Nonnormative discounting: There is more to cue interaction effects than controlling for alternative causes

KELLY M. GOEDERT

Pacific Lutheran University, Tacoma, Washington

and

BARBARA A. SPELLMAN

University of Virginia, Charlottesville, Virginia

Several experiments on human causal reasoning have demonstrated “discounting”—that the presence of a strong alternative cause may decrease the perceived efficacy of a moderate target cause. Some, but not all, of these effects have been shown to be attributable to subjects’ use of conditional rather than unconditional contingencies (i.e., subjects control for alternative causes). We review experimental results that do not conform to the conditionalizing contingency account of causal judgment. In four experiments, we demonstrate that there is “nonnormative discounting” above what is accounted for by conditionalization, that discounting may depend on the nature of the question put to the subjects, and that discounting can be affected by motivation. We conclude that because nonnormative discounting occurs for summary presentations as well as trial-by-trial presentations of information and because nonnormative discounting depends on motivation, it is not a necessary result of cue competition during the contingency learning process.

Everyday causal inference is burdened by information about competing causes. For example, suppose a physician prescribed an antihistamine to relieve a patient’s allergies. The patient would be anxious to learn whether, in fact, the antihistamine worked. But the world does not readily reveal her causal secrets: A symptom-free day could be due to the drug, a reduction in the allergen, or to the presence of some competing allergy inhibitor. How then does one make a causal inference when faced with multiple candidate causes?

A number of researchers have demonstrated that when the effect of a moderately contingent cue is learned in the presence of a strongly contingent cue, causal judgments of the moderate cue are reduced relative to a situation in which the strongly contingent cue is absent (e.g., Baker, Mercier, Vallée-Tourangeau, Frank, & Pan, 1993; Price & Yates, 1993). This phenomenon has been referred to as discounting (e.g., Vallée-Tourangeau, Baker, & Mercier, 1994), cue competition (e.g., Busemeyer, Myung, & McDaniel, 1993), and cue interaction (e.g., Tangen & Allan,

2003). These terms have been used to describe a wide range of phenomena in which causal judgments of a target cue change in response to information about an alternative cue. Such cue interaction effects have been observed in a variety of experimental paradigms, including forward blocking (when a completely predictive cue is learned about first, e.g., A^+ , and then presented in compound with a redundant predictive cue, AT^+), backward blocking (a completely predictive cue compound is presented first, AT^+ , and then a completely predictive element of that compound is presented, A^+), and simultaneous blocking or relative validity (two cues that differ in their predictiveness are learned about simultaneously). These experimental paradigms are various ways of *obtaining* cue interaction effects and, until now, terms referring to cue interaction effects have been used to describe both situations in which it is statistically normative to reduce one’s causal judgment of a target given the presence of a strong alternative (e.g., Price & Yates, 1995; Tangen & Allan, 2003) and situations in which it is not normative (e.g., Busemeyer et al., 1993). In this paper, using a simultaneous blocking paradigm, we provide evidence for two distinct types of cue interaction effects that rely on different underlying cognitive processes and are distinguished on the basis of whether the effect conforms to statistically normative inferences. We label those cases in which cue interaction effects conform to statistical inferences *conditionalization* and those cases in which they do not *nonnormative discounting*.¹

This research was supported by an NIMH grant to B.A.S. We thank Lorraine Allan, Karen Schmidt, Jason Tangen, and an anonymous reviewer for helpful comments on an earlier version of this article. We are also grateful to Heather Alexander, Clyner Antalón, Kyle Littman, Jessica Logan, Nadine Meyers, and Jennifer Samuelson for help with data collection. Correspondence should be addressed to K. M. Goedert, Department of Psychology, Pacific Lutheran University, Tacoma, WA 98447 (e-mail: goedert@plu.edu).

Both associative (e.g., Rescorla–Wagner model, Rescorla & Wagner, 1972) and statistical models (e.g., probabilistic contrast theory, Cheng & Novick, 1990; Power PC theory, Cheng, 1997) have been proposed to explain human causal attribution based on contingency information (see De Houwer & Beckers, 2002, for a review, and Shanks, Holyoak, & Medin, 1996, for a set of relevant articles). It has been demonstrated, however, that at asymptote these models make very similar, if not identical, predictions regarding causal judgment performance (Chapman & Robbins, 1990; Cheng, 1997). According to statistical models, to judge a single candidate cause, the reasoner should consider how the presence of the potentially causal event changes the probability of the outcome event (e.g., Allan, 1980; Cheng & Novick, 1990; also see Cheng, 1997, for discussion of a covariation model that takes into account the base rate of the effect). The two pieces of information that are critical to this assessment are the probability of the effect given the presence of the candidate cause [$P(E|C)$] and the probability of the effect given the absence of the candidate cause [$P(E|\sim C)$]. The change in probability (ΔP) is then computed as follows:

$$\Delta P = P(E|C) - P(E|\sim C).$$

ΔP varies between -1 and $+1$. A nonzero ΔP indicates a statistical relationship between the candidate cause and the effect.

