Causal Judgments That Violate the Predictions of the Power PC Theory of Causal Induction

Frédéric Vallée-Tourangeau, **Robin A. Murphy**, and **Susan Drew**
Department of Psychology
University of Hertfordshire
Hatfield, Hertfordshire
UNITED KINGDOM, AL10 9AB
psyqfv{psyqram, psyqsd}@herts.ac.uk

Abstract

The causal power theory of the probabilistic contrast model (or power PC theory) of causal induction (Cheng, in press) states that estimates of the causal importance of a candidate cause are determined by the covariation between the cause and the effect and the probability of the effect as indexed by the probability of the effect in the absence of the cause. In two causal induction experiments we tested predictions derived from the equations of the power PC theory. In Experiment 1, the power PC theory predicted equivalent causal estimates in conditions where the probability of the effect given the presence of the cause, P(*effect* | *cause*), equalled 1 and in conditions where P(*effect* | *cause*) equalled 0. Judgments, however, differed significantly within these conditions and conformed to the predictions of a simpler contingency model. These prediction failures might be attributable to the particular values of P(*effect* | *cause*), and thus Experiment 2 set this probability to values other than 1 or 0. Causal judgments again disconfirmed the predictions of the power PC theory and this time significantly failed to conform to the predictions of a simple contingency model.

Introduction

In a causal induction task, reasoners seek to determine the importance of different candidate causes in producing a target effect. Cheng and her collaborators have argued that people base this assessment on the computation of covariation between each cause and the effect over an appropriate set of focal observations (e.g., Cheng & Novick, 1992). An index of covariation is ΔP which is the difference between the conditional probability of the effect e given candidate cause A, $P(e \mid A)$ and the conditional probability of the effect given the absence of the cause, $P(e \mid -A)$. More recently, Cheng's (in press) causal power theory of the probabilistic contrast model (or power PC theory) specified that causal judgments are not simply a function of ΔP but also of the base rate of the effect, as indexed by $P(e \mid -A)$. For facilitative causal relationships high effect base rates may mask the causal power of a candidate cause and reasoners' estimates should thus increase as the effect base rate increases. Cheng formalizes this relationship with the following equation,

$$\text{power}_A = \frac{\Delta P_A}{1 \quad P(e \mid -A)} \qquad (1)$$

For preventive causal relationships low effect base rates may obscure the preventive causal power of a candidate cause and consequently reasoners' estimates should increase as the effect base rate decreases. Cheng formalizes this relationship in Equation (2),

$$\text{power}_A = \frac{-\Delta P_A}{P(e \mid -A)} \qquad (2)$$

Note that these equations apply only when alternative candidate causes, in this experiment represented by the constant cause X (as introduced below), are independent of the target cause A, and indeed in the experiments reported below $P(X \mid A) = P(X \mid -A) = P(X)$.

The experiments reported here tested a number of simple predictions derived from Equations (1) and (2) in a causal induction task in which the covariation information between candidate causes and the effect was presented one trial at a time. There were two candidate causes in the task, one present on every trial (the constant cause) and one present on some trials but absent on others (the variable cause). Data from Experiment 1 showed that ratings of the causal importance of the variable cause violated the predictions of the power PC theory but not those of a simple contingency model based on ΔP alone, while the data from Experiment 2 showed that they violated the predictions of both the power PC theory and the simple contingency model. In turn, ratings of the constant cause in both experiments were a function of the probability of the effect in the presence of the constant cause alone. We show below that on the basis of a number of auxiliary assumptions the power PC theory can predict this result.

Experiment 1

In the first experiment subjects judged the importance of variable and constant candidate causes in producing an effect in four conditions. We will refer to the variable cause as cause A and the constant cause as cause X. In two of the conditions the probability of the effect in the presence of the variable cause (and of the constant cause since it is present on every trial or $P(e \mid A.X)$), equalled 1. In these conditions the power PC theory makes identical predictions, namely very high estimates of the facilitatory power of A regardless of the base rate of the effect (and consequently of the contingency between A and the effect; we will derive these predictions below). In the remaining two conditions, $P(e \mid A.X)$ equalled 0. In these conditions, the power PC theory predicts very high estimates of the preventive power of A again regardless of the base rate of the effect .

