

Interpreting Covariation in Causal Structure Learning

Motoyuki Saito (m-saito@kwansei.ac.jp)

Department of Psychological Science, Kwansei Gakuin University
Hyogo, 662-8501, JAPAN

Tsuneo Shimazaki (shimazaki@kwansei.ac.jp)

Department of Psychological Science, Kwansei Gakuin University
Hyogo, 662-8501, JAPAN

Abstract

Recent studies have shown that people use covariation information to infer causal structure. However, there is little information about how people derive causal directionality from covariation. The present study is designed to provide further evidence about the role of covariation in causal structure learning. In Experiment 1, where covariation between two variables was systematically manipulated, participants were asked to observe the states of bacteria (present or absent) and to infer their causal relationship. We found that judgments of causal structure varied as a function of covariation, and that participants interpreted covariation according to necessity of causation. In Experiment 2, participants who received information about high causal strength interpreted covariation according to sufficiency of causation. These results demonstrate that prior knowledge modulates interpretation of covariation and suggest that domain-general covariation information and domain-specific prior knowledge of causal relations interact in causal structure learning.

Keywords: causal learning; covariation; prior knowledge; necessity; sufficiency.

Introduction

Causal knowledge enables us to explain past events, to control the present environment, and to predict future outcomes. Using this knowledge, we can achieve desired outcomes and avoid undesired consequences. Many psychological studies have investigated how people acquire and use knowledge of causality (Gopnik & Schulz, 2007; Sloman, 2005; see also Holyoak & Cheng, 2011, for a review). Despite the importance of causal knowledge, it is often difficult to determine the casual structure among events. For example, imagine that someone feels unmotivated and makes slow progress on their work. In this situation, it is unknown whether the lack of motivation leads to slow progress, or whether slow progress causes lack of motivation. Furthermore, it is also possible that motivation and work progress are unrelated. Given this ambiguity, how do people learn causal structure?

Hume (1739/2000) argued that causal relations are unobservable and must be induced from observable events. Information about covariation among events serves as a fundamental cue for inferring causal structure. Covariation is represented as the pattern of occurrences and non-occurrences for binary variables. Figure 1 shows a standard contingency table where the letters in each cell (a, b, c, d)

represent the joint frequencies for one value of event X and one value of event Y. It is generally accepted that objective measure of contingency is described by ΔP , as shown in Equation 1 (Jenkins & Ward, 1965).

$$\Delta P = P(Y|X) - P(Y|\neg X) = \frac{a}{a+b} - \frac{c}{c+d} \quad (1)$$

In this equation, $P(Y|X)$ is the probability of Y given the presence of X, and $P(Y|\neg X)$ is the probability of Y given the absence of X. Values of ΔP range from -1 to $+1$. Positive ΔP values indicate a generative causal relation; negative ΔP values indicate a preventive causal relation. When a causal relation exists, strong covariation between the cause and the effect is expected. By contrast, lack of covariation indicates that two variables are unrelated (i.e., ΔP does not differ significantly from zero). Many studies have focused on how people estimate causal strength between the candidate cause and the effect, and the results have shown that people are quite sensitive to covariation information (e.g., Wasserman, Elek, Chatlosh, & Baker, 1993). However, covariation itself is inadequate for inferring a unique causal structure: when event X covaries with event Y, it is difficult to determine whether X causes Y, or vice versa.

When combined with additional information, covariation becomes a more useful cue to causal structure. First, temporal order in which people observe the occurrence of events facilitates learning causal directionality. As causes are often observed prior to their effects, when event X precedes event Y, it is highly probable that X causes Y. For example, if becoming unmotivated precedes making slow

		Event Y	
		Y	$\neg Y$
Event X	X	a	b
	$\neg X$	c	d

Figure 1: A contingency table summarizing the covariation between two binary variables. The letters in each cell indicate frequencies of co-occurrence for the two states of events X and Y.

progress on work, temporal order information suggests that decreased motivation causes slow progress. Second, information about the absence of hidden causes also makes covariation cues more useful. When event X covaries with event Y, three possible causal structures are supposed (i.e., $X \rightarrow Y$, $X \leftarrow Y$, or $X \leftarrow Z \rightarrow Y$). The possibility that both events are caused by a hidden common cause, Z, can be excluded if it is known that there are no hidden causes. If event X exists alone, necessity of causation indicates that X causes Y (i.e., $X \rightarrow Y$). This is because nothing happens without a cause (i.e., $P(\text{Effect}|\neg\text{Cause}) = 0$). Therefore, events that exist alone must be a cause variable, not an effect variable. In contrast to necessity of causation, sufficiency of causation draws the opposite conclusion that event Y causes event X in above situation (i.e., $X \leftarrow Y$). Since sufficiency of causation assumes that causes always accompany their effects (i.e., $P(\text{Effect}|\text{Cause}) = 1$), events that occur alone must be an effect variable, not a cause variable. Given that there is no factor that affects both motivation and work progress and that motivation changes spontaneously, in previous example, necessity of causation suggests that decreased motivation causes slow progress and sufficiency of causation indicates that slow progress causes decreased motivation.

