

Research Article

When Is a Cause the “Same”?

Coherent Generalization Across Contexts

Mimi Liljeholm and Patricia W. Cheng

University of California, Los Angeles

ABSTRACT—Two competing psychological approaches to causal learning make different predictions regarding what aspect of perceived causality is generalized across contexts. Two experiments tested these predictions. In one experiment, the task required a judgment regarding the existence of a simple causal relation; in the other, the task required a judgment regarding the existence of an interaction between a candidate cause and unobserved background causes. The task materials did not mention assessments of causal strength. Results indicate that causal power (Cartwright, 1989; Cheng, 1997) is the mental construct that people carry from one context to another.

Judgments about cause-and-effect relations occur in contexts that are like rivers—one never steps into the same context twice. Generalization from one context to another is therefore paramount. In fact, generalizing from the learning context to whatever new context may come is the *raison d'être* of causal learning. What is it that a reasoner carries from one situation into another? Let us formulate this question more specifically in terms of the scenarios illustrated in Figure 1.

Imagine that you are presented with data from two studies, conducted in different laboratories, that tested the influence of two allergy medicines on headache (a possible side effect). In each study, allergy patients were randomly assigned to two groups: a treatment group and a no-treatment (i.e., control) group. In the first study (Fig. 1a), Medicine A alone was administered. In the second (Fig. 1b), Medicines A and B were administered in combination. Headache (indicated by a frowning face, as opposed to a smiling face) occurred with a different frequency in each of the four groups. What is your best bet, based on the results from both studies, regarding whether or not Medicine B causes headache? Presumably, if you perceive a “change” in the results across treatments (i.e., Medicine A in one study and both medicines in the other), you might attribute

this change to the introduction of Medicine B in the second study. But what constitutes a “change”? To put the question another way, what is assumed to be invariant, and hence to generalize, across contexts?

Although the target question in this scenario concerns Medicine B, the psychological representation of interest concerns Medicine A, the medicine that occurs in both contexts. One could ask a direct question about the generalization of Medicine A across contexts. The very wording of a direct question, however, might bias the answer toward one model or another (for a discussion of the striking influence of wording of causal questions on responses, see Buehner, Cheng, & Clifford, 2003). Introducing Medicine B into the scenario and letting it be the target of the question allows one to assess implicitly and without bias what perceived aspect of Medicine A generalizes across contexts. In the rest of this article, we briefly review two approaches to causal learning and present two experiments that tested hypotheses about generalization across contexts, making use of scenarios such as those in Figure 1.

CAUSAL VERSUS PURELY COVARIATIONAL ACCOUNTS OF CAUSAL LEARNING

Causal relations encapsulate how the world works. A classic problem in the field of artificial intelligence is the frame problem (McCarthy & Hayes, 1969): Given the vast amount of empirical information that is available at each moment in each situation, which kinds of information are the most relevant across time and contexts and therefore should be selected for representation? A prevailing answer is: causal relations (e.g., Pearl, 2000; Spirtes, Glymour, & Scheines, 1993/2000). Concurrently in psychology, causal learning has emerged as an important topic (e.g., Blaisdell, Sawa, Leising, & Waldmann, 2006; Dickinson, Shanks, & Evenden, 1984; Glymour, 2001; Griffiths & Tenenbaum, 2005; Waldmann & Holyoak, 1992). The two dominant approaches to causal learning—the causal approach (e.g., Cheng, 1997) and the purely covariational approach (e.g., Rescorla & Wagner, 1972)—make precise and fundamentally different

Address correspondence to Mimi Liljeholm, Department of Psychology, Franz Hall, University of California, Los Angeles, CA 90095-1563, e-mail: mlil@ucla.edu.