In the allergy example, if on 18 of the 24 days on which the patient took the antihistamine she experienced allergy relief but only on 6 of the 24 days on which she did not take the antihistamine she experienced allergy relief, one could calculate ΔP as $18/24 - 6/24$ and find that the contingency between taking the antihistamine and allergy relief was $.5$. The antihistamine could be considered a moderately successful causal agent.

However, what if there exist multiple candidate causes of an outcome? For example, suppose in addition to prescribing an antihistamine the physician also prescribed a nasal steroid. In a demonstration of cue interaction using a simultaneous blocking paradigm, Baker and his colleagues (Baker et al., 1993) had subjects play a computer game in which their task was to successfully move a tank across a minefield. On each trial, subjects could apply camouflage to the tank, but they did not know whether the camouflage increased or decreased their chances of being detected by the mines. Additionally, a spotter plane occasionally flew over the minefield as the tank was traversing it, but the subjects did not know whether the spotter plane was an ally or an enemy. After many trials, subjects judged the effectiveness of the camouflage in helping the tank cross the minefield. Subjects' judgments of the camouflage depended on the contingency between the presence of the plane and success in crossing the minefield. When the plane's contingency was strong ($\Delta P = 1$), subjects' judgments of the camouflage's effectiveness were lower than when the plane's contingency was zero. Baker's results have been inter-

preted as evidence for cue competition; that is, when multiple potentially causal cues are present, they compete with one another for associative strength. As such, the presence of a strong alternative cue will reduce the causal judgments of a moderately contingent cue.

Cue Interaction Explained by Conditionalization

Proponents of statistical accounts of causal judgment have pointed out that when multiple candidate causes exist, as in the aforementioned case, ΔP is not the normative rule to apply (Cheng, 1997; Cheng, Park, Yarlus, & Holyoak, 1996). According to this account, when there are multiple potential causes of an effect, one should assess causality for each candidate cause conditional on the constant absence or constant presence of the alternative. This conditionalization is in effect what scientists do when attempting to determine the cause of an event—they control for alternative causes.

Controlling for alternative causes involves calculating ΔP in part: either across only those events in which the alternative cause was absent or across those in which it was present. We use the allergy example to illustrate conditionalization on the absence of the alternative.² Figure 1 depicts the outcomes of days on which the nasal steroid and/or the antihistamine were used. The proportions in the cells of Figure 1 represent the number of times allergy relief was experienced with that combination (numerator) over the number of times that combination was used (denominator). The proportions outside the right-hand side of the table show the unconditional contingency between antihistamine use and allergy relief, which was $.5$. However, if we consider only those days when the nasal steroid was not used (right side of table) and take the difference between the probability of allergy relief when the antihistamine was used and that when it was not, then the conditional contingency is zero; the antihistamine has no effect.

In this case of two candidate causes, there are three pieces of covariation information that are relevant: (1) the covariation between antihistamine use and allergy relief, (2) the covariation between nasal steroid use and allergy relief, and (3) the covariation between antihistamine use and nasal steroid use. When the two candidate causes covary, the conditional and unconditional contingencies will differ and an unconditional ΔP is not normative. Several researchers have pointed out that when the two cues are positively correlated, as in Baker's experiment (Baker et al., 1993), cue interaction effects reflect subjects' detection of the conditional contingency between each cue and the outcome (Busemeyer et al., 1993; Shanks, 1995; Spellman, 1996a, 1996b).

Indeed, Spellman (1996b) reanalyzed Baker et al.'s (1993) Experiment 1 to show how a conditionalization account could explain their cue interaction effect (see also Cheng, 1993, and Shanks, 1995). In both conditions, ΔP for the camouflage was $.5$, but the conditional contingency was not always equal to its unconditional ΔP . When the plane was not causal ($\Delta P = 0$), the cam-

		Nasal Steroid (Alternative)			
		Present	Absent		
Antihistamine (Target)	Present	$\frac{18}{18}$	$\frac{0}{6}$	$\frac{18}{24}$	Unconditional Contingency (UC) $18/24 - 6/24 = .5$
	Absent	$\frac{6}{6}$	$\frac{0}{18}$	$\frac{6}{24}$	

Conditional
Contingency (CC)
 $0/6 - 0/18 = 0$

Figure 1. Sample outcomes associated with antihistamine and nasal steroid use. Proportions in the cells represent the number of times the outcome occurred with that combination (numerator) over the number of times that combination was used (denominator). UC, unconditional contingency; CC, conditional contingency.

oufflage did have equal conditional and unconditional contingencies of .5. But, when the plane was a strong alternative cause ($\Delta P = 1$), the occurrence of the camouflage and the plane covaried, and therefore the camouflage had unequal conditional and unconditional contingencies. The unconditional contingency was .5, but the contingency for the camouflage conditional on the absence of the plane was 0. Therefore, a conditional contingency account of causal judgment would predict the results that Baker obtained: Subjects reduced their causal rating of the camouflage in the strong alternative condition relative to the non-causal alternative condition. Similar reanalyses of Price and Yates's (1993) Experiment 1 (Shanks, 1995; Spellman, 1996b) have shown how a conditional contingency account could explain the cue interaction reported in that experiment. Note that a number of associative accounts also make this same prediction at asymptote (e.g., Denniston, Savastano, & Miller, 2001; Rescorla & Wagner, 1972).