Method

Design. Each subject participated in four experimental conditions. Each condition was composed of 48 trials which showed either the presence of both a variable and a constant cause (*AX* trials) or the presence of the constant cause alone (*X* trials) followed with the presence or absence of the effect. Table 1 shows the probability of the effect in the presence and in the absence of the variable cause in each of the four conditions. The top half of Table 2 reports the frequency of presentation of the four trial types.

Predictions. Before examining the predictions of the power PC theory a caveat is in order. The output values of Equations (1) and (2) are not understood to map isomorphically onto a rating scale. The mapping function that binds the theory's prediction to a subject's judgments is not known but is assumed to be monotonic.

A simple contingency model based on ΔP predicts that estimates of the variable candidate cause should be proportional to ΔP, namely more positive in Condition 1 where ΔP is 0.83 than in Condition 2 where ΔP is 0.17 and more negative in Condition 3 where ΔP is -0.83 than in Condition 4 where ΔP is -0.17. The power PC theory makes different predictions from the simple contingency model. Thus, on the basis of Equation (1) the predictions of the power PC theory for the two conditions where ΔP is positive are equivalent. In Condition 1 where ΔP = 0.83 and the probability of the effect in the absence of the variable cause, P(*e* | -*A*.*X*), is 0.17 the predicted power of *A* is,

$$\frac{0.83}{1-0.17} = \frac{0.83}{0.83} = 1$$

and in Condition 2 where ΔP = 0.17 and P(*e* | -*A*.*X*) = 0.83, the predicted power of *A* is also 1,

$$\frac{0.17}{1-0.83} = \frac{0.17}{0.17} = 1$$

For the two negative conditions, the power PC theory predictions are also equivalent. The ΔP term in Equation (2) is negated to yield positive estimates of preventive causal power, but since a negative rating scale was used in this study we have removed the negative. In Condition 3 where ΔP = -0.83 and P(*e* | -*A*.*X*) = 0.83 the predicted power of *A* is,

$$\frac{-0.83}{0.83} = -1$$

For Condition 4 where ΔP = -0.17 and P(*e* | -*A*.*X*) = 0.17, the predicted power of *A* is also -1,

$$\frac{-0.17}{0.17} = -1$$

The left panel and the middle panel of Figure 1 illustrate the predictions of the power PC theory and the simple contingency model respectively. While the simple contingency model predicts a contingency polarity (i.e., positive/negative) by absolute ΔP interaction, the power PC theory predicts no such interaction.

The covariation between the constant cause and the effect is incalculable within the task trials since P(*e* | -*A*.-*X*) is undefined. The power PC theory (as well as a simpler model based on the calculation of ΔP alone) must postulate that subjects redefine the focal set to include observations other than the ones presented during the experimental procedure. Since the causal scenario used in these experiments was fictitious, actual or similar observations outside the laboratory that would lend themselves to a definition of P(*e* | -*A*.-*X*) are rare. Instead, subjects may be argued to set P(*e* | -*A*.-*X*) to some value on the basis of their understanding of the task. As long subjects define P(*e* | -*A*.-*X*) in a systematic way across all experimental conditions, the value of the contrast for the constant cause should be determined only by the probability of the effect in the presence of the constant cause alone or P(*e* | -*A*.*X*). Thus, the power PC theory and a simpler contingency model based solely on the computation of ΔP, with the help of a few auxiliary assumptions, predict that the ratings of the causal importance of a constant candidate cause should be a function of P(*e* | -*A*.*X*).

			Cause *A*			Cause *X*
	P(*e* | *A*.*X*)	P(*e* | -*A*.*X*)	ΔP	PPC	est	est
1	1.00	0.17	0.83	1	85.9 (6.3)	-29.5 (9.1)
2	1.00	0.83	0.17	1	61.8 (9.9)	67.9 (8.2)
3	0.00	0.83	-0.83	-1	-85.4 (8.6)	76.8 (3.4)
4	0.00	0.17	-0.17	-1	-61.5 (8.3)	-21.5 (10.7)

Table 1. Probability of the effect in the presence, P(*e* | *A*.*X*), and in the absence of the variable candidate cause, P(*e* | -*A*.*X*); output of a simple contingency model ΔP and of the power PC theory (PPC); mean terminal estimates for the variable cause (*A*) and the constant cause (*X*) in each of the four conditions of Experiment 1 (standard error in parentheses).