Fig. 1. Illustration of two hypothetical studies testing the influences of Medicines A and B on headache. The illustrations at the top show patients who were not exposed to any medicine; the bottom illustrations show patients who were exposed to either (a) Medicine A alone or (b) Medicines A and B in combination. In these studies, the causal power of the treatment (an estimate of a distal property) remains constant across studies while ΔP (a proximal property) varies. Values of Rescorla and Wagner's (1972) model (RW) and of the causal-support model (Griffiths & Tenenbaum, 2005) are also listed. Values of causal support were generated with a noisy-OR function (see the section on Independent Causal Influence).

predictions regarding what remains invariant across contexts (Cheng, 2000; Cheng, Novick, Liljeholm, & Ford, 2007). At the heart of the debate between these approaches is the issue of whether or not reasoners make (tacit) generic assumptions about causal events in the distal world, events that (as Hume, 1739/1987, noted) are unobservable. Assumptions about unobservable events, *ceteris paribus*, are understandably objectionable. In this article, we argue that the payoff of such assumptions is the capability of making coherent generalizations across contexts.

Although some studies have investigated generalization to novel causal contexts (e.g., Lien & Cheng, 2000; Povinelli, 2000), none have tested the distinct predictions of the two approaches. The experiments reported here tested these predictions.

ΔP and Causal Power: Two Psychological Accounts of Causal-Strength Estimation

Purely covariational models of causal learning attempt to sidestep assumptions regarding unobservable distal events. Ac-

cording to a well-established model of covariation, the ΔP rule (Jenkins & Ward, 1965), adapted to apply to subsets of events in which causes other than the candidate cause remain constant (Cheng & Holyoak, 1995), reasoners contrast two probabilities: $P(e|c)$, the probability of a target effect e given the presence of a candidate cause c , and $P(e|\bar{c})$, the probability of e given the absence of c :

$$\Delta P_c = P(e | c) - P(e | \bar{c}) \quad (1)$$

These conditional probabilities are directly estimable by the observable relative frequencies of the relevant events. Both e and c are binary variables. Depending on whether ΔP is greater than, less than, or not noticeably different from 0, c is assumed to be a generative cause, preventive cause, or noncause of e , respectively. Under a set of conditions in which learning is at asymptote, the ΔP rule is equivalent to Rescorla and Wagner's model (1972; see Danks, 2003), and both accounts have been adopted to model causal strength (e.g., Spellman, 1996). For all conditions in our experiments, these accounts make the same ordinal predictions.

In Figure 1, ΔP of the treatment increases from .25 in the first study to .75 in the second study. Therefore, if judgments are based on ΔP , people should perceive a change in the results across treatments and, accordingly, infer that Medicine B causes headache.

In contrast, according to the *causal-power theory of the probabilistic contrast model* (Cheng, 1997; Novick & Cheng, 2004; the *power PC* theory for short), the reasoner explains observable probabilistic contrast (e.g., the covariation between cigarette smoke and lung cancer) by postulating unobservable causal relations in the distal world (e.g., lung cancer being caused by cigarette smoke, background causes, or both). The unobservable causal power (i.e., strength; Cartwright, 1989) of these relations is estimable under a set of default generic causal assumptions. For example, two of these assumptions (which can be revised in light of evidence) are that candidate c influences e independently of *background* causes (i.e., all causes of e other than c occurring in the context), and that the latter do not prevent e . It can be shown that under this set of assumptions, if, in addition, background causes of e are believed to occur independently of c , then when ΔP is greater than or equal to 0, the *generative power* of c with respect to e —the (ideally invariant) unobservable probability with which c produces e —can be estimated as follows:

$$\text{generative power of } c = \frac{P(e|c) - P(e|\bar{c})}{1 - P(e|\bar{c})} \quad (2)$$

The preventive analogue of this equation applies when ΔP is less than or equal to 0 (see Cheng, 1997).

For both studies in Figure 1, the causal power of the treatment is .75, according to Equation 2. Thus, if change is defined by causal-power values, people should judge that Medicine B does not cause headache. Critical to this prediction is the assumption that c and the background causes exert independent influences. If the influence of a cause depends on (i.e., the cause interacts with) context-specific and potentially unobserved background causes, knowledge acquired in one context will not apply in another.

