Rational Assessments of Covariation and Causality

Gustaf Gredebäck (Gustaf.Gredeback@psyk.uu.se)
Department of Psychology, Uppsala University
Box 1225, SE-751 42, Uppsala, Sweden

Anders Winman (Anders.Winman@psyk.uu.se)
Department of Psychology, Uppsala University
Box 1225, SE-751 42, Uppsala, Sweden

Peter Juslin (Peter.Juslin@psy.umu.se)
Department of Psychology, Umeå University
SE-901 87, Umeå, Sweden

Abstract

Are human contingency judgments based on associationistic principles such as cue competition or on normative principles as specified by rational-cognitive models? In this study, participants learned to predict an outcome from several simultaneously presented cues. They were asked to judge the cues in regard to causal power or statistical concepts such as probability or relative frequency. Uniform application of associationistic principles implies cue-interaction effects of blocking (Experiment 1) and conditioned inhibition (Experiment 2) for all judgments. A rational-cognitive framework predicts cue-interaction effects for causality judgments, but not for probability and relative frequency judgments. The results support the rational-cognitive framework on all accounts.

Introduction

The ability to detect causal relations in the environment is of utter importance to all organisms. Fortunately, at first glance, at least, we seem to adjust well to such demands. We readily formulate hypotheses about plausible causal relationships and contingencies. Research deriving from the associationist tradition, however, suggests that this optimistic view is unwarranted. It is proposed that, because of cue-interaction effects, our representations are distorted. This study pits a more optimistic view of human contingency judgment based on the metaphor of the mind as an "intuitive scientist" against this associationist view. It does so by comparing causality judgments with probability and relative frequency judgments in an inductive contingency judgment task.

Imagine that you suffer from an allergic reaction, which you believe originates from eating shellfish. It seems reasonable to assume that this hypothesis of causality originates in your recollection of similar events, in this case your memory of eating shellfish and suffering from allergic reactions. Of course, most meals that you have eaten contained neither shellfish nor resulted in allergic reactions. Still your memory tells you that the allergic reaction on numerous occasions co-occurred with dishes that included shellfish. You also recall dinners which included shellfish but which did not lead to allergic reaction, and times when the allergy sprung up in the absence of shellfish. One popular notion is that memories of previous events are categorized in what resembles a 2×2 matrix. In this *contingency matrix* the presence and absence of a predictor event and an outcome event constitute the two axes.

Table 1: A contingency matrix.

	Outcome	No outcome
Predictor present	*A*	*B*
Predictor absent	*C*	*D*

According to a normative model, judgments of covariation are based on conditional contingency: that is, on the probability of the outcome (e.g., allergy) in the presence and absence of the predictor (e.g., shellfish). Formally, this can be expressed by the Δp *algorithm*:

$$\Delta p = \left(\frac{a}{a+b}\right) - \left(\frac{c}{c+d}\right), \qquad (1)$$

where a, b, c, and d are the number of times that events A, B, C, and D in the contingency matrix have occurred. A positive contingency is perceived if Cell A and D contain more occurrences than Cell B and C ($\Delta p > 0$). Similarly, a negative contingency is indicated by a negative Δp and a zero Δp indicates no contingency. Perceptions of causal relations are thought to be based on observed covariation registered in a form equivalent to the contingency matrix, and computed by the Δp algorithm. Chapman and Robbins (1990) examined two effects that in a conspicuous way violate the normative Δp model: the cue interaction effects referred to as *blocking* and *conditioned inhibition*.

