

Judgmental overshadowing: Further evidence of cue interaction in contingency judgment

PAUL C. PRICE and J. FRANK YATES
University of Michigan, Ann Arbor, Michigan

We investigated a phenomenon called *judgmental overshadowing*. Subjects predicted whether each of several patients had a disease on the basis of whether or not the patient had each of two symptoms. For all the subjects, the presence of the disease was moderately contingent on the presence of one of the symptoms (S1). In Condition 1 of our first experiment, the presence of the disease was highly contingent on the presence of the other symptom (S2). In Condition 2, the presence of the disease was independent of S2. Judgmental overshadowing occurred in that the S1-disease contingency was judged to be stronger in Condition 2 than in Condition 1. Subsequent experiments showed that judgmental overshadowing depends little on the form of the judgment, is not due to a response bias or contrast effect, and does not depend on subjects' actively diagnosing each patient. These results are consistent with, and are generally predicted by, an associative-learning model of contingency judgment.

People often judge the degree to which the occurrence of one event depends on, or is contingent on, the occurrence of another event. For example, through experience, a doctor may conclude that the presence of Disease D is contingent on the presence of Symptom S. Furthermore, the doctor may judge the *degree* to which D is contingent on S, which in turn may determine the diagnosis and treatment of a patient exhibiting S. Such *contingency judgments* have been a topic of interest to human judgment researchers for a number of years (e.g., Arkes & Rothbart, 1985; Beyth-Marom, 1982; Busemeyer, 1991; Chapman, 1991; Lipe, 1990; Smedslund, 1963; Ward & Jenkins, 1965).

In the typical contingency judgment experiment, subjects are presented with several successive trials in which they are informed of the occurrence or nonoccurrence of each of two events. The events could be a symptom and a disease (Smedslund, 1963), the seeding of rain clouds with dry ice and subsequent rainfall (Ward & Jenkins, 1965), or a plant's being a certain color and native to a certain region (Yates & Curley, 1986). The subjects are then asked to judge the degree to which the occurrence of one of the events is contingent on the occurrence of the other. These judgments typically take the form of ratings. For example, Ward and Jenkins had subjects rate, on a 100-point scale, the degree of "control" that cloud

seeding exerted over subsequent rainfall. In some experiments, however, the judgments have been either less constrained (e.g., "Do you think there is a relationship between Symptom A and Diagnosis F? Tell me about it!"; Smedslund, 1963), or more constrained ("Suppose a rhododendron plant is from Region A. What is the chance that it is light in color rather than dark?"; Yates & Curley, 1986).

The traditional psychological account of contingency judgment is that it is based on information analogous to that contained in a simple 2×2 contingency table, but is stored in the memory of the judge (e.g., Beyth-Marom, 1982; Busemeyer, 1991; Shaklee & Mims, 1982). Table 1, for example, is a contingency table containing information about the occurrence and nonoccurrence of each of two events, Event 1 and Event 2. In Cell A of the table is the number of occasions in which both events occurred; in Cell B is the number of occasions in which Event 1 occurred but Event 2 did not; in Cell C is the number of occasions in which Event 1 did not occur but Event 2 did; and in Cell D is the number of occasions in which neither event occurred. Judgment of the degree to which the occurrence of Event 2 is contingent on the occurrence of Event 1, then, is based on some arithmetic combination of this frequency information. A sophisticated judge, for instance, might base his or her judgment on the ΔP statistic: $\Delta P = a/(a+c) - b/(b+d)$, where a , b , c , and d are the numbers of cases in Cells A, B, C, and D of the contingency table, respectively (Allan, 1980).

Cue-Interaction Effects

One problem for this traditional account of contingency judgment is the existence of cue-interaction effects, which arise when people judge the degree of contingency between a single *outcome event* and each of multiple *predictor events*. In such situations, it has often been found that judgment of the degree of contingency between the

Thanks to Jerry Busemeyer, Gregory V. Jones, Robert Proctor, and an anonymous reviewer for their helpful comments on earlier drafts of this manuscript. We would also like to thank the members of the University of Michigan's Judgment and Decision Laboratory—and to remember fondly Carolyn Young—for their support throughout the development of this research. This material is based on work supported under a National Science Foundation Graduate Fellowship to the first author. Address correspondence to P. C. Price at the Department of Psychology, Judgment and Decision Laboratory, University of Michigan, 330 Packard Road, Ann Arbor, MI 48104-2994.

Table 1
A 2×2 Contingency Table

Event 1	Event 2	
	Occurs	Does not occur
Occurs	A $a=20$	B $b=15$
Does not occur	C $c=5$	D $d=10$

outcome event and one predictor event depends, in systematic ways, on the actual degree of contingency between the outcome event and each of the other predictor events.

Chapman (1991; Chapman & Robbins, 1990) has provided good examples of the cue-interaction effect known as *blocking* (see also Shanks & Dickinson, 1987; Waldmann & Holyoak, 1992). In one set of experiments (Chapman & Robbins, 1990), the outcome event was a rise in a fictitious stock market and the predictor events were rises in the stock prices of each of four fictitious corporations. During the first 36 trials of the experiment (Phase 1), a rise in the price of Stock P was always accompanied by a rise in the market, and a rise in the price of Stock N was never accompanied by a rise in the market. There were also trials in which neither stock rose and the market did not rise. During the next 72 trials (Phase 2), two types of compound trials were presented. There were trials in which a rise in the price of Stock P was accompanied by a rise in the price of a new stock, B, and this compound was followed by a rise in the market. There were also trials in which a rise in the price of Stock N was accompanied by a rise in the price of another new stock, C, and this compound was followed by a rise in the market. Additionally, there were trials in which none of the stocks rose and the market did not rise. Thus, although a rise in the market was equally contingent on Stocks B and C, it was judged to be less contingent on a rise in the price of Stock B than on a rise in the price of Stock C. The traditional account of contingency judgment cannot explain these results because it has no mechanism whereby one contingency can affect the perception and judgment of another.

Associative-Learning Account of Contingency Judgment

In the past few years, a number of researchers have proposed that cue-interaction effects can be explained by abandoning the traditional account of contingency judgment in favor of an associative-learning account (Baker, Mercier, Vallée-Tourangeau, Frank, & Pan, 1993; Chapman, 1991; Chapman & Robbins, 1990; Gluck & Bower, 1988; Shanks & Dickinson, 1987). According to the associative-learning explanation of contingency judgment, pairs of events become psychologically associated in essentially the same manner as when a nonhuman animal learns to associate a conditioned stimulus with an unconditioned stimulus or an operant response with an outcome. Judgment of the degree to which one event is

contingent on another, then, is based directly on the strength of the association that has been formed between the two events.

