

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Induction of Casual Chains

Permalink

<https://escholarship.org/uc/item/40x2p690>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 22(22)

Authors

Ahn, Woo-kyoung
Dennis, Martin J.

Publication Date

2000

Peer reviewed

Induction of causal chains

Woo-kyoung Ahn

(woo-kyoung.ahn@vanderbilt.edu)

Department of Psychology, Vanderbilt University; 111 21st Avenue, Nashville, TN 37240 USA

Martin J. Dennis

(martin.dennis@yale.edu)

Department of Psychology, Yale University; 2 Hillhouse Avenue, New Haven, CT 06520 USA

Abstract

The current study examined one way in which people learn complex causal relations from covariation. When participants were presented with covariation information between X and Y and covariation information between Y and Z only, they were willing to infer a causal relationship between X and Z, although it is not warranted by the evidence. Furthermore, the perceived strength of the terminal relationship was tied to the perceived strength of the intermediate relationships, as manipulated through the order of evidence. These results imply that people do not follow normative, contingency-based theories, but instead carry out a hypothesis-testing process and combine piecemeal relationships into an overarching causal induction.

Introduction

Learning causal relations between two events is a fundamental cognitive activity. Although there are many ways to acquire causal knowledge (Ahn & Kalish, in press), the current study examines one way of learning causal relations, namely, through covariation.

In a simple case of covariation involving two events, a possible cause is either present (X) or absent ($\sim X$), and the target effect to be explained is either present (Y) or absent ($\sim Y$) as shown in Figure 1. One way to define covariation between two factors is to calculate an index $\Delta P = P(Y/X) - P(Y/\sim X)$, the difference between the probability that the effect occurs, given that the cause is present, and the probability that the effect occurs, given that the cause is absent (e.g., Cheng & Novick, 1992; Jenkins & Ward, 1965). In the example shown in Figure 1, ΔP is 0.2. Numerous studies have demonstrated positive correlations between objective ΔP and the perceived causal strengths between events (e.g., Wasserman, Chatlosh, & Neunaber, 1983).

	Y	$\sim Y$
X	6	4
$\sim X$	4	6

Figure 1. Example contingency between X and Y. Numbers show example frequencies of each evidence type.

In addition to the simple causal relations examined in these previous studies, people have knowledge about com-

plex causal mechanisms (e.g., Ahn, Kalish, Medin, & Gelman, 1995). For instance, in explaining why Kim had a traffic accident, one might refer to a mechanism of drunk driving rather than just the fact that there is a positive covariation between a traffic accident and drunk driving. Most people understand the mechanism underlying the effect of drunk driving on a traffic accident to be that when drunk, a person's motor responses are uncoordinated, in which case a person might not stay in the road, and so on. The question is, how do we acquire understanding of these causal mechanisms?

In answering this question, it is important to understand first how our knowledge about causal mechanisms might be represented. One useful tool is conditional dependencies or Bayesian networks¹. The idea is that mechanisms can be represented in terms of a complex web of covariation, or more specifically, as a directed graph in which nodes representing variables are connected with arrows indicating causal directions (Glymour, 1998; Glymour & Cheng, 1998; Pearl, 1996; Spirtes, Glymour, & Scheines, 1993; Waldmann & Martignon, 1998). For instance, a mechanism underlying the covariation between drunk driving and a traffic accident might be represented as follows:

drink alcohol \rightarrow uncoordinated motor responses \rightarrow traffic accident

Glymour (1998) proposes that B is a mechanism for a correlation between A and C, if, conditional on B, the correlation of A and C goes to zero. In the above example, one observes that as drunk driving increases, the number of

¹ Although we agree that Bayesian networks are a useful tool for representing people's causal mechanism knowledge, we do not endorse the view that conditional dependencies are all there is to that knowledge. Ahn and Kalish (in press) state that conditional dependencies are consistent with causal mechanisms because people's ideas about mechanisms support patterns of association. For instance, if someone believes that getting sneezed on causes illness via the mechanism of the transmission of germs, they should expect that the covariation between sneezing and illness is conditional on the transmission of germs. However, Ahn and Kalish disagree with Glymour (1998), who argues that patterns of covariation *are* mechanisms, and not just evidence for them. That is, a pattern of covariation might be one useful piece of evidence for identifying a relation as causal (i.e., addressing an epistemic question), but they are not what people mean by causation (i.e., addressing a metaphysical question). The current study addresses the epistemic question rather than the metaphysical question. Thus, we do not focus on this debate and instead assume that mechanisms can be represented in terms of conditional dependencies.