Recent studies on causal structure learning have revealed the importance of covariation (e.g., Deverett & Kemp, 2012; Mayrhofer & Waldmann, 2011; Rottman & Keil, 2012; Saito & Shimazaki, 2012). For instance, Saito and Shimazaki (2012) demonstrated that people judge simple causal structure on the basis of covariation information, but the use of covariation is modulated by task complexity. The experimental task was to observe the states of bacteria and to infer their causal relationship. Participants were instructed that temporal order was unreliable and that there were no hidden causes. In the simple causal structure condition, covariation was favored over temporal order as the basis for inferring causal structure; in contrast, temporal order was more influential in the complex causal structure condition. In addition, Mayrhofer and Waldmann (2011, Experiment 1) reported that people can differentiate common cause models (e.g., $X \leftarrow Z \rightarrow Y$) from common effect models (e.g., $X \rightarrow Z \leftarrow Y$) on the basis of covariation information. Although these recent studies show the ability to infer causal directionality from covariation, how people use covariation to induce causal directionality is still not well-understood. Therefore, it is valuable to study how people make structure judgments according to covariation.

The purpose of the present study is to investigate how people interpret covariation information in causal structure learning. In Experiment 1, we systematically manipulated covariation between two variables and asked participants to make causal structure judgments. Although causal structure between two variables is not determined by covariation alone, this situation enables us to examine whether participants have some sort of tendency in inferring causal directionality from covariation. In Experiment 2, we gave participants different information about causal relations and

investigated whether prior knowledge changed their interpretation of covariation.

Experiment 1

Experiment 1 investigated how people interpret covariation when judging causal structure. The experimental task was to observe the occurrence of two fictitious bacteria and to infer their causal relationship. We manipulated covariation information by varying the number of occurrences and non-occurrences of each bacterium.

Method

Participants and Design Forty-three undergraduates from Kwansai Gakuin University received course credit for participating in this experiment. Two additional participants were excluded from the analyses due to misunderstanding of the instructions. Excluding these participants did not alter the general pattern of results.

Covariation information was systematically manipulated within participants. There were 15 covariation conditions (see Table 1) based on the combinations of five levels (1.00, .75, .50, .25, .00) of the conditional probabilities $P(Y|X)$ and $P(Y|\neg X)$. The difference between $P(Y|X)$ and $P(Y|\neg X)$ for each condition yielded five levels of nonnegative ΔP values (1.00, .75, .50, .25, .00). Each participant completed the causal learning task for all covariation conditions.

Instructions Participants received verbal and written instructions in Japanese, and were asked to confirm that they understood the instructions. An English translation of outlines of the instructions was provided below:

Imagine that you are a scientist attempting to reveal a causal relationship between two types of newly discovered bacteria (These bacteria have the same shapes but different colors to conjure up an image of cell divisions). The term “causal relationship” means a relationship where one bacterium propagates the other bacterium (i.e., generative causal relationship). It is unknown whether one bacterium propagates the other, or whether these bacteria are unrelated. To investigate the relationship between the bacteria, you are going to observe the appearance of the bacteria. The states of the bacteria should help you consider the causal relationship between them.

Your task is to observe the occurrences and non-occurrences of these bacteria and to infer their causal relationship. Note that the experimental task does not require any knowledge of biology. (The remaining instructions describe how to progress through the learning phase and test phase.)

Learning Phase Participants observed the states of bacteria (present or absent) to infer their causal relationship. On each trial, a button labeled “NEXT” was displayed on the screen. After clicking the button, information about the states of both bacteria X and Y was provided. The presence of a

bacterium was indicated by the appearance of the bacterium; in contrast, the absence of a bacterium was represented by the appearance of the bacterium labeled with a cross mark. The screen was returned to its primary state (i.e., “NEXT”) 1s after the bacteria appeared.

