

Seeing the Unobservable – Inferring the Probability and Impact of Hidden Causes

York Hagmayer (york.hagmayer@bio.uni-goettingen.de)

Department of Psychology, University of Göttingen, Gosslerstr. 14
37073 Göttingen, Germany

Michael R. Waldmann (michael.waldmann@bio.uni-goettingen.de)

Department of Psychology, University of Göttingen, Gosslerstr. 14
37073 Göttingen, Germany

Abstract

The causal impact of an observable cause can only be estimated if assumptions are made about the presence and impact of possible additional unobservable causes. Current theories of causal reasoning make different assumptions about hidden causes. Some views assume that hidden causes are always present, others that they are independent of the observed causes. In two experiments we assessed people's assumptions about the occurrence and statistical relations involving a hidden cause. In the experiments, participants either only observed a cause or actively manipulated it. We assessed participants' assumption online after each learning trial and at the end of the learning phase. The results show an interesting dissociation. Whereas there was a tendency to assume negative correlation in the online judgments, the final judgments tended more in the direction of an independence assumption. It could also be shown that the judgments were generally coherent with the learning data. These results are consistent with normative theories that drop independence as the default assumption.

Introduction

Most events are causally influenced by more than a single cause. Unfortunately, very often these other causes are unknown or cannot be easily observed. Therefore we often have to rely on the observed statistical relationship between cause and effect when assessing causal strength. For example, whenever a new influenza virus invades East Asia, health representatives try to estimate its health risks, well aware of the fact that many other factors determine whether a patient will die or not. The question of how to assess causal strength when there are hidden causes has challenged normative theories of causality and psychological theories of causal reasoning for some time. A number of different accounts have been proposed analyzing how the causal impact of an observed factor can be accurately estimated if certain assumptions are made about potential hidden causes. In this report we will first give a brief overview of how two current theories of causal reasoning handle hidden causes. In the second part of the report we will present two experiments in which we assessed the assumptions of learners about the impact and probability of hidden causes. In the final section we will discuss potential theoretical implications of these findings.

Theoretical Accounts of Hidden Causes

We are going to focus on a simple causal structure consisting of a single observable cause C and one possible hidden cause A both influencing a joint observable effect E . The two observable events C and E are statistically related. Cause C is neither sufficient nor necessary for the effect, $P(e|c) < 1$ and $P(e|\sim c) > 0$. How can the causal impact of the observed and - if possible - the impact of the hidden cause be assessed in such a situation?

Associative Theories and the Constant-Background Assumption

Associative theories, such as the Rescorla-Wagner theory (Rescorla & Wagner, 1972), would model this task as learning about the association between a cue representing the observed cause and an outcome representing the effect. Along with the cause cue a second background (or context) cue would be part of the model. This background cue is assumed to be always present and to represent all other factors that might also generate the outcome. Thus, the background cue would play the role of representing the hidden cause A in the outlined causal model. According to the Rescorla-Wagner rule, only weights of cues that are present in a certain trial are being updated. Therefore the permanently present background cue will generally compete with the cause cue in cases in which the cause cue is present. If the outcome is also present the associative weights of both cues will be raised, if the outcome is absent, the weights will be lowered. However, in cases in which the cause cue is absent, only the weight of the background cue will be altered. At the asymptote of learning the associative weight of the observed cause will equal the contingency (i.e., $\Delta P = P(e|c) - P(e|\sim c)$) of the cause cue and the outcome. The associative weight of the background cue will correspond to the probability of the outcome in the absence of the cause cue. Thus, the more often the outcome (=effect) occurs on its own, the higher the associative weight of the background cue will be.

Power PC Theory and the Independence Assumption

Cheng's (1997) Power PC analysis of the causal impact of a single cause can be viewed as a special case of a causal Bayes net in which two causes independently influence a joint common effect (Glymour, 2001, Tenenbaum & Griffiths, 2003). The theory states that the occurrence of the effect E is a consequence of the causal powers of the observed cause C and a hidden cause A (p_c and p_a), and of their base rates $P(c)$ and $P(a)$. Formally the probability of the effect equals the sum of the base rates of the two causes multiplied by their causal power minus the intersection of the causes multiplied by both causal powers:

$$P(e) = P(c) \cdot p_c + P(a) \cdot p_a - P(c) \cdot P(a) \cdot p_c \cdot p_a.$$