	X	~X
Y	6	4
~Y	4	6

	Y	~Y
Z	6	4
~Z	4	6

Figure 2. Example contingency between events X and Y and contingency between events Y and Z.

traffic accidents increases. Conditional on the number of uncoordinated motor responses, however, the covariation between drunk driving and traffic accidents would be greatly reduced. Thus, uncoordinated motor responses serve as a mechanism for this covariation.

Little is known about how people actually learn these complex patterns of covariations. Waldmann and Martignon (1998), who make use of a Bayesian network to represent mechanism knowledge, admit that it is improbable that humans learn such networks bottom-up, as instantiated in some computational models (e.g., Spirtes et al., 1993). For instance, Hashem and Cooper (1996) generated nine sets of relatively simple causal networks (e.g., $A \rightarrow B \rightarrow C$, or $A \leftarrow B \rightarrow C$) instantiated as diseases. Second and third year medical students were instructed to ask for any conditional probabilities among the three variables in each network, and to estimate the causal strength between B and C after receiving answers to their questions. Even from these simple causal networks, their estimates significantly deviated from the normative answers. The results suggest that it is unlikely that people can keep track of all conditional probabilities necessary for acquiring causal networks.

A simpler way of acquiring mechanism knowledge is by combining piecemeal causal relations. In this study, we attempt to show that upon learning that X sometimes causes Y and Y sometimes causes Z, people conclude (albeit erroneously) that X sometimes cause Z. Such inference is non-normative in that even if there is a contingency between X and Y and a contingency between Y and Z, it does not guarantee that there will be a positive contingency between X and Z. A normative conclusion would be that no inference about the relationship between X and Z can be made.

Consider Figure 2 which shows 40 individual cases, 20 of which depict the covariation between X and Y and 20 of which depict the covariation between Y and Z ($\Delta P = 0.2$ in both cases). Each of these 40 cases represents a different observation. On the left, we know, for instance, that there are six cases in which X and Y co-occurred, but we do not know what might have happened about Z in these six cases. Depending on this unknown information, the contingency

between X and Z can vary widely. To demonstrate this point, the top half of Figure 3 shows three possible distributions of these patterns of co-occurrence within the different levels of X and Y. For instance, in the six cases in Figure 3a. where both X and Y occur (the upper left-hand cell of the first contingency table), in two of the cases Z occurs and in the other four Z does not occur. The resulting co-occurrence patterns between X and Z are shown in the bottom half of the figure. Note that, not only is the pattern of co-occurrence between X and Y identical in each example, but the pattern between Y and Z is also identical (i.e., $\Delta P = 0.2$). However, the contingency between X and Z varies widely. Thus, a normative answer given covariation information about X and Y and covariation information about Y and Z only is that contingency between X and Z cannot be determined.

We propose that people would not make such normative judgments, and instead they would frequently assume that they can estimate the relationship between X and Z only from the covariation between X and Y (X-Y covariation) and covariation between Y and Z (Y-Z covariation). The reason for this is two-fold. First, as discussed earlier, keeping track of multiple conditional dependencies seems to be beyond the capacity of human cognition, but people have complex causal mechanism knowledge that can be represented in terms of conditional dependencies. Thus, people must have acquired this knowledge through other means. Second, in real-life situations, constituent covariations are oftentimes revealed in different cases. For instance, one might observe that eating a lot of food high in fat increases one's cholesterol level, and one might also observe that *other* people with high cholesterol die of a heart attack (not knowing whether these people had high-fat diets when alive). Therefore, it is adaptive, although non-normative, to make unwarranted inferences about unobserved covariations based on piecemeal covariations.

Specifically, we propose that people carry out a sort of syllogistic reasoning in this situation (Goldvarg & Johnson-Laird, 1999). Given that X causes Y and Y causes Z, people would subsequently conclude that X causes Z. We also propose that the stronger the perceived intermediate

	a. $\Delta P = -0.2$		b. $\Delta P = 0$		c. $\Delta P = 1$	
	X	~X	X	~X	X	~X
Y	6 (2 Z, 4 ~Z)	4 (4 Z)	6 (5 Z, 1 ~Z)	4 (1 Z, 3 ~Z)	6 (6 Z)	4 (4 ~Z)
~Y	4 (2 Z, 2 ~Z)	6 (2 Z, 4 ~Z)	4 (4 ~Z)	6 (4 Z, 2 ~Z)	4 (4 Z)	6 (6 ~Z)
Z	4	6	5	5	10	0
~Z	6	4	5	5	0	10