According to the conditionalization account, and explicitly predicted by Spellman (1996b), cue interaction should occur when, and only when, the two candidate causes covary (i.e., the unconditional contingency is not equal to the two identical conditional contingencies). A recent simultaneous blocking experiment by Tangen and Allan (2003) supports these predictions. In this experiment, subjects were shown multiple trials in which various combinations of two chemicals were applied to a petri dish containing bacteria. Subjects were to determine the effect of the chemicals on the bacteria's survival. Consistent with the predictions of Spellman (1996b), they found cue interaction when the two chemicals covaried, but not when the two chemicals were independent.

Nonnormative Discounting: Cue Interaction Not Explained by Conditionalization

Not all cue interaction effects, however, can be explained by the conditional contingency account. For example, in Baker et al.'s (1993) Experiments 1 and 2, the unconditional and conditional contingencies for the discounted cause were equal. In Experiment 1, the camouflage had a contingency of 0, the plane had either a contingency of 1 or a contingency of 0, and the camouflage and the plane were independent. Yet subjects rated the camouflage as less causal when the plane was a strong alternative cause than when it was not. Similar results were obtained in Experiment 2 when the contingency for camouflage was 0, the contingency for the plane was either .8 or 0, and the two cues were independent. In another set of experiments using these same contingencies, cue interaction effects were found among candidate causes using a noncausal cover story regarding the presence or absence of geometric figures (Vallée-Tourangeau et al., 1994). Additionally, when the levels of the cues differ in intensity rather than in a binary present or absent fashion, subjects demonstrate cue interaction effects between two independent cues (Bussemeyer et al., 1993). These results are not predicted by the conditionalization account, nor do they concur with asymptotic associative predictions. Thus, in simultaneous blocking paradigms, we call the cue interaction that remains after considering conditionalization *nonnormative discounting*.

In the present paper, we offer evidence that when subjects recognize the existence of a strong alternative cause, they exhibit nonnormative discounting in their causal ratings of a moderately contingent cue. We be-

lieve that this nonnormative discounting is not the result of competition during an associative learning process. We argue for a nonassociative account of nonnormative discounting for three reasons. First, as mentioned previously, the nonnormative discounting effect is inconsistent with most asymptotic associative predictions. Second, we demonstrate that nonnormative discounting is obtained not only when contingency information is encoded on a trial-by-trial basis but also when contingency information is presented to the subjects in a summary format—a situation to which the associative accounts do not apply. Finally, we show that when the importance of accurately assessing the effectiveness of the moderately contingent cue is increased, discounting decreases. This latter finding in particular implies that nonnormative discounting is not a necessary product of the contingency learning or causal judgment process, but is perhaps the result of a heuristic or an inference process that can be flexibly applied.

Our Experiments

In the following experiments, we demonstrate the independent existence of both conditionalization and nonnormative discounting (contradicting the prediction of Spellman, 1996b) using trial-by-trial and summary information encoding tasks (Experiment 1). Furthermore, we replicate Tangen and Allan's (2003) failure to find nonnormative discounting and demonstrate how an aspect of their causal cover story may underlie that finding (Experiments 2 and 3). Finally, we demonstrate that subjects can refrain from discounting when given a reason to fully evaluate the moderately contingent cue (Experiment 4). A summary of our procedures and results is shown in Table 1.

In all experiments, nonnormative discounting was assessed by comparing causal judgments across two conditions. In both conditions, the contingency for the target cause was .33 and the target and alternative causes were independent (i.e., the unconditional and conditional

contingencies were equal). What varied was the causal strength of the alternative: in the high-alternative independent (HA-Ind) condition it was .67; in the low-alternative independent (LA-Ind) condition, it was 0. If subjects rate the target cause lower when the alternative is highly contingent than when it is noncontingent, that would be evidence of nonnormative discounting.

EXPERIMENT 1

In Experiment 1, we wanted to examine whether conditionalization could explain all of the cue interaction we expected to find. To assess what could be explained by conditionalization, we had two "high-alternative" conditions (see Figure 2). In the HA-Independent (HA-Ind) condition, the unconditional and conditional contingencies for the alternative and the target cause were equal, with a contingency of .67 for the alternative and .33 for the target; in the HA-covariation (HA-Cov) condition, the alternative had an unconditional contingency of .67, but a conditional contingency of .80, and the target had an unconditional contingency of .33, but a conditional contingency of $-.20$. If the subjects rate the target as less causal in the covariation than in the independent condition, that would be evidence of conditionalization. To determine whether there were cue interaction effects that could not be explained by conditionalization, we assessed nonnormative discounting by comparing the two independent conditions as previously described.