Therefore the probability of the effect E given that the observed cause C has occurred is

$$P(e|c) = p_c + P(a|c) \cdot p_a - P(a|c) \cdot p_c \cdot p_a \quad [1],$$

and the probability of the effect given that the observed cause is absent is

$$P(e|\sim c) = P(a|\sim c) \cdot p_a \quad [2].$$

Equations [1] and [2] offer an account for hidden causes irrespective of whether they are dependent or independent of the observed cause C . However, if they happen to covary, the power of the causes cannot directly be estimated by the observable data because there are four unknown parameters to be estimated by two observable conditional probabilities. Therefore Power PC theory makes the assumption that the observed and the hidden causes are independent, $P(a|c) = P(a|\sim c) = P(a)$. Based on this assumption the causal power of the observable cause can be calculated by

$$p_i = (P(e|c) - P(e|\sim c)) / (1 - P(e|\sim c)).$$

The independence assumption of Power PC theory implies that the probability of the hidden cause stays the same regardless of whether the observed cause has occurred or not. If this assumption holds, Equation [2] defines lower boundaries for the base rate and the causal strength of the hidden cause. The causal power of the hidden cause and its base rate have to be at least as big as the probability of the effect in the absence of the observed cause, $p_a \geq P(e|\sim c)$ and $P(a) \geq P(e|\sim c)$. Equation [2] also defines a coherence criterion for estimates about hidden causes. In order to be compatible with the observed data, the estimates must honor this equation.

It is important to note that even if independence is not assumed, Equations [1] and [2] still hold and have implications for the unobservable cause. The power of the hidden cause and its probability in the absence of the observed cause are still determined by Equation [2]. Therefore estimates for both values should be constrained by $P(e|\sim c)$. Moreover, Equation [1] provides constraints for the admissible probabilities of the hidden cause in the presence of the observable one. However, this constraint is fairly complex and does not provide the same straightforward implications as Equation [2].

Summary

Both theories consider hidden causes. Associative theories assume that a hidden cause (i.e., the constant background) is always present. In contrast, Power PC and other causal

Bayes net theories assume that the hidden cause is independent of the observed cause and that its probability is constrained by the data. The probability of the effect in the absence of the cause marks its lower boundary. These theories also permit to model statistical dependence between the observed and the hidden causes.

Both theoretical accounts agree that $P(e|\sim c)$ is to a certain degree indicative of the causal strength of the hidden cause. But whereas associative theories generally regard this probability as a valid indicator, Power PC and other causal Bayes net theories view this conditional probability as a lower boundary of the causal impact of the hidden cause.

Experiments

The following two experiments explore what assumptions participants make about the presence and impact of a hidden cause in a trial-by-trial learning task, and whether these assumptions conform to the predictions of any of the discussed theoretical models. Thus far very little research has been conducted about naïve participants' assumptions about hidden causes. An exception is a study by Luhmann and Ahn (2003). They found that participants judged the impact of a hidden cause to be higher if $P(e|\sim c)$ was 0.5 than if it was zero. The experiments presented in this report will go beyond these findings. In addition to causal strength estimates, we collected assessments of the probability of the hidden cause using different kinds of measures. We also varied the learning conditions.

In both experiments participants learned about the causal relation between an observable cause and a single effect. Additionally participants were told that there was one other possible but unobservable cause of the effect. The statistical relation between the observable cause and the effect was manipulated in the two experiments while either keeping the contingency (Experiment 1) or the causal power (Experiment 2) constant. In Experiment 1 participants could only passively observe the cause, which occurred at its natural base rate, in Experiment 2 participants were allowed to manipulate the cause. A number of dependent variables were collected to assess participants' estimates of the probability of the hidden cause and the impact of both the observed and the unobserved causes. Participants were asked to guess the presence of the hidden cause on each trial during learning, and they were asked to give summary estimates after learning was completed. In one condition ("prediction before effect") participants were first informed about the presence or absence of the cause in each trial, and then they were asked to guess the presence of the hidden cause without receiving feedback about this alternative cause. Finally they were informed whether the effect has occurred at this particular trial or not. Predictions of the hidden cause prior to effect information can only be guesses based on observed frequencies of the effect in past trials. Based on normative theories (e.g., Power PC theory) we expected participants to generate independence between the causes. In the second condition ("prediction after effect") participants received information about the presence of both the cause and the effect and then had to predict the hidden cause. As before no feedback was provided about the hidden cause. In this situation participants had complete information about the cause