Cue Interaction Effects

The allergic reaction used in the example exemplifies a simple causal relation where one potential cause is evaluated with respect to its assumed effect ($C \rightarrow E$). However, such occurrences of unambiguous sole cues seldom arise beyond the realms of clinical test scenarios. Instead, we often evaluate complex situations where multiple potential causes may produce an effect (C_1, $C_2 \ldots C_n \rightarrow E$). In Experiment 1 of Chapman and Robbins (1990) participants examined the relationship between the change in prize of four individual stocks (predictors) and the stock market as a whole (outcome). In the first phase, either stock A increased in value, followed by an increase in the value of the market, or stock C increased without a market increase. In the second phase, either of pairs of stocks AB or

CD increased, in both cases followed by a market increase. On one third of the trials in Phases 1 and 2 no stocks increased and the market remained unchanged. After each phase, participants rated the extent to which an increase in each stock predicts a market increase on a scale from -100 to 100. (See Table 2.)

Table 2: The design in Chapman and Robbins (1990).

Phase 1	Phase 2	Test 2
A+, *C*-, ∅-	*AB*+, *CD*+, ∅-	*B*<*D*

Note. (+) = outcome, (-) = no outcome, ∅ = predictor absent.

If the normative Δp model is correct the frequency of occurrences between each predictor and the outcome is mapped within separate contingency matrices. This means that each predictor is evaluated in isolation with contingency judgments based on the recollection of frequencies. Because the Stocks *B* and *D* appear an identical number of times, always in the presence of the outcome, these two stocks should receive identical ratings of predictability in Test 2.

As it turns out, the ratings for individual predictors *interact*. In this case the interaction is *blocking*: of two cues with identical contingencies with the outcome, systematically lower ratings are given to the cue presented with a previously established predictor than to the cue presented with a non-predictor. In a second experiment, Chapman and Robbins examined *conditioned inhibition*. In this design, a predictive cue is followed by the outcome, except when it occurs with a second cue. The second cue, referred to as the inhibitor, is rated lower in predictability than a control cue that has the same objective contingency with the outcome. Cue-interaction such as blocking and conditioned inhibition casts doubt on the Δp model as a descriptive model of human judgment of causation and covariation. The question is whether these results are best explained by an *associationist* or, what we refer to as, a *rational-cognitive* framework.

Associationist Framework

Cue-interaction is routinely observed in studies of animal learning. This has been taken to indicate that human behavior in contingency judgment tasks is best described by a "grand theory" of learning based on the principles that apply to animal conditioning. Associationistic models therefore adhere to a strong analogy between conditioned (*CS*) and unconditioned (*UCS*) stimuli in classical conditioning and predictors and outcomes in human contingency judgments. Contingency judgments are seen as being based on the associative strength of the relationship between conditioned stimuli (predictors) and the unconditioned stimuli (outcomes). In a multiple-cue task, each *CS-UCS* association is based on the informative strength of CS_x with respect to the *US*, in competition with all the *CSs* present.

The model most often called upon to explain cue interaction effects is the *Rescorla-Wagner model* (Rescorla & Wagner, 1972), hereafter referred to as the R-W model.

Formally, the model states that:

$$\Delta V_x^{n+1} = \alpha_x \beta_1 (\lambda_1 - V_{total}^n), \qquad (2)$$

where ΔV_x^{n+1} is the change in associative strength (V) of CS_x as the result of pairing it with UCS_1 on trial n+1, α_x is a learning rate parameter representing the associability of CS_1, and β_1 is the corresponding parameter for UCS_1, λ_1 is the maximum associative strength that the *UCS* can support, called the asymptote ($\lambda_1 = 1$ in the presence of the UCS_1, 0 in its absence), and V_{total}^n is the total associative strength of all *CSs* on trial n+1.

Equation 2 describes the change in associative strength of a *CS* as a function of the current associability of the *UCS* and *CS*, in relation to the remaining associability of the *UCS* ($\lambda_1 - V_{total}^n$). A consequence of Equation 2 is that learning will occur only when an outcome is unexpected or surprising in the light of one's expectations.