One version of the associative-learning account of contingency judgment is based on Rescorla and Wagner's (1972) model of classical conditioning (Chapman, 1991; Chapman & Robbins, 1990), which is formally equivalent to the simple adaptive-network model of Gluck and Bower (1988). One strength of this model is that it applies not only to the simple two-event scenarios traditionally studied by judgment researchers, but also to situations in which a single outcome event is contingent on the occurrence of each of multiple predictor events, as in the case of blocking described above. Specifically, the model indicates that the degree of association between predictor event i and the outcome event (the associative strength of predictor event i) changes from trial to trial according to the following rule:

$$\Delta V_i = \alpha\beta(\lambda_j - \Sigma V_k).$$

Here, V_i is the associative strength of predictor event i , and ΔV_i is the change in the associative strength on the current trial. α and β are parameters, taking on values between 0 and 1, which determine the rate at which the associative strength changes. The value of α depends on the salience of predictor event i and the value of β depends on the salience of the outcome event. λ_j represents the status of the outcome event on the current trial; it takes a value of 1 if the outcome event occurs, and a value of 0 if the outcome event does not occur. The quantity ΣV_k is the sum of the associative strengths of each of the k predictor events that occurred on the current trial, which can also be thought of as the degree to which the outcome event is expected on the basis of the occurrence of the k predictor events. Thus, the trial-by-trial change in the associative strength of predictor event i is proportional to the difference between the actual status of the outcome event and the expected status of the outcome event. When the outcome event occurs but was unexpected, all the present predictor events gain some associative strength. When the outcome event does not occur but was expected, all the present predictor events lose some associative strength. When the actual outcome corresponds perfectly with the expected outcome, no learning takes place.

This version of the associative-learning account of contingency judgment explains the results of Chapman and Robbins's (1990) blocking experiment as follows. In Phase 1 of that experiment, subjects learned to expect a rise in the market on the basis of a rise in the price of Stock P alone. In Phase 2, a rise in the price of Stock P was always accompanied by a rise in the price of Stock B, but because a rise in the market was already expected on the basis of a rise in the price of Stock P, a rise in the price of Stock B acquired little associative strength. Similarly, in Phase 1, the subjects learned to expect a failure to rise in the market on the basis of a rise in the price of Stock N. In Phase 2, a rise in the price of Stock C was always accompanied by a rise in the price of Stock N,

and because a rise in the market was unexpected on the basis of a rise in the price of Stock N, Stock C acquired considerable associative strength.

Judgmental Overshadowing

The associative account of contingency judgment implies another cue-interaction effect, *judgmental overshadowing*, which has received relatively little study (but see Baker et al., 1993; see also Wagner, 1969, for an animal conditioning example). Suppose, for instance, that an outcome event is moderately contingent on Predictor Event 1 (PE1) and strongly contingent on Predictor Event 2 (PE2). On a given trial, then, the status of the outcome event likely will be expected on the basis of the occurrence or nonoccurrence of PE2. Therefore, the degree of association between the outcome event and PE1 will tend to change little from trial to trial. On the other hand, suppose the same outcome event is again moderately contingent on PE1 and weakly contingent on PE2. On a given trial, then, the status of the outcome event likely will not be expected on the basis of the occurrence or nonoccurrence of PE2. Therefore, the degree of association between the outcome event and PE1 will change considerably from trial to trial. This implies that eventually, after several trials, the degree of contingency between the outcome event and PE1 should be judged stronger in the latter condition than in the former.

Before proceeding, it is worth noting that from the associative-learning perspective, the difference between blocking and overshadowing is mainly procedural. In blocking, the relatively strong contingency between the outcome event and one predictor event is established prior to the introduction of the second predictor event. In overshadowing, the contingencies between the outcome event and both predictor events are established simultaneously. According to the associative account, this difference should result in qualitatively similar cue-interaction effects, albeit of greater magnitude in the blocking situation. From another more intuitive viewpoint, however, exactly when the relatively strong contingency is established in relation to the relatively weak contingency should make a considerable difference. If cue-interaction effects result from the conscious application of statistical or causal reasoning (see Waldmann & Holyoak, 1992), then the blocking situation, in which one contingency is firmly established before the other, may produce cue-interaction effects, whereas the overshadowing situation may not. For this reason, an investigation of overshadowing, as a phenomenon distinct from blocking, appears warranted.

There is some prior evidence of the existence of judgmental overshadowing. In a study by Gluck and Bower (1988), subjects learned to predict whether each of several fictional patients had either a rare disease or a mutually exclusive common disease, on the basis of whether the patient exhibited each of four symptoms. Then, for each symptom, the subjects judged the probability that each disease would be present in a patient exhibiting that symptom. In a cross-experiment comparison, Gluck and

Bower found that increasing the degree of contingency between the disease variable and two of the symptoms attenuated the subjects' judgments of the probability of the rare disease, given each of the other two symptoms. Taking these results as suggestive, we decided to perform a more conclusive test of the existence of judgmental overshadowing in contingency judgment by making within-experiment comparisons analogous to those made by Gluck and Bower.

Additionally, in the research reported here, we have tried to address some methodological issues related to the dependent variables used in past studies of cue interaction in contingency judgment. In most of the previous research on contingency judgment, the dependent variable was a rating of the degree of contingency between pairs of events. Chapman and Robbins (1990), for instance, had subjects rate—on a scale of -100 to 100 —the “predictive value” of each stock. They were told that a rating of 100 indicated that the predictor event was a “perfect positive predictor” of the outcome event, a rating of -100 indicated that the predictor event was a “perfect negative predictor,” and a rating of 0 indicated that the predictor event was a “random nonpredictor.” One problem with such ratings is that, from a naive subject's perspective, they have no necessary correspondence with any objective measure of the degree of contingency, although it is possible that a statistically sophisticated subject might consider them to correspond to correlation coefficients. Therefore, the subject is free to use such scales in a variety of ways.

Combined with task demands, the use of such rating scales could result in cue interaction without associative learning. In the experiment by Chapman and Robbins (1990), for example, the subjects may have felt some demand to rate Stocks B and C differently. They may then have reasoned that because Stock B was redundant with another predictive stock but Stock C was not, a lower rating should be given to Stock B. An unanchored rating scale would not prevent the subjects from giving a lower rating to Stock B, but a rating scale that is tied to some objective measure of contingency—and is familiar to the subjects—may prevent them from doing so. Hence the assessment approach employed in the present experiments. For each predictor event, our subjects judged the conditional probability of the outcome event's occurrence, given the predictor event's occurrence. This does not solve all potential response-mode problems, however. Gluck and Bower (1988), for example, had subjects judge the probability of a patient's having a disease, given that the patient exhibited a symptom. But for there to be a contingency between two events, as described earlier, one event must be more likely in the presence of the other than in the absence of the other. Therefore, judgments of the probability of one event in the presence of the other are insufficient to establish subjects' beliefs about the degree of contingency. This is why our subjects also judged the conditional probability of the outcome event's occurrence, given the predictor event's failure to occur.