Procedure. The causal induction task was couched in a fictitious scenario involving samples of laboratory mice having been exposed to radiation. The mice were subsequently injected with certain hormones (e.g., *Curie3CT*, *AsloX7*), and one month later cell growth was analysed. Some mice were injected with two hormones, others with only one. The task instructions informed the subjects that the hormones could either potentiate cancerous cell growth or prevent it. There were four groups of 48 mice, each group corresponding with one of the four experimental conditions. The record for each mouse was presented one at a time on a computer screen and informed the subjects whether it had been injected with one or two hormones. Subjects were then asked to predict the health of that mouse, which they did by entering "1" for "abnormal cell growth" or "0" for "healthy cell growth". One second later feedback appeared on the screen, informing the subjects about the health of that mouse. The feedback stayed on the screen until subjects depressed the spacebar to see the next record. Four sets of labels for the fictitious hormones were used and their assignment counterbalanced across conditions and across hormone types (i.e., variable or constant).

The order of presentation of each of the four conditions was counterbalanced across subjects. The trial presentation was randomized but constrained to match exactly the

conditional probabilities of the effect in the presence and absence of the variable cause shown in Table 1 after each block of 24 trials. After each block of 24 trials, subjects were asked "what is your estimate of the relationship between *<hormone name>* and abnormal cell development?" which they answered by entering a number between -100 and 100. Subjects were told that "a positive estimate indicates that you believe the hormone to CAUSE" and "a negative estimate indicates that you believe the hormone to PREVENT abnormal cell development" Since the trials main effect was not reliable in any of the analyses conducted on the causal ratings, only the terminal ratings will be reported.

Experiment 1				
	Conditions			
Trial Type	1 (0.83)	2 (0.17)	3 (-0.83)	4 (-0.17)
AX → E	24	24	0	0
AX → no E	0	0	24	24
X → E	4	20	20	4
X → no E	20	4	4	20
Experiment 2				
Trial Type	1 (0.50)	2 (0.10)	3 (-0.50)	4 (-0.10)
AX → E	10	18	10	2
AX → no E	10	2	10	18
X → E	0	16	20	4
X → no E	20	4	0	16

Table 2. Frequency of presentation of each of the four trial type in each condition of Experiment 1 (top half) and of Experiment 2 (bottom half). For each condition ΔP for A is given in parentheses. A = variable cause, X = constant cause, E = effect.

Subjects. Twenty-four undergraduates from the University of Hertfordshire received course credits for their participation.

Results and Discussion

Variable Cause. The mean terminal ratings of the causal importance of the variable candidate cause are shown in the sixth column of Table 1 and are plotted in the right-most panel of Figure 1. It seemed that the ratings were influenced by ΔP in that they were more positive when ΔP increased from 0.17 to 0.83 in the two positive conditions, and they were more negative when ΔP decreased from -0.17 to -0.83 in the two negative conditions. Hence, the ratings did not conform to the power PC theory predictions. The statistical analysis supported this observation (the rejection criterion was set at .05 for all analyses unless indicated otherwise). In a two-way repeated measures analysis of variance (ANOVA), the main effect of the absolute value of ΔP (i.e., the |0.17| conditions versus the |0.83| conditions) was not reliable because of the symmetry in the estimates, the main effect of contingency polarity (i.e., the positive versus the negative conditions) was reliable, $F(1, 23) = 262$, and importantly, the interaction between these two factors was reliable, $F(1, 23) = 6.44$. Thus, while the power PC theory predicted no interaction between these two factors (see the left-most panel of Figure 1), there was a significant interaction.