Method

Subjects

One hundred twenty undergraduate students participated in partial fulfillment of a course requirement.

Design

We manipulated two between-subjects factors: encoding task (summary information vs. trial-by-trial) and the contingencies be-

Table 1
Summary of Conditions and Results (Based on Causal Ratings for the Target Cause) in Experiments 1–4

Experiment	Cover Story	Encoding	Comparison Results
1	Plants/fertilizer	Summary	Conditionalization: HA-Cov > HA-Ind Discounting: HA-Ind > LA-Ind
		Trial-by-trial	Conditionalization: HA-Cov > HA-Ind Discounting: HA-Ind > LA-Ind
2A	Patient/medicine	Summary	No Discounting
		Trial-by-trial	Discounting: HA-Ind > LA-Ind
2B	Bacteria/chemicals	Summary	Discounting: HA-Ind > LA-Ind
		Trial-by-trial	No Discounting
3	Modified Bacteria/chemicals	Trial-by-trial	Discounting: HA-Ind > LA-Ind
4	Patient/alt. cheap	Trial-by-trial	Discounting: HA-Ind > LA-Ind
	Patient/alt. expensive	Trial-by-trial	No Discounting

Note—In all experiments, the contingency for the target cause was .33. The contingency for the alternative cause was .67 in the HA-Ind condition and 0 in the LA-Ind condition. Experiment 1 is the only one with an HA-Cov condition. In that condition, the unconditional contingency for the target cause was .33 but the conditional contingency was $-.20$. Also in that condition, the unconditional contingency for the alternative cause was .67 but the conditional contingency was .80.

Figure 2. (A) Event proportions in the high-alternative independent (HA-Ind) condition. (B) Event proportions in the low-alternative independent (LA-Ind) condition. (C) Event proportions in the high-alternative covariation (HA-Cov) condition. Proportions in the cells represent the number of times the outcome occurred with that combination (numerator) over the number of times that combination was used (denominator). UC, unconditional contingency; CC, conditional contingency. These contingency conditions, LA-Ind, HA-Ind, and HA-Cov, are equivalent to the baseline, no discounting, and discounting conditions of Spellman (1996b) and the baseline, equal, and unequal conditions of Tangen and Allan (2003), respectively.

tween the causes and effect (low-alternative independent, LA-Ind; high-alternative independent, HA-Ind; and high-alternative covariation, HA-Cov). See Figure 2 for the contingencies. These contingency conditions, LA-Ind, HA-Ind, and HA-Cov, are equivalent to the baseline, no discounting, and discounting conditions of Spellman (1996b) and the baseline, equal, and unequal conditions of Tangen and Allan (2003), respectively.

Materials and Procedure

Subjects in the summary information condition read a cover story asking them to imagine that they were attempting to figure out if two liquids found in their landlady's garage were fertilizers or plant killers (see the Appendix for complete cover stories for all experiments). Subjects then read a summary account of the contingency information. For example, subjects in the LA-Ind condition read:

She has 72 plants in the greenhouse.

18 got both the RED and BLUE liquid. 6 bloomed.

18 got only the BLUE liquid. 0 bloomed.

18 got only the RED liquid. 6 bloomed.

18 got neither the RED liquid nor the BLUE liquid. 0 bloomed.

Subjects rated the effectiveness of each liquid on a scale from -100 to $+100$ (-100 = total flower inhibitor; $+100$ = total flower stimulator; 0 = no effect).

Subjects in the trial-by-trial condition were tested on Macintosh computers. They read the same cover story and were instructed on how to use the scale described above. They then completed six blocks of 12 prediction trials for a total of 72 trials. Each block of trials contained one sixth of the event frequencies depicted in Figure 2 so that the stated contingencies held for each block. Presentation of the trials was randomized within blocks. Each prediction trial began with a warning screen presented for 1,000 msec. Subjects then saw some combination of red and blue liquids pouring onto a plant without a bloom. This screen remained visible until subjects responded whether they thought the plant would bloom by hitting either the Y (yes) or N (no) key. After responding, subjects learned whether the plant had bloomed on that trial. This feedback was visi-

Table 2
Initial Causal Ratings (After 32 Trials) for Target and
Alternative in Trial-by-Trial Conditions of Experiments 1–4;
Means, With Standard Errors in Parentheses