According to the R-W model, the blocking described by Chapman and Robbins (1990) is due to cue competition. Once a CS_1-*US* association is established, any new CS_2 that is presented with the previous association ($CS_1, CS_2 \rightarrow$ *US*) will not become associated with the outcome. Equation 2 states that the change in associative strength on trial n+1 is defined by the difference between the asymptote (λ) and the total associative strength (V_{total}^n) on trial n. Since CS_1 already predicts the outcome there is no room for CS_2 to become associated with the *US*. In comparing CS_2 with another stimulus CS_3 that has an identical outcome contingency the R-W model thus predicts that CS_2 will be valued as less associative than CS_3.

Conditioned inhibition is seen as the opposite of excitation. According to the R-W model this phenomenon is expressed by ($\lambda_1 - V_{total}^n < 0$). Since the asymptote itself can never be negative, the expression will only be true when V_{total}^n is larger than zero, that is, when some excitation already has occurred. Say that stimulus CS_1 leads to an outcome *E*, while stimulus compound CS_1, CS_2 leads to absence of the outcome. If later tested individually, CS_1 receives an exhibitory value, while CS_2 receives an equally strong inhibitory (negative) value, making the total associative value equal to zero. A number of alternative models have surfaced within the associationistic tradition (e.g., Gluck & Bower, 1988; Pearce, 1994; Van Hamme & Wassermann, 1994). Siegel and Allan (1996) singled out the R-W model as the most successful model and, for the purpose of this paper, the R-W model will represent the associationist framework.

Rational-Cognitive Framework

A number of theoreticians (Waldmann & Holyoak, 1992; Cheng, 1997) propose that the observation of cue interaction effects does not pose a threat to models based on the Δp algorithm. These results merely indicate the inapplicability of the contingency matrix model to situations involving multiple causes. Imagine that someone claims that alcohol consumption causes lung cancer. In support for this claim it is noted that consumers of alcohol more often suffer from lung cancer than others. With the knowledge that alcohol consumption often is accompanied by smoking this line of reasoning may strike you as odd. Instead it

seems reasonable to assume that smoking is the real cause for the increased risk of cancer.

The difference is that your antagonist is considering a simple causal relationship (unconditional contingency) while you apply a more complex analysis of causal relationships, taking into account the *conditional contingency*. This means that the causal relationship is viewed in the context of the presence and absence of alternative causes. In the smoking example, it is nearly impossible to incorporate all of the potential alternative causes. However, the idea is that, like "intuitive scientists" (Kelley, 1967), we have the capacity to control, at least, for "likely" alternative causes (see also Spellman, 1996).

According to a cognitive-rational approach people store information about events in frequency format. This information is available for different forms of analyses by application of cognitive algorithms. Complex causal relationships require a more sophisticated algorithm than the simple Δp rule (Cheng & Novick, 1990; Cheng, 1997). In *Power PC theory* (Cheng, 1997), for example, the strength of a causal factor is estimated by the conditional contingency: that is, the contingency when other potential causes are controlled for[1]. According to this view, cue interaction is viewed as a consequence of the *participants' attempts to control for alternative causes*. In Chapman and Robbins (1990), Experiment 1, for example, the participants may have arrived at the conclusion that the causal relationship between Stock *B* and the outcome event (*E*) is uncertain, since its effect is nullified once control for Stock *A* is performed by a conditional contrast; ($p(E|AB) - p(E|A) = 0$). Applying the same algorithm to Stock *D* indicates that *D* is a strong causal factor ($p(E|CD) - p(E|C) = 1$). The controlling for alternative causes would thus lead to results that coincide with the blocking effect predicted by the R-W model.

The same line of reasoning is applicable to conditioned inhibition. Participants conclude that a stimulus has negative causal power (i. e. power to prevent an effect) due to its conditional contingency with respect to other causes (for a discussion, see Cheng, 1997). Because people are assumed to act like scientists in applying rational arguments and interpreting patterns of covariation in terms of unobservable causes, we refer to this as the rational-cognitive framework.