There were 16 trials for each covariation condition. Bacterium X was present on eight trials and was absent on eight trials (i.e., $P(X) = .50$). Two conditional probabilities, $P(Y|X)$ and $P(Y|\neg X)$, were set to one of five levels in each condition (Table 1) The difference between these probabilities yielded five levels of nonnegative ΔP s ($\Delta P = P(Y|X) - P(Y|\neg X)$). Each condition was described through the difference between the two conditional probabilities. For example, in the .75-.00 condition (i.e., $P(Y|X) = .75$, $P(Y|\neg X) = .00$), bacteria X and Y were both present on six trials. Bacteria X and Y were both absent on eight trials, and on two trials bacterium X was present and bacterium Y was absent. The order of trials and conditions was randomized within participants. To familiarize participants with the procedure, several practice trials were performed prior to the learning phase. Participants were informed that the information in the practice trials was irrelevant to the learning phase.

Test Phase After observing 16 cases, participants were asked two yes/no questions about the causal structure. They were asked whether bacterium X caused bacterium Y, and whether bacterium Y caused bacterium X. Then, after a brief delay, participants began the learning and test phases for the next covariation condition. They were instructed that judgments should be made independently of their answers on prior problems.

Results and Discussion

Combining the answers on the two test questions yields four types of causal models: (1) X causes Y, (2) Y causes X, (3)

bidirectional, and (4) independent. The percentage of responses in each condition is shown in Table 1. Although causal structure could not be uniquely determined in all conditions, participants’ judgments varied greatly. A log-linear model analysis on the 15 (covariation conditions) \times 4 (causal models) cross table revealed a significant interaction between covariation condition and causal model, $\chi^2(42) = 291.37, p < .001$.

In order to explore the interaction between covariation information and causal judgments in greater detail, we conducted a correspondence analysis. The contributions of dimensions 1 and 2 are 63.75% and 20.29%, respectively, and their cumulative contribution is 84.05%. Therefore, we created scatter plots in two dimensions (Figure 2). As can be seen from Figure 2, each judgment is closely related to specific conditions. Most participants concluded that X caused Y in the .25-.00, .50-.00, and .75-.00 conditions and that Y caused X in the 1.00-1.00, 1.00-.75, 1.00-.50, and 1.00-.25 conditions. Bidirectional causal relationships were only inferred in the 1.00-.00 condition. X and Y were judged to be independent in the other conditions.

Participants’ judgments are explained in terms of necessity of causation (cf. Pearl, 2000). Necessity represents the degree to which the cause is necessary for the effect; in contrast, sufficiency is the degree to which the cause is sufficient for the effect. Pearl (2000) introduced three indices that assess causality: the probability of necessity (PN), the probability of sufficiency (PS), and the probability of necessity and sufficiency (PNS). These indices are easily calculated when the covariation information given does not include both the case where the effect is present in the absence of the cause (i.e., cell b), and the case where the effect is absent despite the presence of the cause (i.e., cell c). There also cannot be any common factors that have an

Table 1: Details of conditions, results, and interpretations in Experiment 1

Covariation conditions			Causal models (% of participants)				Interpretations		
$P(Y X)$	$P(Y \neg X)$	ΔP	X→Y	X←Y	X→Y & X←Y	X Y	PN	PS	PNS
1.00	.00	1.00	0.00	2.33	53.49	44.19			
1.00	.25	.75	20.93	48.84	6.98	23.26	X←Y	X→Y	X←Y
.75	.00	.75	53.49	18.60	6.98	20.93	X→Y	X←Y	X←Y
1.00	.50	.50	27.91	53.49	0.00	18.60	X←Y	X→Y	X←Y
.75	.25	.50	11.63	13.95	18.60	55.81			
.50	.00	.50	51.16	23.26	2.33	23.26	X→Y	X←Y	X←Y
1.00	.75	.25	27.91	60.47	2.33	9.30	X←Y	X→Y	X←Y
.75	.50	.25	9.30	20.93	11.63	58.14			
.50	.25	.25	13.95	13.95	16.28	55.81			
.25	.00	.25	46.51	23.26	2.33	27.91	X→Y	X←Y	X←Y
1.00	1.00	.00	20.93	51.16	0.00	27.91			
.75	.75	.00	11.63	27.91	13.95	46.51			
.50	.50	.00	11.63	4.65	11.63	72.09			
.25	.25	.00	16.28	13.95	6.98	62.79			
.00	.00	.00	23.26	16.28	0.00	60.47			