The difference between these two probability judgments is a subjective version of the ΔP statistic, which we call $\Delta P'$ (Yates & Curley, 1986). A second advantage of using conditional probability judgments and $\Delta P'$ is that they allow us to assess judgment accuracy in a straightforward manner by comparing the observed $\Delta P'$ values with the objective ΔP values.

EXPERIMENT 1

We chose to study judgmental overshadowing in the context of a simulated medical decision-making task—a popular method of studying contingency judgment and related topics (e.g., Chapman, 1991; Gluck & Bower, 1988; Medin & Edelson, 1988; Smedslund, 1963). For each of several patients, subjects were told whether or not each of two symptoms was present. After predicting whether or not each patient had a fictional disease, subjects were told whether or not the patient actually had the disease. Finally, the subjects judged the degree to which the presence of the disease was contingent on the presence of each of the symptoms.

To study overshadowing, the degree of contingency between the disease and one of the symptoms was kept moderate, whereas the degree of contingency between the disease and the second symptom was varied across conditions. In one condition, the degree of contingency between the disease and the second symptom was greater than the degree of contingency between the disease and the first symptom. In another condition, the degree of contingency between the disease and the second symptom was less than the degree of contingency between the disease and the first symptom. Therefore, due to judgmental overshadowing, contingency judgments for the first symptom should be lower when the disease is strongly contingent on the second symptom than when the disease is weakly contingent on the second symptom.

Method

Subjects. The subjects were 34 undergraduates at the University of Michigan, who participated in the experiment to fulfill an introductory psychology course requirement.

Stimuli. The stimuli were the "medical charts" of 40 fictional hospital patients, presented on a video display. Each chart indicated whether or not the patient had each of two symptoms, rash and fever.¹ A cover story described the subjects' task as that of a physician who must predict whether or not each patient had the disease "chronitis" on the basis of the information in the medical charts.

Each patient was drawn randomly from a distribution with the following characteristics. The probability of having rash, the probability of having fever, and the probability of having chronitis were all .5. For all the subjects in the experiment, the degree of contingency between chronitis and fever was moderate. The probability of a patient's having chronitis given that the patient had fever was .7, whereas the probability of a patient's having chronitis given that the patient did not have fever was .3 ($\Delta P = .4$). For the subjects in the strong rash context, the degree of contingency between chronitis and rash was strong. The probability of a patient's having chronitis given that the patient had rash was .9, whereas the probability of a patient's having chronitis given that the patient did not have rash was .1 ($\Delta P = .8$). For the subjects in the weak rash context,

the degree of contingency between chronitis and rash was as weak as possible. The probability of a patient's having chronitis given that the patient had rash was .5, as was the probability of a patient's having chronitis given that the patient did not have rash ($\Delta P = 0$). Table 2 shows the relative frequencies with which the patients with each of the eight possible symptom-disease combinations (i.e., rash present, fever present, chronitis present; rash present, fever not present, chronitis present; etc.) were presented to the subjects in the strong and weak rash contexts.

Procedure. The subjects were randomly assigned to either the strong or weak rash context, and were run in small noninteracting groups, with each subject seated at a separate personal computer. Stimulus presentation and response collection were controlled via a program written using the Micro Experimental Laboratory software system (Schneider, 1988).

On each trial, the medical chart of one patient was presented. The subjects indicated whether they thought the patient had chronitis by pressing one of two keys on the computer keyboard, and were provided with immediate feedback. A short message—"Correct (Incorrect), this patient DOES (does NOT) have chronitis"—appeared at the bottom of the screen. The screen was cleared after 3 sec, and the next medical chart was presented when the subject pressed the "enter" key.

After diagnosing each of the 40 patients, the subjects made a series of probability judgments. First, they judged the conditional probability that a patient would have chronitis given that the patient had rash, in the absence of knowledge of whether or not the patient had fever [$P'(C|R)$]. Their instructions read as follows:

Imagine that you encounter a randomly selected patient, W. You know that patient W has RASH. You DO NOT KNOW whether patient W has FEVER. Question: What is the probability that W has chronitis?

The subjects were informed that a judgment of 0 meant that the patient was certain not to have chronitis, a judgment of 50 meant that the patient was as likely as not to have chronitis, and a judgment of 100 meant that the patient was certain to have chronitis. They were also told that their judgments could fall anywhere in the interval from 0 to 100. The subjects then judged the conditional probability that a patient would have chronitis given that the patient did not have rash, in the absence of knowledge of whether or not the patient had fever [$P'(C|-R)$]. Next, they made the analogous judgments with respect to fever [$P'(C|F)$ and $P'(C|-F)$].

After making their probability judgments, the subjects rated each symptom in terms of its ability to predict the presence of chronitis. They did so first for rash by using the right and left arrow keys to move an indicator (an "R") along the length of a horizontally oriented rating scale. The scale was 20 spaces long on the computer screen, with tick marks every five spaces. It was labeled with the words "Not Predictive" on the left, "Moderately Predictive" in the center, and "Perfectly Predictive" on the right. The indicator was initially displayed at the left end of the scale, and the dependent

Table 2
Relative Frequencies of Each Possible Rash-Fever-Chronitis Combination for Both the Strong Rash and Weak Rash Contexts for Experiments 1, 2, and 4

Chronitis	Symptom Combinations			
	Rash/ Fever	Rash/ No Fever	No Rash/ Fever	No Rash/ No Fever
Strong Rash Context				
Present	.325	.125	.025	.025
Absent	.025	.025	.125	.325
Weak Rash Context				
Present	.175	.075	.175	.075
Absent	.075	.175	.075	.175

variable was the number of spaces to the right that the subjects moved the indicator. After rating the predictiveness of rash, the subjects rated the predictiveness of fever. They did so by moving a second indicator (an "F") along the length of the same horizontally oriented rating scale. The first indicator, the "R," continued to be displayed in its final position while the second rating was made.

Results and Discussion

The means and standard deviations of the subjects' probability judgments, implied contingency indexes, and ratings are shown separately for each condition in Table 3. To test for the presence of judgmental overshadowing, we compared judgments about the degree of contingency between chronitis and fever across the two conditions. First, we compared judgments of the probability that a patient with fever would have chronitis [$P'(C|F)$], and found that they were significantly higher in the weak rash context than in the strong rash context [$t(32) = 2.23, p < .05$], indicating a form of judgmental overshadowing (cf. Gluck & Bower, 1988). Recall, however, that this does not necessarily imply that the subjects in the weak rash context believed that the degree of *contingency* between chronitis and fever was stronger than did the subjects in the strong rash context. Therefore, we compared the implied contingency indexes [$\Delta P' = P'(C|F) - P'(C|-F)$] across the two conditions, and found that they too were significantly higher for the weak rash context [$t(32) = 2.85, p < .01$]. In addition, we compared ratings of the usefulness of fever as a predictor of chronitis across the two conditions, and again found them to be significantly higher for the weak rash context than for the strong rash context [$t(32) = 6.70, p < .005$]. To summarize, judgmental overshadowing did occur in Experiment 1, and it did not depend on the mode of judgment. This latter point is taken up further in Experiment 2.