Independent Causal Influence

Let us consider what independent causal influence means. If two or more causes *independently influence* an effect, then when they are jointly present, each operates on the effect as if the other cause or causes were not there. We use Figure 1 to illustrate this concept for Medicine A and the background cause (omit Medicine B from the second study for our purpose here). Suppose that Medicine A is administered to the patients in both control groups (top halves of the two studies illustrated in Fig. 1) and that the medicine causes headache with a probability of 3/4 in each individual, operating in the same way in both studies (i.e., as if the background causes were not there). The resulting frequencies of headache would be consistent with what appears in

the treatment groups (bottom halves of the two studies in the figure), in accordance with the definition of independence in probability theory.

For the treatment group in Figure 1a (bottom left), the probability of “headache being caused independently by both the medicine and the background cause,” an unobservable event, should be 1/2, the product of the probabilities of the constituent events: the probability that headache is caused by the medicine (3/4) and the probability that headache is caused by the background cause (2/3; estimated by reasoning that, given the random assignment of patients to the treatment and control groups in Fig. 1a, if the treatment group had *not* received Medicine A, the proportion of patients who would have a headache anyhow would be similar to the proportion of patients in the control group who do have a headache). (Visualize the bottom left panel as showing the result of the superimposition of the panels at the upper left and bottom right. The overlap between the latter two panels corresponds to this unobservable event.) Given that the probability of the overlap is 1/2, it follows that the probability of headache being caused by the medicine, the background cause, or both is 11/12—the probability of the union of two independent events is the sum of the constituent probabilities minus their product (i.e., the *noisy-OR rule*; Glymour, 2001). And indeed, that is the observed outcome in the bottom left panel. By applying the product definition of independence to postulated causal events, as just explained, the power PC theory (Cheng, 1997) interprets the noisy-OR rule as an implication of independent causal influence (in contrast to related interpretations in the artificial intelligence literature; e.g., Pearl, 1988, pp. 184–185).

The covariational view’s attempt to skirt assumptions about an unobservable event—in this case, the event of headache being independently produced by both c and the background cause—merely results in assuming that the event does not occur, in other words, that c and the background cause have mutually exclusive influences; there is no escape from assumptions regarding some unobservable events (Cheng et al., 2007). Figure 2, with Medicine B omitted in the second study, illustrates this predicament of the ΔP model. Across the two studies in this figure, ΔP of the treatment is unchanged: .25 (in each study, given the treatment, 6 additional patients have a headache). Thus, according to ΔP , the figure represents the independent influences of Medicine A and the background cause. However, for the strength of Medicine A as defined by ΔP to remain constant across the two studies, those 6 additional patients in the bottom half of Figure 2b must not overlap with the 16 who are estimated to have headache due to the background cause; that is, the causes have mutually exclusive influences (i.e., the probability of both causes producing headache in a patient is 0). The summation term in Rescorla and Wagner’s (1972) model carries the same implication.

Note that mutual exclusivity as a definition of independence is self-contradictory. Under this definition, which patients the




Fig. 2. Illustration of two hypothetical studies testing the influences of Medicines A and B on headache. The illustrations at the top show patients who were not exposed to any medicine; the bottom illustrations show patients who were exposed to either (a) Medicine A alone or (b) Medicines A and B in combination. In these studies, the causal power of the treatment varies across studies while ΔP remains constant. Values of Rescorla and Wagner's (1972) model (RW) and of the causal-support model (Griffiths & Tenenbaum, 2005) are also listed. Values of causal support were generated with a noisy-OR function (see the section on Independent Causal Influence).

medicine affects would depend critically on which patients are affected by the background cause; in other words, the medicine cannot operate as if the other causes are not there. In sum, covariational measures are proximal rather than distal, in that they do not map onto any interpretation of causal strength that is coherent across contexts.

Griffiths and Tenenbaum (2005; Tenenbaum & Griffiths, 2001) showed that both ΔP and causal power are Bayesian maximum-likelihood estimates of causal strength, given the existence of the causal relation, under two alternative generating functions: linear and noisy-OR. The former generating function corresponds to ΔP , whereas the latter is the generating function adopted by the power PC theory, and follows from its assumptions (Cheng, 1997). Griffiths and Tenenbaum interpreted each generating function as reflecting a different possible causal mechanism. Their view of ΔP as a viable interpretation of causal strength overlooks the incoherence of ΔP just explained (see Cheng & Novick, 2005, for an alternative interpretation of ΔP , not as strength, but as the probability with which e would not have occurred without c). Our experiments addressed this issue: Do reasoners adopt an interpretation of causal strength that supports coherent generalization across contexts?