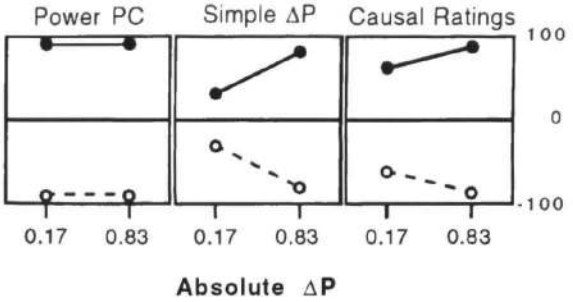

Figure 1. Predicted causal ratings in the two positive contingency conditions (black circles) and in the two negative contingency conditions (white circles) derived from the power PC theory (left panel) and the simple contingency model (middle panel) for Experiment 1; mean observed terminal causal ratings (right panel).

Constant Cause. The mean terminal ratings of the constant cause are reported in the right most column of Table 1. The ratings appeared determined by the likelihood of the effect occurring in the absence of the variable cause (see the $P(e \mid -A.X)$ column of Table 1). In the two positive conditions this probability was low only when ΔP was highly positive, while in the two negative conditions this probability was high only when ΔP was highly negative. We thus expected to obtain a reliable interaction in a two-way repeated measures ANOVA. In such an analysis, neither the main effect of absolute ΔP nor the main effect of contingency polarity was reliable (again because of the symmetry in the estimates), but the interaction between these two factors was, $F(1, 23) = 89.2$. Thus, subjects were willing to rate the importance of a constant candidate cause even if its contingency with the effect was incalculable in the focal set delimited by the task trials (see also Vallée-Tourangeau, Murphy, & Baker, 1996). Moreover, these estimates appeared to be systematically determined by the frequency with which the effect occurred in the absence of the variable cause.

Experiment 2

Subjects' ratings of the variable candidate cause in Experiment 1 were in line with the predictions of a simple contingency model. In this respect, this result is not genuinely embarrassing for the power PC theory as long as auxiliary assumptions are developed that would specify the circumstances in which the computation of ΔP is or is not further weighted by an effect base rate index. One such auxiliary assumption might have something to do with extreme predictions. That is, in Experiment 1 the power PC predictions were either perfect facilitatory or perfect preventive causal power by setting $P(e \mid A.X)$ at 1.00 or 0.00. These extreme values might represent some uninteresting boundaries to the application of the power PC theory. Experiment 2 thus set $P(e \mid A.X)$ at values other than 1.00 and 0.00, and investigated moderate predictions of the power PC theory, namely, predicted causal power of 0.50 or -0.50.

			Cause A			Cause X
	$P(e \mid A.X)$	$P(e \mid -A.X)$	ΔP	PPC	est	est
1	0.50	0.00	0.50	0.5	24.4 (10.3)	-60.6 (10.6)
2	0.90	0.80	0.10	0.5	43.8 (7.5)	80.8 (3.8)
3	0.50	1.00	-0.50	-0.5	-10.2 (10.0)	90.5 (2.2)
4	0.10	0.20	-0.10	-0.5	-49.8 (8.2)	-42.3 (10.2)

Table 3. Probability of the effect in the presence, $P(e \mid A.X)$, and in the absence of the variable candidate cause, $P(e \mid -A.X)$; output of a simple contingency model ΔP and of the power PC theory (PPC); mean terminal estimates for the variable cause (A) and the constant cause (X) in each of the four conditions of Experiment 2 (standard error in parentheses).

Method

Design and Procedure. As in Experiment 1, each subject participated in four experimental conditions. Each condition was composed of 40 trials. The different contingencies between the variable cause and the effect necessitated different trial numbers from Experiment 1 to maintain equal frequencies of the different trial types in both blocks of trials in each condition. Table 3 shows the probability of the effect in the presence and in the absence of the variable cause in each condition. The frequency of the trial types for each condition is shown in the bottom half of Table 2. The same procedure as in Experiment 1 was employed.

