UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Comparison Then Computation: A Model of Independent Causal Efficacy

Permalink
https://escholarship.org/uc/item/98t8c0nd

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 30(30)

ISSN
1069-7977

Author
Goedert, Kelly M.

Publication Date
2008

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Comparison Then Computation: A Model of Independent Causal Efficacy

Kelly M. Goedert (goederke@shu.edu)
Department of Psychology, Seton Hall University, 400 South Orange Ave.
South Orange, NJ 07079 USA

Abstract

This paper introduces the Comparison Then Computation (CTC) model, which describes how individuals assess the causal efficacy of a target event when simultaneously learning of multiple potential causes of the same outcome. The CTC is a two step model that entails an initial comparison between the probability of an outcome given the target alone and that given the joint occurrence of the target and other explicitly represented alternatives. The first step determines the computation to be performed in the second step. This computation heavily weights sufficiency information over necessity information. Two experiments testing novel predictions of the CTC model favored that model over causal power (Cheng, 1997) and weighted averaging (White, 2008). Strengths and weaknesses of the CTC model will be discussed.

Keywords: causal inference; cue competition; contingency learning; covariation.

Introduction

The ability to detect which events in the environment occur together and to make accurate inferences regarding the causes of events from that covariation information are essential for an individual's survival and well-being. Everyday causal inference is, however, burdened by covariation information from multiple competing causes. If a patient with Seasonal Affective Disorder (SAD) experiences improved mood on a given winter day, is it because she used her intensive light therapy or because she went jogging?

How then do individuals determine the efficacy of a target cause in producing an outcome when there are multiple potential causes of that outcome? Here I present a new model of how individuals make this determination: the Comparison Then Computation (CTC) model. First, however, I review the explanation that has garnered the most attention to date.

The current favored explanation is that individuals calculate an estimate of causal efficacy over a subset of events in which alternative causes are held constant – that is, they conditionalize their estimates on the constant presence or constant absence of an alternative cause (Cheng & Novick, 1992; Cheng, 1997; Spellman, 1996a). In particular, when evaluating the generative strength of a candidate cause it is assumed that individuals prefer to derive their estimate from a focal set in which alternative causes of the outcome are absent (Cheng, 1997; Cheng, Park, Yarlas, & Holyoak, 1996). For example, in determining the effectiveness of exercise on relief of SAD symptoms, the aforementioned patient could focus on those days that she did not use the light therapy (i.e., hold the alternative cause constant) and then contrast how often she experienced symptom relief after exercising with how often she experienced symptom relief when she failed to exercise. These frequencies are represented in third column of Table 1. Formally, this contrast is the calculation of ΔP, the change in the probability of the outcome. ΔP is calculated as the difference between the probability of the outcome given the presence of the target cause [$P(O|T)$] and the probability of the outcome given the absence of the target cause [$P(O|\sim T)$]. In this instance, the calculation is performed controlling for the occurrence of the light therapy and thus is a *conditional* ΔP. Calculating this conditional ΔP yields 0/6 – 0/18 or 0. Thus, controlling for light therapy, exercise is not associated with a change in SAD symptoms. Were we to calculate the unconditional ΔP for exercise, we would disregard the light therapy and consider only the presence or absence of exercise (as represented in the row marginals). Evidence indicates, however, that individuals do not use unconditional contingencies (Spellman, 1996b).

Table 1: Illustrative contingency table with hypothetical event frequencies.

	Light Therapy		
Exercise	Present	Absent	Marginals
Present	18/18	0/6	18/24
Absent	6/6	0/18	6/24

Note: Cell ratios are the number of times the outcome occurred over the number of times that combination of candidate causes occurred.

The explanation that individuals base their estimates of causal effectiveness on a focal set of events in which alternative causes are absent has received support from its ability to account for some forms of cue competition. In particular this conditionalization account predicts cue competition when potential causes do not occur independently of one another (Shanks, et al. 1996; Spellman, 1996a). The frequencies in Table 1 illustrate this case. Exercise and light therapy are more often simultaneously present or absent than they are uniquely present or absent and calculation of the conditional ΔP for exercise yields a lower estimate of causal effectiveness than calculation of its unconditional ΔP.