EXPERIMENT 1: A STRUCTURAL DECISION ABOUT A SIMPLE CAUSAL LINK BASED ON MULTIPLE CONTEXTS

The goal of Experiment 1 was to identify the property of a causal relation that is assumed to be invariant, and hence generalize, across contexts. For this experiment, we used the scenarios in Figures 1 and 2 but presented individual trials in each scenario (corresponding to individual faces in the figures) in random sequential order, with trials in corresponding positions in the control and treatment panels paired in a before-and-after design.

Method

Subjects

Fifty undergraduates at the University of California, Los Angeles (UCLA), participated to obtain credit in an introductory psychology course. They were randomly assigned to two groups.

Design

All subjects were presented with data from two separate studies (i.e., two data sets). One group of subjects, the power-constant group, was presented with the two data sets illustrated in Figure 1,

in which causal power was constant while ΔP varied. A second group, the ΔP -constant group, was presented with the data sets illustrated in Figure 2, in which causal power varied while ΔP remained constant.

The critical measure was a judgment about whether Medicine B is a cause of headache. We constructed this measure on the basis of the assumption that reasoners expect causal influences to remain invariant across contexts. To avoid biasing subjects' answers (e.g., toward a decision based on causal power or ΔP), we did not request evaluation of any aspect of the individual studies.

If subjects spontaneously generalized across studies on the basis of causal power, those in the power-constant group would have judged Medicine B to be noncausal. Conversely, if subjects generalized on the basis of ΔP , those in the ΔP -constant group would have been the ones to judge Medicine B to be noncausal.

Materials and Task

The task was presented on a computer. First, subjects were presented with the following cover story:

A pharmaceutical company is investigating if two allergy medicines (Medicines A and B) might produce headache as a side effect. The company has conducted two experiments that test the influence of these medicines, and you will see the results from both experiments. The two experiments were conducted in different labs, so the number of allergy patients who have a headache before receiving any medicine may vary across experiments. After reviewing the results from both experiments, you will be asked about the influence of the medicines on headache.

Subjects were further informed that data from the two studies would be presented separately and that patients in the first study received Medicine A only, whereas those in the second study received Medicines A and B in combination.

On each trial, subjects were presented with a picture of an individual allergy patient indicating that patient's state (with or without headache) before receiving any medicine (i.e., an enlarged version of a face in one of the top panels in Fig. 1 or 2). They were then asked to predict whether the patient had a headache after receiving the medicine or medicines. They were instructed to press the "Y" key to indicate "yes" and the "N" key to indicate "no." After making a prediction, subjects were given feedback, which consisted of a picture indicating the patient's state after taking the medicine or medicines (i.e., an enlarged version of a face in one of the bottom panels in Fig. 1 or 2) and a statement reporting whether the subject's prediction was right or wrong. The before-and-after format was adopted to encourage the assumption that background causes occurred independently of the treatment, a prerequisite for estimating causal power.

The data sets were separated by a screen indicating that the data from the first study had ended and that the data from the second study would follow. Subjects were reminded that they would be asked about the results from both studies. After viewing all the trials, subjects were asked this question:

Based on the information from BOTH experiments, what is your best bet on whether or not Medicine B causes headache?

Subjects pressed the "Y" key to indicate "yes, it does" and the "N" key to indicate "no, it does not." Note that our question concerned neither estimates of causal strength nor estimates of the influence of Medicine A, even though the predictions being tested rested on the causal strength of that medicine.

Results

Our results clearly indicated that change across contexts is defined with respect to causal power, rather than ΔP . Whereas more than two thirds of the power-constant group (18 out of 25) responded that Medicine B was noncausal, only one fifth of the ΔP -constant group (5 out of 25) did so, $\chi^2(1, N = 50) = 13.6, p < .001$.