Figure 3. Example frequencies of co-occurrence between X and Z, holding constant the co-occurrence between X and Y, and between Y and Z. Note: ΔP 's show contingency between X and Z.

causal relations are, the stronger the perceived causal relationship between the two terminal events would be. Thus, if causal relations between X and Y, and Y and Z are weak, one would infer a weak causal relation between X and Z. If so, any manipulation that increases the perceived intermediate causal strengths should also increase the perceived causal strength between the terminal events. One such manipulation is presented in Dennis and Ahn (in press) who manipulated the order of evidence supporting a positive causal relationship versus evidence supporting a negative causal relationship. Because the current study utilized the same manipulation, we will first describe this study in detail, and then return to the issue of deducing overarching causal relations from piecemeal covariations.

Order manipulation

Consider four cells in Figure 1 again. Cells XY and $\sim X\sim Y$ serve to confirm that a positive causal relationship exists between X and Y. Henceforth, we will call these two cells positive evidence. Cells X $\sim Y$ and $\sim XY$ serve to confirm that there is a negative causal relationship between X and Y (negative evidence, henceforth). Participants in Dennis and Ahn (in press) observed a sequence of trials, each of which described presence or absence of two events, and judged the causal strength between the two events at the end of the sequence. Participants in one condition observed the bulk of positive evidence followed by the bulk of negative evidence (positive-first condition). In the other condition, participants observed the bulk of negative evidence followed by the bulk of positive evidence (negative-first condition). Although the order was different, all participants observed an identical covariation between X and Y, namely zero, in their experiment. The three possible results from this experiment were; (1) no effect of order, (2) a recency effect in which the negative-first condition leads to more positive causal estimates than the positive-first condition, and (3) a primacy effect in which the positive-first condition leads to more positive causal estimates than the negative-first condition. Existing models of causal induction predict either no effect (Cheng, 1997) or a recency effect (Rescorla & Wagner, 1972; see Dennis & Ahn, in press for more details of this prediction.) However, the results showed a strong primacy effect. This result was obtained even with the prospect of receiving reward for accurate judgments, indicating that the results are unlikely to be due to a fatigue effect.

Dennis and Ahn (in press) proposed that the primacy effect is obtained because causal learning occurs through a process of belief formation and updating. In this view, the information that a person receives at the beginning is used to construct an initial hypothesis about possible causal relationships. This initial belief then helps to provide an anchor point for future adjustments (Hogarth & Einhorn, 1992). However, as shown by Tversky and Kahneman (1974), people do not sufficiently adjust their initial anchor, resulting in the primacy effect.

Order effect in the two-step causal chain

In addition to showing that people frequently infer unwarranted overarching causal relations from constituent covariations, the second goal of the current study is to examine whether the order effect is obtained when judging X to Z causal strength based on X-Y covariation and Y-Z covariation. In the positive-first condition, participants observed a bulk of positive evidence for X and Y, and for Y and Z, followed by a bulk of negative evidence for X and Y, and for Y and Z. In the negative-first condition, participants observed identical contingencies with the negative evidence preceding the positive evidence. It is hypothesized that compared to the negative-first condition, the positive-first condition will lead to more positive causal estimates for the relationship between X and Z.

Method

Overview of Methods

In general, there were four phases: instructions, a learning phase, a test phase, and a follow-up phase. In the learning phase participants observed a series of trials providing X-Y covariation, and Y-Z covariation. The test phase required that participants make judgments about the causal relationship between events X and Z. The main experimental manipulations which occurred during the learning phase were the order in which participants received a bulk of positive or negative evidence for X-Y and Y-Z covariation. This manipulation was a within-subject variable, so that each participant actually saw two sets of learning and test phases. In the follow-up phase, participants described their thought processes. Each phase is explained below.

The Instruction and the learning phases were presented on iMac computers, using Microsoft PowerPoint 98 ®. The test and follow-up phases were presented as a paper-and-pencil task. Participants were 39 undergraduates at Vanderbilt University.

Procedure

Instruction phase In order to make participants get acquainted with the format of events, participants first received ten example learning trials with animations in which a person either does or does not eat a fictional plant called Ablex, and the same person subsequently does or does not exhibit a fictional physical reaction called Burlosis. The face of the person in each trial varied in order to have participants familiarized with the fact that each trial dealt with different cases.

