

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Causal Induction: The Power PC Theory versus the Rescorla-Wagner Model

Permalink

<https://escholarship.org/uc/item/4tm70424>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Buehner, Marc J.

Cheng, Patricia W.

Publication Date

1997

Peer reviewed

Causal Induction: The Power PC Theory versus the Rescorla-Wagner Model

Marc J. Buehner

Universität Regensburg

Glockengasse 3

93047 Regensburg, Germany

marc.buehner@psychologie.uni-
regensburg.de

Patricia W. Cheng

Department of Psychology

University of California, Los Angeles

Los Angeles, CA 90095

cheng@lifesci.ucla.edu

Abstract

Two experiments compared the influence of the probability of the effect given the absence of the candidate cause on the causal judgments of candidate causes with the same ΔP , defined as the difference between the probability of the effect in the presence of a candidate cause and that in its absence. Our results strongly support the power PC theory (Cheng, 1997) but contradict the Rescorla-Wagner model (1972) and the traditional ΔP model.

Introduction

Causal induction allows humans and other intelligent systems to explain, predict, and control their environment, making it a critical tool in an otherwise chaotic world. How do people discover the causes of events? A long-standing proposal (e.g., Jenkins & Ward, 1965) for candidate causes and effects that can be represented by binary variables is that causal judgments are based on

$$\Delta P = P(efc) - P(e|c), \quad (1)$$

where c is a candidate cause, e is the effect in question, $P(efc)$ is the probability of e given the presence of c , and $P(e|c)$ is the probability of e given the absence of c . ΔP is a measure of the extent to which c and e covary (i.e., are both present or both absent), and has variously been called the *contingency* or *contrast*. The conditional probabilities are estimated by the respective relative frequencies of the events for which e occurs in the presence and in the absence of c . If ΔP is noticeably positive, c is a *generative* or facilitatory cause, and if it is noticeably negative, c is a *preventive* or inhibitory cause. If ΔP does not noticeably differ from zero, c is independent of e and is noncausal.

Rescorla-Wagner Model

An influential connectionist model that under some conditions asymptotically computes a variant of ΔP is the Rescorla-Wagner (1972) model (RWM). Although this model was originally proposed for describing Pavlovian conditioning, it has often been adopted to explain causal induction (e.g., Dickinson, Shanks, & Evenden, 1984; Wasserman, Elek, Chatlosh, & Baker, 1993). Using the RWM to explain causal induction implies reducing causal learning to associative learning. According to this model, learning proceeds by changes in the strength of a connection

between a conditioned stimulus c (e.g., a flash of light) and an unconditioned stimulus e (e.g., a shock). In causal terms, c is a candidate cause and e is the effect. This change in the strength of c , ΔV_c , is specified by the learning rule

$$\Delta V_c = \alpha_c \beta_e (\lambda_e - \sum V), \quad (2)$$

where α_c and β_e are rate parameters that, respectively, depend on the salience of c and e , and λ_e is the actual outcome of the trial. If e is present, this value is typically set to 1; if e is absent, it is set to 0. $\sum V$ is the expected outcome, which is the sum of the strengths of all candidate causes present on that trial. Learning thus occurs by reducing the discrepancy between the actual outcome λ_e and the expected outcome $\sum V$ until this discrepancy ($\lambda_e - \sum V$) approximates zero. Note that ΔP from the contingency model and ΔV in RWM do not both represent causal strength: whereas ΔP in the contingency model represents causal strength, ΔV in RWM is the change in strength on a trial. When ΔV approximates zero (when learning has reached asymptote), the weight of the link from a candidate cause to the effect represents its causal strength. These asymptotic weights are traditionally obtained by computer simulation, but they can also be obtained mathematically for many experimental designs (see Cheng, 1997).