Before accepting the results of Experiment 1 as support for the associative-learning account of contingency judgment, however, we should consider two relatively simple alternative explanations. First, recall that we asked

our subjects to judge the probability that a patient with or without each symptom would have chronitis, in the absence of any knowledge of whether the other symptom was present. For example, we asked them to judge the probability that a patient with fever, but with an unknown status with respect to rash, would have chronitis. It is possible, however, that instead, the subjects judged the probability that a patient with fever, but *without* rash, would have chronitis. They may have done so either because they misunderstood our instructions, or because they assumed that missing information about a symptom implied that the symptom was not present. This too would produce apparent judgmental overshadowing.²

The results of Experiment 1, however, lend little support to this hypothesis. Figures 1 and 2 show the means of the subjects' judgments of the probability of chronitis given the presence and absence of each symptom (the black bars). For comparison, we have plotted the objective probabilities that we intended the subjects to judge (the gray bars in Figures 1 and 2), along with the objective probabilities of chronitis given the presence and absence of each symptom, assuming the absence of the other symptom (the white bars). It is clear from the figures that the subjects' judgments were not entirely accurate with respect to the objective probabilities that we intended them to judge. It is also clear, however, that their judgments were not entirely accurate with respect to the objective probabilities of chronitis given each symptom, assuming the absence of the other symptom. We calculated the mean squared deviation (*MSE*) of the subjects' mean judgments from both sets of objective probabilities plotted in Figures 1 and 2. Their judgments were slightly closer to the objective probabilities we intended them to judge for the weak rash context (*MSE* = .027 vs. .056), but were slightly farther from those same objective probabilities for the strong rash context (*MSE* = .014 vs. .010). There is no reason to conclude, therefore, that the judgmental overshadowing found in Experiment 1 was the result of

Table 3
Means and Standard Deviations of Subjects' Probability Judgments (%),
Implied Contingency Indexes ($\Delta P'$), and Ratings for Each Condition
in Experiment 1: Contingency Judgment With Predictor Events
Presented Simultaneously and With Predictions Demanded

Judgment	Condition			
	1: Strong Rash Context		2: Weak Rash Context	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
P' (Chronitis Rash)	78.12	17.66	45.06	11.57
P' (Chronitis No Rash)	21.24	21.82	31.47	22.41
$\Delta P'$ (Rash)	56.88	34.37	13.59	25.95
Rating (Rash)	15.94	2.25	7.64	3.14
P' (Chronitis Fever)	44.47	21.53	60.00	18.96
P' (Chronitis No Fever)	42.82	18.83	33.53	18.09
$\Delta P'$ (Fever)	1.65	19.37	26.47	30.20
Rating (Fever)	3.53	2.45	10.82	3.76

Note—Condition 1: Moderate chronitis-fever contingency ($\Delta P = .4$); strong chronitis-rash contingency ($\Delta P = .8$). Condition 2: Moderate chronitis-fever contingency ($\Delta P = .4$); weak (null) chronitis-rash contingency ($\Delta P = 0$).

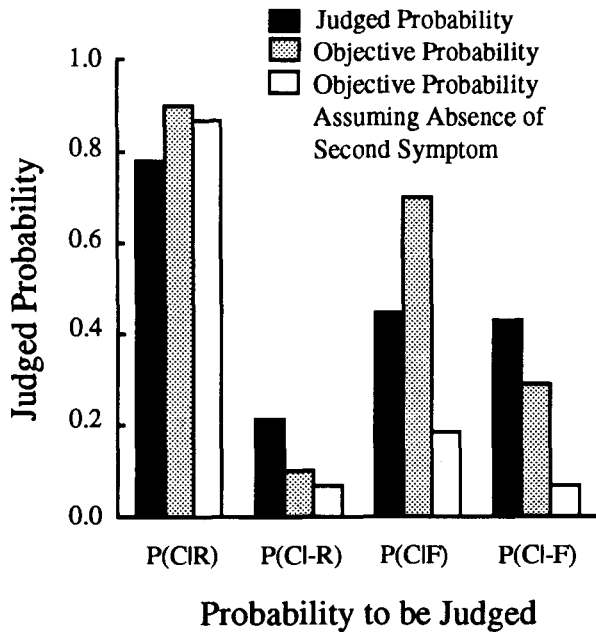


Figure 1. Experiment 1 (strong rash context). Judged probabilities of chronitis given the presence and absence of each symptom (black bars), compared with the objective probabilities (gray bars), and the objective probabilities assuming the absence of the other symptom (white bars).

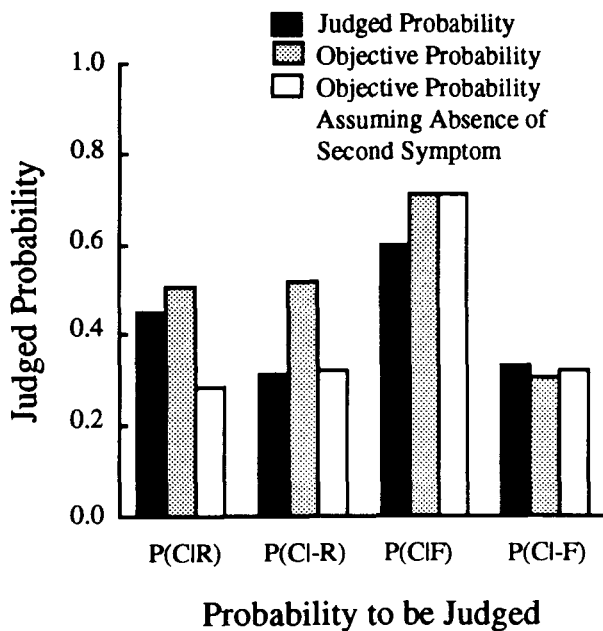


Figure 2. Experiment 1 (weak rash context). Judged probabilities of chronitis given the presence and absence of each symptom (black bars), compared with the objective probabilities (gray bars), and the objective probabilities assuming the absence of the other symptom (white bars).

the subjects' judging the probabilities of chronitis given the presence and absence of each symptom, assuming the absence of the other symptom.