influence on both events X and Y. Under these conditions, the probability of necessity, PN , is calculated according to the following equation:

$$PN = \frac{P(Y|X) - P(Y|\neg X)}{P(Y|X)} \quad (2)$$

When event X generates event Y (i.e., $\Delta P > 0$), values of PN become higher as the probability of Y given the absence of X, $P(Y|\neg X)$, decreases. This reflects the fact that necessity of causation assumes that the base rate of the effect is low (i.e., $P(\text{Effect}|\neg\text{Cause}) = 0$). In contrast, the probability of sufficiency, PS , is based on the assumption that causes always accompany their effects (i.e., $P(\text{Effect}|\text{Cause}) = 1$) and is defined as follows:

$$PS = \frac{P(Y|X) - P(Y|\neg X)}{1 - P(Y|\neg X)} \quad (3)$$

The probability of necessity and sufficiency, PNS , takes both necessity and sufficiency aspects of causal relations into account:

$$PNS = P(Y|X) - P(Y|\neg X) \quad (4)$$

These indices are calculated on the basis of the causal direction from event X to event Y. Therefore, indices based on inverse direction are calculated by interchanging the rows and columns of the 2×2 contingency table. When index values based on the direction from X to Y (e.g., PN from X to Y) are compared with those based on the direction from Y to X (e.g., PN from Y to X) and causal directionality is inferred by higher agreement with the conception (i.e., higher values), the three indices lead to different interpretations (see Table 1). For example, in

the .25-.00, .50-.00, and .75-.00 conditions, PN predicts that X causes Y. In contrast, in the 1.00-.75, 1.00-.50, and 1.00-.25 conditions, PN makes the opposite prediction. On the basis of responses in the conditions where the three indices are defined, we classified participants into one of five clusters: necessity, sufficiency, necessity and sufficiency, random, and unclassified. These classifications were made by rates of agreement between judgments and index predictions (1 for predicted judgments, 0 for unpredicted judgments). When participants had the same rate of agreement for different clusters, they were included in the unclassified cluster. As a result, more than half of the participants (55.81%) were classified to the necessity cluster and 27.91% of participants were classified in the sufficiency cluster. There were few participants in the other clusters (6.98% in the necessity and sufficiency cluster; 2.33% in the random cluster; 6.98% in the unclassified cluster). This suggests that most people interpret covariation information according to necessity.

However, these results are inconsistent with recent work suggesting that people judge causal relations on the basis of sufficiency (Mayrhofer & Waldmann, 2011). Sufficiency of causation assumes that causal relations are deterministic (i.e., $P(\text{Effect}|\text{Cause}) = 1$). According to sufficiency, the presence of event X in the absence of event Y is interpreted as an indication that X does not cause Y. Therefore, when covariation information includes such cases, it is suggested that Y causes X (i.e., $X \leftarrow Y$). On the other hand, necessity of causation assumes that all events have a cause (i.e., $P(\text{Effect}|\neg\text{Cause}) = 0$) and such cases are taken as evidence that Y does not cause X. In contrast to sufficiency, necessity indicates that X causes Y (i.e., $X \rightarrow Y$) in the situation described above. Thus, judgments of causal directionality between two variables depend on the interpretation of covariation. In a second experiment, Mayrhofer and Waldmann (2011) had participants observe communications between two mind-reading aliens, and asked them to infer causal directionality. Covariation information included the case where two aliens X and Y thought the same thing, and the case where only one alien thought something (e.g., X) and the other alien thought nothing (e.g., $\neg Y$). Whereas sufficiency would suggest that alien Y transferred his thought to alien X (i.e., $X \leftarrow Y$), necessity favors the opposite conclusion (i.e., $X \rightarrow Y$). More participants concluded that Y caused X, suggesting that people judge causal relations on the basis of sufficiency.