Afterwards, participants were told to estimate "the extent to which Ablex plants cause Burlosis" on a scale from -100 (i.e., Ablex plants may prevent Burlosis) to 100 (i.e., Ablex may be a strong cause of Burlosis). Participants received instructions about the scale and examples of some of the scores. In addition, participants were instructed that

XY YZ $Y \sim Z$ $\sim X \sim Y$ $\sim Y \sim Z$ XY YZ $\sim X \sim Y$ $\sim Y \sim Z$ $X \sim Y$ XY YZ $\sim X \sim Y$ $\sim Y \sim Z$
 $\sim XY$ XY YZ $\sim YZ$ $\sim X \sim Y$ $\sim Y \sim Z$ XY YZ $\sim X \sim Y$ $\sim Y \sim Z$ (Block2) $Y \sim Z$ $X \sim Y$ $\sim XY$
 XY YZ $Y \sim Z$ $\sim XY$ $\sim YZ$ $Y \sim Z$ $X \sim Y$ $\sim YZ$ $\sim XY$ $X \sim Y$ $\sim X \sim Y$ $\sim Y \sim Z$ $\sim YZ$

Figure 4. The sequence used for the positive-first condition. Note: The sequence should be read from left to right. The trials in outline are negative evidence. (Block 2) indicates where the positive block ends.

"You may also decide that you cannot determine an estimate, given the information presented. In this case, you should give an estimate of 'NA.'" Participants wrote down their estimate from the practice trials on the sheet provided.

Learning Phase Upon completing the practice trials, participants were told that in the actual experimental trials, they would see descriptions of three events; the possible application of a fictitious fertilizer, the possible increase in the level of a fictitious chemical in the soil, and the possible blooming of a fictitious flower. During each learning phase, they were told what these three events were; they were presented with animations that would accompany each event. (See the material section.)

Participants were specifically told that they will have only two pieces of information available during learning (e.g., "whether it [i.e., the plot] had increased levels of the chemical compound alizene and whether the plant Lanya subsequently bloomed on it, or whether it received the fertilizer Yerban and whether it subsequently had increased levels of alizene"). They were also explicitly told that they would never receive information about both the fertilizer and the plant. This instruction was added to prevent any false memory of having observed the covariation between the fertilizer and the plant. That is, if participants did not select "NA" in estimating the causal strength between the fertilizer and the plant, it cannot be due to the fact that they misremembered what covariation information they had seen. In addition, we attempted to reduce participants' cognitive load during the learning phase by instructing them what their task is in advance. Thus, participants were told that their task was, for instance, "to judge the causal relationship between Yerban and Lanya." After these instructions, participants were presented with 40 learning trials. (See the material section for more detail.)

Test Phase After observing the entire sequence of trials in a learning phase, participants provided causal strength ratings for the effect of the fertilizer on the plant's blooming. Following Wasserman, Elek, Chatlosh, and Baker (1993), participants were asked, for instance: "To what extent does the fertilizer Yerban cause the plant Lanya to bloom?" Participants wrote a number between -100 and 100. They were also reminded to write "NA" if they "cannot determine an answer from the evidence given."

Follow-up Phase When participants were done with the learning and test phases for two sets of materials, they were asked to rate how much thought they put into each judgment on a 5-point scale where 1 indicated "no" and 5 indi-

cated "very much." Finally, they were asked to write about their "thought process in performing the experimental task" such as "Were there any strategies in particular you used while observing the experimental trial? How did you interpret each type of evidence?"

Design and Materials

During the learning phase, participants received 40 trials, in which 20 provided X-Y covariation information and 20 provided Y-Z covariation information. For both, ΔP was 0.2 as in Figure 2.

Two experimental conditions were defined by the order in which covariation information was presented during the learning phase. In order to construct the experimental sequences, two different blocks (positive and negative blocks) were created. The positive block had 24 trials, 20 of which were positive evidence (i.e., XY , YZ , $\sim X \sim Y$, or $\sim Y \sim Z$). The negative block had 16 trials, 12 of which were negative evidence (i.e., $X \sim Y$, $\sim X Y$, $Y \sim Z$, or $\sim Y Z$). Within each block (positive or negative), the trials were randomly ordered except that XY was always followed by YZ , and $\sim X \sim Y$ was always followed by $\sim Y \sim Z$ ². This random order was fixed across participants.