A second alternative explanation of the judgmental overshadowing found in Experiment 1 is that the subjects made comparative rather than absolute judgments. For example, after judging the probability of chronitis given rash to be quite high, the subjects in the strong rash context may have reported the probability of chronitis given fever to be quite low, simply to reflect what they felt to be a large difference or contrast between the two contingencies. Similarly, after judging the probability of chronitis given rash to be quite low, the subjects in the weak rash context may have reported the probability of chronitis given fever to be quite high in order to reflect a similar contrast in the other direction. Note that this is essentially the problem we intended to eliminate by having our subjects make conditional probability judgments, in addition to ratings of the degree of contingency. However, considering the fact that probability judgments do not always obey the constraints normally placed on probabilities (e.g., Wright & Whalley, 1983; Yates, 1990), this seemed like an alternative that was worth further exploration. We began to address this issue in Experiment 2.

EXPERIMENT 2

The results of Experiment 1 provided some evidence that judgmental overshadowing and, most likely, other cue-interaction effects do not depend heavily on the mode of judgment. Whether we considered predictiveness ratings, conditional probability judgments, or judgments of $\Delta P'$, judgmental overshadowing occurred. The possibility remained, however, that alternative subjective contingency indexes may have failed to reflect the phenomenon. One such alternative index is based on χ , the product of the conditional proportion of Event 1 occurrences given the occurrence of Event 2 and the conditional proportion of Event 2 occurrences given the occurrence of Event 1 (Jones, 1985). The index χ is particularly interesting because it has been argued that people's intuitive notions of contingency or association may, in some cases, be captured better by χ than by ΔP (Jones, 1983). Accordingly, in Experiment 2, we tested for judgmental overshadowing by using an implied association index: $\chi' = P'(C|F) \times P'(F|C)$.

A second addition to the procedure of Experiment 2 is that we counterbalanced the order in which the symptoms were judged. Recall that in Experiment 1, the subjects were always asked first about the degree of contingency between chronitis and rash, and then about the degree of contingency between chronitis and fever. In Experiment 2, however, only half the subjects were asked about the contingencies in the same order as in Experiment 1; the rest received the same questions in the reverse order. This allowed us to test the possibility that judgmental overshadowing in Experiment 1 was the result of a re-

sponse bias or contrast effect that depends on judging the chronitis-fever contingency after judging either the strong or weak chronitis-rash contingency. In other words, it may be that when the subjects in the strong and weak rash contexts judge the strength of the chronitis-fever contingency without first having judged the strength of the chronitis-rash contingency, there is no difference between the two conditions.

Method

Subjects. The subjects were 84 undergraduates at the University of Michigan, who participated to fulfill an introductory psychology course requirement.

Procedure. The procedure of Experiment 2 was the same as that of Experiment 1, but with the following additions. After diagnosing the 40 patients, half of the subjects judged the probability of chronitis given rash and the probability of chronitis given no rash, as well as the probability of rash given chronitis. Then they judged the probability of chronitis given fever, the probability of chronitis given no fever, and the probability of fever given chronitis. Note that the objective value of κ was .49 for the fever-chronitis association, .81 for the rash-chronitis association in the strong rash context, and .25 for the rash-chronitis association in the weak rash context. Next, the subjects rated the effectiveness of rash as a predictor of chronitis, and then the effectiveness of fever as a predictor of chronitis. The other half of the subjects made the same judgments, but they were always asked about the chronitis-fever contingency before being asked about the chronitis-rash contingency.

Results and Discussion

The means and standard deviations of the subjects' probability judgments, implied contingency indexes, implied association indexes, and ratings are shown separately for each condition in Table 4. Analyses of variance (ANOVAs), with condition and order as between-subjects factors, were run on each dependent variable. There was a significant main effect of condition on judgments of the probability that a patient with fever would have chronitis

[$F(1,80) = 11.84, p < .001$], a significant main effect of condition on the implied contingency index [$F(1,80) = 5.34, p < .05$], and a significant main effect of condition on ratings of the usefulness of fever as a predictor of chronitis [$F(1,80) = 12.69, p < .001$]. That is, the primary results of Experiment 1 were replicated in Experiment 2. In addition, there was a significant main effect of condition on the implied association index [$F(1,80) = 7.52, p < .01$]. We can therefore conclude that judgmental overshadowing is a robust phenomenon in that it does not depend on the contingency index used to assess subjects' beliefs. The implications of the apparent robustness of judgmental overshadowing will be taken up further in the General Discussion.

The only significant effect involving the order in which the two contingencies were judged was the main effect on ratings of the usefulness of fever as a predictor of chronitis [$F(1,80) = 4.99, p < .05$]. None of the probability judgments, and neither the implied contingency indexes nor the implied association indexes, differed significantly across the two judgment orders. The ratings may have been affected either because they are somewhat arbitrary—unlike the probability estimates—or because they were made on a graphic scale, which may encourage this kind of effect. The lack of an order effect in the probability judgments suggests that judgmental overshadowing was not the result of a response bias that depended on judging the chronitis-fever contingency after judging either the strong or weak chronitis-rash contingency.

These results do not, however, rule out the possibility that judgmental overshadowing is due to a similar response bias that does not depend on the order in which the contingencies are judged. That is, subjects may judge one contingency relative to the other, regardless of whether they have previously made an explicit judgment about the other contingency. The results of Experiment 2 also do

Table 4
Means and Standard Deviations of Subjects' Probability Judgments (%), Implied Contingency Indexes ($\Delta P'$), Implied Association Indexes (κ'), and Ratings for Each Condition in Experiment 2: Contingency Judgment With Predictor Events Presented Simultaneously and With Predictions Demanded

Judgment	Condition			
	1: Strong Rash Context		2: Weak Rash Context	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
P' (Chronitis Rash)	71.05	18.47	46.95	13.36
P' (Chronitis No Rash)	19.71	15.30	44.17	21.21
P' (Rash Chronitis)	82.93	19.36	58.55	21.86
$\Delta P'$ (Rash)	51.33	28.92	2.79	24.90
κ' (Rash)	59.87	23.73	28.58	14.71
Rating (Rash)	13.29	6.38	6.29	4.06
P' (Chronitis Fever)	38.43	22.78	53.93	17.53
P' (Chronitis No Fever)	35.48	26.16	38.00	18.74
P' (Fever Chronitis)	60.45	22.03	63.00	24.22
$\Delta P'$ (Fever)	2.95	27.05	15.93	25.12
κ' (Fever)	23.33	17.40	35.10	21.32
Rating (Fever)	5.33	3.98	8.55	4.51

Note—Condition 1: Moderate chronitis-fever contingency ($\Delta P = .4, \kappa = .49$); strong chronitis-rash contingency ($\Delta P = .8, \kappa = .81$). Condition 2: Moderate chronitis-fever contingency ($\Delta P = .4, \kappa = .49$); weak (null) chronitis-rash contingency ($\Delta P = 0, \kappa = .25$).

not rule out the possibility that judgmental overshadowing is the result of what might be termed a *true contrast effect*, rather than a mere response bias. In other words, the subjective strength of the chronitis-*fever* contingency may in fact differ across contexts, not as the result of an associative-learning process, but as the result of a contrast effect analogous to perceptual contrast effects. Substantial evidence that such contrast effects occur in non-perceptual judgment has been provided by Manis, Nelson, and Shedler (1988). Experiment 3 was intended to rule out the possibility that either a response bias or a true contrast effect is the cause of judgmental overshadowing.