These conflicting findings could be due to differences in prior knowledge about causal relations. Sufficiency of causation requires high causal strength, whereas necessity of causation requires the low base rate of the effect. If participants expect the effect's base rate to be low before the learning phase, covariation information is likely to be interpreted according to necessity. In contrast, prior knowledge about high causal strength might lead to an interpretation based on sufficiency. Indeed, it is difficult to imagine that the effect bacterium could occur in the absence of the cause bacterium in the bacteria story. That is,

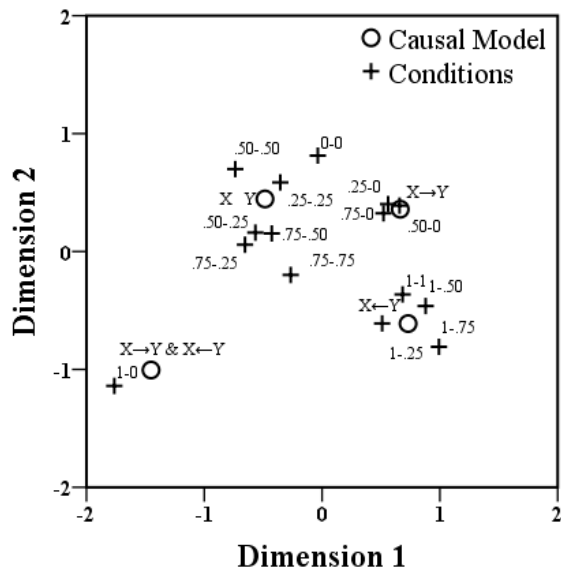


Figure 2: Results of correspondence analysis.

participants will believe that bacteria do not arise spontaneously (i.e., $P(\text{Effect}|\neg\text{Cause}) = 0$) and therefore, think the bacterium that exists alone must be a cause. In the alien story, however, such a situation is more plausible: an alien has the potential to think spontaneously, regardless of whether a cause alien is present (i.e., $P(\text{Effect}|\neg\text{Cause}) > 0$). Participants will assume multiple causes in the alien cover story and regard the single-occurrence of the thought as an effect. Since necessity and sufficiency differ in their assumption about the base rate of effect and causal strength, differences in prior knowledge about these parameters might result in different judgments of causal structure based on covariation. We test this hypothesis in Experiment 2.

Experiment 2

Experiment 1 demonstrated that participants made different judgments as a function of covariation. Whereas the results of Experiment 1 indicate that people interpret covariation information according to necessity of causation, Mayrhofer and Waldmann (2011) suggest that people interpret covariation according to sufficiency of causation. Experiment 2 was designed to investigate the effect of prior knowledge on interpretation of covariation information. The experimental procedure was similar to that of Experiment 1, but participants received different instructions about causal relations. We expected that additional instructions about high causal strength would lead to interpretations based on sufficiency of causation, and that participants who were not given additional instructions would infer causal structure according to necessity of causation.

Method

Participants and Design Twenty-four undergraduates from Kwansei Gakuin University participated in the experiment and received course credit. None of them took part in Experiment 1. They were randomly assigned to either the sufficiency instruction or control group.

Procedure Each participant observed the states of bacteria (present or absent) and inferred their causal relationship. The procedure was the same as Experiment 1, with the following exceptions. First, participants in the sufficiency instruction group received additional instructions that emphasized the sufficiency of causation. In addition, to ensure that participants remembered this additional information, they were allowed to re-read the instructions during the learning and test phases. Finally, covariation information was manipulated within a context where

inferences could be uniquely identified as being made according to necessity or sufficiency.

In the instructions, the cover story was explained and participants were told to determine the causal relationship between two newly discovered bacteria. For participants in the sufficiency instruction group, instructions stated that the cause bacterium always accompanied the effect bacterium when one bacterium propagates the other bacterium (i.e., $P(\text{Effect}|\text{Cause}) = 1$), and that there are other causes in the environment that can produce the bacteria (i.e., $P(\text{Effect}|\neg\text{Cause}) > 0$). This information was not provided for participants in the control group.

In the learning phase, participants observed the states of bacteria on 16 trials. Six covariation conditions (.25-.00, .50-.00, .75-.00, 1.00-.75, 1.00-.50, and 1.00-.25) were used to determine whether participants interpreted covariation according to necessity (*PN*) or sufficiency (*PS*). Participants performed each condition twice in order to counterbalance the role of bacteria.

In the test phase, participants were told to judge the causal relationship in the same way as in Experiment 1. After a brief delay, participants completed the learning and test phases for the next condition. They were instructed that judgments should be made independently of their answers on prior problems.

Results and Discussion

Participants' responses were analyzed in a manner similar to Experiment 1. First, judgments were categorized as one of four types of causal models. Next, we classified participants into one of five clusters (i.e., necessity, sufficiency, necessity and sufficiency, random, and unclassified) according to whether their judgments were predicted by *PN*, *PS*, or *PNS*. Table 2 shows the number of participants assigned to each cluster. Participants in the sufficiency instruction group were largely divided into the necessity cluster and sufficiency cluster. In contrast, almost all participants in the control group were assigned to the necessity cluster, replicating Experiment 1 where the majority of participants interpreted covariation on the basis of necessity. Fisher's exact test confirmed that there were significantly more judgments according to sufficiency of causation for participants in the sufficiency instruction group than the control group ($p < .05$). Although some participants still interpreted covariation according to necessity, these results indicate that prior knowledge modulated the interpretation of covariation information.