Predictions. The simple contingency model predicts that causal ratings should be less positive in Condition 2 (ΔP = 0.10) than in Condition 1 (ΔP = 0.50). It also predicts that the ratings should be less negative in Condition 4 (ΔP = -0.10) than in Condition 3 (ΔP = -0.50). The power PC predictions are different: on the basis of Equation (1) the predicted ratings should be equivalent in the positive conditions. In Condition 1 where ΔP = 0.50 and the probability of the effect in the absence of the variable cause, $P(e \mid -A.X)$, equals 0, the predicted power of A is,

$$\frac{0.50}{1-0.00} = \frac{0.50}{1} = 0.50$$

and in Condition 2 where ΔP = 0.10 and $P(e \mid -A.X)$ = 0.80, the predicted power of A is the same, namely

$$\frac{0.10}{1-0.80} = \frac{0.10}{0.20} = 0.50$$

For the two negative conditions, the power PC theory predicts equivalent causal ratings as well. In Condition 3 where ΔP = -0.50 and $P(e \mid -A.X)$ = 1.00 the predicted power of A is,

$$\frac{-0.50}{1.00} = -0.50$$

In Condition 4 where ΔP = -0.10 and $P(e \mid -A.X)$ = 0.20, the predicted power of A is also -0.50, namely

$$\frac{-0.10}{0.20} = -0.50$$

The far left panel and the middle panel of Figure 2 illustrate the predictions of the power PC theory and the simple contingency model respectively. As in Experiment 1 the power PC theory predicts no interaction between contingency polarity and absolute ΔP, whereas the simple contingency model does.

As in Experiment 1, both the power PC theory and the simple ΔP model can predict how the constant cause will be rated by postulating that subjects calculate its contingency by redefining the focal set. As long as subjects redefine the focal set in the same way in the four conditions, then these models predict that ratings of the constant cause will be determined by the probability of the effect in the presence of the constant cause alone.

Subjects. Twenty four undergraduates from the University of Hertfordshire received course credits for their participation.

Results and Discussion

Variable Cause. The mean terminal ratings of the variable candidate cause are reported in the sixth column of Table 3 and are plotted in the right-most panel of Figure 2. It seemed that the estimates were in line with the predictions of neither the power PC theory nor of the simple contingency model. Interestingly, in the positive conditions, ratings were more positive when ΔP equalled 0.10 than when it equalled 0.50, and in the negative conditions, ratings were less negative when ΔP equalled -0.50 than when it equalled -0.10. Hence, there seemed to be a strong interaction in the ratings but of a kind opposite to the one predicted by the simple contingency model. In a two-way repeated measures ANOVA, the main effect of the absolute value of ΔP (i.e., the |0.50| conditions versus the |0.10| conditions) was not reliable again because of the symmetry in the estimates, the main effect of contingency polarity (i.e., the positive versus the negative conditions) was reliable, $F(1, 23) = 47.8$, and importantly, the interaction between these two factors was reliable, $F(1, 23) = 6.9$. These data undermine the claim that the power PC theory proffers an adequate description of causal judgments and invalidate the argument that subjects were still sensitive to some elements of the power PC theory equations, namely ΔP, since ratings were significantly different from those predicted by a simple ΔP model.

Constant Cause. The mean terminal ratings of the constant candidate cause are reported in the right-most column of Table 3. As in Experiment 1, the estimates appeared to be a function of $P(e \mid -A.X)$. In a two-way ANOVA on these estimates the main effect of absolute ΔP was not reliable, the main effect of the variable cause's contingency polarity was reliable, $F(1, 23) = 15.8$ and the interaction between these two factors, as expected, was highly significant, $F(1, 23) = 166$. The contingency polarity main effect was reliable because $P(e \mid -A.X)$ in the two negative conditions (viz., 1.00 and 0.20 in Conditions 3 and 4 respectively, see Table 3) was larger than in the two

positive conditions (viz., 0.00 and 0.80 in Conditions 1 and 2 respectively). This is further evidence that subjects' ratings of the constant cause closely reflected the frequency with which the effect occurred in the absence of the variable cause.

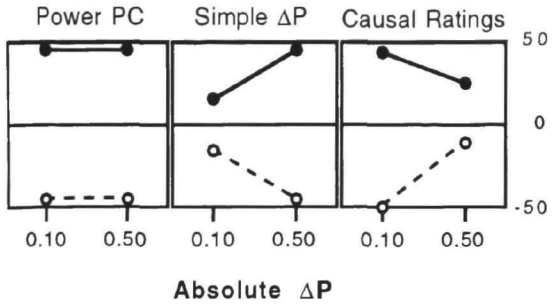

Figure 2. Predicted causal ratings in the two positive contingency conditions (black circles) and in the two negative contingency conditions (white circles) derived from the power PC theory (left panel) and the simple contingency model (middle panel) for Experiment 2; mean observed terminal causal ratings (right panel).