Conditionalization, however, fails to account for all forms of cue competition. Even in cases in which two causes occur independently of one another, individuals reduce their effectiveness judgments of a target cause that is moderately contingent with the outcome when it is learned about in the presence of highly contingent alternative (e.g., Busemeyer,

Myung, & McDaniel, 1993; Baker et al., 1993; Goedert, Harsch, & Spellman, 2005; Goedert & Spellman, 2005). This effect is not predicted by the conditionalization account.

An additional drawback of the conditionalization / focal set account is that it discards information relevant to the assessment of causal efficacy when there are multiple possible causes explicitly under consideration. Here, I present the CTC model of independent causal influence. This new model predicts cue competition regardless of the relation between the potential causes. Furthermore, the model highlights the information to be gained from the consideration of events in which alternative causes are present.

Comparison Then Computation (CTC) Model

CTC is a descriptive model of causal inference from covariation information for situations in which multiple potential causes of an outcome are explicitly under consideration. Table 2 depicts a generic contingency table for the multiple cause case. In this table and elsewhere in this paper, *P* refers to observed probability, *O* to outcome, *T* to target cause, and *A* to alternative cause. When there are two potential causes of a common outcome we can obtain the probability of the outcome in each of four cases: the joint occurrence of the target and alternative [P(O|TA)], the presence of the target, but absence of the alternative [P(O|T~A)], the absence of the target, but presence of the alternative [P(O|~TA)], and the joint absence of the target and alternative [P(O|~T~A)].

Table 2: Generic contingency table for the multiple cause case.

		Alternative	
Target		Present	Absent
	Present	P(O|TA)	P(O|T~A)
	Absent	P(O|~TA)	P(O|~T~A)

Note: P = probability; O = outcome; T = target; A = Alternative; ~ indicates absence.

According to the CTC model, to assess the independent causal efficacy of the target one performs a two-step process: a comparison and then a computation (see Table 3). In the first step, a comparison between the probability of the outcome given the target alone [P(O|T~A)] and the probability of the outcome given the joint occurrence of the target and alternative [P(O|TA)] determines the nature of other explicitly represented candidate causes and thus determines the relevant information for and the form of computation to be performed in the second step of the process. The computation in the second step determines the causal efficacy (*CE*) and nature of the target cause. *CE* values greater than zero indicate the target is a generative cause of the outcome. Those less than zero indicate it is a preventive cause; those at zero indicate it is not causal.

Table 3: Summary of the Comparison Then Computation (CTC) model.

Step 1: Comparison
(a) If P(O|TA) ≤ P(O|T~A) then perform computation 1.
(b) If P(O|TA) > P(O|T~A) then perform computation 2.
Step 2: Computation
(1) CE = P(O|T~A) – (.4)[P(O|~T~A)]*
(2) CE = P(O|T~A) – (.2)[P(O|~TA)] – (.2)*[P(O|~T~A)]*

Note: P = probability; O = outcome; T = target; A = Alternative; ~ indicates absence.

The Comparison Table 4 depicts the possible results of the initial comparison and the information gained given the result. In the initial comparison, if the outcome is *less* likely to occur when both the target and alternative are present than when the target alone is present, it suggests the alternative is a preventive cause of the outcome. If the outcome is equally likely in the two situations (row 2 of Table 4), it suggests that *either* the target is the sole cause of the outcome *or* there is another cause of the outcome, which may be the explicitly represented alternative or an unobserved cause. Finally, if the outcome is *more* likely to occur when both the target and alternative [P(O|TA)] are present than with the target alone [P(O|T~A)], it suggests the alternative is a generative cause of the outcome.

Table 4: Possible results and information gained from Step 1 comparison.