The Effect of Judgment Type

If the assumption of veridical representation of frequencies in models like power PC-theory is correct, we should not expect interaction effects for judgments of probability or relative frequency (at least, to the extent that that probability ratings are based on representations of relative frequencies). This is a crucial difference between the R-W model and power-PC theory. An orthodox interpretation of the R-W model, which presupposes that the same processes underlie judgments of causality and covariation, predicts that there should be no difference between conceptually distinct ways of probing for the relationship between cue and outcome. The judgments are mapped from associative strengths and we expect interaction effects for judgments of causality, probability, and frequency alike. In contrast, the rational-cognitive framework presupposes correct representations of environmental frequencies, which deviate from causality ratings in predictable ways. In short: With power PC-theory there is a distinction between judgments of causality and covariation, with the R-W model there is not.

Cue interaction has occasionally been reported also with judgments of probability or frequency (Chapman, 1991, Price & Yates, 1995). Nevertheless, the main body of empirical findings on blocking and conditioned inhibition rests on assessments of often vaguely defined judgment scales (e.g., predictability). We know of no study of conditioned inhibition with relative frequency or probability judgments. In this study, we thus compared judgments of a) the *causal power* of the predictor on the outcome, b) the *probability* of the outcome given the event, and c) the *relative frequency* of the outcome given that the predictor was present on the trial, in a between-subjects design. The associationist account suggests cue-interaction effects with all three judgments. The rational-cognitive approach implies cue interaction effects for causality judgments, but an absence of this effect with the other two judgments.

Experiment 1: The Blocking Effect

Experiment 1 of Chapman and Robbins (1990) examined the blocking effect in a multiple cue task involving stock market predictions (Table 2). Participants were asked to rate the "predictability" of each stock on a scale from -100 to 100. The present experiment applies the design from Chapman and Robbins' Experiment 1, with the addition of three new groups. In addition to predictability (to replicate their results), the participants judged either explicit causality, explicit probability, or relative frequency.

Method

Participants. Participants were 64 undergraduate students from Uppsala University. They received either course credit or a movie ticket in exchange for their participation.

Materials and procedure The experiment was divided in two learning-phases (L_1, L_2) and two test-phases (T_1, T_2)[2], appearing in the order L_1, T_1, L_2, T_2. In each learning-phase participants were to assess whether the stock market as a whole would change in value based on the individual movement of four fictional stocks (see Table 2). L_1 contained 36 trials (12×*A*+, 12×*C*-, and 12×∅-). L_2 contained 76 trials (24×*AB*+, 24× *CD*+, 24×∅-). In each test phase, participants assessed the relationship between the increase of each separate stock and an increased stock market. One group rated the stocks according to their *predictability* on

[1] It is important to note, however, that covariation is merely one component of the process of causal induction in power PC theory. Another important component is an a priori framework for interpreting input in terms of causal mechanisms (Cheng, 1997).

[2] Only the most central results, that is, those of the second test phase are included in the present paper.

a scale from -100 to 100, in accordance with Chapman and Robbins (1990). The other three groups rated either relative frequency (*In what percentage of the occasions on which the stock X increased, did the outcome occur?*), probability (*Given that stock X increases, what is the probability of the outcome?*), or causality (*To what degree does stock X cause the outcome?*) with respect to an increased market on a scale from 0 to 100 (a bi-directional scale does not apply to frequency and probability). The judgments were varied between groups.

Results

The results from Experiment 1 are presented in Table 3. A blocking index was calculated by subtracting ratings to predictor *B* from ratings given to predictor *D*, where a positive score indicates a blocking effect. As illustrated in Figure 1A, both the mean predictability blocking index 27.2 and the mean causality blocking index 24.3 show significant blocking. In contrast, the mean probability blocking index 5.7 and the mean frequency blocking index -15.2 show no sign of the blocking effect (the latter is even negative).

Table 3: The average rating of stimuli *A* through *D* during test Phase II. 95% confidence intervals within parentheses.