When there is only one varying candidate cause in a context, the RWM asymptotically computes ΔP as the measure of associative strength when β_e is assumed to remain constant across trials on which e does and does not occur (Chapman & Robbins, 1990). However, if β_e is assumed to be *larger* on trials with e present, RWM makes predictions different from the contingency model: RWM then predicts the (absolute) magnitudes of the judged causal strengths to be smaller as $P(e)$ increases for any fixed positive or negative ΔP (see Wasserman et al., 1993). It is difficult to motivate why β_e might be *smaller* on trials with e present (and researchers applying the RWM hardly if ever make this assumption), but if this assumption is made, RWM would predict the opposite trend: the (absolute) magnitudes of the judged causal strengths should be larger as $P(e)$ increases for any fixed positive or negative ΔP . Regardless of assumptions about β_e , RWM predicts the same influence of $P(e)$ on the absolute causal strengths of candidates with same positive ΔP as for those with the same negative ΔP . Also regardless of assumptions about β_e , RWM predicts judged causal strength to be zero when $\Delta P = 0$.

Power PC Theory

As many have noted, covariation is insufficient as a criterion for inferring causality. For example, one would not infer from the covariation between a drop in the barometric reading and the approach of a storm that a barometric drop causes the approach of storms. This example fails to satisfy one of the putative boundary conditions for inferring causality from covariation -- that alternative causes of the effect (e.g., changes in atmospheric pressure) are controlled. To explain this and other boundary conditions for inferring causality from covariation, Cheng (1997) proposed the power PC theory (a causal power theory of the probabilistic contrast model of Cheng and Novick, 1990). According to this theory, the causal reasoner's goal is to optimally estimate the unobservable causal power of a candidate cause from observable events. This theory assumes that *the reasoner believes that there are such things in the world as causes that have the power to generate an effect and causes that have the power to prevent an effect, and that only such things influence the occurrence of an effect* (cf. Kant, 1781/1965). Applying this assumption in a probabilistic form to explain Equation 1, this theory shows (1) the conditions under which ΔP_c provides an estimate of causal power and (2) how well it does so under those conditions. Cheng (1997) reviews a diverse range of phenomena supporting this theory over other models. One of the mathematical consequences of this theory is that when causes alternative to the candidate cause c are controlled and ΔP_c as defined by Equation 1 is nonnegative, p_c , the causal power of candidate c to generate effect e is

$$p_c = \frac{\Delta P_c}{1 - P(e|\bar{c})}. \quad (3)$$

Analogously, when causes alternative to the candidate cause c are controlled and ΔP_c is nonpositive, the causal power of candidate c to prevent effect e is

$$p_c = \frac{-\Delta P_c}{P(e|\bar{c})}. \quad (4)$$

The predictions of the theory are assumed to be only ordinal.

Thus, whereas the traditional contingency model (Equation 1), and the RWM under the assumption that β_c remains constant across trials on which the effect does and does not occur, each predicts ΔP to be the sole source of judged causality, the power PC Theory predicts that equal levels of ΔP , with different values of $P(e|\bar{c})$ (and hence also different values of $P(elt)$) should *not* yield equal causal judgments. When alternative causes are controlled, Equation 3 applies and predicts that candidate causes with the same positive ΔP_c should be judged to have increasingly *large* generative power as $P(e|\bar{c})$ increases, but does not equal 1. When $P(e|\bar{c})=1$, the generative power of the candidate cause is undefined. That is, a reasoner cannot draw any conclusion about the causal strength of c generating e if e occurs all the time regardless of whether c occurs.

In contrast, Equation 4 predicts that as $P(e|\bar{c})$ increases, candidate causes with the same negative ΔP_c should be judged to have increasingly *small* preventive power. Note that (1) $P(e|\bar{c})$ influences the (absolute) magnitude of estimated generative and preventive powers in opposite

directions, and (2) the direction of these influences are not dependent on any parameter settings. The power PC theory and the RWM differ on both of these points. When $P(e|\bar{c})=0$, the preventive power of the candidate cause is undefined according to Equation 4. That is, one cannot make inferences about the strength of c preventing e , if e never occurs in the first place.

