Learning & Behavior
2003, 31 (2), 192-204

Assessing power PC

LORRAINE G. ALLAN
McMaster University, Hamilton, Ontario, Canada

In a recent theoretical paper, Cheng (1997) presented a new causal model, power PC. She argued that power PC was able to account for data in the literature that raised problems for associative models—notably, the Rescorla–Wagner model (Rescorla & Wagner, 1972). The purpose of the present paper is threefold: (1) to show that, overall, the data in the literature, which Cheng relied on to make her case, do not in fact provide support for power PC, (2) to show that, overall, the experiments reported in the literature since the publication of Cheng, designed specifically to evaluate the predictions of power PC, also do not provide support for power PC, and (3) to suggest that Cheng's assessment of associative models was too narrowly defined.

Although everyone agrees that organisms must be able to appreciate the relationships among events in their environment in order to survive, there is less agreement about the mechanism that detects such relationships (see Allan, 1993; Shanks, 1993; Shanks, Holyoak, & Medin, 1996). Table 1 presents the standard 2 × 2 contingency matrix for the generic laboratory task used to study how human observers make judgments about the relationship between two binary variables (see Allan, 1980). In such tasks, the cue is either present (C) or absent (~C), and the outcome is either present (O) or absent (~O). The letters in the cells (a, b, c, and d) represent the joint frequency of occurrence of the four possible cue–outcome combinations. After a series of trials on which each of the four cue–outcome combinations are presented with a predefined probability, the observer is asked about the relationship between the cue and the outcome. The models that have been proposed to account for data generated in such binary judgment tasks are often categorized either as covariational or as causal.[1] The main distinction between these two categories of models is that covariation models do not assign cause and effect meanings to the cue and the outcome, whereas causal models do. Moreover, causal models postulate that judgments are influenced by knowledge about how causes are related to effects. The models are also categorized either as statistical or as associative. Statistical models represent observers as intuitive statisticians who extract covariation information by applying a rule to integrate frequencies or probabilities of events over time. In contrast, associative models postulate that judgments are determined by associative links or connections that are formed between contiguously presented cues and outcomes.

In a recent theoretical paper, Cheng (1997; see also Cheng, Park, Yarlas, & Holyoak, 1996) presented a new causal model, power PC. She argued that power PC was able to account for data in the literature that raised problems for associative models—notably, the Rescorla–Wagner (RW) model (Rescorla & Wagner, 1972). The purpose of the present paper is threefold: (1) to show that, overall, the data in the literature, which Cheng relied on to make her case, do not in fact provide support for power PC,[2] (2) to show that, overall, the experiments reported in the literature since the publication of Cheng, designed specifically to evaluate the predictions of power PC, also do not provide support for power PC, and (3) to suggest that Cheng's assessment of associative models was too narrowly defined.

POWER PC

One statistical measure of the covariation or contingency between the cue and the outcome is ΔP, which is the difference between two independent conditional probabilities (see Allan, 1980). With reference to Table 1,

$$\Delta P = P(\mathrm{O}|\mathrm{C}) - P(\mathrm{O}|{\sim}\mathrm{C}) = \frac{a}{a+b} - \frac{c}{c+d}. \qquad (1)$$

Many of the early studies of covariation judgments concentrated on determining whether humans could accurately judge the size and the sign of the contingency between two binary variables. Although most of these studies reported a high correlation between judgments and ΔP, systematic departures from ΔP were frequently noted (see Allan, 1993; Shanks, 1993; Shanks et al., 1996). For

The preparation of this paper was supported by a grant to L.G.A. from the Natural Sciences and Engineering Research Council of Canada. From January 2001 to May 2001, a group of faculty and graduate students at McMaster met weekly to discuss the various influences of associative learning theories on current cognitive research. My task was to summarize the causal learning literature, including the debate between researchers who support an associative account of causal learning and those who do not. The other members of the group encouraged me to put my verbal critique of Cheng (1997) to paper, and this manuscript is the outcome. I thank Lee Brooks, Jason LeBoe, Bruce Milliken, and Jason Tangen for the exciting and thought-provoking discussions that occurred during those weekly meetings. The manuscript benefited greatly from the insightful suggestions made by Ralph Miller, David Shanks, and Edward Wasserman on an earlier draft. Correspondence should be addressed to L. G. Allan, Department of Psychology, McMaster University, Hamilton, ON, L8S 4K1 Canada (e-mail: allan@mcmaster.ca).

Table 1
Standard 2 × 2 Contingency Matrix

		Outcome O	Outcome ~O	
Cue	C	a	b	$a + b$
	~C	c	d	$c + d$
		$a + c$	$b + d$	

$$\Delta P = P(\mathrm{O}|\mathrm{C}) - P(\mathrm{O}|\sim\mathrm{C}) = \frac{a}{a+b} - \frac{c}{c+d}$$

example, there have been many reports of an outcome *density bias*: Judgments of contingency are not constant for a fixed ΔP but increase with the frequency of the outcome. Also, in experiments in which there are multiple cues, judgment of the relationship between each cue and the outcome is influenced by the copresence of the other cues and by the pairing history of the cues.

Cheng and Novick (1990, 1992) developed the probabilistic contrast model (PCM) to account for departures in judgments from ΔP—notably, those that occur when multiple cues are presented. For the standard or unconditional ΔP rule (Equation 1), $P(\mathrm{O}|\mathrm{C})$ is based on all C trials, $P(\mathrm{O}|\sim\mathrm{C})$ is based on all ~C trials, and other cues are ignored in the determination of ΔP. In PCM, the ΔP rule applies across a *focal set* of trial types, rather than across all trials. Specifically,

$$\Delta P = P(\mathrm{O}|\mathrm{CA}) - P(\mathrm{O}|\sim\mathrm{CA}), \tag{2}$$

where A represents all other cues that are the same on C and ~C trials. That is, focal set ΔP is based on trials that are identical except for the presence or the absence of C. Consider the situation in which there are two cues, C_1 and C_2, and a common outcome, O. A focal set for Cue C_1 can be delineated by the presence or absence of Cue C_2; that is, the value of ΔP for C_1, conditional on the presence of C_2, is

$$\Delta P_{\mathrm{C}_1|\mathrm{C}_2} = P(\mathrm{O}|\mathrm{C}_1\mathrm{C}_2) - P(\mathrm{O}|\sim\mathrm{C}_1\mathrm{C}_2), \tag{3a}$$

and the value of ΔP for C_1, conditional on the absence of C_2, is

$$\Delta P_{\mathrm{C}_1|\sim\mathrm{C}_2} = P(\mathrm{O}|\mathrm{C}_1 \sim\mathrm{C}_2) - P(\mathrm{O}|\sim\mathrm{C}_1 \sim\mathrm{C}_2). \tag{3b}$$

PCM, like the ΔP rule, is concerned with covariation, not causation. Cheng (1997) elaborated PCM to encompass causal power—the causal power theory of the probabilistic contrast model, or more simply, power PC. According to power PC, when alternative causes to C are controlled, the generative causal power of C for $\Delta P \geq 0$ is

$$p_\mathrm{c} = \frac{\Delta P}{1 - P(\mathrm{O}|\sim\mathrm{C})} \text{ for } P(\mathrm{O}|\sim\mathrm{C}) \neq 1, \tag{4}$$

and the preventive causal power of C for $\Delta P \leq 0$ is

$$p_\mathrm{c} = \frac{-\Delta P}{P(\mathrm{O}|\sim\mathrm{C})} \text{ for } P(\mathrm{O}|\sim\mathrm{C}) \neq 0, \tag{5}$$

where the predictions are assumed to be only ordinal. According to Equations 4 and 5, reaching a conclusion about the causal status of C depends not only on ΔP, but also on $P(\mathrm{O}|\sim\mathrm{C})$.[3] Specifically, (1) generative p_c is undefined for $P(\mathrm{O}|\sim\mathrm{C}) = 1$, and preventive p_c is undefined for $P(\mathrm{O}|\sim\mathrm{C}) = 0$; (2) $p_\mathrm{c} = \Delta P$ for $\Delta P = 0$, unless p_c is undefined; and for a constant ΔP ($\Delta P \neq 0$), as $P(\mathrm{O}|\sim\mathrm{C})$ increases, generative p_c becomes larger, and preventive p_c becomes smaller (see Figure 1).