Result	Information Gained
P(O|TA) < P(O|T~A)	alternative preventive
P(O|TA) = P(O|T~A)	either target, alternative or unobserved generative
P(O|TA) > P(O|T~A)	alternative generative

Note: P = probability; O = outcome; T = target; A = Alternative; ~ indicates absence.

According to the model, knowing the nature of alternative causes is essential for determining which events will be useful in estimating causal efficacy. If the alternative is a preventive cause, then when evaluating the efficacy of the target it would be desirable for the reasoner to exclude cases in which the alternative was present, because in those cases the alternative may mask the efficacy of the target. Therefore, given a preventive alternative (result *a* in Table 3), the focal set of events chosen for computation are those in which the alternative cause is absent and the computational choice is a weighted conditional ΔP (calculation 1 in Table 3).

The same set of focal events and computation are performed when P(O|TA) is equal to P(O|T~A). This case suggests that only one generative cause is at work. Thus, if the target is causal, there are no other generative causes present. As such, the calculation task is simplified to include only those cases in which alternative causes are absent.

Hence, the weighted conditional ΔP is calculated (calculation 1 in Table 3).

In the preventive and sole generative cause cases, the focal set of events used for the computation of causal efficacy does not differ from that favored in previous accounts: the focal set is the set of events in which alternative causes are absent (Cheng, 1997; Cheng et al, 1996). The CTC model diverges from previous accounts when the result of the initial comparison suggests that the alternative is a generative cause of the outcome (result *b* in Table 3). In this case, the reasoner weighs the evidence that uniquely confirms the efficacy of the target [P(O|T~A)] against all of the evidence that disconfirms the efficacy of the target [P(O|~TA) and P(O|~T~A)] by performing calculation 2 in Table 3. The choice of events to include in this calculation is Bayesian in nature in that evidence favoring the alternative as a generative cause is counted against the efficacy of the target. This property will be elaborated in the next section.

The Computation In both calculations depicted in Table 3, evidence reflecting whether the target cause is sufficient to produce the outcome (i.e., the probability of the outcome given the target alone) is given an implicit weight of 1.0. This sufficiency evidence is weighted more heavily than evidence reflecting whether the target is necessary to produce the outcome (i.e., the probability of the outcome in the absence of the target), which is given a weight of 0.4. When the alternative is a generative cause of the outcome, this weight is distributed equally over the target absent events in which the alternative is present [P(O|~TA)] and those in which it is absent [P(O|~T~A)].

The choice of the 1.0 and 0.4 weights for the sufficiency and necessity information in the current model was guided by a review of studies that have derived how individuals weight the cells of the contingency table in the case of a single candidate cause. In the single cause case cells C and D (i.e., evidence regarding necessity) are weighted on average a little less than half that of cells A and B (i.e., evidence regarding sufficiency; Kao & Wasserman, 1993; Mandel & Lehman, 1998; Perales & Shanks, 2008). As such, the choice of the weight for the necessity information in the current calculations is a little less than half the weight given to the sufficiency information.

In causal support (Griffiths & Tenenbuam, 2005), a Bayesian model of causal inference, the likelihood that a particular candidate cause is responsible for producing the outcome is determined by weighing the evidence in favor of that cause against evidence favoring other causes of the same outcome. Although the CTC model does not perform an explicit Bayesian analysis, the model captures this property of the Bayesian model in calculation 2 of Table 3. When the initial comparison reveals that the alternative cause is generative, not only is the probability of the outcome given the absence of both causes [P(O|~T~A)] subtracted from the unique evidence favoring the target [P(O|T~A)], but the unique evidence favoring the alternative cause [P(O|~TA)] is also subtracted from the evidence favoring the target. In this way, cue competition is built into the CTC model. While the causal support model of Griffiths and Tenenbaum (2005) assesses the causal structure of events, causal power (Cheng, 1997) is a Bayesian estimate of causal strength, or a parameterization, of that structure. Notably, the CTC model described here predicts cue competition in situations not predicted by causal power. Causal power is assumed to be calculated over the set of events in which alternative causes are absent (Cheng, 1997). As described earlier, this choice of focal set leads to cue competition when causes are correlated, but it fails to predict cue competition when causes are independent. The CTC model predicts cue competition regardless of the relation between the causes. In particular, when the alternative is generative, the CTC model takes unique evidence favoring the target as a cause and subtracts from that *all evidence against* the target.