Condition	Rating of Target	Rating of Alternative
Experiment 1		
HA-Cov	–36.8 (8.6)	81.8 (4.8)
HA-Ind	–2.4 (10.5)	43.8 (7.2)
LA-Ind	37.2 (9.9)	–45.8 (10.9)
Experiment 2A		
HA-Ind	10.4 (8.7)	51.2 (7.2)
LA-Ind	23.6 (6.9)	–35.0 (8.6)
Experiment 2B		
HA-Ind	–0.4 (12.4)	14.5 (12.9)
LA-Ind	–7.7 (16.7)	–15.6 (13.8)
Experiment 3		
HA-Ind	10.8 (9.6)	57.5 (9.1)
LA-Ind	20.4 (15.1)	1.3 (16.7)
Experiment 4		
Expensive		
HA-Ind	17.1 (7.3)	62.3 (5.9)
LA-Ind	31.4 (6.8)	–36.5 (9.1)
Same		
HA-Ind	9.7 (8.0)	50.2 (7.0)
LA-Ind	18.7 (5.8)	–37.5 (7.4)

Note—In all experiments, the contingency for the target cause was .33. The contingency for the alternative cause was .67 in the HA-Ind condition and 0 in the LA-Ind condition. Experiment 1 is the only one with an HA-Cov condition. In that condition, the unconditional contingency for the target cause was .33 but the conditional contingency was –.20. Also in that condition, the unconditional contingency for the alternative cause was .67 but the conditional contingency was .80.

ble for 2,500 msec. After the 36th and 72nd trials, subjects were reminded of how to use the scale and they made their ratings for each liquid separately by typing in a number between –100 and 100.

Results

Because we were primarily interested in directly comparing the trial-by-trial and summary encoding conditions and because subjects in the trial-by-trial conditions did not have the same amount of information as those in the summary until after 72 trials, the primary analyses and conclusions for all experiments are based on subjects' final causal ratings.³ We do, however, report briefly on the initial causal ratings in the trial-by-trial conditions. Table 2 contains the initial causal ratings made after 36 trials for subjects in the trial-by-trial encoding conditions for all experiments. In addition to an omnibus analysis of variance (ANOVA), for all experiments we carried out theoretically driven planned comparisons, which are depicted in Table 1.⁴

Target Liquid

Final causal ratings for the target liquid are critical to the assessment of conditionalization and nonnormative discounting. The ratings appear in Figure 3. An ANOVA on the final target ratings with encoding task (summary information, trial-by-trial) and contingency (HA-Cov, HA-Ind, LA-Ind) as between-subjects factors revealed a main effect of contingency [$F(1,114) = 53.319$, $MS_e =$

1,067, $p < .001$], no effect of task ($F < 1$), and no interaction [$F(2,114) = 1.95$, $p = .147$].

Assessing discounting: Independent conditions.

Comparison of the two independent conditions of Figure 3 indicates that subjects rated the target as less causal when there was a strong alternative cause present than when there was not. Indeed, planned comparisons revealed that subjects demonstrated nonnormative discounting in both the trial-by-trial [$t(38) = -2.17$, $p = .036$] and the summary [$t(38) = -3.84$, $p < .001$] encoding conditions. This pattern of nonnormative discounting was also observed in the initial trial-by-trial ratings [$t(38) = -2.74$, $p = .009$].

Assessing conditionalization: High-alternative conditions.

Comparison of the two high-alternative conditions of Figure 3 indicates that subjects rated the target as less causal in the covarying condition than in the independent condition. Planned comparisons confirmed that subjects' causal ratings varied with the conditional contingency in both the trial-by-trial [$t(38) = -3.59$, $p < .001$] and the summary [$t(38) = -4.63$, $p < .001$] encoding conditions. Conditionalization was also observed in the initial trial-by-trial ratings [$t(38) = -2.53$, $p = .016$].

Alternative Liquid

Although the causal ratings of the alternative liquid are not critical to the assessment of nonnormative discounting and conditionalization, they do bear on the interpretation of the ratings for the target liquid (i.e., we must ensure that subjects are sensitive to the varying contingency of the alternative). In both encoding tasks, subjects' ratings of the alternative liquid varied appropriately with the conditional contingency. An ANOVA on the ratings of the alternative with encoding task and contingency as between-subjects factors yielded an effect of contingency [$F(2,114) = 180.21$, $MS_e = 932$, $p < .01$], no effect of task [$F(1,114) = 1.22$, $p = .272$], and no interaction ($F < 1$). Subjects rated the alternative as more causal in the HA-Cov ($M = 78.8$, $SE = 4.3$) than in the HA-Ind ($M = 60.5$, $SE = 2.9$) condition [$t(78) = 3.53$, $p < .001$]. Likewise, the ratings in the HA-Ind condition were higher than those in the LA-Ind ($M = -41.5$, $SE = 6.6$) condition [$t(78) = 14.21$, $p < .001$]. The initial trial-by-trial ratings of the alternative showed the same pattern: Subjects rated the alternative as more causal in the HA-Cov condition than in the HA-Ind condition [$t(38) = 4.37$, $p < .001$], and more causal in the HA-Ind condition than in the LA-Ind condition [$t(38) = 6.86$, $p < .001$].