Finally, Equations 3 and 4 both predict that when $\Delta P_c = 0$, the power of c should remain at 0 and be uninfluenced by $P(e|\bar{c})$ as long as the denominator in the relevant equation is not 0.

Consider a concrete illustration of the power PC theory's predictions. Suppose a researcher wants to evaluate the preventive power of a new drug against headaches. In a study involving 16 participants, 8 receive treatment with the drug (candidate present), and 8 receive a placebo (candidate absent). Now suppose neither the 8 participants who received drug treatment nor the 8 participants who received the placebo report headaches. In this case ΔP_c equals zero. Yet the researcher would not infer that the drug is ineffective (i.e., noncausal). Since headaches did not occur even in the control group, how then could a preventive candidate prevent them in the drug group? The researcher simply cannot draw any causal inferences. This intuition is captured by the special case in which preventive power in Equation 4 is undefined.

Now consider two studies that involve the same nonzero ΔP_c but different $P(e|\bar{c})$. In the first study, all of the 8 participants in the control group have headaches, whereas only 6 of the 8 participants who received the drug have headaches. In this case, $\Delta P_c = P(elt) - P(e|\bar{c}) = .75 - 1.00 = -.25$. Assuming that all alternative causes producing or preventing headaches are constant across the two groups, the researcher would assume that if not for the drug, all 8 participants in the drug group, just as in the control group, would have had headaches. The drug therefore has a small preventive power, preventing headaches with a probability of .25. In a second study, 4 of the 8 participants in the control group report headaches and 2 of the 8 participants receiving drug treatment report headaches (i.e., $P(elt) - P(e|\bar{c}) = .25 - .50 = -.25$). Again, assuming that the composite of alternative causes remains constant between groups, the composite would have produced headaches in 4 of the 8 participants in the drug group, just as in the control group. The drug therefore prevents headaches in 2 of the remaining 4 participants, yielding a probability of .50. Thus, although $\Delta P_c = -.25$ here as in the preceding study, the researcher would attribute a higher preventive power to the latter candidate. Equation 4 formalizes this intuition.

Analogous intuitions about when the generative power of a candidate cause is undefined and about the causal power of candidate causes with the same positive ΔP_c but different $P(e|\bar{c})$ are captured by Equation 3.

The present paper presents a preliminary report of some experimental tests of the traditional contingency model, the Rescorla-Wagner model, and the power PC theory. We report the results of two experiments that discriminate between the predictions of these three accounts. We tested candidates with positive, negative, and zero ΔP s to evaluate the full range of predictions of the power PC theory. An

additional reason for doing so was to cover all possible consistent parameter settings of the RWM. Recall that although the RWM can predict causal judgments for candidates with the same non-zero ΔP that are either positive or negative functions of $P(e|c)$, or not a function of $P(e|c)$, it cannot predict a positive function in one experimental condition and a negative function in another condition unless its parameter settings are reversed across conditions.

A test of the predictions of the three accounts would involve varying levels of $P(e|c)$ and $P(e|\bar{c})$. Although many previous experiments have manipulated these variables, and the available results support the power PC theory (see review in Cheng, 1997), these experiments either tested too few levels of these variables to evaluate this prediction systematically (e.g., Allen & Jenkins, 1983; Anderson & Sheu, 1995; Baker, Berbrier, & Vallée-Tourangeau, 1989; Dickinson et al., 1984; Shanks, 1987), or they used an effect that occurred in continuous time rather than in discrete entities (e.g., Wasserman et al., 1993), in which case the power PC theory does not directly apply. In the experiments we report, we therefore tested combinations of many levels of $P(e|c)$ and $P(e|\bar{c})$ and used a binary effect that occurs in discrete trials. We adapted Wasserman et al.'s (1993) design (the most comprehensive study of this nature to date), in which 5 levels of $P(e|c)$ and of $P(e|\bar{c})$ are independently combined within-subject. Our first experiment presented subjects with a relatively small number of individual events, and the second experiment presented a larger number of events in a visual summary format. Within each experiment, conditions involving negative ΔP s were tested separately from those involving positive ΔP s, and the former conditions are reported before the latter. Both parts of each experiment included conditions with zero ΔP s.