Extant Evidence Accounted for by CTC

CTC predicts forms of cue competition not accounted for by conditionalization (e.g., Baker et al., 1993; Goedert & Spellman, 2005). For example, Goedert and Spellman's Experiment 1 consisted of two conditions in which two potential causes occurred independently of one another: a moderately contingent target was paired with either a strongly contingent or non-contingent alternative. In a third condition, the occurrence of the target and alternative causes covaried. In that experiment, as predicted by the conditionalization account, participants rated the target less causal in the covarying condition than in the independent, strong-alternative condition. However, they also rated the target as less causal in the strong-alternative condition than in the non-contingent alternative condition – an effect not predicted by conditionalization. Conversely, the CTC model yields estimates of causal efficacy consistent with the ordering of participants' ratings across the three conditions: *CE* is .33 in the non-contingent alternative condition, .19 in the strong-alternative condition, and -.24 in the covarying condition.

CTC also accounts for the finding that participants sometimes give negative ratings to an event whose conditional ΔP is 0, and thus, whose causal power is 0. For example, across three conditions in Spellman's (1996b) Experiment 2 participants gave negative ratings to a target whose conditional ΔP was 0. The CTC model, however, predicts this pattern of negative ratings across the three conditions (*CE* = -.25 in each).

Experiments Testing Novel Predictions of CTC

Two novel predictions of the CTC model were tested here. Experiment 1 tested the hypothesis that when the alternative cause is generative, both the probability of the outcome given the alternative only [P(O|~TA)] and the probability of the outcome given the joint absence of the target and alternative [P(O|~T~A)] count against the causal efficacy of the target. Experiment 2 tested the hypothesis that the form

of computation applied to learned covariation information depends on the relation between the proportion of times the outcome occurs given only the target event [P(O|T~A)] and the proportion of times the outcome occurs given the joint presence of the target and the alternative [P(O|TA)]. In both experiments the predictions of the CTC model were compared to the predictions of the causal power model (Cheng, 1997) and the recently proposed weighted averaging model (White, 2008).

General Method

Thirty six participants completed seven conditions designed to test two sets of hypotheses, described separately as Experiment 1 and Experiment 2. Testing in groups of up to 20, participants worked through a booklet containing the seven experimental conditions, which appeared in different random orders for different participants. The cover story, adapted from Spellman (1996b), directed participants to imagine that they worked as a chemical engineer for DOW chemicals and that they were attempting to discover new chemicals that could be sold either as fertilizers or as weed killers. Each page of the experiment booklet summarized the results of a single study simultaneously assessing the effects of two chemicals in plant blooming. For example, in one condition participants read the following:

80 plants were assigned to one of the following four treatment conditions:

20 received both 2ZY and UT62. 10 of these bloomed.
20 received only 2ZY. 10 of these bloomed.
20 received only UT62. None of these bloomed.
20 received neither chemical. 10 of these bloomed.

After reading the summary information, participants judged the effectiveness of each of the causes on a scale from -100 to +100, with -100 indicating that the chemical completely prevents plant blooming and +100 indicating that the chemical completely causes plant blooming.