Undefined p_c

Cheng (1997) described the following anecdotal example to illustrate a judgment situation in which $P(\mathrm{O}|\sim\mathrm{C}) = 1$. Suppose that you think that you are allergic to certain foods. The doctor makes a grid of scratches on your skin, puts food samples (C) on some of the scratches, and leaves other scratches untouched (~C). You observe that there is an allergic reaction (O) at every scratch, those scratches with food samples and also those scratches without food samples—that is, $P(\mathrm{O}|\mathrm{C}) = P(\mathrm{O}|\sim\mathrm{C}) = 1$ and $\Delta P = 0$. According to a statistical covariation model, the observers should judge the cue and the outcome as unrelated, but according to power PC, the observers would be uncertain about the causal status of C. Cheng also described the following anecdotal example to illustrate a judgment situation in which $P(\mathrm{O}|\sim\mathrm{C}) = 0$. Suppose that you are a medical researcher who has developed a drug for relieving headaches. You randomly assign participants to two groups, administering the new drug to one group (C) and a placebo to the other group (~C). No participant in either group reports a headache (O); that is, $P(\mathrm{O}|\mathrm{C}) = P(\mathrm{O}|\sim\mathrm{C}) = 0$ and $\Delta P = 0$. According to a statistical covariation model, the observers should judge the cue and the outcome as unrelated, but according to power PC, observers would realize that no firm conclusion could be drawn about C.

Cheng (1997) acknowledged that there were no published studies in which observers were given the option of explicitly expressing uncertainty about the causal status of C and, therefore, there were no published data directly relevant to these differential predictions of the covariation and the power approaches. The only research reported in Cheng that apparently showed that observers who were given the opportunity did express uncertainty about the causal status of C was attributed to Fratianne and Cheng but has not yet been published.

Since the publication of Cheng (1997), Wu and Cheng (1999) have provided data from an experiment designed to evaluate power PC predictions. Observers were presented with a descriptive scenario regarding a cue and an outcome. There were six types of scenarios, resulting from crossing two types of power and three levels of $P(\mathrm{O})$. In the generative scenario, C was intended to increase the outcome, and in the preventive scenario, C was intended to decrease the outcome. For both types of power, the outcome always occurred [$P(\mathrm{O}) = 1$], the outcome never occurred [$P(\mathrm{O}) = 0$], or the outcome sometimes occurred [$0 < P(\mathrm{O}) < 1$]. The observers were asked to decide (1) whether C was causal, (2) whether C was noncausal, or (3) whether the information in the scenario was uninformative with regard to the evaluation of the effectiveness of C. According to power

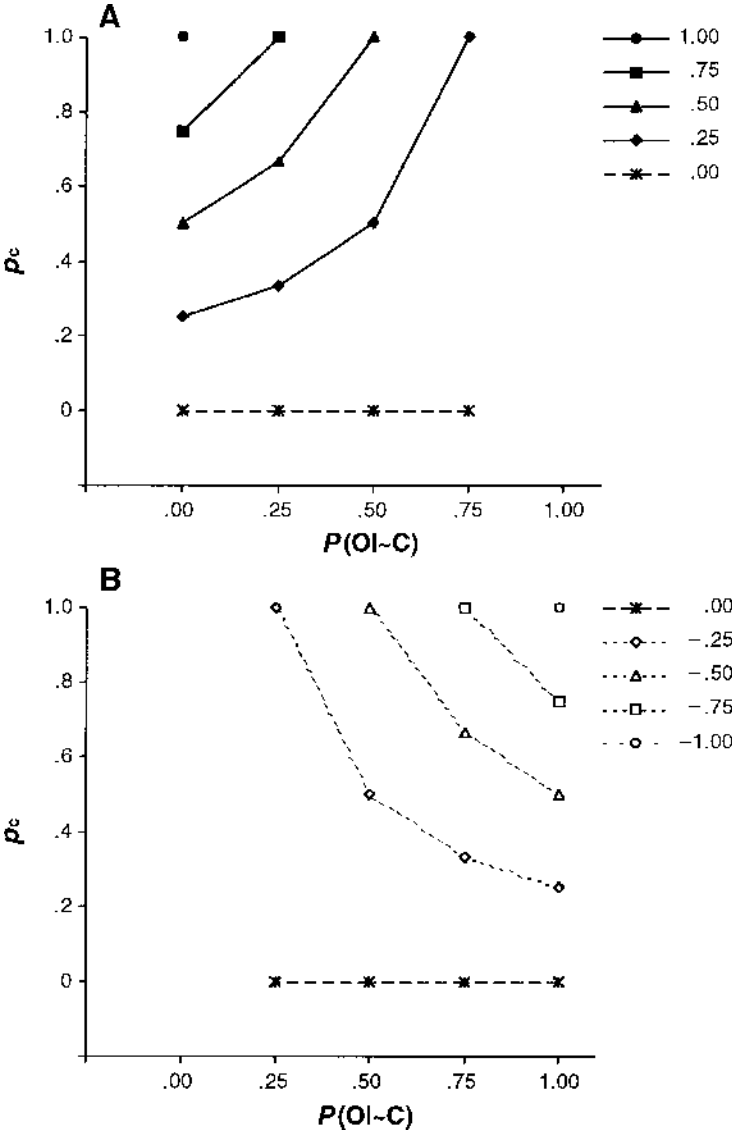

Figure 1. p_c as a function of P(O|~C). Each curve describes the relationship for a constant value of ΔP. Panel A shows the relationship for generative p_c (ΔP = 0, .25, .50, .75, and 1.0), and Panel B shows the relationship for preventive p_c (ΔP = 0, −.25, −.50, −.75, and −1.0). As P(O|~C) increases, generative p_c increases, and preventive p_c decreases.

PC, the observers should select the uninformative option for the generative scenario when P(O|~C) = 1 and that for the preventive scenario when P(O|~C) = 0. Although the majority of the observers in these two conditions did select the uninformative option, about 30% selected the noncausal option. Also, according to power PC, observers should select the noncausal option for the remaining four conditions, but many selected the uninformative option. In fact, Wu and Cheng noted that "a moderate majority of the participants responded in accord with the power PC theory. One possible interpretation is that this theory does only moderately well" (p. 95). However, rather than accepting the interpretation that the theory does only moderately well, Wu and Cheng suggested that their data were noisy because they were generated as part of a long questionnaire given to large groups of observers.[4]

Buehner and Cheng (1997, Experiment 2) also gave their observers the opportunity not to provide a rating if they thought the information was inadequate. For preventative causal power, 35% of the observers indicated that they were unable to give a rating when ΔP = 0 and P(O|~C) = 0; for generative causal power, 22% of the observers indicated that they were unable to give a rating when $\Delta P = 0$ and $P(\text{O}|{\sim}\text{C}) = 1$. Thus, the majority of the observers were quite willing to provide ratings.