and the effect which should allow them to make more informed guesses about the hidden cause, especially if the observed cause is absent: If in this case the effect is present, participants should conclude that the hidden cause is also present. However if the effect is absent, they should have the intuition that the hidden cause is absent. Predictions based on the presence of the observed cause are more difficult. If in this case the effect is absent, participants should infer that the hidden cause is more likely to be absent than present; if the effect is present the hidden cause should also be given a higher probability of being absent. Based on the theories outlined above, we expected that participants in both conditions would generate independence between the causes in their trial-by-trial predictions. A third control condition left out the trial-by-trial predictions. In this condition participants rated the causal strength of the observed and the hidden cause as well as the probability of the hidden cause in the presence and in the absence of the observed cause after the learning phase. Again we expected participants to rate the causes to be independent. We also expected that the strength ratings for the observed cause would be based on causal power, and that the ratings for the hidden cause would be influenced by $P(e|\sim c)$.

Experiment 1

With Experiment 1 we pursued two goals. The first was to investigate whether participants would assume independence between the observable and unobservable cause. The second goal was to find out whether the power estimates for the unobservable cause would be influenced by the probability of the effect in the absence of the observed cause. Participants were given the task to assess the causal relation between a fictitious microbe (“colorophages”) and the discoloration of certain flowers. In addition they were told that there was only one other possible cause of the effect, an infection with another fictitious microbe (“mamococcus”), which was currently not observable. Participants were then directed to a stack of index cards providing information about individual flowers. The front side of each index card showed whether the flower was infected by colorophages or not, and the backside informed about whether the flower was discolored or not. Then participants were instructed about the specific learning procedure in their condition. The learning conditions were manipulated as a between-subjects factor. In Condition 1 (“prediction before effect”) participants were first shown the front side of the card, then they had to guess whether the flower was also infected by the other microbe, and finally the card was turned around by the experimenter revealing whether the flower was in fact discolored or not. In contrast, in Condition 2 (“prediction after effect”) the card was first turned around and then the participant made her guess about the hidden cause. Guesses were recorded without giving feedback. In the third, control condition cards were simply shown and turned around by the experimenter.

As a second factor the statistical relation between the observed microbe and discoloration was manipulated. Three different data sets consisting of 20 cases each were constructed. Table 1 summarizes the statistical properties of the three data sets. As the table shows, the contingency ΔP was

constant across the data sets, whereas both $P(e|\sim c)$ and causal power were rising. All three data sets were shown to every participant in a within-subjects design. Different data sets were introduced as data from different species of flowers. It was pointed out to participants that the effectiveness of the microbes might vary depending on the species. The order of the presented data sets was counterbalanced.

Table 1: Statistical properties of data sets shown in Experiment 1

	Data Set 1	Data Set 2	Data Set 3
$P(c)$	0.50	0.50	0.50
$P(e c)$	0.60	0.80	1.00
$P(e \sim c)$	0.10	0.30	0.50
ΔP	0.50	0.50	0.50
Power p_c	0.56	0.71	1.00

After each learning phase participants were asked to rate the causal influence of the observed and the hidden cause on a scale ranging from 0 (“no impact”) to 100 (“deterministic impact”). Participants were also asked to estimate how many of ten flowers that were infected with the observed microbe were also infected with the other microbe, and how many of ten flowers that were not infected with the observed microbe were instead infected with the other microbe. No feedback was provided about these assessments.

36 students from the University of Göttingen were randomly assigned to one of the learning conditions. Figure 1 shows the mean ratings of the impact of the observed and the hidden causes.




Figure 1: Mean ratings of causal impact for the observed cause (left) and the unobserved cause (right) in Exp. 1.