EXPERIMENT 2: A DECISION ABOUT WHETHER A CANDIDATE CAUSE INTERACTS WITH BACKGROUND CAUSES

Most causes are complex and are likely to interact with unobserved background factors. Thus, the assumption that causal influences are independent would often be too strong (see Cheng, 2000, for derivations of weaker conditions that mimic independent causal influence). This assumption, however, provides a criterion for its own demise: If the candidate's causal power changes across contexts, one can infer that the assumption has been violated. This violation provides a signal for seeking a more complex explanation. In Experiment 2, we explored the criterion subjects use to decide that a simple causal hypothesis needs revision.

Experiment 2 used the same method as Experiment 1, except that subjects in Experiment 2 were presented with three hypothetical studies in which only one medicine was tested. Subjects were asked to judge whether that medicine interacted with background causes that might have varied across the three studies. As in Experiment 1, we avoided requesting evaluations of causal influence for individual studies.

Method

Thirty UCLA undergraduates participated to obtain credit in an introductory psychology course. They were randomly assigned to two groups.

Design

Each subject was presented with three studies. Causal power remained constant across the data sets for one group of subjects, the power-constant group, but varied across the studies for the other group, the power-varying group. In other words, the assumption of independent causal influence for the candidate and background causes held in the data presented to the power-constant group, but not in the data presented to the power-

TABLE 1

Relative Frequencies of Headache for the Three Hypothetical Studies in Experiment 2

Subject group	Study 1		Study 2		Study 3	
	elno A	eIA	elno A	eIA	elno A	eIA
Power-constant	16/24	22/24	8/24	20/24	0/24	18/24
Power-varying	0/24	6/24	0/24	12/24	0/24	18/24

Note. A = administration of Medicine A; e = effect (i.e., headache).

varying group. Unlike in Experiment 1, ΔP varied across the studies for both groups, with the variation held constant across groups. Tables 1 and 2 list the event frequencies and values of causal power and ΔP for the three studies presented to each group. Within each group, the order of the three studies was counterbalanced across subjects in a Latin-square design.

The critical measure was a judgment of whether Medicine A interacted with unobserved background causes. If subjects generalized across contexts on the basis of causal power, those in the power-varying group would have inferred an interaction, but those in the power-constant group would have inferred no interaction. In contrast, if subjects generalized on the basis of ΔP values, subjects in the two groups would have inferred an interaction equally often.

Materials and Task

Except for modifications to convey that a single medicine was tested and that three studies would be presented, the materials and task were identical to those in Experiment 1. After viewing all trials, subjects were asked:

Based on the results from ALLTHREE experiments, do you think that Medicine A interacts with some factor that varies across experiments, or do you think that the medicine influences the patients in different experiments in the same way?

YES: I think that the medicine interacts with some factor that varies across experiments.

NO: I think that the medicine has the same influence across experiments.

Note that we did not specify what “influences . . . in the same way” involved.

Results

Our results corroborated those of Experiment 1, demonstrating that generalizations are coherent, that is, in accord with causal power rather than ΔP . Whereas only one third of the subjects in the power-constant group (5 out of 15) responded “yes” to the interaction query, most subjects (13 out of 15) in the power-varying group did so, $\chi^2(1, N = 30) = 8.89, p < .005$.

A Group (2) \times Order (3) between-subjects analysis of variance indicated that there was an influence of group, $F(1, 24) = 10.67, p < .005$, but no influence of order or interaction between group and order, $F_s(2, 24) < 1$.

GENERAL DISCUSSION

Both experiments indicate that people tacitly adopt generic assumptions regarding unobservable causal events so that coherent generalization across contexts is possible. Although the experimental task concerned a structural decision (regarding the existence of a simple causal link or of an interaction with background causes), people spontaneously responded on the basis of estimates of causal strength; more specifically, they judged according to causal power, a coherent estimate of the unobservable distal property of causal strength, rather than according to ΔP , a proximal measure of covariation. The same pattern of results was obtained with variations of our experimental materials: scenarios involving a between-subjects rather than before-and-after design, simultaneous rather than sequential presentation of trials, other values of causal power and ΔP , and different sample sizes for the scenarios (Liljeholm, 2006).