Experiment 1

Experiment 1 compared three conditions in which there was a generative alternative cause and the confirming evidence for the target was held constant while the disconfirming evidence for the target varied. These conditions are depicted in Table 5. In the first condition (1), the outcome never occurred when the target was absent [P(O|~TA) = 0 and P(O|~T~A) = 0]. In the second condition (2), the outcome never occurred when the target and alternative were jointly absent [P(O|~T~A) = 0], but occurred half of the time when the alternative alone was present [P(O|~TA) = .5]. In the third condition (3), the outcome never occurred when the alternative alone was present [P(O|~TA) = 0], but occurred half of the time when the target and alternative were jointly absent [P(O|~T~A)]. For all three conditions, step 1 of the CTC model reveals that the alternative is a generative cause and causal efficacy is computed using calculation 2 from Table 3.

Applying this calculation to the probabilities in Table 5 yields the CTC predictions in the third column of Table 6. The CTC model predicts that participants will rate the target the same in conditions 2 and 3, but these will be less than the ratings given in condition 1. This prediction differs from the causal power (*CP*; Cheng, 1997) and weighted averaging (*WA*; White, 2008) predictions, which are also depicted in Table 6.

Table 5: Conditions compared in Experiment 1.

Con	P(O|TA)	P(O|T~A)	P(O|~TA)	P(O|~T~A)
(1)	1.0	.5	0	0
(2)	1.0	.5	.5	0
(3)	1.0	.5	0	.5

Table 6: Model predictions and results for Experiment 1.

Con	*CP*	*WA*	*CTC*	Ratings (SE)
(1)	.5	.65	.50	59.2^{a} (3.2)
(2)	.5	.70	.40	48.9^{b} (4.5)
(3)	0	.51	.40	53.2^{b} (3.0)

Results and Discussion Participants' ratings of the target are depicted in the far right column of Table 6. Conditions denoted with different superscripts differed significantly from each other. I performed one set of planned Helmert contrasts to test the predictions of the CTC model. The first contrast compared condition 1 to combined conditions 2 and 3. The second contrast compared condition 2 to condition 3. This set of comparisons confirmed the predictions of the CTC, but not the other, models. Condition 1 differed from conditions 2 and 3, $F(1, 35) = 6.13$, $p < .05$, which did not differ from each other, $F(1, 35) = 1.07$, $p = .31$.

Experiment 2

Experiment 2 tested the hypothesis that the computation chosen in step 2 depends on the result of the initial comparison between P(O|T~A) and P(O|TA). In both Experiment 2a and 2b the probability of the outcome given the joint presence of the target and alternative was manipulated [P(O|TA)] while all other probabilities were held constant. In Experiment 2a the base rate probability of the outcome given the absence of the target and alternative was .5; whereas, in Experiment 2b this probability was 0. Table 7 depicts the probabilities across conditions in Experiments 2a and 2b.

In condition 1 of Experiments 2a and 2b, P(O|TA) is less than P(O|T~A); therefore, computation 1 from Table 3 is applied to these probabilities. In condition 2 of Experiments 2a and 2b, P(O|TA) and P(O|T~A) are equal; therefore, computation 1 is again applied. Finally, in condition 3 of Experiments 2a and 2b, P(O|TA) is greater than P(O|T~A) and thus, computation 2 is applied. Applying these computations to the probabilities depicted in Table 7 yields the CTC model predictions depicted in the third column of Table 8. In Experiment 2a, the CTC model predicts that

participants will rate the target the same in conditions 1 and 2, but that their ratings in these conditions will be less than that in condition 3. Notably, the CTC model predicts that the ratings will be reliably greater than 0. These predictions differ from those of the causal power (*CP*; Cheng, 1997) and weighted averaging (*WA*; White, 2008) models, which are also depicted in Table 8.

Table 7: Conditions compared in Experiments 2a and 2b.