Experiment 1a

Method

Subjects. 13 male and 44 female students in undergraduate psychology classes at the University of California, Los Angeles, served as subjects in exchange for course credit.

Design and Procedure. Subjects were given a cover story in which they were asked to pretend they were virologists testing several new vaccines against viruses. They were asked to evaluate the outcome of studies performed on laboratory rats which had all been infected with certain viruses. For every subject, there was one practice condition and 15 experimental conditions, with each condition consisting of 16 laboratory records (i.e., learning trials). Each record provided information about whether one particular rat was vaccinated prior to virus exposure and whether the rat developed the disease related to the virus. The 15 different experimental conditions represented 15 independent studies on different viruses and vaccines. Because Experiment 1a examined the evaluation of preventive causal power, 5 levels of the two conditional probabilities, $P(e|c)$ and $P(e|\bar{c})$ -- 1.00, .75, .50, .25, and

.00 -- were combined to yield 5 levels of nonpositive ΔP s: -1.00 , $-.75$, $-.50$, $-.25$, and $.00$. These combinations yielded a total of 15 conditions (see Figure 1).

The 16 laboratory records in each condition consisted of 8 for which c was present (the rat was vaccinated) and 8 for which c was absent (no vaccination given). The 15 conditions and the laboratory records within each condition were both presented in random order. After studying the 16 records in each condition, subjects were asked to evaluate the effectiveness of the studied vaccine at preventing the disease related to the virus in question. They were asked to give a rating on a scale from 0 to 100, where 0 meant that the vaccine does not prevent the disease at all and 100 meant that the vaccine prevents the disease every time.

Results and Discussion. Figure 1 displays subjects' mean ratings of the preventive power of the candidate cause. On the abscissa are the 5 levels of $P(e|\bar{c})$. For Figure 1 and subsequent analogous figures, the lines connect data points with the same level of ΔP , except for the point with the zero ΔP that has undefined causal power according to the power PC theory (and intuition). This point is displayed separately. The corresponding values for $P(e|c)$ follow from each combination of ΔP and $P(e|\bar{c})$. The power PC theory predicts decreasing estimates of preventive power with increasing $P(e|\bar{c})$ for candidates with equal negative ΔP s. A trend analysis was performed over each set of conditions with the same ΔP . A linear negative trend was highly reliable for each level of ΔP , $t(56) = 4.1$, $p < .001$ for $\Delta P = -.75$, $t(112) = 9.5$, $p < .001$ for $\Delta P = -.50$, $t(168) = 13.3$, $p < .001$ for $\Delta P = -.25$, and $t(224) = 7.6$, $p < .001$ for $\Delta P = 0.0$.

The reliable negative trends for candidates with the same negative ΔP s contradict the traditional contingency model. They also contradict the RWM if β_e is assumed to be either constant across trials or smaller on trials in which e occurs than those in which e does not occur. They support the RWM if β_e is assumed to be larger on trials in which e




Figure 1: Mean judgments of preventive causal power for each ΔP level in Experiment 1a.

occurs. They also support the power PC theory.

The reliable trend for $\Delta P = .00$ replicates the observation in many previous studies, but contradicts the asymptotic predictions of all three accounts considered in this paper. The single data point for $\Delta P = -1.00$ is irrelevant to our evaluation of alternative accounts.

Experiment 1b

Experiment 1b is the analogue of Experiment 1a for the evaluation of generative causal power.

Method

Subjects. 23 male and 29 female students in undergraduate psychology classes at the University of California, Los Angeles, participated in exchange for course credit.

