

# Differences between Observation and Intervention in Causal 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

Previous studies have suggested that learning is improved when people actively intervene rather than when they passively observe in causal structure learning tasks. Two experiments were conducted to investigate whether a facilitative effect will occur in the judgment of causal strength. In Experiment 1, participants were asked to learn causal strength in a situation where the target cause and context independently produced the effect. The intervention group could manipulate the state of the cause, which was later presented to the observation group (i.e., yoked-control procedure). The results demonstrated that participants made similar evaluations for the target cause, but not for the context. Experiment 2 was designed to examine whether different estimations were because of facilitation or bias in which participants undervalue other causes. The results provide support for a facilitative effect, but suggest that the improvement with intervention may be limited to the estimation of weak causal strength.

**Keywords:** causal reasoning; causal inference; intervention; causal power; yoked control procedure.

## Introduction

The ability to learn causal relations is essential for adapting to complex environments. Causal knowledge enables people to explain past events, to control present situations, and to predict future consequences. When someone catches a cold, for example, this might be attributed to viruses and lack of sleep. The person might take medicine to control this condition, with the expectation of getting better. It has been recognized that both children and adults easily represent causal relations (Gopnik & Schulz, 2007; Sloman, 2005; see also Holyoak & Cheng, 2011 for a review).

Information about causal relations can be acquired by passive observation and active intervention. Whereas learning by observation includes an observation of a system's autonomous behavior, learning by intervention involves an exogenous manipulation to the causal system that changes the state of the variable in some way (Pearl, 2000; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). In the above example, observation corresponds to seeing your friend take medicine, while intervention corresponds to making your friend take medicine. Both observation and intervention provide the state (i.e., presence or absence) of the cause and its effect, but they differ in how this information is conveyed.

Differences between observation and intervention have been well documented in structure learning tasks where participants are required to infer causal directions among events. Steyvers et al. (2003) investigated people's ability to infer causal structure from both observation and intervention. The results demonstrated that performance improved when people were allowed to intervene in the causal system. Similar results were obtained in Lagnado and Sloman (2004). The benefit of intervention is explained by the fact that interventions can discriminate Markov equivalent models which mere observation cannot. For instance, observations about co-occurrences imply statistical relations between two events X and Y; however, they cannot differentiate whether X causes Y or Y causes X. If event X is manipulated and nothing happens, the possibility that X causes Y is eliminated. The relation that Y causes X is confirmed by the fact that a manipulation of Y changes the state of X. Sobel and Kushnir (2006) revealed that learners were better at inferring causal structure when they intervened of their own free will, rather than observing data from another's intervention. They discussed that causal learning is facilitated because learners who intervened freely were able to receive the data in a desired order, and concluded that decision making is an important part of causal learning from intervention.

In contrast to many studies about causal structure learning, only a few studies have focused on differences between observation and intervention when people make judgments of causal strength. Whether intervention facilitates estimation of causal strength remains unknown. Hattori and Oaksford (2007) investigated the differences between observation and intervention when assessing causal strength. Participants were asked to learn the strength of the causal relation between using a particular type of fertilizer and plants blooming, by either observation or intervention. Their results demonstrated that people tended to ignore the information that both cause and effect were absent (i.e., cell  $d$  information in a  $2 \times 2$  contingency table) when learning with observation. However, the study is limited in two ways. First, they did not make comparisons between the causal ratings of the observation group and those of the intervention group, nor between participants' ratings and normative values, such as causal power (Cheng, 1997), because the experiment was conducted to establish the criteria for meta-analysis. In addition, covariation information for each group was programmed separately,

resulting in a difference in the presented information between the observation and intervention groups. Therefore, the measured difference between observation and intervention might result from the cause density effect, in which estimated causal strength becomes higher as the probability of the cause increases (e.g., Perales & Shanks, 2007), or from the outcome density effect, in which estimated causal strength becomes higher as the probability of the outcome increases (e.g., Shanks, 1985). This problem can be resolved by using a yoked control procedure. In this procedure, the same sequence of trials generated by a participant in the intervention group is presented to a yoked participant in the observation group. No previous research has provided direct evidence about the role of intervention in judgments of causal strength with a yoked control procedure.