Experiment 2a				
Con	P(O|TA)	P(O|T~A)	P(O|~TA)	P(O|~T~A)
(1)	0	.5	0	.5
(2)	.5	.5	0	.5
(3)	1.0	.5	0	.5
Experiment 2b				
Con	P(O|TA)	P(O|T~A)	P(O|~TA)	P(O|~T~A)
(1)	0	.5	0	0
(2)	.5	.5	0	0
(3)	1.0	.5	0	0

In Experiment 2b, both the CTC model and the causal power model predict that participants will rate the target similarly in all three conditions. Assessing performance across this set of conditions is important, however, because it can rule out the possibility that any differences observed in Experiment 2a are due solely to changes in P(O|TA) rather than to differences in the choice of computation applied. In Experiment 2a, the same computation leads to different estimates of causal efficacy, but in Experiment 2b different computations lead to the same estimates of causal efficacy.

Table 8: Model predictions and results for Experiments 2a and 2b.

Experiment 2a				
	CP	*WA*	*CTC*	Ratings (SE)
(1)	0	.03	.30	27.6^{a} (5.4)
(2)	0	.29	.30	33.9^{a} (4.0)
(3)	0	.51	.40	48.9^{b} (4.5)
Experiment 2b				
	CP	*WA*	*CTC*	Ratings (SE)
(1)	.50	.14	.50	52.5^{a} (3.3)
(2)	.50	.42	.50	55.9^{a} (3.5)
(3)	.50	.65	.50	59.1^{a} (3.2)

Experiment 2a Results and Discussion Participants' causal effectiveness ratings appear in the far right column of Table 8. Conditions denoted with different subscripts differed significantly from each other. As in Experiment 1, I performed a set of planned Helmert contrasts to test the predictions of the CTC model. The first contrast compared condition 3 to combined conditions 1 and 2; the second contrast compared condition 1 to condition 2. This set of comparisons confirmed the predictions of the CTC model. Condition 3 differed from conditions 1 and 2, $F(1, 35) = 13.54$, $p < .01$ and conditions 1 and 2 did not differ significantly from each other, $F(1, 35) = 1.77$, $p = .19$.

Experiment 2b Results and Discussion A within-groups MANOVA with condition as a factor revealed no reliable effect, $F(2, 34) = 1.10$, $p = .34$. Hence, in Experiment 2b, despite differences in the probability of the outcome given the joint occurrence of the target and alternative causes, participants rated the cause similarly and in a manner commensurate with the CTC model predictions.

Discussion

In the current paper I introduced and provided evidence for the Comparison Then Computation (CTC) model, which describes how individuals assess the causal efficacy of a target event when simultaneously learning of multiple potential causes of the same outcome. The previously favored account of this process – the hypothesis that individuals choose focal sets in which alternative causes are absent (Cheng, 1997; Cheng et al., 1996) – failed to account for cue competition effects that occur when individuals learn about two candidate causes whose occurrence is uncorrelated (Goedert & Spellman, 2005). Additionally, this focal set account discarded information relevant to the assessment of the causal efficacy of the target.

The new CTC model predicts cue competition regardless of the relation between the occurrence of the two candidate causes and thus accounts for extant demonstrations of cue competition effects (e.g., Baker et al., 1993; Goedert & Spellman, 2005; Price & Yates, 1993; Spellman, 1996b). Furthermore, the model assumes a special, initial role for information from events in which two candidate causes are jointly present. Computation of causal efficacy, however, focuses on the unique evidence for the target cause. Two experiments demonstrated the unique ability of the CTC model to account for patterns of causal judgment when the relation between P(O|TA) and P(O|T~A) is systematically varied (Experiments 2a and 2b) and when the form of disconfirming evidence is systematically varied (Experiment 1).

Limitations There are several limitations of the CTC model. One limitation is that it assumes the reasoner has acquired all of the covariation information. Therefore, the model fails to account for learning phenomena such as order effects (Dennis & Ahn, 2001). A second limitation is that it fails to account for observed dissociations in forms of cue competition (Goedert et al., 2005). Finally, the model fails to account for some important phenomena known to influence causal judgment in the single cause case: sample size and outcome density. Individuals' causal judgments may increase as sample size increases (White, 2002). Likewise, individuals' causal judgments increase as the overall probability of the outcome increases (i.e., the outcome density effect; Lober & Shanks, 2000).