$\Delta P = 0$ and $P(\text{O}|{\sim}\text{C}) \neq 0, 1$

According to Equations 4 and 5, $p_c = 0$ (or at least is constant) when $\Delta P = 0$ and $P(\text{O}|{\sim}\text{C}) \neq 0, 1$. Cheng (1997) did admit that there was ample evidence in the literature that judgments are not constant across variations in $P(\text{O}|{\sim}\text{C})$ when $\Delta P = 0$. Many studies have reported an outcome density bias in which judgments increase with the probability of the outcome [$P(\text{O})$]. Cheng concluded that although this density bias "might appear to contradict the power PC theory . . . this 'bias' is consistent with the power PC theory . . . if (a) there are variations in the objective value of ΔP from participant to participant (with a mean of zero across participants) or (b) participants are likely to misperceive an objective ΔP of 0" (p. 389). That is, although the value of ΔP is programmed to be zero over the complete series of trials, the actual presented value might not be zero because of the random nature of the generator of the trial events. If the actual presented value of ΔP is not zero, then according to Equations 4 and 5, p_c will depend on $P(\text{O}|{\sim}\text{C})$. Cheng also noted that in most studies observers made judgments about many different contingencies. She argued that such within-subjects designs might bias the observer to report a contingent relationship when $\Delta P = 0$. That is, preceding contingent relationships might induce an observer to judge a subsequent noncontingent relationship as contingent.

Vallée-Tourangeau, Murphy, Drew, and Baker (1998) evaluated Cheng's (1997) arguments for the deviation of data from the predictions of power PC when $\Delta P = 0$ and $P(\text{O}|{\sim}\text{C}) \neq 0, 1$. They noted that presented ΔP would approach the programmed value of zero with increasing trials and, therefore, the dependence of p_c on $P(\text{O}|{\sim}\text{C})$ when $\Delta P = 0$ should be transient. They presented evidence that the effect of $P(\text{O}|{\sim}\text{C})$ on judgments was present over trials (i.e., the density bias was maintained). They also provided a summary of the density bias literature, which indicated that a density bias had been reported in studies in which a between-subjects design had been used. Moreover, they found a density bias in their data even when they restricted their analysis to the first judgment an observer made. Lober and Shanks (2000) also noted that the patterns of results in their experiments were similar when they restricted their analysis to the first set of trials presented to an observer.

p_c and $P(\text{O}|{\sim}\text{C})$

As is shown in Figure 1, for a constant ΔP, as $P(\text{O}|{\sim}\text{C})$ increases, generative p_c increases, and preventive p_c decreases. Cheng (1997) noted that several published articles had reported that cues with the same nonzero ΔP but different values of $P(\text{O}|{\sim}\text{C})$ were, in fact, judged differently. She specifically referred to the data reported by Wasserman, Elek, Chatlosh, and Baker (1993) and by Allan and Jenkins (1983). Wasserman et al. (1993) varied the size and the sign of ΔP and the value of $P(\text{O}|{\sim}\text{C})$. Allan and Jenkins (1983) also varied $P(\text{C})$.

Wasserman et al. (1993). In Experiment 1 in Wasserman et al. (1993), observers were required to judge whether tapping a key influenced the presentation of a light. The light was programmed to occur with various probabilities at the end of 1-sec sampling intervals. Specifically, if the observer responded at least once during the 1-sec interval, $P(\text{O}|\text{C})$ was probed at the end of the interval; otherwise, $P(\text{O}|{\sim}\text{C})$ was probed. If the probe was positive, the light occurred, whereas if the probe was negative, the light did not occur. Wasserman et al. (1993) varied both $P(\text{O}|\text{C})$ and $P(\text{O}|{\sim}\text{C})$, resulting in values of ΔP that varied from -1.0 to $+1.0$. The observers were asked to rate the effect that tapping the key had on producing the light, on a scale ranging from -100 to $+100$. The data from Wasserman et al. (1993, Experiment 1) are reproduced in Figure 2A, which plots ratings as a function of $P(\text{O}|{\sim}\text{C})$. The lines connect ratings for a constant ΔP value. In Wasserman et al. (1993), the rating scale ranged from -100 to $+100$, whereas p_c is always positive for preventive power (Equation 5), as well as for generative power (Equation 4) and, therefore, ranges from 0 to $+1.0$. To facilitate the visual comparison of the predictions of power PC with the data in Wasserman et al. (1993), the predictions of power PC shown in Figures 1A and 1B were replotted in Figure 2B, where preventive values of p_c have a negative sign.

Figure 2A indicates that for both positive and negative ΔP values, ratings tend toward zero as $P(\text{O}|{\sim}\text{C})$ increases. Although this trend toward zero is consistent with the predictions of power PC for negative ΔP values, it is not consistent with the predictions for positive ΔP values. Cheng (1997), however, concluded that the deviation of the positive ΔP data did not raise problems for power PC, because Wasserman et al. (1993) varied rate, not probability. She argued that whereas probabilities have a lower and an upper bound, rates have a lower bound (not at all), but not an upper bound. For this reason, the preventive judgments in Wasserman et al. (1993) should have decreased with increases in $P(\text{O}|{\sim}\text{C})$, whereas the generative judgments should not have changed as a function of $P(\text{O}|{\sim}\text{C})$.

Cheng (1997) did not acknowledge that Wasserman, Chatlosh, and Neunaber (1983) examined the validity of partitioning their continuous trial procedure into discrete intervals in order to calculate probabilities. In Experiment 2 in Wasserman et al. (1983), three sampling intervals were used: a 1-sec interval with 240 intervals, a 4-sec interval with 60 intervals, and a 1-sec interval with 60 intervals. In Experiment 3, a fixed sampling interval of 3-sec was compared with an average sampling interval of 3-sec, where on any trial the sampling interval could be 1, 3, or 5 sec. In both experiments, variations in sampling interval did not influence either the probability of a response or the relationship of the contingency judgments to ΔP. Wasserman et al. (1983) concluded that their continuous trial procedure could be partitioned into discrete intervals in order to calculate probabilities.

