to reveal the true causal structure. The participants in the previous experiments were not guaranteed to receive all of these interventions and, moreover, may have received nonessential data that could have influenced what particular models they considered possible at any given time.

GENERAL DISCUSSION

Several researchers have emphasized the importance of causal learning from interventions over learning from simply observing data (Gopnik et al., 2004; Kushnir et al., 2003; Lagnado & Sloman, 2004; Steyvers et al., 2003;

Waldmann & Hagmayer, 2005). Why do interventions benefit learning? Beyond providing the learner with necessary conditional probability and temporal information, we propose that deciding what interventions to generate provides learners the opportunity to recognize the relationship between interventions and outcomes—and to recognize data that are critical for learning a causal structure.

Experiment 1 showed that generating interventions benefited learning more than did observation of the same data, even when care was taken to provide the learners in the observation-of-intervention condition with the temporal information inherent to those interventions. Experiment 2 expanded these findings to a condition in which the learners were forced to make a set of interventions. This ruled out the possibility that kinesthetic effects, salience, or reinforcement effects were responsible for this difference and further supported the hypothesis that decision making influences causal learning.

Experiment 3 was an attempt to introduce decision making into an observation-of-intervention condition. When learners were shown data in a structured manner, so that they could potentially determine what interventions they would want to make, and then were shown those data, learning was facilitated, in comparison with cases in which the same models were learned from the same data, presented in a random order. This suggests that causal learning is facilitated if learners can identify what data are necessary to observe in order to disambiguate the data they have already observed (and then observe those data).

One limitation with the present experiments is that our learning environment restricted learning to a particular case, in which no temporal information about the events existed that would distinguish causes from effects. This situation is not unlike many cases of causal learning (such as many cases involving light switches, remote controls, and other electrical devices). But there are certainly other examples of causal learning in which temporal information between events provides clues to causal structure. For instance, even young children recognize that an effect is produced by a potential cause that precedes it, rather than by one that follows it (Bullock, Gelman, & Baillargeon, 1982). Such temporal information clearly plays a role in causal learning (Lagnado & Sloman, 2004; White, 2000).

Another limitation with the present learning environment was that events could not occur spontaneously. Lagnado and Sloman (2004) found that learners had difficulty with fully probabilistic learning environments, and one motivation for the present study was to improve learners' overall accuracy. Sobel (2005) converted the present learning environment into one in which learners could observe the effects of their own interventions, as well as observe events occur spontaneously. In this paradigm, accuracy on similar self-intervention conditions did not show a great drop-off as compared with the present data.

What are the implications of these experiments for describing a mechanism of causal learning? Many data-driven mechanisms of causal learning (e.g., those that

involve computing associative strength or measures of conditional probability) do not naturally account for the differences we find among the learning conditions. It is possible that decision making makes data more salient, which might motivate greater interest in the problem of learning (e.g., Renninger & Wozniak, 1985). But we feel that a more rational account of causal learning involves a top-down learning mechanism, in which deciding what interventions to make allows learners to reflect on whether those data help disambiguate possible representations of the causal structure.

This is not to say that all top-down strategies are consistent with these results. Shaklee and Fischhoff (1982) have described three top-down strategies for causal learning: parallel search, in which all possible hypotheses (in this case, causal models) are considered together; serial search, in which each hypothesis is tested individually but exhaustively; and truncated search, in which hypothesis testing stops once a particular hypothesis is confirmed. A parallel search strategy seems inconsistent with the present findings. Shaklee and Fischhoff found that adult learners preferred a truncated search strategy. The present data are consistent with this possibility.

One way to investigate these issues further would be to construct a computer simulation of what the learning process would look like given different hypothesis-testing strategies. As part of their investigation, Steyvers et al. (2003) attempted to model the strategies of their learners. They used a Bayesian learning algorithm with an active learning component, in which the hypotheses under consideration determined what intervention the simulation would generate. They found that human learners generated interventions in which they could observe data that provided the maximum contrast among their hypotheses, in agreement with the predictions of their simulation. Whether this model generalizes to other data—including the present data set—is uncertain, but we are presently examining this question.

The field of causal learning has been populated by a variety of computational models that focus on learning through associations and interventions. The present study suggests that in addition to the temporal and conditional probability information inherent in interventions, theories of causal learning must take the ability to make decisions about interventions into account. Some top-down learning algorithms do seem to incorporate this information. Specifying the nature of such an algorithm, however, remains a topic for further investigation.