	Predictor			
	A	*B*	*C*	*D*
Predictability	53.1 (21.4-84.8)	35.9 (10.2-61.7)	-10.6 (-35.3-14.0)	63.1 (35.9-90.2)
Causality	89.3 (78.133-100.4)	49.3 (34.7-63.8)	27.1 (10.1-44.1)	73.6 (55.4-91.8)
Probability	69.7 (55.1-84.2)	52.3 (43.4-61.3)	36.7 (19.5-53.8)	58.0 (43.2-72.8)
Frequency	87.5 (70.8-104.2)	98.4 (95.4-101.5)	71.6 (52.0-91.1)	83.2 (67.0-99.5)

We replicate the findings by Chapman and Robbins (1990) with blocking effects for predictability judgments, and the effect is even more consistent with the causality scale. These results support the idea that the judgment labeled predictability is interpreted as a mix of causality and covariation. Within the same settings we fail to observe cue interaction effects for probability and relative frequency judgments. This predicted pattern is significant as concluded from a planned comparison of means analysis, F (1, 57) = 9.8, $p<.01$[3].

The results thus support theories such as power PC theory with respect to its analysis of causal reasoning, as well as the more general notion in rational-cognitive models that (roughly) veridical representations of event frequencies are preserved. At the same time the results are in opposition to the R-W model. To extend and validate the result from the first experiment, a second experiment was conducted in order to examine conditioned inhibition.

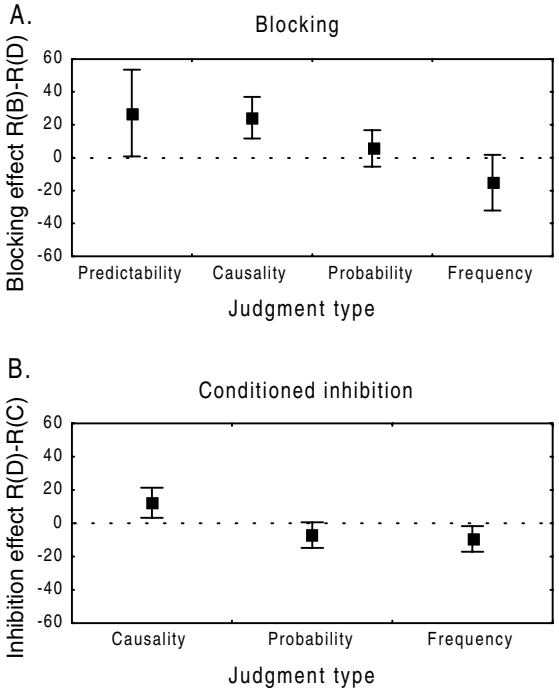

Figure 1: Panel A: Blocking effect in Experiment 1 as a function of judgment. Panel B: Conditioned inhibition in Experiment 2 as a function of judgment.

Study 2: Conditioned Inhibition

Experiment 2 of Chapman and Robbins (1990) was the first study to examine the conditioned inhibition effect in humans. The results showed clear signs of conditioned inhibition. Williams, Sagness, and McPhee (1994) report several failures to replicate Chapman and Robbins' (1990) cue interaction effects. They reasoned that these failures might be due to the way participants approach the task. People can either interpret a stimulus compound (*A*, *B*) as one distinct stimulus (*AB*) or as the combination of the two stimuli (*A* and *B*). The former is termed a *configural encoding* while the later is referred to as an *elemental encoding*. In a multiple cue task conditioned inhibition can not be obtained with configural encoding: the effect demands that the participants view each stimuli in isolation. Williams et al. therefore tried to experimentally encourage participants to engage in elemental encoding.

After a couple of fruitless attempts to replicate the effect we therefore abandoned the design of Chapman and Robbins (1990). Inspired by Williams et al.'s (1994) Experiment 2, we attempted to promote elemental strategies in favor of configural strategies. In order to take every measure to obtain an effect we made some additional changes. In the original task the outcome always occurs with the positive predictor alone, but never with this predictor in conjunction with the inhibitor. The deterministic design may promote the learning of explicit rules, which

[3] Because the predictability group used a bi-directional scale whereas the other groups used unidirectional scales, the scores for all groups were standardized within each condition before they were entered into the planned comparison.

might diminish a true effect. Furthermore, floor effects might mask a real effect since both the inhibitor and the control cue can be expected to be rated at the lower end of the scale.