Compelling intuition, as well as our findings, contradicts the purely covariational approach. The ΔP model makes the anomalous and counterintuitive prediction that a cause can be predestined to interact with others. Let us return to Figure 1b, but with Medicine B omitted from the treatment in the second study. In this scenario, Medicine A has a ΔP value of .75. Now consider testing this medicine on the patient groups in Figure 1a. Not only does the covariational approach make the inaccurate prediction that reasoners would infer that the medicine interacts with the background causes in view of the results for these additional patients, but it predicts that they would infer that this interaction is inevitable: Given the top left panel (the

TABLE 2

Values of Causal Power, ΔP , Rescorla and Wagner’s (1972) Model (RW), and Causal Support for the Three Hypothetical Studies in Experiment 2

Subject group	Causal power			ΔP and RW			Causal support ^a		
	Study 1	Study 2	Study 3	Study 1	Study 2	Study 3	Study 1	Study 2	Study 3
Power-constant	.75	.75	.75	.25	.50	.75	2.0	5.6	15.3
Power-varying	.25	.50	.75	.25	.50	.75	2.0	7.7	15.3

^aValues of causal support (Griffiths & Tenenbaum, 2005) were generated with a noisy-OR function (see the section on Independent Causal Influence).

control group for that study), it is impossible for the medicine to have a ΔP value of .75. Therefore, one need not conduct the study to find out if the medicine interacts with background causes. The ΔP values for a given cause across contexts are like two-dimensional snapshots that do not cohere into a three-dimensional representation.

We have made use of only causal power and ΔP to illustrate the rather abstract distinction between a distal and a proximal property (cf. Gallistel's, 1990, distinction between representational and associative learning). Another proximal concept is Griffiths and Tenenbaum's (2005) Bayesian *causal support*, a normative measure of the amount of evidence that a sample provides in favor of the existence of a causal relation. This measure and chi-square are highly correlated (Tenenbaum & Griffiths, 2001). Both are properties of the sample; both increase with sample size, other things being equal. Is causal support, as a proximal quantity, something that reasoners carry from one context into another? The clear answer is "no," as in our various experimental conditions, values of causal support were controlled in tandem with ΔP (see Figs. 1 and 2 and Tables 1 and 2). Note, however, that our results do not pose a problem for the Bayesian approach per se. Bayesian models appropriate for the structural decisions in our experimental scenarios can be developed (e.g., see Jaynes, 2003) and indeed have been developed (T. Griffiths, personal communication, January 15, 2006).

The distinction between distal and proximal properties is a generalization of the psychophysical distinction between distal and proximal stimuli. A proximal stimulus consists of the patterns of energy impinging on receptors, whereas a distal stimulus consists of the patterns of energy that emanate from an object. It is a fundamental fact that people have no access to the distal world except through proximal stimulation. But useful distal properties are often not directly represented in the proximal stimuli. Objects have viewpoint-invariant three-dimensional shapes that are not directly represented in the two-dimensional viewpoint-specific inputs in retinal images. Likewise, invariant causal relations are not directly represented in the observed frequencies of events. For example, each panel in Figure 1 shows a different number of patients with headache, even though we constructed the influence of the treatment to be invariant across contexts. Thus, for organisms to operate intelligently in the distal world, their cognitive processes must have the goal of recovering distal properties (e.g., the constant shape of a three-dimensional object) from proximal input (e.g., viewpoint-specific two-dimensional retinal images).

As cognitive scientists have learned, the problem of recovering the distal stimulus is typically unsolvable unless constraints are introduced (see Marr, 1982, for examples). Indefinitely many distal stimuli can give rise to the same proximal stimuli. The theoretical construct of causal power serves a purpose analogous to the tacit assumption in the human visual system that the world is three-dimensional—both enable parsimonious solutions to the recovery problem. From this per-

spective, our results indicating that causal power is the feature people carry from one context into another should come as no surprise. Neither should it be a surprise that coherence matters. How else can one learn about the distal world that presumably exists, if not by assuming the coherence of disparate proximal stimuli?