Recent studies on causal learning from intervention have suggested that people can select the more informative situation to assess causal strength (e.g., Barberia, Baetu, Sansa, & Baker, 2010; Green & Over, 2009). Barberia et al. (2010) demonstrated that participants preferred to test generative causes in the low base rate context where the alternative causes have little effectiveness. In such contexts, the effectiveness of the target cause is not confounded by the effectiveness of the alternative causes. Learning by intervention includes the opportunity for informative selection; in contrast, learning by observation does not. Although the yoked control procedure equalizes the presented information, receiving information in a desired order would serve as a facilitator. Therefore, we hypothesized that intervention leads to more accurate estimation of causal strength.

The purpose of the present study is to investigate whether intervention facilitates evaluation of causal strength. In Experiment 1, participants were asked to learn causal strength in the situation where the target cause and context independently produced the effect. In Experiment 2, another cause was added in order to highlight the difference between observation and intervention. If intervention improves judgments of causal strength, ratings of the intervention group should be closer to the normative values than those of the observation group.

## Experiment 1

In order to investigate the difference between observation and intervention, we used a common-effect causal structure in which two causes independently produced the effect. One cause is the target cause which participants could manipulate; the other cause is the context which is always present regardless of the state of the target cause (i.e.,  $P(\text{context}) = 1$ ). Therefore, when participants observe the presence of the effect, they cannot be sure whether it occurred because of the target cause or because of the context. In the cover story of the experiment, the target cause and context corresponded to the missiles and minefield respectively. Participants were asked to assess the likelihood that the missile caused the tank to blow up. The

common-effect structure provides the state of the target cause with different informative values. When the target cause is present, participants receive information about the joint influence of the cause and context. In contrast, the absence of the target cause conveys information about the causal strength of the context. In order to precisely assess the causal power of the target cause, participants have to grasp the causal power of the context. The intervention group could select the information in their desired order.

The causal power of the target cause,  $q_{(\text{target cause})}$ , is the likelihood that the cause will generate an effect in a context with no alternative causes (Cheng, 1997). When the target cause and context influence the occurrence of the effect independently,  $q_{(\text{target cause})}$  is calculated as follows:

$$q_{(\text{target cause})} = \frac{P(\text{effect} \mid \text{cause, context}) - P(\text{effect} \mid \neg\text{cause, context})}{1 - P(\text{effect} \mid \neg\text{cause, context})}$$

In this equation,  $P(\text{effect} \mid \text{cause, context})$  is the probability of the effect given the presence of the cause and context.  $P(\text{effect} \mid \neg\text{cause, context})$  is the probability of the effect given the absence of the cause and the presence of the context. The causal power of the context,  $q_{(\text{context})}$ , is defined as follows:

$$q_{(\text{context})} = P(\text{effect} \mid \neg\text{cause, context})$$

These indices are easily calculated on the basis of the presented information. Participants were asked to estimate the causal power of the target cause and context. The distinctive feature of the present study lies in using a yoked control procedure and a common-effect causal model. The hypothesis predicts that people become better at inferring the causal strength when they are allowed to intervene in the target cause.

## Method

**Participants and design** Twenty-four undergraduates from Kwansei Gakuin University received course credit for taking part in this experiment. They were randomly assigned to one of two groups resulting from the manipulation of the type of learning (observation or intervention).

**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 researcher in military facilities attempting to determine the effectiveness of new anti-tank missiles. The term “effectiveness” means how likely the missile will blow up the tanks.

When a missile hits a tank, it causes the tank to blow up. However, the tank will not be destroyed if the missile fails to hit it. In addition, the target tank

runs through a minefield. Thus, when a tank explodes in the minefield, you cannot be sure whether it exploded because of the missile or because of a mine. Of course, if you do not launch a missile at the tank and it explodes, then you know that it must have exploded because of a mine. Note that the target tank always runs in a different area of the minefield and therefore a single explosion by a mine does not imply that the tank will never explode because of a mine.