An analysis of variance revealed a significant increase in impact ratings for the observed cause, $F(2,66)=12.7$, $MSE=296.6$, $p<.01$, supporting the predictions of Power PC theory. The same analysis for the hidden cause resulted also in a significant main effect of the factor data set, $F(2,66)=4.92$, $MSE=408.1$, $p<.05$, which indicates that with increasing $P(e|\sim c)$ participants tended to assume a stronger impact of the hidden cause. This result is in accordance with the predictions of all theoretical accounts. However, the interaction between data sets and learning condition also

turned out to be significant, $F(4,66)=4.55$, $MSE=408.1$, $p<.05$. The observed increase was strongest in the ‘prediction after effect’ condition followed by the control condition. This interaction might be due to the learning procedure. In the ‘prediction after effect’ condition participants were sensitized to the possible presence and impact of the hidden cause more than in the other two conditions. Being informed about the occurrence of the effect in the absence of the observable cause is a strong cue pointing to the presence of the hidden cause.

Table 2 shows the results concerning participants’ assumptions about the dependence between the causes. The online predictions of the hidden cause in the presence and absence of the target cause were transformed into conditional frequencies, and combined into subjective contingencies, $\Delta P=P(a|i) - P(a|\sim i)$. On the left side of Table 2 the generated contingencies underlying online predictions are listed, the right hand side shows the corresponding contingencies based on the final probability ratings.

Table 2: Mean estimates of dependence between observed and unobserved cause. Numbers indicate contingencies (possible range:-100 to +100).

	Generated Dependence			Estimated Dependence		
	Data Set 1	Data Set 2	Data Set 3	Data Set 1	Data Set 2	Data Set 3
Before Eff.	33.3	8.3	-21.7	30.0	5.8	-3.3
After Eff.	17.0	17.5	-28.8	0.8	0.0	-22.6
Control	-	-	-	14.4	5.8	-4.2

An analysis of variance of the generated contingencies yielded a significant trend from positive to negative assessments which proved independent of learning condition, $F(2,44)=22.1$, $MSE=749.8$, $p<.01$. The estimated contingencies showed a similar trend, $F(2,66)=4.2$, $MSE=753.9$, $p<.05$. The mean contingencies in the different data sets varied slightly across learning conditions, $F(2,33)=2.8$, $MSE=1646.4$, $p<.10$. The generated contingencies were significantly above zero if $P(e|\sim c)$ was zero, and significantly below zero if $P(e|\sim c)$ was 0.5. The estimated contingencies showed a similar but only marginally significant pattern. Thus, there was a hint of a dissociation between online and post hoc assessments which will be followed up in Experiment 2.

These results do not conform to the theoretical assumptions of the discussed theories. Participants did not assume that the hidden cause was always present or that the two causes were independent.

A closer analysis of the conditional probabilities revealed that the negative trend was due to an increase in the generated and estimated probability of the hidden cause in the absence of the observed cause, whereas the subjective probability of the hidden cause in the presence of the observed cause remained relatively stable. This pattern is in part consistent with the analysis outlined in the introduction, $P(a|\sim c)$ seems directly constrained by $P(e|\sim c)$. In contrast, the con-

straint for $P(a|c)$ is more complex, which may be the reason why participants had more difficulties honoring it.

Even if participants’ answers did not conform to the independence assumption, their answers still might be coherent with the observed data. Both Power PC theory and Bayesian models can model dependence between observed and hidden causes. Although precise power estimates might be impossible, the data still yields constraints on plausible estimates. The most important constraint is that the product of the causal power (or strength) of the hidden cause and the probability of the hidden cause in the absence of the observed cause must equal the probability of the effect in the absence of the observed cause. To find out whether participants honor this constraint we used their ratings to recalculate the probability of the effect when cause C was absent:

$$P(e|\sim c)_{rec} = \text{Rating Impact } A \cdot \text{Rating } P(a|\sim c).$$

The results are shown in Figure 2. It can be seen that the recalculated probabilities in the ‘prediction after effect’ condition were surprisingly close to the actually observed probabilities. In contrast, the recalculated probabilities in the other two conditions were inaccurate. Apparently, participants had to be sensitized by the learning procedure to the presence and impact of the hidden cause to be able to derive coherent estimates. Learning that the effect is present in the absence of the target cause apparently provided the necessary information to make educated guesses about the hidden cause. Without this information the guesses showed some systematicity but did not conform very well to the observed data.