Summary

Experiment 1 demonstrates that even when subjects do not have statistical reason to reduce their causal judgments of a moderately contingent cue in the presence of a strongly contingent cue, they do so. This nonnormative discounting is evidenced not only when subjects encode information on a trial-by-trial basis, but also when they are given summary information regarding the contin-

Figure 3. Causal ratings for the target liquid in Experiment 1 (plant cover story). HA-Cov, high-alternative covarying condition; HA-Ind, high-alternative independent condition; LA-Ind, low-alternative independent condition. In both LA-Ind and HA-Ind conditions, the conditional contingency for the target was .33. In the HA-Cov condition it was $-.20$. Error bars represent $\pm 1 SE$.

gencies. As such, this nonnormative discounting does not appear to depend on a competitive associative learning process. The results of Experiment 1 are inconsistent with the predictions of Spellman (1996b) and inconsistent with the findings of Tangen and Allan (2003).

EXPERIMENT 2

Why were the results of Experiment 1 inconsistent with the findings of Tangen and Allan (2003)? We had two hypotheses: the different cover stories and the different procedures. Among the procedural differences were that Tangen and Allan manipulated contingency condition within subjects whereas we did it between subjects, and their subjects made causal judgments more frequently (after 36, 54, and 72 trials) than our subjects (after 36 and 72 trials). Therefore, in Experiment 2 we attempted to replicate our nonnormative discounting results with a new cover story (one in which the subjects learned about potential cures for a disease) and to replicate the findings of Tangen and Allan by using their bacteria cover story with our procedure.

Method

Subjects

One hundred seventy-nine undergraduates participated in partial fulfillment of a course requirement: Experiment 2A had 98; Experiment 2B had 81.

Design and Procedure

The design and procedure were the same as those of Experiment 1 with exceptions here. Subjects read one of two different

cover stories (Experiment 2A = patients/medicine; Experiment 2B = bacteria/chemicals). As in Experiment 1, some subjects got summary statistics, whereas others learned information trial by trial. Because we were primarily interested in nonnormative discounting, we excluded the HA-Cov condition, leaving only the HA-Ind and LA-Ind conditions.

In Experiment 2A, subjects were asked to evaluate the effectiveness of two potential cures for a deadly virus. In the summary conditions, subjects read information about how many patients with the deadly virus were cured after receiving various combinations of the medicines. In the trial-by-trial conditions, subjects saw trials depicting a patient with the disease who had received some combination of the two potential treatments. Subjects predicted whether the patient would be cured, and on each trial received feedback regarding whether that patient was cured.

In Experiment 2B, subjects were told that they were to determine the effect of two different chemicals on the survival of bacteria in a petri dish. (The cover story was based on Tangen & Allan, 2003). In the summary conditions, subjects read information about whether bacteria in petri dishes survived after having various combinations of the two chemicals added to them. Subjects in the trial-by-trial conditions saw trials depicting various combinations of the two chemicals added to different petri dishes containing the bacteria. Subjects predicted whether the bacteria would survive and on each trial received feedback regarding the bacteria's survival. The complete stories are in the Appendix.

Results: Experiment 2A Disease Cover Story

Causal Ratings of Target Cure

Consistent with Experiment 1, subjects overall discounted the target cure (Figure 4A). An ANOVA on the target ratings with encoding task (summary, trial-by-trial) and contingency (HA-Ind, LA-Ind) as between-subjects factors yielded a main effect of contingency [$F(1,94) = 6.56$, $MS_e = 576$, $p = .012$], a main effect of encoding task [$F(1,94) = 9.73$, $p = .002$], and no interaction ($F < 1$). The main effect of encoding task reflects the tendency for subjects in the trial-by-trial condition to give lower causal ratings overall. Planned comparisons revealed significant discounting with the trial-by-trial but not the summary encoding [$t(48) = -2.28$, $p = .027$, and $t(46) = -1.29$, $p = .203$, respectively]. The initial trial-by-trial ratings of the target did not show nonnormative discounting [$t(48) = -1.19$, $p = .242$].

Causal Ratings of Alternative Cure

Subjects' ratings of the alternative cure varied appropriately with the contingency of that cure. An ANOVA on the ratings of the alternative with encoding task and contingency as between-subjects factors yielded an effect of contingency [$F(1,94) = 147.39$, $MS_e = 1,162$, $p < .001$], an effect of encoding task [$F(1,94) = 4.35$, $p = .039$], and no interaction ($F < 1$). Once again, the main effect of encoding task reflects the tendency for subjects in the trial-by-trial condition to give lower causal ratings overall ($M = 27.4$, $SE = 7.2$, for the summary condition and $M = 13$, $SE = 8.1$, for the trial-by-trial condition). As expected, however, subjects' ratings of the alternative were lower in the LA-Ind condition ($M = -21.8$, $SE = 6.4$) than in the HA-Ind condition ($M = 61.9$, $SE = 2.7$). The initial trial-by-trial ratings

Figure 4. (A) Causal ratings for the target in Experiment 2A (disease cover story). (B) Causal ratings for the target in Experiment 2B (bacteria cover story). (C) Causal ratings for the target in Experiment 3 (modified bacteria cover story). HA-Ind, high-alternative independent condition; LA-Ind, low-alternative independent condition. In both LA-Ind and HA-Ind conditions, the conditional contingency for the target was .33. Error bars represent ± 1 SE.

reflected the same sensitivity to the contingency [$t(48) = 7.72, p < .001$].