Your task is to observe whether missiles cause the tank to blow up and to judge how effective the missiles are. Note that the experimental task does not require any knowledge of missiles and tanks. (The remaining instructions describe how to progress through the learning phase and test phase.)

**Learning Phase** The learning phase consisted of 40 trials that presented information about the launch of missiles and the tank's explosion. Participants were requested to observe the states of the missiles and tanks, and to infer causal relationships between them. They were randomly assigned to either the intervention or the observation group. Participants in the intervention group were told to decide whether to launch a missile. First, a tank was displayed on a screen. There were also two buttons on the screen: one for launching the missile, and the other for not launching the missile. After they made a choice, the state of the missile was presented. The presence of the missile was indicated by the picture of the missile. The absence of the missile was represented by the appearance of the missile labeled with a cross mark. At the same time, a button labeled "NEXT" was displayed on the screen. After clicking the button, the outcome of launching the missile was shown. The destruction of the tank was indicated by the appearance of an explosion; in contrast, the tank remained unchanged when the missile failed to hit. The screen was returned to its primary state 1.5s after the result was displayed.

The same procedure was used for participants in the observation group, except that they did not decide whether to launch the missiles. First, a tank and the state of the missile (presence or absence) were shown on the screen. After participants clicked the "NEXT" button, the state of the tank was provided. For each participant in the observation group, the states of the missiles and the subsequent outcomes were yoked to those of a participant in the intervention group. Thus, the covariation information between the two groups was identical. Each participant completed four contingency conditions of 40 trials. The different conditions were .25-.25, .75-.75, .75-0, and .75-.25 (Table 1). The first term refers to the probability of the tank blowing-up on a trial with a launched missile (i.e.,  $P(\text{effect}|\text{target cause, context})$ ) and the second term to the probability of the tank's destruction on a trial with no launched missile (i.e.,  $P(\text{effect}|\neg\text{target cause, context})$ ). In the .25-.25 condition, for example, the tank exploded 25% of the times, regardless of whether the missile was launched. That is, the missile had no causal power [ $q_{(\text{target cause})} = 0$ ] and

Table 1: Causal Power of Target Cause and Context in Each Condition in Experiment 1

Causal power	Contingency conditions			
	.25-.25	.75-.75	.75-0	.75-.25
$q_{(\text{target cause})}$	0	0	.75	.67
$q_{(\text{context})}$	.25	.75	0	.25

*Note.* The name of conditions represent  $P(\text{effect}|\text{target cause, context})$ ,  $P(\text{effect}|\neg\text{target cause, context})$  in order.

the mine had weak causal power [ $q_{(\text{context})} = .25$ ]. Similarly, the causal power of the missile was zero in the .75-.75 condition [ $q_{(\text{target cause})} = 0$ ]. The outcome density in the .75-.75 condition was higher than that in the .25-.25 condition because of the high causal strength of the mine [ $q_{(\text{context})} = .75$ ]. In contrast, the missile had strong causal power and the mine had no causal power in the .75-0 condition [i.e.,  $q_{(\text{target cause})} = .75$ ,  $q_{(\text{context})} = 0$ ]. In the .75-.25 condition, both missiles and mines caused the tank to blow up. Whereas the missile had strong causal power [ $q_{(\text{target cause})} = .75$ ], the mine had weak causal power [ $q_{(\text{context})} = .25$ ]. Thus, the causal power of the missile and mine differed in each contingency condition. In addition, the actual contingency varied slightly from participant to participant because the participants' actions were out of our control. To reduce the potential variance of the actual contingency, 90% of the outcome states were determined in a manner that converged with the programmed values and the remaining 10% were determined at random. The order of the contingency conditions was counterbalanced with the constraint that two ineffective missile conditions (i.e., .25-.25 condition, .75-.75 condition) did not follow each other.