Design and Procedure. The procedure was the same as in Experiment 1a. This time the cover story asked the subjects to imagine they were microbiologists studying how ray exposure influences the mutation of viruses. The subjects studied laboratory records that provided them with information about whether or not a petri dish with viruses was exposed to certain rays and whether or not mutation occurred. The design was exactly symmetrical to that in Experiment 1a, yielding only nonnegative ΔP s with values .00, .25, .50, .75, and 1.00. After studying the laboratory records relevant to a certain virus and ray combination, subjects gave a rating of how strongly they thought the particular rays cause mutation. The scale ranged from 0 to 100, where 0 meant that the rays do not cause the virus to mutate at all and 100 meant that the rays cause the virus to mutate every time.

Results and Discussion. Subjects' mean ratings of causal power are shown in Figure 2. As before, the data points with equal levels of ΔP are connected. Recall that for the evaluation of generative causal power, the power PC theory predicts a positive linear trend for candidates with equal levels of positive ΔP with increasing $P(e|c)$. As before, a trend analysis was performed to check for positive linear trends in equal levels of ΔP . For $\Delta P = .75$ the positive linear trend fell just short of significance, $t(51) = 1.99, p = .052$. For $\Delta P = .50$ the positive trend was highly significant, $t(102) = 3.72, p < .001$; as was the linear positive trend for $\Delta P = .25, t(153) = 3.91, p < .001$.

As for the evaluation of preventive power, both RWM with constant β and the traditional ΔP model predict flat lines connecting candidates with equal levels of ΔP . While assigning larger values to β for trials where e occurs allows the RWM to explain the negative trends found in Experiment 1a, this modification also predicts negative trends for positive ΔP s. The observed positive trends in Experiment 1b contradict this prediction of the RWM but support the power PC theory.

For zero ΔP s, all three accounts considered predict noncausal judgments. However, subjects did not rate zero ΔP s as noncausal: We unexpectedly also found sloping trends in the zero ΔP conditions. Why were there such trends?

A possible explanation is that our subjects conflated reliability with causal strength. For example, in the generative experiment, in the 0/8-0/8 condition ($P(eltc) =$




Figure 2: Mean judgments of generative causal power for each ΔP level in Experiment 1b.

$P(e|c) = 0$) e neither occurred in any of the 8 trials given the presence of c nor in those given its absence. Assuming that alternative causes were constant, it follows that c had 8 trials to "prove" its generative power, and it failed to do so. In the 6/8-6/8 condition ($P(eltc) = P(e|c) = .75$), however, e would be expected to be produced by alternative causes on 6 of the 8 trials when c was present, just as when c was absent. This leaves only 2 out of the 8 trials for c to prove its causal power, and it failed to do so. Because there were more trials on which c could have but in fact failed to prove its power in the former condition than in the latter, subjects might be more confident of a noncausal judgment in the former condition, leading to a rating closer to 0 in that condition. Providing subjects with a constant number of trials across conditions thus yields varying reliability of the information presented.

An analogous explanation applies to the zero ΔP conditions of the preventive experiment.

This explanation does not apply to non-zero ΔP s. First, note that unlike noncontingent candidates, which all produce e with the same causal power (i.e., probability) of 0, candidates with the same non-zero ΔP do not produce e with the same causal power. For example, for preventive candidates that all have $\Delta P = -.5$, in the 0/8-4/8 condition, $p_c = 1.0$; in the 2/8-6/8 condition, $p_c = .67$; and in the 4/8-8/8 condition, $p_c = .5$. Their causal ratings therefore would not reflect reliability alone. Second, if one were to ignore the varying causal powers, and consider reliability alone, then the 3 conditions just mentioned should have increasingly greater reliability: The 0/8-4/8 condition had only 4 trials to reveal preventive power, the 2/8-6/8 condition had 6 trials, and the 4/8-8/8 condition had 8 trials. Therefore, given that these candidates have non-zero preventive power, reliability predicts increasing preventive ratings for these 3 conditions, contrary to the observed ratings, which were in the decreasing order predicted by their causal powers. That is, reliability cannot explain the trends observed for candidates with equal non-zero ΔP s.