Acknowledgments—Preparation of this article was supported by National Institutes of Health Grant MH64810. The article is based in part on a Ph.D. dissertation completed by the first author under the advisorship of the second. We thank James Cutting, David Danks, Clark Glymour, Tom Griffiths, and Keith Holyoak for their valuable comments.

REFERENCES

- Blaisdell, A.P., Sawa, K., Leising, K., & Waldmann, M.R. (2006). Causal reasoning in rats. *Science*, 311, 1020–1022.
- Buehner, M.J., Cheng, P.W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1119–1140.
- Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford, England: Clarendon Press.
- Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- Cheng, P.W. (2000). Causality in the mind: Estimating contextual and conjunctive causal power. In F. Keil & R. Wilson (Eds.), *Explanation and cognition* (pp. 227–253). Cambridge, MA: MIT Press.
- Cheng, P.W., & Holyoak, K.J. (1995). Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In J.-A. Meyer & H. Roitblat (Eds.), *Comparative approaches to cognition* (pp. 271–302). Cambridge, MA: MIT Press.
- Cheng, P.W., & Novick, L.R. (2005). Constraints and nonconstraints in causal learning: Reply to White (2005) and to Luhmann and Ahn (2005). *Psychological Review*, 112, 694–707.
- Cheng, P.W., Novick, L.R., Liljeholm, M., & Ford, C. (2007). Explaining four psychological asymmetries in causal reasoning: Implications of causal assumptions for coherence. In M. O'Rourke (Ed.), *Topics in contemporary philosophy: Vol. 4. Explanation and causation* (pp. 1–32). Cambridge, MA: MIT Press.
- Danks, D. (2003). Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology*, 47, 109–121.
- Dickinson, A., Shanks, D.R., & Evenden, J.L. (1984). Judgment of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, 36A, 29–50.
- Gallistel, C.R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical models in psychology*. Cambridge, MA: MIT Press.
- Griffiths, T.L., & Tenenbaum, J.B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51, 334–384.
- Hume, D. (1987). *A treatise of human nature* (2nd ed.). Oxford, England: Clarendon Press. (Original work published 1739)
- Jaynes, E.T. (2003). *Probability theory: The logic of science*. Cambridge, England: Cambridge University Press.
- Jenkins, H.M., & Ward, W.C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs: General and Applied*, 79(1, Whole No. 594).

- Lien, Y., & Cheng, P. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, 40, 87–137.
- Liljeholm, M. (2006). *Structure learning, parameter estimation and causal assumptions*. Unpublished doctoral dissertation, University of California, Los Angeles.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- McCarthy, J., & Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In D. Michie (Ed.), *Machine intelligence* (pp. 463–502). New York: Elsevier.
- Novick, L.R., & Cheng, P.W. (2004). Assessing interactive causal influence. *Psychological Review*, 111, 455–485.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Mateo, CA: Morgan Kaufman.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, England: Cambridge University Press.
- Povinelli, D.J. (2000). *Folk physics for apes*. New York: Oxford University Press.
- Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Spellman, B. (1996). Conditionalizing causality. In D.R. Shanks, K.J. Holyoak, & D.L. Medin (Eds.), *Advances in the psychology of learning and motivation: Vol. 34. Causal learning* (pp. 167–206). New York: Academic Press.
- Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction and search* (2nd ed.). Cambridge, MA: MIT Press. (Original work published 1993)
- Tenenbaum, J.B., & Griffiths, T.L. (2001). Structure learning in human causal induction. In T.K. Leen, T.G. Dietterich, & V. Tresp (Eds.), *Advances in neural information processing systems 13* (pp. 59–65). Cambridge, MA: MIT Press.
- Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, 121, 222–236.

(RECEIVED 12/14/06; REVISION ACCEPTED 1/31/07;
FINAL MATERIALS RECEIVED 3/2/07)