Results: Experiment 2B Bacteria Cover Story

Causal Ratings of Target Chemical

Inconsistent with the results of Experiments 1 and 2A, subjects receiving summary information discounted and those receiving trial-by-trial information did not (Figure 4B). An ANOVA on the target ratings with encoding and contingency as between-subjects factors yielded no overall effect of contingency condition [$F(1,77) = 3.25, p = .075, MS_e = 1892$], no effect of task ($F < 1$), and no interaction [$F(1,77) = 2.10, p = .151$]. Further planned comparisons revealed that, as is apparent in Figure 4B, subjects in the summary information condition discounted [$t(38) = -2.52, p = .016$], whereas those in the trial-by-trial condition did not [$t(39) = -0.23, p = .819$]. Nor was there discounting apparent in the initial trial-by-trial ratings of the target [$t(39) = 0.23, p = .818$]. The lack of nonnormative discounting in the trial-by-trial condition replicates Tangen and Allan (2003).

Causal Ratings of Alternative Chemical

Subjects' ratings of the alternative chemical varied appropriately with the contingency of that chemical. An ANOVA on the ratings of the alternative with encoding task and contingency as between-subjects factors yielded an effect of contingency [$F(1,77) = 34.81, MS_e = 2,047, p < .001$], and no other effects ($F_s < 1$). As expected, subjects' ratings of the alternative were lower in the LA-Ind condition ($M = -20.8, SE = 11.9$ and $M = -24.2, SE = 8.6$, for the trial-by-trial and summary information conditions, respectively) than in the HA-Ind condition ($M = 28.6, SE = 11.2$ and $M = 45.0, SE = 7.9$, for the trial-by-trial and summary information conditions, respectively). However, the initial trial-by-trial ratings of

the HA-Ind and LA-Ind conditions did not differ [$t(39) = 1.38, p = .176$].

Summary

In Experiment 2A, we found nonnormative discounting using yet another cover story, that of evaluating two potential cures for a disease. Although this nonnormative discounting was reliable only in the trial-by-trial encoding conditions, there was an overall effect, indicating discounting across the two types of encoding tasks. Such discrepant results may reflect a less robust discounting effect with the disease cover story than with the plant cover story of Experiment 1. Indeed, it took more exposure to the task for nonnormative discounting to emerge in the trial-by-trial conditions of Experiment 2A (discounting was not present in the initial ratings of the target).

In Experiment 2B, we replicated Tangen and Allan's (2003) results using their bacteria cover story with our procedure; we, too, found an absence of nonnormative discounting in the trial-by-trial encoding condition. We did, however, find nonnormative discounting with the bacteria cover story in the summary condition (note that Tangen and Allan did not have a summary encoding condition). Thus, something must be different between the two cover stories that elicit nonnormative discounting (plants and disease) and the one that does not (bacteria). Because nonnormative discounting results from reducing the judgment of a moderately contingent cue in the presence of a strongly contingent cue, it is imperative that subjects recognize the strongly contingent cue as such. In our replication of Tangen and Allan using the bacteria cover story (Experiment 2B trial-by-trial encoding condition), the strong alternative, which had a contingency of .67, was rated only 28.6—far weaker than in the cover stories in which nonnormative discounting was found. With the other cover stories (trial-by-trial

versions), the strong alternative was rated 68.7 (plant) and 69.5 (disease). The failure to find nonnormative discounting with the bacteria cover story therefore may result from subjects' failure to recognize the strong alternative as such. Indeed, when given the bacteria cover story, subjects assigned less causality to either cue than they did with other cover stories (e.g., compare the magnitude of the causal ratings of the target cue in Figure 4A, the disease cover story, with that in Figure 4B, the bacteria cover story).

A structural difference in the causal scenarios becomes apparent when we take a closer look at the question that is being asked of subjects in each of these experiments. This structure is illustrated in Figure 5. On each trial in the plant story, subjects are initially shown a plant *without* a bloom and must answer the question of whether the plant *will* bloom, and then receive feedback about whether the plant bloomed. After many trials, subjects evaluate the effectiveness of each liquid on plant blooming. On each trial in the disease story, subjects are shown a person with the disease, make a judgment about whether the patient will be cured, and then receive feedback as to whether the patient was cured. In both the plant and disease stories, an increase in the positive effectiveness of the interventions (i.e., cues) is associated with a change in the status quo: A plant without a bloom now has a bloom; a person with a disease is now disease free. In the bacteria story, however, an increase in the positive effectiveness of the interventions is associated with a maintenance of the status quo. On each trial in the bacteria story, subjects are shown a petri dish *with* the