EXPERIMENT 3

Experiment 3 was designed so that judgmental overshadowing would not be expected on the basis of the associative account of contingency judgment, but would still be expected if it were the result of either a response bias or a true contrast effect. In each trial, a subject was informed of *either* the patient's status with respect to rash *or* the patient's status with respect to fever. Then, on the basis of this one piece of information, the subject tried to determine whether the patient had chronitis. After many such trials, the subject judged the degree to which the presence of chronitis was contingent on each symptom.

According to the associative account of contingency judgment, this experiment should not produce judgmental overshadowing. To see why, consider only the trials in which the subject was informed that the patient had fever. Because the subject had no information about rash in these trials, the degree to which chronitis was expected depended only on the degree of contingency between chronitis and fever. Therefore, the degree of association between fever and chronitis was independent of the degree of contingency between chronitis and rash, and judgments of the degree of contingency between chronitis and fever were independent of the degree of contingency between chronitis and rash. On the other hand, if judgmental over-

shadowing is simply the result of a contrast effect, it should still occur in Experiment 3. This is because, as in Experiments 1 and 2, the subjects would still consider two contingencies simultaneously and therefore would judge each one relative to the other.

Method

Subjects. The subjects were 50 undergraduates at the University of Michigan, who participated to fulfill an introductory psychology course requirement.

Stimuli and Procedure. The stimuli were similar to those used in Experiments 1 and 2, but each medical chart provided information about only one of the two symptoms—either rash or fever. The probability that a given medical chart contained information about rash was .5, as was the probability that it contained information about fever. The probability of a patient's having chronitis was also .5. As in Experiment 1, the probability of chronitis given the presence of fever was .7 and the probability of chronitis given the absence of fever was .3. The probabilities of chronitis given the presence of rash were .9 (strong rash context) and .5 (weak rash context), and the probabilities of chronitis given the absence of rash were .1 (strong rash context) and .5 (weak rash context). To ensure that each symptom was presented approximately as many times as in Experiments 1 and 2, the number of patients was doubled from 40 to 80.

The procedure, except as noted above, was the same as in Experiment 1.

Results and Discussion

The means and standard deviations of subjects' probability judgments, implied contingency indexes, and ratings are shown separately for each condition in Table 5. Again, to test for the presence of judgmental overshadowing, we compared judgments about the degree of contingency between chronitis and fever across the two conditions. In this case, however, there were no significant differences between the judgments of the subjects in the two conditions. Neither the judgments of the conditional probability of chronitis given fever, the implied contingency indexes, nor the ratings of the usefulness of fever as a predictor of chronitis differed significantly across conditions [$t(48) = 0.09$, n.s.; $t(48) = .10$, n.s.; $t(48) = 0.14$,

Table 5
Means and Standard Deviations of Subjects' Probability Judgments (%), Implied Contingency Indexes ($\Delta P'$), and Ratings for Each Condition in Experiment 3: Contingency Judgment With Predictor Events Presented Independently and With Predictions Demanded

Judgment	Condition			
	1: Strong Rash Context		2: Weak Rash Context	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
P' (Chronitis Rash)	71.00	20.31	53.56	19.57
P' (Chronitis No Rash)	35.60	29.17	38.20	15.56
$\Delta P'$ (Rash)	35.40	32.34	15.36	24.31
Rating (Rash)	15.12	2.45	9.20	3.44
P' (Chronitis Fever)	53.20	19.52	53.68	19.93
P' (Chronitis No Fever)	36.40	15.65	37.56	17.44
$\Delta P'$ (Fever)	16.80	26.18	16.12	20.87
Rating (Fever)	9.24	5.21	9.44	4.64

Note—Condition 1: Moderate chronitis-*fever* contingency ($\Delta P = .4$); strong chronitis-rash contingency ($\Delta P = .8$). Condition 2: Moderate chronitis-*fever* contingency ($\Delta P = .4$); weak (null) chronitis-rash contingency ($\Delta P = 0$).

n.s.]. These results are consistent with the associative account of contingency judgment. They also indicate that the judgmental overshadowing found in Experiments 1 and 2, as well as the cue-interaction effects demonstrated by others (Chapman, 1991; Chapman & Robbins, 1990; Gluck & Bower, 1988), are unlikely to be due to either a response bias or a true contrast effect.

EXPERIMENT 4

The procedure of Experiments 1–3, in which subjects predicted the occurrence or nonoccurrence of one event on the basis of the occurrence or nonoccurrence of another, is atypical of research on contingency judgment. More often, subjects have simply been presented with information about the occurrence or nonoccurrence of two events simultaneously on each of a series of trials (e.g., Smedslund, 1963; Ward & Jenkins, 1965). Presumably, this procedure is like many, if not most, real-world contingency judgment tasks, in which we are called on to make judgments about contingencies without having made explicit predictions. This raises the question of whether judgmental overshadowing would still occur in the absence of subjects making such predictions. If judgmental overshadowing is in fact based on the product of a trial-by-trial, associative-learning process, then having to make predictions may foster the formation of the expectations and associations that lead to judgmental overshadowing. On the other hand, not having to make predictions—and being informed about the predictor and outcome events simultaneously—may render the formation of expectations and associations unnecessary, and thus judgmental overshadowing may not occur.

We performed a fourth experiment, therefore, in which subjects did not predict whether each patient had chronicity on the basis of whether the patient had rash and fever. Instead, the subjects were presented simultaneously with information about rash, fever, and chronicity. This manipulation not only makes our research more comparable with previous research on contingency judgment, but also al-

lows us to test the hypothesis that explicit predictions are necessary for judgmental overshadowing to occur.

Method

Subjects. The subjects were 46 undergraduates at the University of Michigan, who participated to fulfill an introductory psychology course requirement.

Procedure. The procedure was identical to that of Experiment 1, except for the following change. The subjects were informed whether each patient had chronicity without first having to make a prediction. As soon as each medical chart was presented, feedback about whether the patient did or did not have chronicity appeared at the bottom of the screen. The screen was cleared after 4 sec, and the subject pressed the space bar to present the next patient.