**Test phase** After 40 trials had been completed, participants were asked to estimate the causal strength of the target cause and context. A rating scale was presented on the screen together with the question, "To what extent does the missile cause the tank to blow up?" A rating was made on a scale from 0 (the missile does not cause the tank to blow up at all) to 100 (the tank causes the tank to blow up every time). Subsequently, participants were requested to infer the causal power of the context in a similar manner. Then, after a brief delay, participants began the learning and test phases for the next contingency condition. They were instructed that their judgments should be made independently of their answers in prior conditions.

## Results and Discussion

Figure 1 shows the mean ratings for target cause and context in each condition. Separate analyses of variance for target cause and context with the type of learning (observation, intervention) as a between-participants factor, and the contingency condition (.25-.25, .75-.75, .75-0, .75-.25) as a within-participants factor were conducted. The analysis for the target cause revealed a significant effect of the contingency condition,  $F(3, 66) = 35.21$ ,  $MSE = 340.01$ ,  $p < .001$ ,  $\eta^2_G = .51$ , but no effect of the type of learning,  $F(1, 22) = 0.21$ ,  $MSE = 567.95$ ,  $p = .648$ ,  $\eta^2_G = .003$ . The

interaction between the type of learning and the contingency condition was not significant,  $F(3, 66) = 0.58$ ,  $MSE = 340.01$ ,  $p = .631$ ,  $\eta^2_G = .02$ . Individual comparisons of the contingency condition revealed significant differences among the four conditions ( $ps < .01$ ), except for the comparison between the .75-0 and .75-.25 conditions ( $p = .31$ ). These results suggest that both the intervention and observation group roughly differentiated the causal strength of the target cause in each contingency condition. Although the target cause had no causal power in the .25-.25 and .75-.75 conditions [i.e.,  $q_{(\text{target cause})} = 0$ ], causal ratings in .75-.75 condition were higher than that in .25-.25 condition. This pattern of results clearly demonstrates the outcome density bias, which suggests that the probability of the outcome positively affects judgments of causal strength (e.g., Shanks, 1985).

The parallel analysis for the context yielded significant main effects of the type of learning,  $F(1, 22) = 7.11$ ,  $MSE = 398.98$ ,  $p = .014$ ,  $\eta^2_G = .09$ , and the contingency condition,  $F(3, 66) = 59.93$ ,  $MSE = 282.15$ ,  $p < .001$ ,  $\eta^2_G = .65$ , but no interaction between the type of learning and the contingency condition,  $F(3, 66) = 0.35$ ,  $MSE = 282.15$ ,  $p = .790$ ,  $\eta^2_G = .01$ . Individual comparisons revealed significant differences among all conditions ( $ps < .001$ ), except for the comparison between the .25-.25 condition and the .75-.25 condition ( $p = .35$ ). As the causal power of the context can be obtained from the probability of the effect being present given the absence of the target cause (see Table 1), participants in the intervention group made more accurate estimations than those in the observation group.

Although the results of Experiment 1 are consistent with the hypothesis that intervention leads to more accurate estimation of causal strength, there is still an alternative interpretation to consider: the lower estimation of causal strength of the context might be due to the possibility that the intervention causes participants to undervalue alternative causes besides those in which they intervened. Indeed, Kushnir, Wellman, and Gelman (2009) reported the self-agency bias in which people weigh their own action as more

effective than the action of others. There remains a possibility that participants in the intervention group exhibited a self-agency bias and undervalued the causal power of the context. Experiment 2 was conducted to assess this possibility.

## Experiment 2

Experiment 1 demonstrated that participants made similar evaluations of the target cause regardless of whether they intervened or observed, but made different estimates of the causal power of the context. The results were interpreted as an indication that interventions lead to more accurate inferences for context. However, there is an alternative interpretation that intervention results in the decreased evaluations of the other causes due to a self-agency bias. Experiment 2 was designed to investigate these two interpretations. The experimental procedure was similar to that of Experiment 1, but another cause was introduced in addition to the target cause and context. Adding another cause enabled us to differentiate whether interventions facilitate the estimation of causal strength or decrease estimations of other causes. If people accurately estimate for causal strength with interventions, the deviations from normative values in the intervention group should be smaller than those in the observation group. If interventions simply make people undervalue other causes, participants in the intervention group should make light of the causal strength of another cause and context.