We therefore made the task probabilistic. The outcome occurred with a probability of .95 in the presence of the positive predictors alone, and with a probability of .3 in the presence of the positive predictor in conjunction with the inhibitor. This modification deals with both of the unfortunate characteristics of the Chapman and Robbins (1990) design. In addition, separate single presentations of the negative cue which have been found to increase an effect (Williams, 1995) were added. Finally, the content of the task was changed. In the original task, the content consists of stocks that change. It could be argued that this content does not encourage inhibition since it may be hard to create a mental model of a causal mechanism of how a particular stock hinders the outcome to occur. We used a task of evaluating experimental fertilizers (cf. Spellman, 1994) where it is easier to form a model of how a particular substance may hinder growth[4]. To summarize: Study 2 was designed to investigate whether conditioned inhibition effects will occur also for probability and relative frequency judgments or if a dissociation will be observed between these and judgments of causality.

Method

Participants Seventy-five undergraduates from Uppsala University took part in the study. They received a movie ticket or course credit in exchange for their participation

Materials and procedure The experiment was divided in two tasks. The first (pretraining) involved one learning phase and one test phase, the second (main experiment) was divided in two learning-phases (L_1, L_2) and two test-phases (T_1, T_2), appearing in the order L_1, T_1, L_2, T_2. Both the pretest and the experiment involved the prediction of whether a plant would produce flowers or not after an observation of which fertilizers that had been added to an irrigation fluid. The learning phase of the pretraining was identical to the *Explicit Condition* of Williams et al. (1994). It consisted of 12×*X*+, 12×*XY*+, 12×*Y*- and 12×*Z*- trials. When finished, participants were asked to rate each fertilizer with respect to the outcome according to either causality, probability, or relative frequency[5] on a scale from 0-100. The purpose of the pretreatment was to encourage an elemental encoding and results of this phase were not investigated further. In Phase L_1 of the main experiment either of fertilizers *A* through *E* were added to the liquid 20 times each. In the case of fertilizers *A* and *B* the plant produced flowers in 19/20 (95%) of the occasions. Fertilizers *C*, *D*, and *E* were coupled with the outcome on 6/20 (30%) occurrences. In L_2 three constellations of fertilizers were followed by the outcome with a base rate of 19/20. These were *A*, *B*, and *AB*[6]. Fertilizer *E*, as well as fertilizer combinations *AC* and *DE* was coupled to the outcome with a base rate of 6/20. Thus, in this design, predictor *C* is the inhibitor and predictor *D* is the control, with exactly the same contingency with the outcome and number of occurrences. Table 4 describes the conditioned inhibition design in the experiment. After the learning phases, participants rated the relationship between each fertilizer and the outcome based on causality, probability, or frequency with the same scale as described in Study 1. Throughout Study 2, each participant only made judgments for one of the three scales.

Table 4: Conditioned inhibition design in Experiment 2.

Phase 1	Phase 2	Test 2
A+ B+ C- D- E-	*A+ B+ E- AB+ AC- DE-*	*C<D?*

Note. (+) indicates a probability of outcome of .95, (-) indicates a probability of outcome of .3.