General Discussion

In the experiments reported here, estimates of the causal importance of a variable candidate cause disconfirmed the predictions of the power PC theory. However, in the first experiment, these estimates conformed to the predictions of a simpler model based simply on the computation of ΔP. Even this simpler model had to be rejected in the light of the data obtained in the second experiment where the causal ratings were inversely proportional to ΔP. Thus, the ratings of the variable candidate cause provided strong evidence that the process of causal induction is not describable in terms of the equations developed in the power PC theory.

Subjects also rated the causal importance of a constant candidate cause. These ratings were a function of the probability of the effect in the absence of the variable cause, or P(e | -A.X). If indeed the computation of contrasts underlies causal induction then subjects asked to rate the causal important of a constant cause were facing an 'impasse' because the second probability needed for this computation, namely P(e | -A.-X) was undefined. The power PC theory must therefore assume that subjects redefine the focal set to include observation from outside the experimental procedure. Assuming that this new focal set fixed P(e | -A.-X) at the same value in all conditions, ratings of the constant cause should be a function of P(e | -A.X), which indeed they were. A direct test of these assumptions has not yet been undertaken.

A class of simple connectionist architectures originally developed to model Pavlovian conditioning offer an alternative account of causal induction. Associative models of causal induction such as the Rescorla-Wagner model (Rescorla & Wagner, 1972; henceforth R-W model) and Pearce's (1987) model of stimulus generalization do not postulate that reasoners compute contrasts or redefine focal sets. Rather, they take the form of simple two-layer networks with an input layer coding for the candidate causes and an output layer coding for the target effect. Both use error reduction learning algorithms similar to the delta rule. The R-W model stipulates that the change in the weight of any cause-effect connection on any given trial is a function of Equation (3)

$$\Delta w = \alpha\beta(\lambda - \sum_{1}^{k} w) \quad (3)$$

where α and β are learning parameters, λ is the target activation on that trial and Σw represents the actual activation of the output node (determined by the sum of the weights along the k active connections (a connection is active when the corresponding candidate cause is present; the activation function of the output unit is linear, not sigmoidal).

The input layer of a Pearce network codes for the configuration of causes experienced by the subjects, namely AX and X. On any given trial, the weight for the AX configuration is updated using the following rule

$$\Delta w_{AX} = \beta(\lambda - [\bar{E}_{AX} - \bar{I}_{AX}]) \quad (4)$$

where β is a learning parameter and λ is the target activation of the output node on that trial. The actual activation of the output node reflects the sum of the strength of two types of connection, one excitatory ($\bar{E}$) and one inhibitory ($\bar{I}$; again the activation function is linear). Critically, these measures of excitatory and inhibitory associative strength include the generalized excitatory (or inhibitory) associative strength of candidate causes that are similar to the compound AX, in this case X (see Pearce, 1987, Equations (6) and (8)). The change in weight for the constant cause is determined by the learning rule,

$$\Delta w_X = \beta(\lambda - [\bar{E}_X - \bar{I}_X]) \quad (5)$$

The causal ratings for the constant cause are a function of the cumulative weight changes as specified in Equation (5).

The predictions concerning the ratings of the constant candidate cause derived from the R-W model and the Pearce model are plotted in the left half of Figure 3. Both associative models predict that ratings of X will be a function of P(e | -A.X) (shown in Tables 1 and 3). While the power PC theory formulated the same predictions it did so by postulating that subjects 1. redefined the focal set to include real or imagined 'observations' outside the task trials and 2. computed a contrast.

The predictions of the three models fare much poorer for the ratings of the variable candidate cause. As Figures 1 and 2 showed, the predictions of the power PC theory were disconfirmed in both experiments. The predictions of the R-W model and the Pearce[1] model are shown in the right half of Figure 3. The R-W model predicts the development of

[1] Since the variable cause was never experienced on its own, the Pearce model stipulates that its causal estimation is a function of the cumulative weight changes for the AX configuration as specified in Equation (4) multiplied by a coefficient that indexes the similarity between A and AX. Since A constitutes half of the AX configuration, the Pearce model assumes that that coefficient equals .5 (see Pearce, 1987, Equation (3)).