The two different experimental conditions were constructed by manipulating the order of these two blocks, so that, in the positive-first order condition, the positive block came before the negative block. This pattern is shown in Figure 4 where the positive and the negative blocks are separated (Block 2). Although lines separate the positive and the negative blocks in this figure, the entire sequence was presented to the participants without any indication of blocks. In the negative-first order condition, the negative block came before the positive block, which can be seen by switching the two blocks in Figure 4. Each participant went through both experimental conditions; the order of conditions was counterbalanced across participants.

The actual events used for X was the application of a fertilizer called Yerban or Zertax, Y was a change in level of a chemical called Alizene or Banizon, and Z was the blooming of a plant called Lanya or Hyaeth. We used animations to show spraying of fertilizer, increasing chemical level, and blooming plant. These animations were intended to keep participants' attention and to reduce their cognitive load by visualizing the events, so that participants would not make NA responses simply because they were over-

² This constraint could not have limited our interpretation of the effect of order because the same constraint was used for both order conditions. In another experiment where this constraint was not imposed, the number of NA responses was approximately the same as in the current experiment.

whelmed with too many combinations of presence and absence of events. Finally, each trial had a unique plot number displayed at the top of the screen, so that it was clear that each observation was separate.

To summarize, after receiving general instructions, each participant observed a series of trials about covariation between X and Y, and covariation between Y and Z in either the positive-first or the negative-first order, and made a causal strength judgment about X and Z. Afterwards, they observed another series of trials about three new events in the other condition, and then made a second judgment. Finally, they wrote about their thought processes.

Results

We first examined the number of NA responses. In order to be truly valid NA responses, a participant should have given NA responses in both the positive-first and the negative-first conditions. Only one out of 39 participants did so. This participant's explanation also agreed with the true justification for doing so, "It was very difficult to reason without seeing all three factors together...."

Overall, 20.5% of responses across the two conditions were NA responses. There are a number of reasons to believe that these NA responses were unlikely to indicate a response of "indeterminate," but rather were a way to indicate a lack of causal relation between the two events. First, as reported earlier, only one of these subjects gave NA responses in both conditions. Second, the other participants' reasons for giving an NA response are consistent with this interpretation. For instance, one participant stated, "... no causal relationship, or lack thereof, could be estimated because every relationship that was shown had another that contradicted it..."; another participant stated, "...There seemed to be no relationship between any..." Third, most interestingly, there were more NA responses from the negative-first condition (35.9% of participants) than from the positive-first condition (7.7% of participants), $\chi^2(1, N=39) = 8.1, p < .01$, McNemar's test (McNemar, 1947). As we shall see below, those who did not give NA responses gave lower estimates in the negative-first condition than in the positive-first condition. Thus, more NA responses in the negative-first condition seemed to reflect participants' belief in weaker causal strengths.

Finally, we examined the mean estimates for each condition. The mean rating in the positive-first condition were 32.5 whereas that in the negative-first condition was only 5.8. For a statistical analysis, we excluded data of those who gave at least one NA response in either condition. With the remaining 23 participants, a dependent *t*-test showed that the mean rating in the positive-first condition (22.6) was reliably higher than that in the negative-first condition (4.1), $t(22) = 3.71, p = .001$. Thus, although participants saw identical contingencies between X and Y, and between Y and Z, their estimated causal strength between X and Z was stronger when they first saw positive evidence

for these two contingencies than when they first saw negative evidence.

Discussion

The experiment reported here suggests that people are willing to make overarching causal inductions from constituent covariations. The bulk of participants in our experiment were willing to infer a causal relationship between the two terminal events in a proposed causal chain, even though they did not see the actual covariation between the two events. Of those people who were not so willing to make that inference, the majority seemed not to understand the normatively correct reason for a response of "indeterminate," instead using such a response as a proxy for a perceived lack of causal relationship. This willingness to make overarching inductions seems to be a sensible thing to do, given that people rarely have the luxury in the real world of observing a complete set of covariation patterns between multiple events.

When people make these overarching inductions, they seem to first infer that X causes Y and Y causes Z. Based on these inferences, they conclude that X causes Z. Some participants' explanations for their responses supports this. For example, one participant wrote, "I tried to find the patterns; for example, that A caused B, and B caused C, so A probably causes C" Another wrote, "I tried to see the relationship between the plant and the compound, and compound and fertilizer separately first. From there I tried to determine whether or not the presence or absence of fertilizer yielded the presence or absence of plant..." In other words, it appears that people may try to integrate the relative strengths of the intermediate relationships to estimate the strength of the relationship between the terminal events.