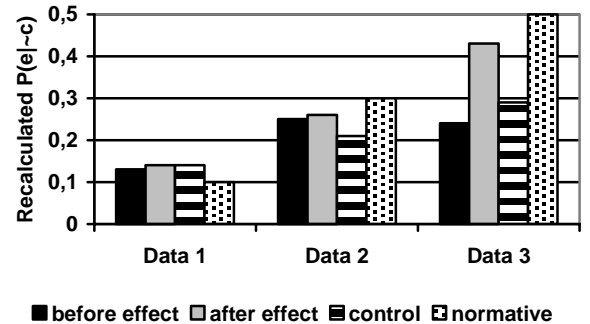


Figure 2: Mean recalculated probabilities of the effect in the absence of the observed cause (Experiment 1).

Experiment 2

In Experiment 1 we used a scenario in which the observable cause could only be passively observed. Therefore a dependence of the observed and unobserved cause was possible and maybe for some participants plausible. In Experiment 2 we allowed participants to arbitrarily manipulate the observable cause. Since these random interventions cannot be based on the presence or absence of the hidden cause, they should make the independence between the alternative causes more salient than in the observation context. Thus, we expected that participants would now assume the causes to be independent in all conditions of the present experiment.

Participants were instructed to imagine being a captain on a pirate ship firing his battery at a fortress. A second ship, which cannot be seen, was also firing at the fortress. Participants had a certain number of shells available and had to decide on each trial whether to fire or not. This procedure ensured that all participants saw the same data despite the fact that they set the cause themselves. The three learning conditions of Experiment 1 were used again. Participants had to guess whether the other ship currently fires either before they were informed about the occurrence of an explosion in the fortress (“prediction before effect”), or they had to predict the other ship’s action after they had learned whether the fortress was hit (“prediction after effect”). In a third, control condition no predictions were requested.

Three data sets consisting of 60 cases each were constructed. Table 3 shows the statistical properties of the data. In contrast to Experiment 1 the contingency between the observed cause and the effect decreased across the data sets, whereas the causal power remained stable. Participants learned about all the three data sets with order being counterbalanced.

Table 3: Data shown in Experiment 2

	Data Set 1	Data Set 2	Data Set 3
$P(c)$	0.50	0.50	0.50
$P(e c)$	0.70	0.80	0.90
$P(e \sim c)$	0.00	0.33	0.67
ΔP	0.70	0.47	0.23
Power p_c	0.70	0.70	0.70

60 students from the University of Göttingen were randomly assigned to one of the three learning conditions. The same dependent variables as in Experiment 1 were collected.

Figure 3 shows the results for the estimates of the causal impact of the observed and the hidden cause.




Figure 3: Mean ratings of causal impact for the observed cause (left) and the unobserved cause (right) in Experiment 2.

An analysis of variance of the impact ratings for the observed cause yielded no significant effects, which is in line with the predictions of Power PC theory. As in Experiment 1 the estimated impact of the hidden cause rose significantly

across the data sets, $F(2,114)=65.7$, $MSE=408.2$, $p<.01$. This finding is consistent with the predictions of all discussed theories. There was also a significant difference between learning conditions, $F(2,57)=4.06$, $MSE=591.8$, $p<.05$. Participants in the ‘prediction after effect’ condition rated the impact of the hidden cause to be higher than in the other two conditions. This results points in the same direction as the results of Experiment 1 indicating that predictions with effect information may have sensitized participants to the role of the hidden cause.

Table 4: Mean estimates of dependence between observed and unobserved causes. The numbers express contingencies.

	Generated Dependence			Estimated Dependence		
	Data Set 1	Data Set 2	Data Set 3	Data Set 1	Data Set 2	Data Set 3
Before Eff.	-23.8	-33.2	-29.2	-4.1	8.2	-8.2
After Eff.	-3.9	-20.4	-35.4	12.8	-10.8	-1.0
Control	-	-	-	9.5	9.2	-6.0

Table 4 shows the results concerning the assumed dependence between the two causes. Although the random interventions were expected to increase the salience of independence, participants generated a negative dependence between the two causes which rose across the data sets, $F(2,76)=6.97$, $MSE=510.1$, $p<.01$. The interaction also turned out to be significant, $F(2,76)=3.57$, $MSE=510.1$, $p<.05$. The negative ratings decreased more strongly when participants had received effect information than in the contrasting condition (“prediction before effect”). As in Experiment 1 this trend can be traced to an increase in the generated probability of the hidden cause in the absence of the observed cause. In contrast, the estimated dependencies did not statistically differ. The results are consistent with an independence assumption. Thus, in this experiment there was a clear dissociation between online and posthoc judgments.