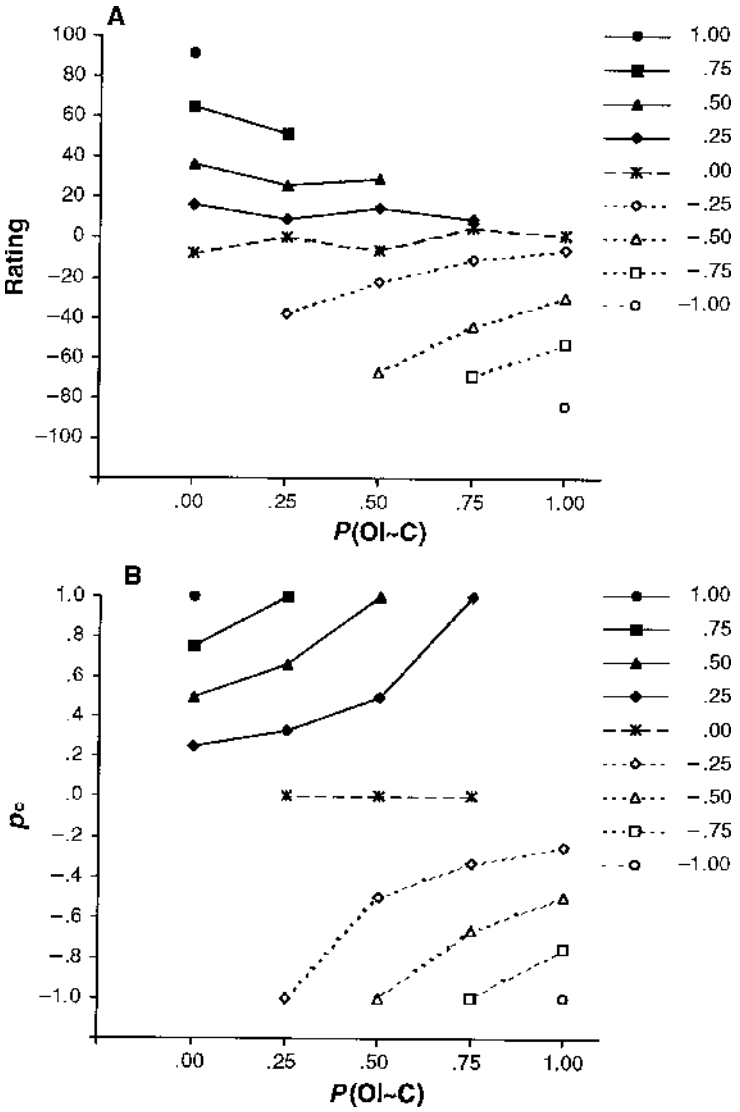

Figure 2. (A) Ratings as a function of $P(\mathrm{O}|{\sim}\mathrm{C})$. Each curve describes the relationship for a constant value of ΔP. Data are replotted from Wasserman, Elek, Chatlosh, and Baker (1993, Experiment 1). (B) p_c as a function of $P(\mathrm{O}|{\sim}\mathrm{C})$. Each curve describes the relationship for a constant value of ΔP. Preventive p_c is signed as negative.

In addition, Cheng (1997) did not reference other papers by Wasserman and his colleagues (e.g., Kao & Wasserman, 1993; Levin, Wasserman, & Kao, 1993; Wasserman, Dorner, & Kao, 1990), which avoided the rate versus probability issue. In these studies, either the trials were discrete or the cue–outcome information was summarized in a 2 × 2 matrix. These papers were concerned with the relative influence of the four cells of the contingency matrix on judgments. The data clearly showed that the four cells were weighted differentially. Specifically, they were weighted in the order cell $a >$ cell $b >$ cell $c >$ cell d. Thus, for any constant ΔP, judgments were not constant but depended on the cell frequencies. Since these studies predated Cheng, the relationship between the ratings and p_c was not examined. For many of the studies, the information needed to calculate p_c is available, as are the observer ratings. The data from these studies do not support power PC. As an example, the 25 problems in Experiment 2 in Wasserman et al. (1990) and the mean rating for each problem are reproduced in Table 2. It is clear that the ratings are far from constant for a constant p_c.

Since the publication of Cheng (1997), Buehner and Cheng (1997)[5] reported data from an experiment that was similar to Experiment 1 in Wasserman et al. (1993), but with discrete trials, thereby avoiding the rate versus probability issue. They also explicitly identified causal type

Table 2
Problems Used in Wasserman, Dorner, and Kao (1990, Experiment 2)

a	b	c	d	P(O|C)	P(O|~C)	ΔP	p_c	Rating
				Preventive				
10	10	20	10	.500	.667	−.167	.250	−0.162
10	10	10	5	.500	.667	−.167	.250	0.018
10	20	10	10	.333	.500	−.167	.333	−0.304
5	10	10	10	.333	.500	−.167	.333	−0.183
10	10	20	5	.500	.800	−.300	.375	−0.118
10	20	10	5	.333	.667	−.333	.500	−0.330
5	10	20	10	.333	.667	−.333	.500	−0.341
5	10	10	5	.333	.667	−.333	.500	−0.272
5	10	20	5	.333	.800	−.467	.583	−0.365
5	20	10	10	.200	.500	−.300	.600	−0.510
5	20	10	5	.200	.667	−.467	.700	−0.478
5	20	20	5	.200	.800	−.600	.750	−0.541
				$\Delta P = 0$				
10	10	10	10	.500	.500	.000	.000	0.033
				Generative				
10	10	10	20	.500	.333	.167	.250	0.113
10	10	5	10	.500	.333	.167	.250	0.124
20	10	10	10	.667	.500	.167	.333	0.452
10	5	10	10	.667	.500	.167	.333	0.262
10	10	5	20	.500	.200	.300	.375	0.138
20	10	5	10	.667	.333	.333	.500	0.573
10	5	10	20	.667	.333	.333	.500	0.282
10	5	5	10	.667	.333	.333	.500	0.389
10	5	5	20	.667	.200	.467	.583	0.335
20	5	10	10	.800	.500	.300	.600	0.548
20	5	5	10	.800	.333	.467	.700	0.678
20	5	5	20	.800	.200	.600	.750	0.693

(generative or preventive). Buehner and Cheng found that when $\Delta P \neq 0$, judgments varied systematically with $P(O|{\sim}C)$ for a constant ΔP. Specifically, as $P(O|{\sim}C)$ increased, generative p_c increased, and preventive p_c decreased. Buehner and Cheng concluded that this pattern was consistent with power PC. For $\Delta P = 0$, neither generative nor preventive ratings were constant (as predicted by Equations 4 and 5). Rather, as with $\Delta P \neq 0$, generative ratings increased and preventive ratings decreased as $P(O|{\sim}C)$ increased. To explain the deviation of the data from power PC for $\Delta P = 0$, Buehner and Cheng suggested that observers conflated reliability with causal strength. Consider the generative conditions in which $P(O|C) = P(O|{\sim}C) = 0$ and $P(O|C) = P(O|{\sim}C) = .75$. When $P(O|C) = P(O|{\sim}C) = 0$, there are more trials on which C could have but failed to prove its causal power than when $P(O|C) = P(O|{\sim}C) = .75$. The observer, therefore, would be more confident of a noncausal rating in the former condition than in the latter condition, leading to a causal rating closer to zero. That is, providing observers with a constant number of trials across the $\Delta P = 0$ conditions yielded varying reliability of the information presented.

Buehner and Cheng (1997) suggested that increasing the number of trials might unconfound reliability and causal strength, since there would be more opportunities for C to prove its causal power. In Experiment 2, they increased the number of trials by presenting the cue–outcome information in a summary (matrix) format, rather than in a trial format, and informing the observer that the summary data were based on 100 trials. Although the influence of $P(O|{\sim}C)$ was smaller in the summary format than in the trial format, the linear trend was significant, for both generative and preventive power, across the four values of $P(O|{\sim}C)$ for which p_c was defined.

Lober and Shanks (2000) replotted the data from Buehner and Cheng (1997) to explicitly show the relationship between judgments and p_c. According to Equations 4 and 5, judgments should be constant for a constant p_c. The replotted data clearly showed systematic changes in judgment when p_c was constant. Buehner (2001) attributed these deviations to a conflation of reliability with causal strength. He did admit, however, that "the experimental design reported here cannot shed light on whether this hypothesis is correct or not" (p. 51).

Lober and Shanks (2000) reported a series of six experiments explicitly designed to evaluate whether ratings were better related to ΔP or to p_c. In Experiments 1–3, the information about the cue–outcome relationship was presented in a trial format. To address the concern of Buehner and Cheng (1997) over confounding reliability and power, Lober and Shanks used a large number of trials (56 or 60 trials, as compared with 16 in Buehner and Cheng) and obtained ratings after blocks of 10 or 20 trials. The strategy in these experiments was to keep p_c constant and vary ΔP or to keep ΔP constant and vary p_c. In Experiments 1

Table 3
Problems Used in Allan and Jenkins (1983, Experiment 3)

P(O|C)	P(O|~C)	ΔP
.1	.1	0
.3	.3	0
.5	.5	0
.7	.7	0
.9	.9	0
.3	.1	.2
.9	.7	.2
.5	.1	.4
.9	.5	.4
.9	.1	.8

and 2, Lober and Shanks found that, contrary to power PC, ratings increased with ΔP even though p_c was constant. The data from Experiment 3, however, were consistent with power PC, in that ratings for a constant ΔP increased as p_c increased. Experiments 4–6 were similar to Experiments 1–3, except that the cue–outcome relationship was presented in summary format. For the summary format, the data from all three experiments were at variance with power PC, in that judgments tracked ΔP, and not p_c. Lober and Shanks concluded that "the power PC theory was unequivocally contradicted by the results obtained in these experiments" (p. 195).