The failure of the model to account for these latter effects is driven by the fact that computations in the model are

carried out on probabilities rather than on event frequencies. A recent review of causal inference in the single cause case highlighted the superiority of models based on weighted frequencies over those based on probabilities in accounting for these effects (Perales & Shanks, 2007). Indeed, a recent model of causal inference in the multiple cause case, the weighted averaging model of White (2008), is based on event frequencies and accounts for sample size and outcome density effects. That model fails, however, to account for patterns of causal judgment observed in the current set of experiments in which sample size was held constant.

Future versions of the CTC model may be able to account for sample size and outcome density by adjusting the computation performed in step 2. This adjustment should be possible while retaining the essential characteristics of the model: the initial comparison step, the heavy weighting of sufficiency information, and the inclusion of all disconfirming evidence when the alternative is a generative cause of the outcome.

Conclusion In sum, the Comparison then Computation (CTC) model captures some important aspects of human causal inference when there are multiple explicitly represented causes of the same outcome. Additionally, the model makes novel predictions regarding how individuals will use the probability of the outcome given the joint occurrence of a target and alternative cause and how individuals use disconfirming evidence when evaluating a target cause in the presence of a generative alternative.

Acknowledgments

The author wishes to thank Stephanie Cox for assistance with data collection and entry.

References

Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory & Cognition, 19,* 414-432.

Busemeyer, J.R., Myung, J., & McDaniel, M.A. (1993). Cue competition effects: Empirical tests of adaptive network learning models. *Psychological Science, 4,* 190-195.

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

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

Cheng, P. W., Park, J., Yarlas, A. S., & Holyoak, K. J. (1996). A causal-power theory of focal sets. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The Psychology of Learning and Motivation: Causal Learning (Vol. 34,* pp. 313-355). San Diego: Academic Press.

Dennis, M. & Ahn, W-K. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, *29*, 152-164.

Goedert, K. M., Harsch, J., & Spellman, B. A. (2005). Discounting and conditionalization: Dissociable cognitive processes in human causal inference. *Psychological Science, 16,* 590-595.

Goedert, K. M., & Spellman, B. A. (2005). Non-normative discounting: There is more to cue interaction effects than controlling for alternative causes. *Learning & Behavior, 33,* 197-210.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology, 51,* 334-384.

Kao, S., & Wasserman, E. A. (1993). Assessment of an information integration account of contingency judgment with examination of subjective cell importance and method of information presentation. *Journal of Experimental Psychology: Learning, Memory & Cognition, 19,* 1363-1386.

Lober, K., & Shanks, D. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review, 107,* 195-212.

Mandel, D. R., & Lehman, D. R. (1998). Integration of contingency information in judgments of cause, covariation, and probability. *Journal of Experimental Psychology: General, 127,* 269-285.

Perales, J. C. & Shanks, D. R. (2007). Models of covariation-based causal judgment: A review and synthesis. *Psychonomic Bulletin and Review, 14,* 577-596.

Price, P. C., & Yates, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition, 21,* 561-572.

Shanks, D R., Lopez, F. J., Darby, R. J., & Dickinson, A. (1996). Distinguishing associative and probabilistic contrast theories of human contingency judgment. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The Psychology of Learning and Motivation: Causal Learning (Vol. 34,* pp. 167-206). San Diego: Academic Press.

Spellman, B. A. (1996a). Conditionalizing causality. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The Psychology of Learning and Motivation: Causal Learning (Vol. 34,* pp. 167-206). San Diego: Academic Press.

Spellman, B. A. (1996b). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science, 7,* 337-342.

White, P. A. (2002). Perceiving a strong causal relation in a weak contingency: Further investigation of the evidential evaluation model of causal judgment. *Quarterly Journal of Experimental Psychology, 55A,* 97-114.

White, P. A. (2008). Accounting for occurrences: A new view of the use of contingency information in causal judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 204-218.