Figure 4: Mean recalculated probabilities of the effect in the absence of the observed cause (Experiment 2).

Figure 4 is based on an analysis of the coherence of the estimates with the constraints from the learning data (using the same method as in Experiment 1). It can again be seen that participants honored the normative constraints.

Conclusions

All current theories of causal reasoning consider hidden causes that may also influence observable effects. In most theories independence between the observable cause and the hidden cause is the default assumption, which is a precondition for giving precise estimates for the causal strength of the observed cause-effect relation. Whereas associative theories create independence by assuming constant presence of alternative causes, Power PC theory and Bayesian models are more flexible. Typically these models assume a varying independent hidden cause. However, these theories can also model situations violating the independence assumption by providing bounds for consistent estimates. All theoretical accounts agree that the impact of the hidden cause has to be at least as high as $P(e|\sim c)$. We found evidence in both experiments that participants honored this constraint. Moreover, our analyses showed that participants' judgments about the probability and impact of the hidden cause were in most conditions coherent with the data.

Furthermore, we assessed participants' assumptions about the statistical relation between the observed and the hidden causes. In Experiment 1 learners passively observed the causal relations. In this experiment participants expressed that the causes were positively correlated when $P(e|\sim c)$ was low but they assumed a negative correlation when $P(e|\sim c)$ was high. The generated and estimated probabilities suggest that participants may have assumed that $P(e|\sim c)$ is an indicator of the probability of the hidden cause in the absence of the observed cause ($P(a|\sim c)$) and an indicator of the impact of the hidden cause (p_a) but that $P(e|c)$ conveys little information about the probability of the hidden cause conditional upon the presence of the observed cause ($P(a|c)$). As a consequence participants only adapted their guesses about $P(a|\sim c)$ to the observed $P(e|\sim c)$ while sticking with the initial assumption about $P(a|c)$ irrespective of $P(e|c)$.

In Experiment 2 we increased the salience and plausibility of independence between the alternative causes by letting participants randomly manipulate the observable cause. And indeed the final estimates expressed the assumption of independence. However, surprisingly participants generated a negative correlation in their trial-by-trial predictions. Using the explanation we gave for Experiment 1, this pattern implies that the initial assumption of $P(a|c)$ was at a relatively low level. People may find it unlikely that two independent actions are performed simultaneously by coincidence. In addition participants may erroneously overapply the 'principle of explaining away' (Pearl, 1988) in this task. This principle states that it is generally true that alternative independent causes are less likely in the subset of events in which the cause and effect are present as compared to the whole set of events in which only the effect has occurred. However, in the overall set of events causes should still exhibit independence regardless of the order in which the causal events are experienced. Another related possible explanation might be that people are reluctant to consider overdetermination of

effects. Since one cause suffices to explain the effect, assuming a second hidden cause is not necessary. Intuition tells us that one cause is enough for the presence of an effect. It is interesting to see that this intuition seems particularly strong when participants consider single trials of cause-effect patterns. In this situation learners have to decide whether one or two causes generated the effect. Looking back at the learning set at the end of the learning phase seems to decrease the salience of these possible cases of overdetermination, which may be the reason for the interesting dissociation between the tendency to assume negative correlations in online judgments but independence in the summary judgments at the end.

Theoretical Implications

Our results contradict the assumption of associative theories that learners assume constant presence of alternative, hidden causes. The results also indicate that independence of varying causes is not the general default assumption. The online predictions revealed a tendency to assume correlations between alternative causes. Both Power PC theory and causal Bayes nets allow modeling this assumption. Although causal power may in these cases not always be numerically identifiable, these theories can provide constraints for plausible estimates. Future research will have to explore the boundary conditions and the generality of people's assumption across different tasks. The observed dissociations in the present studies indicate that a simple account may be unlikely.

References

- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.
- Glymour, C. (2001). *The mind's arrow. Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- Luhmann, C. C., & Ahn, W.-K. (2003). Evaluating the causal role of unobserved variables. *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society*, Mahwah, NJ: Erlbaum.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA: Morgan Kaufmann.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, MA: Cambridge University Press.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II. Current research and theory* (pp. 64-99) New York: Appleton-Century-Crofts.
- Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal inference. *Advances in Neural Information Processing Systems*, 15, 35-42.