Vallée-Tourangeau, Murphy, and Drew (1997) also examined the effect of ΔP and p_c on ratings. They found that judgments varied with changes in ΔP for a constant value of p_c, and they concluded that their data "disconfirmed the predictions of the power PC theory" (p. 779). In a later paper, Vallée-Tourangeau et al. (1998) reported results similar to those of Buehner and Cheng (1997). For a constant ΔP, judgments varied systematically with P(O|~C). Generative p_c increased and preventive p_c decreased as P(O|~C) increased. This was the case for all values of ΔP, including $\Delta P = 0$. As was indicated earlier, Vallée-Tourangeau et al. (1998) were critical of the post hoc accounts provided by Buehner and Cheng for the effect of P(O|~C) on ratings when $\Delta P = 0$.

Allan and Jenkins (1983). Allan and Jenkins (1983) used a discrete trial procedure, and Cheng concluded that, overall, their data were consistent with power PC. This conclusion was based on only nine comparison pairs selected by Cheng (1997) across all the experiments reported by Allan and Jenkins. Within each of these nine pairs, ΔP was constant, and P(O|~C) varied. For seven of these nine comparisons, the generative rating was higher for the larger value P(O|~C). It should be noted that although the majority of the pairs selected by Cheng were consistent with power PC, this result would not be significant according to a binomial test.

In fact, although not noted by Cheng (1997), Experiment 3 in Allan and Jenkins (1983) provided data that were clearly at variance with power PC. In that experiment, observers were required to judge the influence of the movement of a joystick on the movement of a dot. At the beginning of a trial, the joystick was represented in its resting position on the left side of the computer screen, and the dot was present in the middle of the right side of the screen. On each trial, the joystick either moved or remained in its resting position (the cue), and then the dot either moved downward or remained stationary (the outcome). At the end of a series of 50 cue–outcome pairings, the observer rated the influence of joystick position on dot movement on a 40-point scale. Allan and Jenkins varied the contingency between joystick movement and dot movement (ΔP), the probability of dot movement [P(O)],

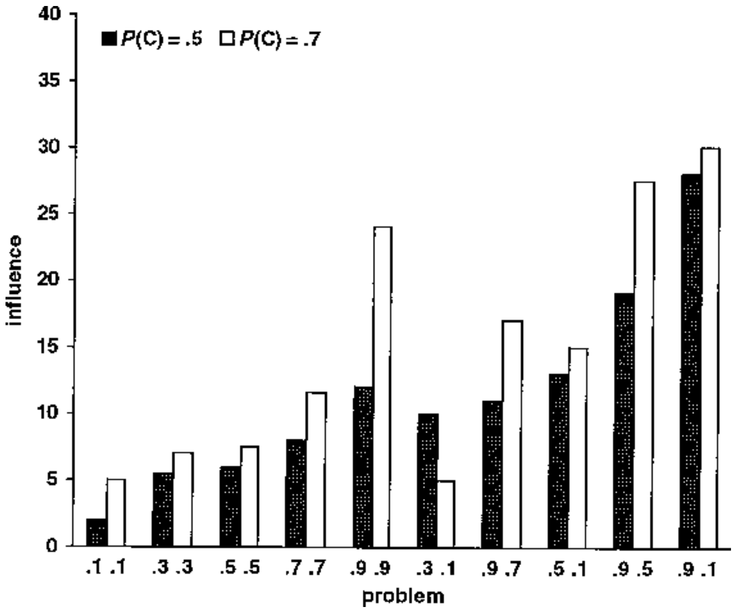

Figure 3. Mean influence for each of the 20 problems in Allan and Jenkins (1983, Experiment 3).

and the probability of joystick movement [P(C)], resulting in 20 different trial sequences or problems (see Table 3). For 10 problems, P(C) = .5, and for 10 problems, P(C) = .7. There were 10 noncontingent ($\Delta P = 0$) problems and 10 contingent ($\Delta P > 0$) problems, and for a fixed value of ΔP, P(O|~C) was varied. Each observer rated each of the 20 problems, which were presented in a random order over two sessions.

The data from Experiment 3 in Allan and Jenkins (1983) are reproduced in Figure 3. Mean influence is shown for each of the 20 problems. Influence judgments increased as ΔP increased but were also dependent on both P(O) and P(C). Two features of the data are inconsistent with power PC. First consider the 10 pairs in which both ΔP and P(O|~C) were constant but P(C) varied (either .5 or .7). Contrary to power PC, judgments were not constant with variations in P(C). Second, consider judgments for ΔP = 0. Again, contrary to power PC, judgments were not constant but were dependent on both P(O|~C) and P(C).

Trial Effects and Trial-Sequencing Effects

Cheng (1997) did not address data from experiments concerned with trial effects. At that time, there was ample evidence in the literature that covariation judgments changed over trials (see Allan, 1993; Shanks, 1993). For example, Shanks (1985a, 1987) probed his observers for judgments a number of times during the sequential presentation of the cue–outcome pairings. He found that judgments became more positive across trials when the contingency was positive and became more negative across trials when the contingency was negative. When the contingency was zero, judgments first increased across trials and then decreased. The size of the deviation of the rating from zero depended on P(O): the higher the value P(O), the greater the deviation. Power PC (and also PCM) is not able to encompass systematic changes in ratings across trials. Although estimates of conditional probabilities become more accurate with increasing sample size, the mean estimate is independent of sample size. Thus, contrary to the data, power PC predicts that mean judgment should not change over trials.

Cheng (1997) also did not address data from experiments concerned with trial-sequencing effects (e.g., Yates & Curley, 1986). In most causal learning studies, information from the four cells of Table 1 are evenly distributed throughout the trial sequence. Yates and Curley, and, more recently, Dennis and Ahn (2001) and López, Shanks, Almaraz, and Fernández (1998) manipulated the sequencing of the four trial types. These trial-sequencing experiments indicate that judgments usually are not based on information integrated over the entire trial sequence. Rather, the order in which the information is presented influences the judgment. For example, Dennis and Ahn found that judgments were more influenced by early trials than by late trials (a primacy effect). Although others (e.g., López et al., 1998) have reported that late trials were more important than early trials (a recency effect), there is general agreement that information is not integrated equally across all trials. Dennis and Ahn concluded that "at this stage, any order effect is beyond the boundary conditions of the power PC theory because the causal strength of an event is calculated over all available trials all at once when enough observations are assumed to have been accumulated" (p. 160).

Summary

The data in the literature, which Cheng (1997) relied on to make her case, do not in fact provide support for power PC. With regard to ratings when p_c is undefined, there were no data available in the published literature to support power PC's prediction that observers would be uncertain. With regard to the effect of P(O|~C) on ratings when $\Delta P = 0$, Cheng's post hoc explanations of the outcome density bias do not provide an adequate account of the deviations of the data from the predictions of power PC. With regard to the effect of P(O|~C) on ratings when $\Delta P \neq 0$, Cheng's criticism of Wasserman et al. (1993) is unwarranted, and Allan and Jenkins's (1983) data, rather than supporting power PC, actually provide strong evidence against the model. Moreover, data available in the literature that were not in accord with the predictions of power PC were not cited (e.g., Kao & Wasserman, 1993; Levin et al., 1993; Wasserman et al., 1990). Finally, putatively relevant data cited by Cheng have not been published.