## Method

**Participants and design** Eighteen undergraduates from Kwansei Gakuin University participated in the experiment and received course credit. None of them took part in Experiment 1. As in Experiment 1, they were randomly assigned to either the observation or intervention group.

**Procedure** The procedure corresponded to the one in Experiment 1, except that another cause was introduced in addition to the target cause and context. In the instructions,

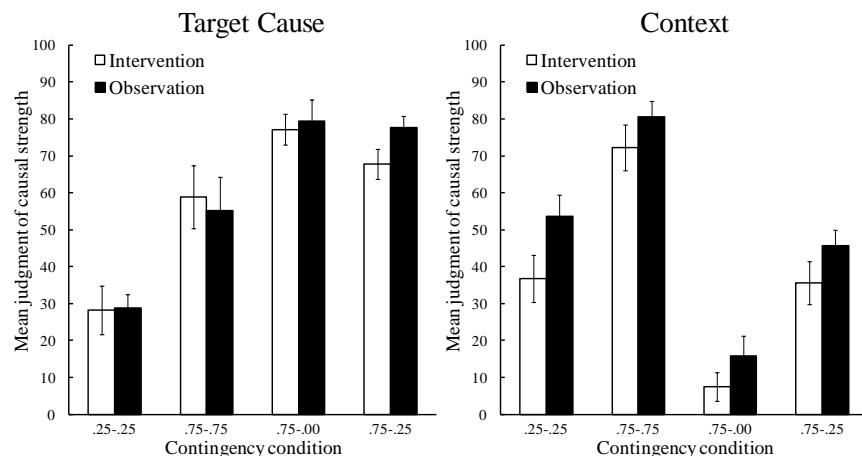


Figure 1. Mean causal ratings of the intervention and observation group in each contingency condition of Experiment 1, for target cause (left panel) and context (right panel). The error bars represent standard errors of the mean.

the cover story was explained and participants were told to infer the influence of the new missile (i.e., target cause) on the tank's explosion. Participants were informed that the tank's explosion could also be caused by the mine (i.e., context) or by another missile also targeted at the tank (i.e., another cause).

In the learning phase, participants received information about the target cause, another cause, and the context on 40 trials. Whereas the context was always present (i.e.,  $P(\text{context}) = 1$ ), another cause was present in the half of the trials (i.e.,  $P(\text{another cause}) = .5$ ). The probability that the target cause occurred in the trial depended on the choices of the intervention group. Participants in the intervention group could choose the state of the target cause with knowledge of whether another cause was present or not; in contrast, participants in the observation group were shown the states of the target cause and another cause. Each participant completed four contingency conditions. The different conditions were .75-.75-.75, .25-.75-.25, .75-.25-.25, and .75-.75-.25 (Table 2). The first term refers to the probability of the tank blowing up on a trial with a launched missile (i.e.,  $P(\text{effect}|\text{target cause}, \neg\text{another cause}, \text{context})$ ), and the second term to the probability of the tank's destruction on a trial with another launched missile (i.e.,  $P(\text{effect}|\neg\text{target cause}, \text{another cause}, \text{context})$ ), and the third term to the probability of the tank's explosion on a trial with no launched missile (i.e.,  $P(\text{effect}|\neg\text{target cause}, \neg\text{another cause}, \text{context})$ ). Thus, each condition differed in the causal power of the target cause, another cause, and the context.

In the test phase, participants were told to judge the causal strength of the target cause, another cause, and the context in the same way as in Experiment 1. After a brief delay, participants completed the learning and test phases for the next contingency condition. The order of the contingency condition was counterbalanced across participants via a Graeco-Latin square design.