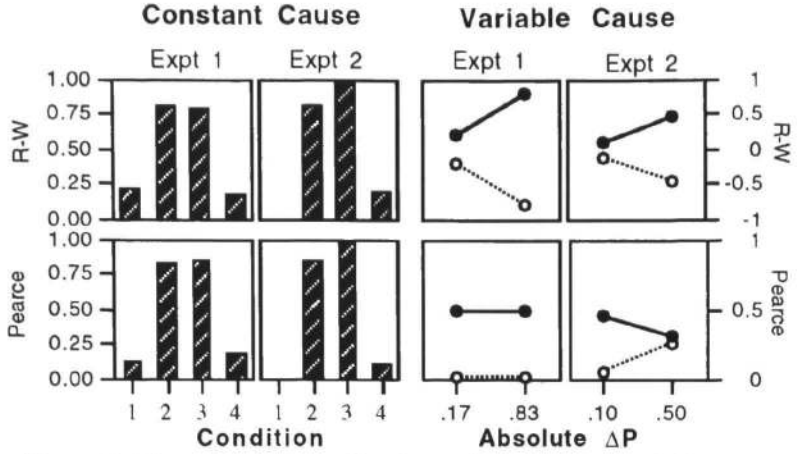

Figure 3. Mean terminal predicted associative strength of the constant candidate cause and the variable candidate cause in the four conditions of each experiment. The α parameter in the R-W algorithm was set to 0.35. The β parameters in both the R-W and the Pearce algorithms were equated and set to 1. The mean terminal associative strength was calculated after 100 training trials over 100 epochs.

associative bonds whose magnitude approximates ΔP at asymptote. Thus, its predictions are the same as those from a simple contingency model. Hence, the R-W can anticipate the results of the first experiment, but not those of the second. In turn, for Experiment 1, the Pearce model predicts, as did the power PC theory, that ratings of the variable cause within the positive conditions and within the negative conditions should be equivalent. The Pearce model, however, is the only model that can anticipate the ratings of the variable cause in Experiment 2 because the associative strength of the constant cause, which increases as the effect base rate increases, generalizes to the variable cause. The notion that the influence of the effect base rate on the ratings of the variable cause may be mediated by the transfer of associative strength from the constant cause may have important implications for associative models of causal induction. For example, Dickinson and Burke (1996) have argued that retrospective revaluation such as backward blocking may be mediated by within compound associations. Their account proposes that the associative strength of a given cause can be influenced by the strength of other causes with which it has been repeatedly paired.

In summary, then, all three models considered here (R-W, Pearce, power PC), can anticipate the ratings of the constant cause. However the associative models do so without postulating that subjects redefine focal sets and compute contrasts (ΔP is not a theoretical primitive in those models). Instead, the magnitude of the causal ratings reflect the acquisition of an association between a cause and an effect, and this developmental process is mediated by the cause-effect *contiguity*. As for the ratings of the variable candidate cause, all three models postulate that they are influenced by both the cause-effect covariation and the base rate of the effect. The local prediction failures of the Pearce model (for Experiment 1) and the R-W model (for Experiment 2) suggest that these associative models have not fully captured the interactive nature of these influences. Of course, it is precisely the integration of these influences that the equations of the power PC theory attempt to formalize. However, the fact that the predictions derived from these equations were disconfirmed in both experiments demonstrate that the formalization is inadequate.

Acknowledgments

We would like to thank Andy Baker and two anonymous reviewers for thoughtful comments on a previous version of this paper.

References

Cheng, P W. (in press). From covariation to causation: A causal power theory. *Psychological review.*

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

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgments. *The Quarterly Journal of Experimental Psychology*, **49B**, 60-80.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, **94**, 61-73.

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.

Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (1996). Judging the contingency of a constant cue: Contrasting predictions from an associative and a statistical model. In G. W. Cottrell (Ed.) *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 457-452). Mahwah, NJ: Lawrence Erlbaum Associates, Publishers.