Since the publication of Cheng (1997), Cheng and her colleagues (Buehner, 2001; Buehner & Cheng, 1997; Wu & Cheng, 1999) have reported new data, and they have concluded that these data are supportive of power PC. In fact, the data in these papers deviated from the predictions of power PC, and the authors relied on ad hoc explanations in an attempt to explain the deviations of the data from the predictions of power PC. Wu and Cheng were critical of the design of their own experiment. Buehner and Cheng (1997; Buehner, 2001) acknowledged that their data deviated from the predictions of Equations 4 and 5 but attributed these deviations to a conflation of reliability with causal strength and to the use of within-subjects designs. These post hoc accounts either do not stand up to scrutiny or have not been empirically investigated.

Other researchers who have conducted experiments explicitly to evaluate power PC have concluded that, overall, their data deviated from the predictions of power PC (Lober & Shanks, 2000; Vallée-Tourangeau et al., 1997; Vallée-Tourangeau et al., 1998). Also the data in the literature concerning trial effects (see Allan, 1993; Shanks, 1993) and trial-sequencing effects (e.g., Dennis & Ahn, 2001; López et al., 1998; Yates & Curley, 1986) cannot be encompassed by power PC.

ASSOCIATIVE MODELS: RESCORLA–WAGNER

Cheng (1997) acknowledged that, for many of the experiments reported in the literature, the predictions of power PC are the same as the asymptotic predictions of RW. She concluded that when the predictions were different, power

PC was superior to RW. It should be emphasized that Cheng's conclusion was based on the asymptotic predictions of RW and, also, on the original version of RW, rather than on more recent modifications of RW (e.g., Van Hamme & Wasserman, 1994).

Preasymptotic Judgments

As was noted earlier, covariation judgments change over trials (see Allan, 1993; Shanks, 1993). One of the strengths of associative models is that they predict both the course of acquisition and the final asymptotic level of judgments. According to RW, for example, a cue gains associative (predictive) strength only to the extent that it provides information about the occurrence of the outcome that was not available from another source. The change in the predictive strength of the outcome by the cue is proportional to the degree to which the outcome is unexpected or surprising given all the alternative cues present on that trial. More precisely, the predictive strength of a cue (V_C) for the outcome will change on each trial in which it is presented according to the standard linear operator equation

$$\Delta V_C = \alpha\beta\left(\lambda - \Sigma V\right), \tag{6}$$

where ΔV_C is the change in predictive strength of the cue, α and β are learning rate parameters that depend on the salience of the cue and the effectiveness of the outcome, respectively, λ is the maximum amount of predictive strength supported by the outcome, and $\sum V$ is the sum of the predictive strengths of all cues present on that trial. V_C is assumed to be linearly related to performance.

The essence of RW is competition: There is a limit to the amount of predictive strength that an outcome can support. This limited amount of predictive strength is allocated among all the cues present on the trial: If one cue acquires more of the predictive strength available, all the other cues that are present at the same time must get less. Since a cue never occurs in isolation but is always compounded with alternative cues, the outcome is associated with the target cue and, also, with alternative cues. Contingency manipulations affect predictive strength of the target cue because of the competition for the predictive strength among the various associations formed, both with the target cue and with alternative cues.

V_C will continue to change over trials until ΔV_C is zero. When V_C is at asymptote, $V_C = \Delta P$ if β_O (outcome present) $= \beta_{\sim O}$ (outcome absent).[6] That is, under some circumstances, the predictive strength of the target cue, as described by RW, is identical to the contingency between the cue and the outcome as described by ΔP (see Chapman & Robbins, 1990). Thus, RW explains how a sensitivity to ΔP emerges from a process that does not explicitly represent ΔP.

RW, like the ΔP rule, stipulates that asymptotic V_C is independent of P(O) and of P(C). These marginal probabilities, however, will affect preasymptotic V_C according to RW. Figure 4 shows simulated RW acquisition functions for two of the zero contingencies in Allan and Jenkins (1983, Experiment 3): [P(O|C) = P(O|~C) = .5] and [P(O|C) = P(O|~C) = .9]. For each of these zero contingencies, curves were generated for P(C) = .5 and P(C) = .7. The simulations were for 50 trials, and the data were averaged over 25 runs, with trials randomly presented on each run. They were generated for an arbitrary set of parameter values: α_C (salience of cue) = .9, α_A (salience of alternative cues) = .35, $\beta_O = \beta_{\sim O} = .2$, $\lambda = 1$ on O trials, and $\lambda = 0$ on ~O trials. It is clear from Figure 4 that RW predicts that V_C is positive on early trials even when $\Delta P = 0$ and that the size of V_C depends on P(O|~C) and on P(C). That is, the RW model predicts a preasymptotic density

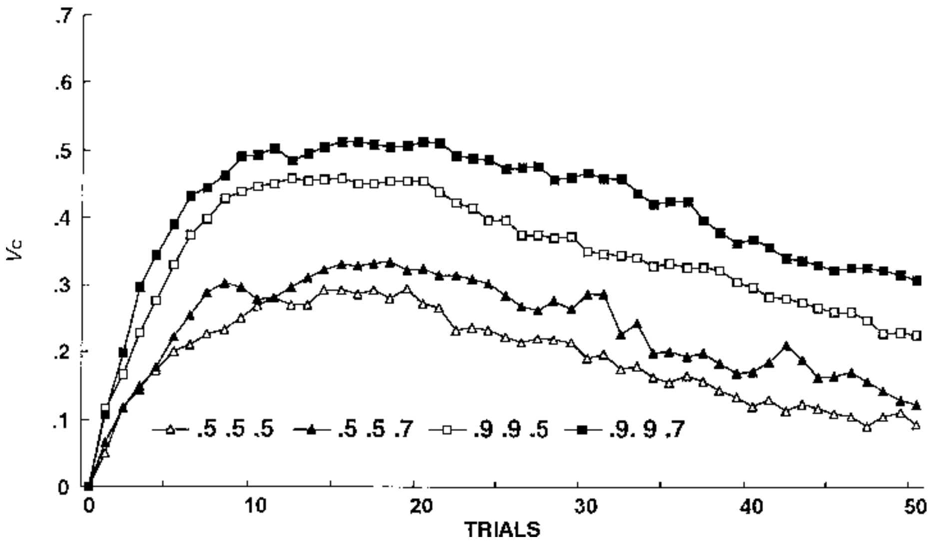

Figure 4. Simulated Rescorla–Wagner acquisition functions for two of the zero contingencies in Allan and Jenkins (1983, Experiment 3): [P(O|C) = P(O|~C) = .5] and [P(O|C) = P(O|~C) = .9]. For each zero contingency, curves were generated for P(C) = .5 and P(C) = .7.

bias. The size of the density bias depends on the values of the parameters.

Thus, RW can readily accommodate the course of judgments over trials. Moreover, it provides an account for the frequently observed density bias. Within the framework of RW, one would expect to observe a density bias if the judgments were preasymptotic. Experiments that have tracked judgments over trials generally have shown that the density bias is present for preasymptotic judgments (see Allan, 1993).