Results

Table 5 presents the results of Experiment 2. To repeat, conditioned inhibition is observed if stimulus *C* is rated lower than stimulus *D*. A significant conditioned inhibition effect was found in the causality group (average difference between *D* and *C* = 12.4: See Figure 1B). In neither of the other groups is there conditioned inhibition. Both groups have higher ratings for the *C* than the *D* predictor (difference between *D* and *C* = -9.5 in the frequency group and -7.2 in the probability group.). In fact, this reversed difference is significant in the frequency group. A planned comparison shows that the predicted difference between the causality group and the probability and frequency groups is significant (F (1, 72) = 13.7, p <. 001). The reason for the significant difference in the opposite direction in the frequency group is unclear, but interestingly the trend was the same in Experiment 1. An explanation, (undeniably speculative) could be that higher level deductive reasoning influence frequency ratings; Maybe participants reason that since predictor *A* occurred often together with the outcome and predictor *C* often occurred in conjunction with predictor *A*, then predictor *C* probably occurred quite often together with the outcome. Note, however, that the observed significance in no way indicates that the frequency ratings are severely distorted. A comparison between true frequencies vs. rated probabilities and frequencies show that these agree approximately (although the ratings are moderately regressive). In neither case are the true frequencies excluded by the confidence intervals for the ratings, making it impossible to

[4] A change of content should not affect the result according to associative models, which imply independence of content. In cognitive-rational models on the other hand content may play an important role since prior causal models potentially can affect interpretation according to these.

[5] Due to the similar results for predictability and causality ratings in Study 1 and the deviating scale for the predictability ratings, the latter were dropped in Study 2.

[6] The conjunction of the positive predictors *AB* was included in order to eliminate the possibility that participants learned a rule implying that a conjunction of any two predictors was followed by a decreased probability of the outcome.

reject the hypothesis that these are made on basis of undistorted representations of the true frequencies. These results are in line with predictions of a rational-cognitive model, with conditioned inhibition effects only for judgments of causality.

Table 5: The average rating of stimuli *A* through *E* during test Phase II (95% confidence intervals within parentheses).

	Predictor				
	A	B	C	D	E
Causality	83.9 (78.1-89.7)	88.2 (82.0-94.5)	23.80 (16.3-31.3)	36.2 (28.8-43.5)	36.0 (27.8-44.2)
Probability	82.0 (73.9-90.1)	90.1 (85.0-95.1)	32.8 (24.3-41.2)	25.6 (17.9-33.2)	33.3 (23.8-42.8)
Frequency	79.8 (71.8-87.7)	87.4 (81.6-93.2)	33.2 (24.9-41.5)	23.7 (16.3-31.2)	29.9 (21.4-38.4)

Discussion

In this paper, we have contrasted two different frameworks for the processes that underlie human contingency judgment. An associationist account which stresses the similarity to the processes derived from learning in animals, as epitomized in the R-W model, and one rational-cognitive account that relies on the metaphor of the mind as an intuitive scientist.

The rational-cognitive account implies that the participants can appreciate a distinction between judgments that concern the causal power of a factor, and judgments that pertain to covariation, such as probability and relative frequency. On this view, blocking and conditioned inhibition arise from appropriate considerations of the confounding between multiple potential causes. This reasoning is compatible with—and indeed presupposes—availability of accurate information about frequencies. An orthodox interpretation of the R-W model, presuming the same process behind judgments of causality and covariation, suggests no effect of the judgment type manipulation. On any account, the model does not provides an explanation for the observed effect. The results from two separate experiments, with fairly disparate designs covering the two most well known cue interaction effects clearly favor the rational-cognitive account. The participants seem to appreciate the distinction between a judgment of causality and judgments of probability and relative frequency.

These results suggest that, functionally the same behavior may be implemented by different mechanisms in different organisms. The same behavior that is computed by associationist processes in lower animals may be the results of high-level reasoning in humans. This conclusion may come as no surprise: Regardless of our ontogenetic sophistication we all share the challenge of dealing with a complex and uncertain environment, and the evolutionary and adaptive pressures we face may thus be very similar in the end.

References

Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, and 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 and Social Psychology*, *58*, 545-567.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General, 117*, 225-244.

Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska Symposium on motivation* (Vol. 15, pp. 192-238). Lincoln: University of Nebraska Press.

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, *101*, 587-607.

Price, P. C., & Yates J. F. (1995). Associative and rule-based accounts of cue interaction in contingency judgment, *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*, 1639-1655.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.

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

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

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

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*, 222-236.

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

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 694-709.