In this study, we used three events that may have reflected prior knowledge about the function of chemical fertilizers. Participants could have judged the strength of the causal relationships based solely on this prior knowledge. But such an interpretation is unlikely, given that the events we used were fictitious ones (and thus, there could not have been prior knowledge about causal strengths among these events), and furthermore, people's causal strength estimates were susceptible to manipulation of the order of evidence. Finally, preliminary results from a new study show that the same effects occur using very abstract events (e.g. squares changing shape or triangles changing color).

Extending Dennis and Ahn (in press), we found an order effect in situations involving three events. As we suggested, we think this order effect occurs because of an anchoring-and-adjustment process. One participant's description precisely illustrates this process: "If a particular pattern kept coming, but one or two trials deviated from the pattern, I would excuse them as flukes." In this case, the adjustments to the initial anchor was not strong enough, leading to biased final estimates of causal strength.

These results also have implications for current, nor-

mative theories of causal learning (Cheng, 1997; Glymour, 1998; Glymour & Cheng, 1998). These theories propose that people's estimates of causal power match those predicted by contingency indices calculated from observed conditional probabilities. However, in the current experiment no such index can be calculated, given the lack of observed co-occurrence between the terminal events. Yet people still were willing to provide judgments of causal strength, suggesting that the normative contingency-based theories are inadequate descriptions of human causal learning.

In contrast, the results are consistent with a causal power view of causal learning. According to this view, people infer causal relationships based on the proposed transfer of some sort of causal force or energy between one object and another. Specifically, the mechanism by which one event brings about another is proposed to be the main focus of causal reasoning (Ahn, et al., 1995; Bullock, Gelman, & Baillargeon, 1982; Harré, 1988). In the case of our experimental results, the presence of a putative mechanism (i.e. the change in soil chemistry) seems to outweigh the absence of the covariation information necessary to draw accurate causal inferences. Furthermore, the perceived strength of the target relationship was tied to the perceived strength of the mechanism, as evidenced by the primacy effect obtained. That is, the current results demonstrate people's reliance on mechanism information in the acquisition of new causal learning.

Acknowledgements

Support for this research was provided by a National Institute of Health Grant (NIH R01-MH57737) to Woo-kyoung Ahn.

References

Ahn, W., & Kalish, C. W. (in press). The role of mechanism beliefs in causal reasoning. In F. Keil & R. Wilson (Eds.), *Explanation and cognition*. MIT Press.

Ahn, W., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299-352.

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.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.

Dennis, M. J., & Ahn, W. (in press). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*.

Glymour, C. (1998). Learning causes: Psychological expla-

nations of causal explanation. *Minds and Machines*, *8*, 39-60.

Glymour, & Cheng, P. W. (1998). Causal mechanism and probability: A normative approach. In M. Oaksford & N. Chater (Eds.) *Rational models of cognition*. Oxford University Press.

Goldvarg, Y., & Johnson-Laird, P. N. (1999). Naive causality: a mental model theory of causal meaning and reasoning. *Proceedings for 21st annual meeting of Cognitive Science Society*, Vancouver, Canada.

Harré, R. (1988). Modes of explanation. In D.J. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality* (pp. 129-144). Brighton, Sussex, UK: Harvester Press.

Hashem, A. I., & Cooper, G. F. (1996). Human causal discovery from observational data. *Proceedings of the 1996 Symposium of the American Medical Informatics Association*.

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, *24*, 1-55.

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, *79*(1, Whole No. 594).

McNemar, Q. (1947). Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, *12*, 153-157.

Pearl, J. (1996). Structural and probabilistic causality. In D. R. Shanks, D. L. Medin, & K. J. Holyoak (Eds.), *Psychology of learning and motivation, Vol. 34: Causal learning*. San Diego, CA: Academic Press.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research*. New York: Appleton-Century-Crofts.

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

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *211*, 453-458.

Waldmann, M. R., & Martignon, L. (1998). A Bayesian Network Model of Causal Learning. *Proceeding of the 20th Cognitive Science Conference*, Hillsdale, NJ: Erlbaum.

Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning and Motivation*, *14*, 406-432.

Wasserman, E. A., Elek, S. M., Chatlosh, D. L., & Baker, A. G. (1993). Rating causal relations: Role of probability in judgments of response-outcome contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 174-188.