COMPARISON OF POWER PC AND RW

As was summarized by Allan (1993) and Shanks (1993), human contingency judgments show cue interaction effects, such as blocking, overshadowing, conditioned inhibition, and relative cue validity.[7] Cheng (1997) acknowledged that for most cue interaction experiments reported in the literature, power PC and RW do not make differential predictions for asymptotic ratings. She specifically directed her attention toward experimental manipulations under which power PC and RW make different predictions. She cited three papers reporting data from her laboratory, which she claimed were at variance with RW but were consistent with power PC. Two of these papers have not been published (Fratianne & Cheng; Park & Cheng). Thus, there was only one relevant published report (Yarlas, Cheng, & Holyoak, 1995).

The experiments reported in Yarlas et al. (1995) were concerned with conditioned inhibition. In these experiments, Cue C was paired with the outcome, and Cue C compounded with Cue X (CX) was presented without the outcome. According to power PC,

$$\Delta P_{\mathrm{X}|\mathrm{C}} = P(\mathrm{O}|\mathrm{CX}) - P(\mathrm{O}|\mathrm{C}\sim\mathrm{X}) \tag{7}$$

will be negative, and therefore, p_c will be preventive. For RW, ΔV_X will be negative, because on CX trials, the outcome is absent and, therefore, $\lambda = 0$ (see Equation 6). Thus, both power PC and RW predict that Cue X would be established as a conditioned inhibitor. Yarlas et al. (1995), as well as others (e.g., Chapman, 1991; Chapman & Robbins, 1990; Williams, 1995), demonstrated that conditioned inhibition can be established in human judgment tasks.

Although both RW and power PC predict the acquisition of conditioned inhibition, they make different predictions regarding the extinction of a conditioned inhibitor. Cheng (1997) discussed two extinction procedures for conditioned inhibition. Under the *direct* procedure, only Cue X is presented, and the outcome is absent. According to power PC, Equation 5 would apply. Since $P(\mathrm{O}|\sim\mathrm{C}\sim\mathrm{X}) = 0$, p_X would be undefined. According to Cheng, the new uninterruptible information obtained about X under the direct extinction procedure would be ignored, and the value of X established during acquisition would not change. In contrast, RW predicts that a conditioned inhibitor would extinguish under the direct extinction procedure. On an extinction trial, $\lambda = 0$ and the parenthetical term in Equation 6 would be positive for the inhibitory Cue X. Thus, X should become less inhibitory, and judgments should become more positive. Contrary to RW, most studies with animal subjects have not been able to extinguish a conditioned inhibitor with the direct procedure (see Miller, Barnet, & Grahame, 1995). Yarlas et al. (1995) reported a similar result for human observers: Judgments of X did not change when X was subjected to the direct extinction procedure.

Under the *indirect* extinction procedure, Cue C is presented alone, without the outcome. According to power PC, X will become less preventative, because ΔP in Equation 6, which had been negative, approaches 0 during indirect extinction. $P(\mathrm{O}|\mathrm{CX})$ remains 0, and $P(\mathrm{O}|\mathrm{C}\sim\mathrm{X})$ shifts from 1 to 0. According to RW, the inhibitory value of X will remain unchanged, since the strength of a cue does not change when the cue is not presented. Contrary to RW, most studies with animal subjects have shown a reduction in the inhibitory control by X after extinction of C (see Miller et al., 1995). Yarlas et al. (1995) reported a similar result for human observers: Judgments of X became more positive after C was subjected to the indirect extinction procedure.

Within the context of associative models, the indirect extinction procedure is an example of retrospective revaluation. In the indirect extinction procedure, C is presented in extinction, and then the observer is asked to revaluate Cue X, which had not been presented during extinction. The original RW model does not provide a means for the predictive strength of an absent cue to change. However, a simple modification to RW, proposed by Van Hamme and Wasserman (1994), provides for revaluation of a nonpresented cue and predicts that a conditioned inhibitor will be extinguished under the indirect extinction procedure. In this modified RW (MRW), the predictive strength of all cues are updated on all trials, even when a cue is not presented. Van Hamme and Wasserman (1994; Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996) provided empirical data that showed that judgments changed over trials and that these changes occurred even on trials in which the cue was not presented.

MRW encompasses other retrospective revaluation data as well, such as backward blocking and recovery from overshadowing (e.g., Chapman, 1991; Dickinson & Burke, 1996; Shanks, 1985b; Wasserman & Berglan, 1998), that raise problems for the original RW model. Since RW is a special case of the MRW, one might have expected Cheng (1997) to have compared power PC with the MRW model. In fact, MRW was not even mentioned by Cheng. Although Cheng did not reference Van Hamme and Wasserman (1994), Cheng et al. (1996) did, and they were highly critical of MRW.

In the original RW model, the predictive strength of a cue does not change on cue-absent trials. That is, α in Equation 6 is positive ($0 < \alpha \leq 1$) on cue-present trials and is zero on cue-absent trials. Van Hamme and Wasserman (1994) allowed α to be negative on cue-absent trials. That is, nonpresented cues have a negative salience. One criticism of MRW put forth by Cheng et al. (1996) is that by allowing α to be negative on cue-absent trials, RW loses its conceptual interpretation. In fact, Van Hamme and Wasserman provided a justification for the negative α values on cue-absent trials. They suggested that the associative strength of a cue could change on trials in which the cue was absent but *expected.* For example, in the first phase of a backward-blocking experiment, two cues (A and B) are presented, whereas in the second phase, only one cue (A) is presented. According to MRW, B is expected on A-only trials by virtue of the association formed between the representations of B and A during Phase 1. In MRW, the direction of change in associative strength of an expected but absent cue is opposite to the direction of the change in associative strength of a cue that is actually presented. A trial analysis of Van Hamme and Wasserman's data provided support for MRW, in that a nonpresented cue did change its associative value in the opposite direction from a presented cue. Later, Dickinson and Burke (1996) and Wasserman and Berglan (1998) provided data indicating that formation of an association between Cue A and Cue B leads to the expectation that B will occur when A is presented. They showed that the value of α on cue-absent trials depended on the strength of the between-cue associations. Contrary to the claim of Cheng et al., RW does not lose its conceptual interpretation when α is allowed to assume a negative value when the cue is absent but expected.

A second criticism put forth by Cheng et al. (1996) was that MRW, unlike the original model, is not interpretable as a connectionist model. Van Hamme and Wasserman (1994; Wasserman et al., 1996), in fact, highlighted that Markman (1989) and Tassoni (1995) incorporated a similar modification within their connectionist models.

Cheng et al. (1996) were also critical of MRW at an empirical level. They correctly noted that MRW does not address the failure of the prediction that a conditioned inhibitor should extinguish under the direct extinction procedure. This difficulty for RW (and MRW) is created by the assumption that excitation and inhibition are symmetrical opposites. The empirical lack of symmetry between conditioned excitation and conditioned inhibition has been widely recognized in the animal learning literature for years (see Miller et al., 1995). Indeed, Wagner and Rescorla (1972) expressed some doubts about their symmetry assumption when they applied the model to conditioned inhibition: "Whether or not a model of Pavlovian conditioning without a special inhibitory process is generally tenable is a larger issue on which we have not committed ourselves" (p. 308). Alternative ways of conceptualizing conditioned inhibition within an associative framework have been proposed that would encompass the empirical data showing that the direct extinction procedure does not result in the extinction of a conditioned inhibitor (e.g., Miller & Matzel, 1988; Rescorla, 1979; Wagner, 1981).

The main focus in Cheng et al. (1996) with regard to empirical failures of MRW was that the predictions were not in accord with Park and Cheng's overexpectation data. As we noted earlier, these data have never been published.