Experiments 2a and 2b

The goal of Experiments 2a and 2b was to reduce the possibility of varying reliability by presenting the information in summary format involving a large number of trials as opposed to a small number of individual trials. We

expect the new format to reduce the slopes for the zero ΔP conditions only.

Method

Subjects. 79 (Experiment 2a) and 74 (Experiment 2b) students in undergraduate psychology classes at the University of California, Los Angeles, participated in exchange for course credit.

Design and Procedure. Subjects were given the same cover stories as in Experiments 1a and 1b. As before, they were asked to evaluate the outcome of studies performed on rats (Experiment 2a) and on petri-dishes of viruses (Experiment 2b). There was one practice condition and 15 within-subject experimental conditions in both Experiment 2a and 2b. The 15 experimental conditions represented 15 independent studies that had to be evaluated by the subjects. The same conditional probabilities and levels of ΔP as in the previous experiments were used to create nonpositive ΔP s (Experiment 2a) and nonnegative ΔP s (Experiment 2b).




Figure 3: Example of Data Presentation in Experiment 2b

The main difference between Experiments 2a and 2b and the previous experiments was that this time subjects did not study individual laboratory reports but were rather presented with visual summaries of the outcome of performed studies. This was realized by presenting subjects with two pie charts containing information about relative frequencies of the effect, given the presence or the absence of the candidate. Subjects were informed that each chart summarizes the results for 100 cases.

Figure 3 represents an example summary as used in Experiment 2b. For each of these summaries, subjects were asked to give a rating for how strongly they think a vaccine prevents the disease (Experiment 2a) or for how strongly they think certain rays cause mutation to a virus

(Experiment 2b). As before, subjects were asked to give ratings from 0 to 100. Additionally, this time subjects also had the opportunity to give no answer when they thought they could not draw a conclusion.

Results and Discussion.




Figure 4: Mean judgments of preventive causal power for each ΔP level in Experiment 2a.

Figure 4 displays subjects's mean ratings of preventive power in Experiment 2a. Visual examination of Figure 4 shows the same trends for the negative ΔP s as in Experiment 1. But this time the line representing the zero ΔP s appears flatter. As predicted by the power PC Theory, subjects were not sure about their causality ratings for the $P(e|c)=P(e|\sim c)=.00$ condition: 35% of the subjects said they were unable to give a rating, compared to at most 14% of the same subjects in the other conditions. We therefore excluded this data point from further analysis.

A trend analysis for $\Delta P = .00$ still shows a significant linear trend $t(129) = 2.6, p < .02$. However, this trend appears to be solely due to the difference between the $P(e|c)=P(e|\sim c)=1.00$ condition and the other 3 conditions (for which causal power is defined). The linear trend for the other 3 zero ΔP conditions alone was not reliable, $t(86) = .85, p > .05$.

All other levels of ΔP yielded highly significant linear trends: $t(76) = 5.5, p < .001$ for $\Delta P = .75$, $t(150) = 8.68, p < .001$ for $\Delta P = .50$ and $t(225) = 10.7, p < .001$ for $\Delta P = .25$.

Figure 5 shows the results of Experiment 2b. Analogous to Experiment 2a, the line representing the zero ΔP s is flatter than that in Experiment 1b. In Experiment 2b, 22% of the subjects claimed to be unable to give a causal rating for the $P(e|c)=P(e|\sim c)=1.00$ condition, which is undefined according to the power PC Theory. As in Experiment 2a, this data point was excluded from the trend analysis. Despite the flatter line, the zero ΔP s still yielded a significant linear trend, $t(195) = 2.67, p < 0.002$. As in Experiment 2a, however, this trend appears to be entirely due to the lower mean rating for the $P(e|c)=P(e|\sim c)=.00$ condition. The trend analysis for the remaining 3 conditions was not reliable, $t(130) = .25, p > .05$.