Results and Discussion

The results of Experiment 4 are presented in Table 6. Again, we compared judgments of the degree of contingency between chronicity and fever across the two conditions. As in Experiments 1 and 2, judgments of the probability that a patient with fever would have chronicity were higher for the weak rash context than for the strong rash context [$t(44) = 2.25, p < .05$], as were the implied contingency indexes [$t(44) = 1.90, p < .05$]. Two subjects' ratings of the usefulness of fever as a predictor of chronicity were not obtained, due to computer failure. These ratings were also higher for the weak rash context than for the strong rash context [$t(42) = 1.88, p < .05$]. This pattern of results overall is very similar to that of Experiments 1 and 2, and it suggests that engaging in a prediction task is not necessary for judgmental overshadowing to occur, although it may indeed heighten the effect somewhat by forcing subjects to attend more closely to the stimuli.

GENERAL DISCUSSION

We have established that judgmental overshadowing does occur, consistent with previous work on overshadowing (Gluck & Bower, 1988), as well as the cue-interaction effects known as blocking and conditioned inhibition (Chapman, 1991; Chapman & Robbins, 1990; Waldmann

Table 6
Means and Standard Deviations of Subjects' Probability Judgments (%), Implied Contingency Indexes ($\Delta P'$), and Ratings for Each Condition in Experiment 4: Contingency Judgment With Predictor Events Presented Simultaneously and With No Predictions Demanded

Judgment	Condition			
	1: Strong Rash Context		2: Weak Rash Context	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
P' (Chronicity Rash)	72.61	17.83	42.17	18.03
P' (Chronicity No Rash)	15.13	14.51	35.70	20.02
$\Delta P'$ (Rash)	57.48	25.69	6.48	25.72
Rating (Rash)	14.71	3.23	6.30	4.38
P' (Chronicity Fever)	40.09	20.56	54.44	22.72
P' (Chronicity No Fever)	31.17	22.49	29.65	19.36
$\Delta P'$ (Fever)	8.91	28.83	24.78	27.86
Rating (Fever)	5.76	3.52	8.39	5.47

Note—Condition 1: Moderate chronicity-fever contingency ($\Delta P = .4$); strong chronicity-rash contingency ($\Delta P = .8$). Condition 2: Moderate chronicity-fever contingency ($\Delta P = .4$); weak (null) chronicity-rash contingency ($\Delta P = 0$).

& Holyoak, 1992). Furthermore, we have demonstrated that judgmental overshadowing—and probably the other cue-interaction effects—does not depend heavily on the mode of judgment. It occurred whether the dependent variable considered was a rating of the degree of contingency, a judgment of the conditional probability of the occurrence of one event given the occurrence of the other, the implied contingency index $\Delta P'$, or the implied association index κ' , as in Experiment 2. In addition, we showed in Experiments 2 and 3 that judgmental overshadowing is the result of neither a response bias nor a true contrast effect, which bolsters the associative-learning account. And, finally, we demonstrated in Experiment 4 that making explicit predictions is not necessary for the effect to occur.

Our work adds to a growing body of literature on cue-interaction effects, including judgmental overshadowing, on contingency and causality judgment. Recent work by Baker et al. (1993), for instance, has demonstrated judgmental overshadowing in a paradigm somewhat similar to our own. In their study, subjects played a video game in which tanks either safely traversed a minefield or were blown up in the attempt. Then the subjects judged the effectiveness of both a camouflaging operation that they were able to initiate on each trial and the presence of a "spotter plane," over which they had no control, in leading to a successful minefield crossing. Consistent with the present results, the camouflaging operation was judged to be more effective when the presence of the spotter plane was uncorrelated with successful minefield crossings than when it was highly correlated with successful minefield crossings. Additionally, some recent work by Busemeyer, Myung, and McDaniel (1993) has demonstrated overshadowing in a very different paradigm. In that study, subjects learned to predict the value of a numerical criterion variable (a plant's height) on the basis of two numerical cue variables (amounts of two different growth hormones applied to the plant), which were uncorrelated with each other, over a number of trials. It was found that a moderately valid cue came to be relied on more heavily when paired with a less valid cue than when paired with a more valid cue. Although Busemeyer et al. assume that this form of overshadowing is produced by the same mechanism as the judgmental overshadowing demonstrated here, the relationship between the two phenomena has yet to be clarified.

One of the unique findings to come from the present research is the apparent robustness of judgmental overshadowing across dependent variables. We should make clear, however, that this robustness is in part illusory. The significant differences across the two conditions between both $\Delta P'$ judgments and κ' judgments were driven mainly by the significant differences in the subjects' judgments of the probability of chronitis given the presence of fever. Their judgments of the probability of chronitis given the absence of fever did not differ significantly across conditions in any of the experiments reported here. Similarly, the subjects' judgments of the probability of fever given chronitis did not differ significantly across conditions in

Experiment 2. Although it is important to know that taking into account judgments of $P(C|F)$ and $P(F|C)$ does not eliminate judgmental overshadowing, it is notable that those judgments themselves did not reflect judgmental overshadowing.

Why might this be? From the associative perspective, the lack of an overshadowing effect on judgments of chronitis given the absence of fever suggests that people do not spontaneously form associations between the non-occurrence of one event and the occurrence of another—even when the nonoccurrence is explicitly acknowledged. Also from the associative perspective, the lack of an overshadowing effect on judgments of the probability of fever given chronitis has two implications. First, it suggests that the association formed between two events is not bidirectional. Otherwise, because $P(C|F)$ and $P(F|C)$ were equal in the present experiments, we would have expected judgments of these two probabilities to be equal and to reflect the same overshadowing effect. Instead, it appears that the subjects recognized the nondependence of these two probabilities.

Thus, although there are a number of associations that might reasonably be formed between pairs of event occurrences and nonoccurrences, there appear to be certain constraints on which of these associations are actually formed and on how they are used. One type of constraint could be attentional limitations. It may be difficult to form associations between both the occurrence and nonoccurrence of an event and the occurrence of a second event, due to our limited capacity to process the incoming information. A more interesting type of constraint may derive from people's higher level knowledge about cause-effect relationships. This idea is consistent with the work of Waldmann and Holyoak (1992), who found that blocking occurs only when the predictor events are interpreted by subjects as potential causes of a single effect. If the predictor events are interpreted as potential effects of a single cause, blocking does not occur. According to Waldmann and Holyoak, this is because rather than simply forming associations, people apply knowledge about the structure of cause-effect relationships to the judgment task. People know, for instance, to beware of spurious causes. Therefore, they hesitate to draw conclusions about the causal status of a predictor event that is redundant with another predictor event that has already been established as a cause of the outcome event, and they respond with lower judgments for the redundant predictor event. People also know, however, that there is no such thing as a spurious effect. That is why blocking does not occur when the predictor events are interpreted as multiple effects of a single cause.³

We should point out that Waldmann and Holyoak (1992) do not see associative learning playing any role at all in contingency judgment. They contend that people base their judgments on an accurate analysis of the relevant conditional proportions, as in the traditional account discussed earlier (see also Cheng & Novick, 1992). Cue-interaction effects, according to this view, are due only to the type of causal reasoning process described above. If this were

true, however, then any contingency judgment situation that did not require causal reasoning would not produce cue-interaction effects. But in our experiments, the subjects simply judged conditional probabilities of the occurrence of one event given the occurrence or nonoccurrence of another. Causal reasoning was irrelevant to the subjects' task, yet judgmental overshadowing still occurred. This fact is difficult to reconcile with the view that people make accurate contingency judgments in multiple-predictor situations on the basis of an analysis of the relevant conditional proportions.