Cheng (1997) indicated that she would concentrate on RW, rather than on more recent associative models, because RW has been the most influential associative model for the past quarter century. As has been documented by Siegel and Allan (1996), the model indeed has been spectacularly successful. However, despite its successes, the model also has failed to account for some of the empirical data that have been reported during the past quarter century (see Miller et al., 1995). Extinction of conditioned inhibition is one such example. One of the successes of the model, which is emphasized by Miller et al., has been its influence on new model development. To ignore these new developments when assessing associative models, as Cheng (1997) did, is to ignore progress.

CONCLUSIONS

The data in the literature that Cheng (1997) relied on to make her case did not in fact provide support for power PC. In the intervening years, Cheng and her colleagues (Buehner, 2001; Buehner & Cheng, 1997; Wu & Cheng, 1999) have reported new data, and they have concluded that these data were supportive of power PC. In fact, the data in these papers deviated from the predictions of power PC, and the authors invoked ad hoc explanations in an attempt to explain the deviations of the data from the predictions of power PC. Other researchers, who have conducted experiments explicitly to evaluate power PC, have concluded that, overall, power PC did not provide an adequate account of their data (e.g., Lober & Shanks, 2000; Vallée-Tourangeau et al., 1997; Vallée-Tourangeau et al., 1998). Also power PC is unable to encompass trial effects (see Allan, 1993; Shanks, 1993) and trial-sequencing effects (e.g., Dennis & Ahn, 2001; López et al., 1998; Yates & Curley, 1986).

Cheng's (1997) claim that power PC provided a better account than did RW relied largely on unpublished data from her laboratory. In her criticisms of RW, she minimized one of the strengths of associative models—namely, their ability to describe not only final judgments, but also the trial-by-trial changes in those judgments. In contrast to power PC, RW predicts the frequently observed density bias for preasymptotic judgments.

Cheng (1997) concluded that "the results uniquely support the power PC theory" (p. 367). Our evaluation indicates that they do not. Although modifications to RW have provided accounts for originally troublesome data, such as backward blocking and the indirect extinction of a conditioned inhibitor, other challenges remain (see Miller et al., 1995). Our contention is simply that Cheng's criticisms of RW ignore important strengths of the model and are based on the original model, rather than on more recent modifications to the model.

REFERENCES

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, **15**, 147-149.

Allan, L. G. (1993). Human contingency judgments: Rule-based or associative? *Psychological Bulletin*, **114**, 435-448.

Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning & Motivation*, **14**, 381-405.

Buehner, M. J. (2001). Inducing causation: Covariation assessment and the assumption of causal power. In M. May & U. Oestermeier (Eds.), *Interdisciplinary perspectives on causation* (pp. 33-58). Norderstedt, Germany: Libri.

Buehner, M. J., & Cheng, P. W. (1997). Causal induction: The power PC theory versus the Rescorla–Wagner model. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 55-60). Hillsdale, NJ: Erlbaum.

Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 837-854.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, **18**, 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 & Social Psychology*, **58**, 545-567.

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

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

Fales, E., & Wasserman, E. A. (1992). Causal knowledge: What can psychology teach philosophers? *Journal of Mind & Behavior*, **13**, 1-27.

Kao, S. F., & 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.

Levin, I. P., Wasserman, E. A., & Kao, S. F. (1993). Multiple methods for examining biased information use in contingency judgments. *Organizational Behavior & Human Decision Processes*, **55**, 228-250.

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

López, F. J., Shanks, D. R., Almaraz, J., & Fernández, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 672-694.

Markman, A. B. (1989). LMS rules and the inverse base-rate effect: Comment on Gluck and Bower (1988). *Journal of Experimental Psychology: General*, **118**, 417-421.

Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla–Wagner model. *Psychological Bulletin*, **117**, 363-386.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51-92). San Diego: Academic Press.

Rescorla, R. A. (1979). Conditioned inhibition and extinction. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume to Jerzy Konorski* (pp. 83-110). Hillsdale, NJ: Erlbaum.

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 research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Shanks, D. R. (1985a). Continuous monitoring of human contingency judgment across trials. *Memory & Cognition*, **13**, 158-167.

Shanks, D. R. (1985b). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, **37B**, 1-21.

Shanks, D. R. (1987). Acquisition functions in contingency judgment. *Learning & Motivation*, **18**, 147-166.

Shanks, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, **84**, 319-354.

Shanks, D. R., Holyoak, K. J., & Medin, D. L. (1996). *The psychology of learning and motivation: Causal learning*. San Diego: Academic Press.

Siegel, S., & Allan, L. G. (1996). The widespread influence of the Rescorla–Wagner model. *Psychonomic Bulletin & Review*, **3**, 314-321.

Tassoni, C. J. (1995). The least mean squares network with information coding: A model of cue learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 193-204.

Vallée-Tourangeau, F., Murphy, R. A., & Drew, S. (1997). Causal judgments that violate the predictions of the power PC theory of causal induction. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 775-780). Hillsdale, NJ: Erlbaum.

Vallée-Tourangeau, F., Murphy, R. A., Drew, S., & Baker, A. G. (1998). Judging the importance of constant and variable candidate causes: A test of the power PC theory. *Quarterly Journal of Experimental Psychology*, **51A**, 65-84.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning & Motivation*, **25**, 127-151.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application to a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning* (pp. 301-336). London: Academic Press.

Wasserman, E. A., & Berglan, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgement: The role of within-compound associations. *Quarterly Journal of Experimental Psychology*, **51B**, 121-138.

Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response–outcome contingencies under free-operant procedures. *Learning & Motivation*, **14**, 406-432.

Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990). The contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 509-521.

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, & Cognition*, **19**, 174-188.

Wasserman, E. A., Kao, S., Van Hamme, L. J., Katagiri, M., & Young, M. E. (1996). Causation and association. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Causal learning* (Vol. 34, pp. 207-263). San Diego: Academic Press.

Williams, D. A. (1995). Forms of inhibition in animal and human learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **21**, 129-142.

Wu, M., & Cheng, P. W. (1999). Why causation need not follow from statistical association: Boundary conditions for the evaluation of generative and preventive causal powers. *Psychological Science*, **10**, 92-97.

Yarlas, A. S., Cheng, P. W., & Holyoak, K. J. (1995). Alternative approaches to causal induction: The probabilistic contrast versus the Rescorla–Wagner model. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 431-436). Hillsdale, NJ: Erlbaum.

Yates, J. F., & Curley, S. P. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, **62**, 293-302.

NOTES

1. A fuller discussion of the philosophical lineage of the covariation (Hume) and causal (Kant) traditions is available in a number of sources, such as Fales and Wasserman (1992), as well as in Cheng (1997).

2. Cheng (1997) also relied heavily on data from her laboratory reported in manuscripts that had been submitted for review prior to 1997 (Fratianne & Cheng, 1995; Park & Cheng, 1995). These manuscripts were not accepted, the data have yet to be published, and the original manuscripts are not available (Cheng, personal communication, April 2001).

3. Cheng (1997) provides the mathematical derivations of these equations.

4. In the analysis of their data, Wu and Cheng (1999) conducted multiple nonorthogonal χ^2 tests and did not correct the experiment-wise α. A more acceptable analysis would be an overall χ^2, followed by partitioning of the overall matrix.

5. Some of the data reported by Buehner and Cheng (1997) also are presented in Buehner (2001).

6. The assumption that β_O (outcome present) = $\beta_{\sim O}$ (outcome absent) is often made in the human judgment task.

7. In the literature, these terms have both a methodological and a theoretical connotation. At the methodological level, they describe experimental procedures. At the theoretical level, they derive from associative models. In this section, the terms are used methodologically, as they were in Cheng (1997).

(Manuscript received July 4, 2002;
revision accepted for publication September 17, 2002.)