Table 2: Causal Power of Target Cause, Another Cause, and Context in Each Condition in Experiment 2

Causal power	Contingency conditions			
	.75-.75-.75	.25-.75-.25	.75-.25-.25	.75-.75-.25
$q_{(\text{target cause})}$	0	0	.67	.67
$q_{(\text{another cause})}$	0	.67	0	.67
$q_{(\text{context})}$	.75	.25	.25	.25

Note. The contingency conditions represent  $P(\text{effect}|\text{target cause}, \neg\text{another cause}, \text{context})$ ,  $P(\text{effect}|\neg\text{target cause}, \text{another cause}, \text{context})$ , and  $P(\text{effect}|\neg\text{target cause}, \neg\text{another cause}, \text{context})$  in order.

## Results and Discussion

Figure 2 shows the mean ratings for target cause, another cause and the context in each condition. Separate analyses of variance for target cause, another cause, and the context with the type of learning (observation, intervention) as a between-participants factor, and the contingency condition (.75-.75-.75, .25-.75-.25, .75-.25-.25, .75-.75-.25) as a within-participants factor were conducted. The analysis for the target cause yielded only a significant main effect of the contingency condition,  $F(3, 48) = 13.26, MSE = 344.45, p < .001, \eta^2_G = .38$  (all other  $F$ s < 1). As in Experiment 1, both the observation and intervention groups demonstrated strong outcome density bias in the .75-.75-.75 condition where the target cause had no causal power. Although comparisons between participants' judgments and causal powers suggest that participants in the intervention group made more accurate estimations than those in the observation group in the .25-.75-.25 condition, the difference was not statistically significant.

The parallel analysis for another cause revealed a significant effect of the contingency condition,  $F(3, 48) = 25.60, MSE = 240.38, p < .001, \eta^2_G = .51$ , but no significant effect of the type of learning,  $F(1, 16) = 1.17, MSE = 377.59, p = .296, \eta^2_G = .02$ . There was no significant interaction,  $F < 1$ . If participants exhibited a self-agency bias in the

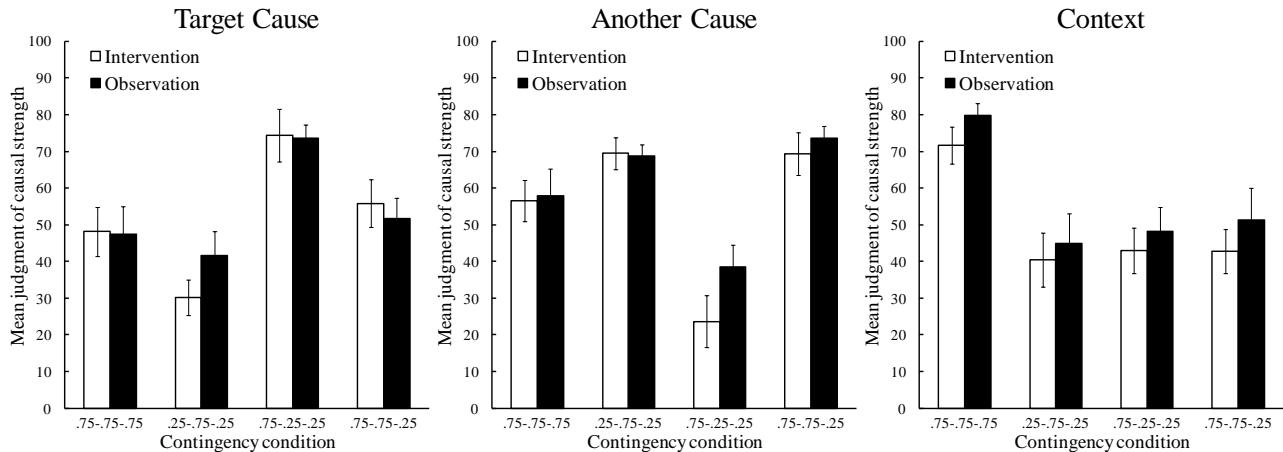


Figure 2. Mean ratings of the intervention and observation group in each contingency condition of Experiment 1, for target cause (left panel), another cause (center panel), and context (right panel). The error bars represent standard errors of the mean.

judgments of another cause, participants in the intervention group would make lower estimations than the observation group when the causal power of another cause was high. However, both groups made similar estimations when another cause had strong causal strength (i.e., .25-.75-.25 and .75-.75-.25 conditions). Instead, participants in the intervention group seemed to make better estimates when another cause had weak causal power (i.e., .75-.25-.25 condition).