REFERENCES

- BACON, F. (1986). *Novum organum*. Paris: Presses Universitaires de France. (Original work published 1620)
- BULLOCK, M., GELMAN, R., & BAILLARGEON, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209-254). New York: Academic Press.
- CHURCH, R. M. (1964). Systematic effect of random error in the yoked control design. *Psychological Bulletin*, *62*, 122-131.
- DICKINSON, A., & SHANKS, D. (1995). Instrumental action and causal representation. In D. Sperber, D. Premack, & A. J. Premack (Eds.),

- Causal cognition: A multidisciplinary debate* (pp. 5-25). Oxford: Oxford University Press, Clarendon Press.
- FIREMAN, G., KOSE, G., & SOLOMON, M. J. (2003). Self-observation and learning: The effect of watching oneself on problem solving performance. *Cognitive Development*, **18**, 339-354.
- FURUHATA, H. (1999). Traditional, natural and TPR approaches to ESL: A study of Japanese students. *Language, Culture & Curriculum*, **12**, 128-142.
- GIBSON, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- 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.
- GOPNIK, A., & MELTZOFF, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- HELM, D. J. (1991). Neuro-linguistic programming: Establishing rapport between the school counselor and the student. *Journal of Instructional Psychology*, **18**, 255-257.
- KLAYMAN, J. (1988). Cue discovery in probabilistic environments: Uncertainty and experimentation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 317-330.
- KLAYMAN, J., & HA, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, **94**, 211-228.
- KUSHNIR, T., GOPNIK, A., SCHULZ, L. E., & DANKS, D. (2003). Inferring hidden causes. In R. Alterman & D. Kirsch (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 699-703). Mahwah, NJ: Erlbaum.
- KUSHNIR, T., & GOPNIK, A. (2005). Young children infer causal strength from probability and intervention. *Psychological Science*, **16**, 678-683.
- LAGNADO, D. A., & SLOMAN, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 856-876.
- LARISH, J. F., & ANDERSEN, G. J. (1995). Active control in interrupted dynamic spatial orientation: The detection of orientation change. *Perception & Psychophysics*, **57**, 533-545.
- MURPHY, G. L., & MEDIN, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, **92**, 289-316.
- PEARL, J. (2000). *Causality*. New York: Oxford University Press.
- REIS, E. M., & WOLF, A. E. (1987). Helping secondary L.D. students read social studies text. *Journal of Instructional Psychology*, **14**, 133-139.
- RENNINGER, K. A., & WOZNIAK, R. H. (1985). Effect of interest on attentional shift, recognition, and recall in young children. *Developmental Psychology*, **21**, 624-632.
- SCHULZ, L. E. (2003, April). *The play's the thing: Interventions and causal inference*. Paper presented at the 2003 biennial meeting of the Society for Research in Child Development, Tampa, FL.
- SHAKLEE, H., & FISCHHOFF, B. (1982). Strategies of information search in causal analysis. *Memory & Cognition*, **10**, 520-530.
- SOBEL, D. M. (2005). [The role of model complexity in causal learning from interventions]. Unpublished raw data.
- SOMMERVILLE, J. A., WOODWARD, A. L., & NEEDHAM, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, **96**, B1-B11.
- SOPHIAN, C., & HUBER, A. (1984). Early developments in children's causal judgments. *Child Development*, **55**, 512-526.
- SPELBER, D., PREMACK, D., & PREMACK, A. J. (1995). *Causal cognition: A multidisciplinary debate*. Oxford: Oxford University Press, Clarendon Press.
- STEYVERS, M., TENENBAUM, J. B., WAGENMAKERS, E. J., & BLUM, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, **27**, 453-489.
- WALDMANN, M. R., & HAGMAYER, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **31**, 216-227.
- WASON, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, **20**, 273-281.
- WELLMAN, H. M., & GELMAN, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, **43**, 337-375.
- WHITE, P. A. (2000). Causal judgment from contingency information: Relation between subjective reports and individual tendencies in judgment. *Memory & Cognition*, **28**, 415-426.
- WOODWARD, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.

(Manuscript received August 28, 2003;
revision accepted for publication March 3, 2005.)