In contrast, the linear trends for two of the nonzero ΔP conditions were highly reliable. $\Delta P = .75$ resulted in a highly significant linear trend, $t(72) = 3.17, p < .003$. $\Delta P = .50$ did not produce a significant trend, with $t(142) = 1.69, p$

< .1. But $\Delta P = .25$ showed a reliable trend, with $t(198) = 2.20, p < .05$.




Figure 5: Mean judgments of generative causal power for each ΔP level in Experiment 2b.

In summary, to reduce the possibility that subjects conflate reliability with strength, Experiments 2a and b presented the same information as Experiments 1a and b, but with a large number of replications across all conditions. We found reliable positive and negative trends for the non-zero ΔP conditions despite this change in procedure. For the zero ΔP conditions, the trends were still reliable, but they appear to be reduced and were due solely to the difference between the condition with an extreme P(e) and the rest of the conditions, which were not influenced by P(e|c).

General Discussion

The two experiments reported in this paper clearly contradict both the traditional contingency model and RWM as descriptions of human causality judgments. In Experiment 1a we showed that subjects's mean judgments of preventive causal power for conditions with equal levels of ΔP showed a significant negative linear trend as P(e|c) increased. The traditional contingency model, which holds ΔP as the only determinant of judged causality, erroneously predicts no trend. RWM with constant values of β across trials also predicts no trend. By assuming higher values of β for trials on which the effect is present than on trials on which it is absent, the RWM can explain the negative trends for candidates with equal negative ΔP s. But this assumption also predicts negative trends for candidates with equal positive ΔP s. Experiment 1b shows that, to the contrary, the trends were positive for such candidates.

The negative trends for candidates with equal negative ΔP s observed in Experiment 1a and the positive trends for candidates with equal positive ΔP s observed in Experiment 1b both support the power PC theory.

Subjects's ratings for zero ΔP s in Experiments 1a and 1b systematically deviated from the predictions of the accounts considered in this paper. A negative trend was found for the evaluation of preventive power; a positive trend was found for the evaluation of generative power. A possible explanation is that subjects conflated reliability with causal strength.

In Experiment 2a and 2b we attempted to eliminate this conflation by presenting the information in a visual format summarizing a large number of trials, thereby increasing

reliability across all conditions. We still found the respective negative and positive trends for candidates with equal negative and positive ΔP s, but we found flat lines for the candidates with zero ΔP s except for the condition with an extreme P(e). Except for this extreme condition, this pattern of results supports the power PC theory but contradicts the RWM and the traditional ΔP model. It is unclear why, contrary to all 3 accounts considered, the extreme conditions received mean ratings different from the other zero ΔP conditions.

In conclusion, the experiments reported in this paper clearly favor the power PC theory over either the RWM or the traditional contingency model as a description of human causal induction.

Acknowledgments

This paper is based on work supported under a Fulbright grant and a Friedrich-Ebert-Stipend awarded to the first author. Sean McAuliffe programmed Experiment 1a and 1b. We thank Michael Gold, Ernest Vina, and Melissa Wu for their assistance and Bruce Burns for valuable comments on an earlier draft.

References

- Baker, A.G., Berbrier, M.W., & Vallée-Tourangeau, F. (1989). Judgments of a 2x2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology*, 41B, 65-97.
- Chapman, G.B., Robbins, S.J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, 18 (5), 537-545.
- Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.
- Cheng, P.W., & Novick, L.R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, 58, 545-567.
- 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.
- Jenkins, H., & Ward, W. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, 7, 1-17.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research*. New York: Appleton-Century Crofts.
- Shanks, D.R. (1987). Acquisition functions in causality judgment. *Learning and Motivation*, 18, 147-166.
- Wasserman, E. A., Elek, S. M., Chatlosh, D.L. & Baker, A. G. (1993). Rating causal relations: Role of probability in judgments of response-outcome contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 174-188.