The parallel analysis for context revealed a significant effect of the contingency condition,  $F(3, 48) = 13.56$ ,  $MSE = 315.13$ ,  $p < .001$ ,  $\eta^2_G = .34$ , but no effect of the type of learning,  $F(1, 16) = 1.27$ ,  $MSE = 613.26$ ,  $p = .276$ ,  $\eta^2_G = .03$ . The interaction between the type of learning and the contingency condition was not significant,  $F < 1$ . These results suggest that participants could differentiate the strong context from the weak context.

In order to investigate the difference between learning by observation and learning by intervention in greater detail, we calculated the root mean square error (RMSE) of the causal judgments (mean square deviation between judgment and causal power). As a result, the RMSEs of the intervention group ( $RMSE_{(target\ cause)} = 33.4$ ,  $RMSE_{(another\ cause)} = 33.3$ ,  $RMSE_{(context)} = 20.7$ ) were lower than those of the observation group ( $RMSE_{(target\ cause)} = 35.4$ ,  $RMSE_{(another\ cause)} = 36.7$ ,  $RMSE_{(context)} = 26.4$ ) in all types of the causal judgments. These results support the hypothesis that intervention leads to more accurate estimation of causal strength.

## General Discussion

The present study demonstrates that learning by intervention leads to more accurate judgments of causal strength than learning by observation. Although these results are consistent with the previous findings in causal structure learning tasks, the improvement with intervention is limited in two ways. First, intervention does not facilitate all judgments, but improves judgments when causal strength is weak. In order to accurately assess the weak causal power, people have to pay attention to information about negative relationships (i.e., cell  $b$  and  $c$  information in a  $2 \times 2$  contingency table). Maldonado, Jiménez, Herrera, Perales, and Catena (2006) reported that people tend to ignore such information in incidental situations. Since the intervention group could select the type of information they received, they might have been more likely to consider negative evidence into account than the observation group. That is, the benefit of the intervention might result from increased attention to information about negative relationships. This possibility can be assessed by asking participants to estimate the number of trials of each type that have been presented in learning phase. Second, participants overestimated the causal strength even though the target cause was irrelevant to the occurrence of the effect (i.e., outcome density bias). Since the effect is often present, there is little chance to disconfirm the positive relationship between action and outcome. Future research will reveal how intervention leads

to judgments that are more accurate and distorts judgments of non-contingent relationships.

## References

- Barberia, I., Baetu, I., Sansa, J., & Baker, A. G. (2010). Choosing optimal causal backgrounds for causal discovery. *Quarterly Journal of Experimental Psychology*, 63, 2413–2431.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- Gopnik, A., & Schulz, L. E. (2007). *Causal learning: Psychology, philosophy, and computation*. New York: Oxford University Press.
- Green, D. W., & Over, D. E. (2009). Valuing intervention and observation. *Quarterly Journal of Experimental Psychology*, 62, 1010–1022.
- Hattori, M., & Oaksford, M. (2007). Adaptive non-interventional heuristics for covariation detection in causal induction: Model comparison and rational analysis. *Cognitive Science*, 31, 765–814.
- 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.
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in preschoolers' causal inferences. *Developmental Psychology*, 45, 597–603.
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 30, 856–876.
- Maldonado, A., Jiménez, G., Herrera, A., Perales, J. C., & Catena, A. (2006). Inattentional blindness for negative relationships in human causal learning. *Quarterly Journal of Experimental Psychology*, 59, 457–470.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, United Kingdom: Cambridge University Press.
- Perales, J. C., & Shanks, D. R. (2007). Models of covariation-based causal judgment: A review and synthesis. *Psychonomic Bulletin & Review*, 14, 577–596.
- Shanks, D. R. (1985). Continuous monitoring of human contingency judgment across trials. *Memory & Cognition*, 13, 158–167.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. New York: Oxford University Press.
- Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, 34, 411–419.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453–489.