In summary, Experiment 2 showed that judgments of

Table 2: Number of participants assigned to each cluster in Experiment 2

	Necessity	Sufficiency	Necessity and Sufficiency	Random	Unclassified
Sufficiency instruction Group	4	6	0	1	1
Control group	10	1	0	0	1

causal structure were largely affected by prior knowledge about causal relations. When participants were informed about high causal strength, they were more likely to infer causal directionality on the basis of sufficiency; in contrast, participants not given additional instructions always judged causal structure according to necessity. These results bridge the gap between the results showing that judgments of causal structure are based on necessity (Experiment 1) and those showing that judgments of causal structure are based on sufficiency (Mayrhofer & Waldmann, 2011).

General Discussion

Recent studies have shown that people use covariation to infer causal directionality. However, there is little information about how people infer causal directionality from covariation. The present study was designed to investigate how people make causal structure judgments on the basis of covariation. Experiment 1 demonstrated that judgments of causal structure vary as a function of covariation, and that participants' answers can be explained in terms of necessity of causation. Experiment 2 showed that prior knowledge about high causal strength led more participants to interpret covariation according to sufficiency. The results of Experiment 2 are consistent with both findings concerning necessity interpretation of covariation (Experiment 1) and sufficiency interpretation of covariation (Mayrhofer & Waldmann, 2011). These results reveal the importance of interpretations of covariation information in causal structure learning.

The results of the present study are closely related to the finding that learners flexibly interpret covariation during causal learning (Luhmann & Ahn, 2011). Luhmann and Ahn (2011) asked participants whether they interpreted single pieces of covariation information as evidence of generative or preventive causal relations. The results showed that observations from Cell A can be interpreted as evidence for either a generative or preventive causal relation. These studies share the view that covariation information is flexibly interpreted, but focus on different aspects of causal learning. Whereas Luhmann and Ahn (2011) focused on learning causal strength, the present study addressed learning causal structure. An intriguing question for future research is to ask participants whether they interpret covariation as evidence for X causes Y or Y causes X.

The difference between a necessity interpretation of covariation in the bacteria story (Experiment 1) and a sufficiency interpretation in the mind-reading aliens story (Mayrhofer & Waldmann, 2011) can be regarded as an interaction between domain-general causal inference and domain-specific knowledge. Whereas covariation is thought to be domain-general information, prior knowledge about causal relations seems to differ between the two stories. The results of Experiment 2 demonstrate that the basis for interpreting covariation can change from necessity to sufficiency when information about high causal strength is provided, but it remains unknown whether there are conditions that will change participants' basis for

interpretation from sufficiency to necessity. Another key question for future research is to investigate whether information about low base rate of the effect encourages a necessity interpretation of covariation in the mind-reading aliens cover story. Future research will provide further evidence about interactions between domain-general covariation and domain-specific prior knowledge.

References

- Devereett, B. & Kemp, C. (2012). Learning deterministic causal networks from observational data. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 288-293). Austin, TX: Cognitive Science Society.
- Gopnik, A., & Schulz, L. E. (2007). *Causal learning: Psychology, philosophy, and computation*. New York: Oxford University Press.
- Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inference as a rational process: The new synthesis. *Annual Review of Psychology*, 62, 135-163.
- Hume, D. (1739/2000). *A treatise of human nature*. Oxford, England: Oxford University Press.
- Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, 79, 1-17.
- Luhmann, C. C., & Ahn, W.-K. (2011). Expectations and interpretations during causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 568-87.
- Mayrhofer, R., & Waldmann, M. R. (2011). Heuristics in covariation-based induction of causal models: Sufficiency and necessity priors. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 3110-3115). Austin, TX: Cognitive Science Society.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, United Kingdom: Cambridge University Press.
- Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, 64, 93-125.
- Saito, M., & Shimazaki, T. (2012). Strategy changes in causal structure learning: The role of task complexity. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 2264-2269). Austin, TX: Cognitive Science Society.
- Slooman, S. A. (2005). *Causal models: how people think about the world and its alternatives*. New York: Oxford University Press.
- 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.