Therefore, we speculate that cue-interaction effects are the result of an associative-learning process that is constrained by attentional limitations and guided by high-level knowledge about cause-effect relationships. This suggests a number of interesting research questions. How, for instance, do people judge the conditional probability of Event 2 given the nonoccurrence of Event 1 if they have not formed the relevant association? Do they do so according to a memory-based judgment strategy (Hastie & Park, 1986), or do they derive this probability on the basis of their judgments of other probabilities, such as the probability of Event 2 given the occurrence of Event 1? Can they form associations between the nonoccurrence of one event and the occurrence of another if instructed to do so? Does judgmental overshadowing still occur if people are not warned ahead of time that they will be making contingency judgments? Does it occur if the events have no apparent causal relationship?

Finally, we should consider the practical implications of judgmental overshadowing. If we think of the subjects' task as one of making judgments in the face of missing information (Jaccard & Wood, 1988), these implications become clear. The following is an example of the possible effects of judgmental overshadowing. A personnel manager in charge of hiring new graduates for a large corporation has been exposed to a strong contingency between an applicant's having a degree in natural science and good job performance, and a moderate contingency between having many extracurricular activities and good job performance. Then the personnel manager receives an application from someone with many extracurricular activities, but whose major field of study is not indicated. Due to judgmental overshadowing, the personnel manager may give less consideration to this particular application than he or she should on the basis of the actual degree of contingency between good job performance having many extracurricular activities. Future research on cue interaction in contingency judgment should explore such practical issues as well as the theoretical issues discussed earlier.

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.
- ARKES, H. R., & ROTHBART, M. (1985). Memory, retrieval, and contingency judgments. *Journal of Personality & Social Psychology*, *49*, 598-606.
- BAKER, A. G., MERCIER, P., VALLÉE-TOURANGEAU, F., FRANK, R., & PAN, M. (1993). Selective associations and causality judgments: The presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *19*, 414-432.
- BEYTH-MAROM, R. (1982). Perception of correlation reexamined. *Memory & Cognition*, *10*, 511-519.
- BUSEMEYER, J. R. (1991). Intuitive statistical estimation. In N. H. Anderson (Ed.), *Contributions to information integration theory: Vol. 1. Cognition* (pp. 187-215). Hillsdale, NJ: Erlbaum.
- BUSEMEYER, J. R., MYUNG, I. J., & MCDANIEL, M. A. (1993). Cue-competition effects: Empirical tests of adaptive network learning models. *Psychological Science*, *4*, 190-195.
- 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., & NOVICK, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.
- GLUCK, M. A., & BOWER, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227-247.
- GLUCK, M. A., & BOWER, G. H. (1990). Component and pattern information in adaptive networks. *Journal of Experimental Psychology: General*, *119*, 105-109.
- HASTIE, R., & PARK, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, *93*, 258-268.
- JACCARD, J., & WOOD, G. (1988). The effects of incomplete information on the formation of attitudes toward behavioral alternatives. *Journal of Personality & Social Psychology*, *54*, 580-591.
- JONES, G. V. (1983). Identifying basic categories. *Psychological Bulletin*, *94*, 423-428.
- JONES, G. V. (1985). Disinterring the Spearman-Thomson coefficient of association, $r_{JK}/\sqrt{r_{JK}}$. *British Journal of Mathematical & Statistical Psychology*, *38*, 122-123.
- LIPE, M. G. (1990). A lens model analysis of covariation research. *Journal of Behavioral Decision Making*, *3*, 47-59.
- MANIS, M., NELSON, T. W., & SHEDLER, J. (1988). Stereotypes and social judgment: Extremity, assimilation, and contrast. *Journal of Personality & Social Psychology*, *55*, 28-36.
- MEDIN, D. L., & EDELSON, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, *117*, 68-85.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.
- SCHNEIDER, W. (1988). Micro Experimental Laboratory: An integrated system for IBM PC compatibles. *Behavior Research Methods, Instruments, & Computers*, *20*, 206-217.
- SHAKLEE, H., & MIMS, M. (1982). Sources of error in judging event covariations: Effects of memory demands. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *8*, 208-224.
- SHANKS, D. R. (1990). Connectionism and human learning: Critique of Gluck and Bower (1988). *Journal of Experimental Psychology: General*, *119*, 101-104.
- SHANKS, D. R., & DICKINSON, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229-261). New York: Academic Press.
- SMEDSLUND, J. (1963). The concept of correlation in adults. *Scandinavian Journal of Psychology*, *4*, 165-173.
- WAGNER, A. R. (1969). Stimulus selection and a "modified continuity theory." In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 3, pp. 1-41). New York: Academic Press.
- 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.

- WARD, W. C., & JENKINS, H. M. (1965). The display of information and the judgment of contingency. *Canadian Journal of Psychology*, **19**, 231-241.
- WRIGHT, G., & WHALLEY, P. (1983). The supra-additivity of subjective probability. In B. P. Stigum & F. Wenstop (Eds.), *Foundations of utility and risk theory with applications* (pp. 233-244). Dordrecht, The Netherlands: Reidel.
- YATES, J. F. (1990). *Judgment and decision making*. Englewood Cliffs, NJ: Prentice-Hall.
- YATES, J. F., & CURLEY, S. P. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, **62**, 293-302.

NOTES

1. Technically, rash and fever are medical *signs* rather than *symptoms*. Because we adopted the lay terminology (referring to rash and fever as symptoms) during our experiments, we decided to use it throughout this paper.

2. This is essentially the same criticism that was directed at the work of Gluck and Bower (1988) by Shanks (1990). See also Gluck and Bower (1990) for their response to this criticism.

3. Waldmann and Holyoak (1992) identify the symptom/disease relationship as one of effect to cause. This seems to put the present results, as well as those of Chapman (1991; Chapman & Robbins, 1990) and Gluck and Bower (1988), at odds with the notion that cue-interaction effects occur only when the predictor events are causes of a single effect. We feel, however, that the symptom/disease relationship is somewhat ambiguous in terms of cause and effect, especially for naive subjects. This may be why cue-interaction effects still occur using a symptom/disease cover story.

(Manuscript received June 19, 1992;
revision accepted for publication March 11, 1993.)