bacteria, make a judgment about whether the bacteria will survive (i.e., whether the petri dish will *continue* to have the bacteria), and then receive feedback about whether the bacteria did survive. This structure requires subjects to pay particular attention to what would have happened in the absence of either intervention (i.e., the base rate). In order to assign causality to either chemical in the bacteria cover story, one must recognize that in the absence of these cues, the bacteria would likely die. This type of base-rate information may be difficult for subjects to use, and it has been shown that subjects think it the least important piece of covariation information for assessing the contingency between events (Kao & Wasserman, 1993; Mandel & Lehman, 1998). Thus subjects reading the bacteria cover stories may have had difficulty assigning significant causality to either cue. According to this analysis, rephrasing the causal question posed to the subjects so that they must determine how effective each chemical is in *killing* the bacteria (i.e., changing the status quo) should lead to greater causal attribution overall, allowing subjects to recognize the strong alternative as strongly contingent. As a result, with the minor reframing of the causal question, we predict that subjects will discount.

EXPERIMENT 3

In Experiment 3, we assessed the possibility that the direction of causal judgment and the wording of the dependent variable caused the lack of nonnormative discounting in Tangen and Allan's (2003) results and our Experiment 2B. The experiment was identical to Experiment 2B,

Figure 5. Structure of the causal question in all experiments. Solid lines indicate maintenance of the status quo, and dashed lines a change from the status quo. Placement of the words *causality question* indicates the type of judgment subjects made regarding the interventions (i.e., whether the interventions helped maintain or helped change the status quo).

with the exception that subjects now rated the effectiveness of the chemicals in killing the bacteria. Experiment 3 used only a trial-by-trial encoding condition.

Method

Subjects

Thirty undergraduates participated in partial fulfillment of a course requirement.

Design and Procedure

We manipulated one between-subjects factor, the contingency between the cause and effect (HA-Ind vs. LA-Ind). Subjects were told that they were to determine whether each of two different chemicals was effective in killing bacteria in a petri dish. Subjects encoded the contingency information on a trial-by-trial basis. On each trial, subjects saw various combinations of the two chemicals added to different petri dishes containing bacteria. Subjects predicted whether the bacteria would die, and on each trial received feedback regarding whether the bacteria died. Subjects rated the effectiveness of each chemical in killing the bacteria.

Results

As predicted, subjects' causal ratings of the target chemical revealed nonnormative discounting (see Figure 4C). Subjects rated the target as less causal in the HA-Ind condition than in the LA-Ind condition [$t(28) = 2.34, p = .024$]. This difference was not apparent in the initial trial-by-trial ratings of the target [$t(29) = 0.56, p = .579$]. As one might expect given the finding of nonnormative discounting, subjects appropriately rated the alternative chemical as more causal in the HA-Ind condition ($M = 54.2, SE = 7.2$) than in the LA-Ind condition ($M = -33.8, SE = 13.4$) [$t(28) = -6.27, p < .001$]. This same pattern was seen in the initial trial-by-trial ratings of the alternative [$t(28) = -3.20, p = .003$].

Summary

The results of Experiment 3 confirmed our hypothesis that Tangen and Allan's (2003) and our own failure to find nonnormative discounting using the bacteria cover story was the result of the structure of the causal scenario. In those experiments, subjects were asked if an intervention would maintain the current status of bacteria in a petri dish and subjects overall attributed little causal effectiveness to either the strong alternative or the target cue. In the present experiment, when subjects were asked to determine whether an intervention effects a *change* in the current status of bacteria in a petri dish, they recognized the strong alternative as such and, as a result, discounted their causal ratings of the moderately contingent target.⁵

EXPERIMENT 4

In our final experiment, we sought to determine whether nonnormative discounting is a necessary product of the contingency learning or causal reasoning process, or whether it can be flexibly applied. We examined whether subjects would no longer discount when given a reason to more accurately assess the moderately contingent cue.

We asked subjects to evaluate the effectiveness of two different cures in treating a deadly virus. The cover story was similar to that of Experiment 2A (patient/medicines), except that we manipulated the cost of producing the strong alternative. In one condition, the strong alternative cure was very expensive to produce and the target cure very cheap to produce. In the other condition, both the cures were very cheap to produce. We expected that if subjects could flexibly apply discounting, then they would no longer discount the target cure when the strong alternative cure was very expensive to produce.

Method

Subjects

One hundred twelve undergraduates participated in partial fulfillment of a course requirement.

Design

We manipulated two between-subjects factors: the relative cost of producing the strongly contingent alternative medication (expensive, same) and the contingency between the cues and outcome (HA-Ind, LA-Ind).

Procedure

As in Experiment 2A, subjects were told that their task was to determine the effectiveness of two potential cures of a deadly virus. Additionally, they were told that they were under some pressure from the World Health Organization to produce an inexpensive cure. Subjects in the "expensive" condition were further told that the one of the potential cures cost \$3,000 and the other \$20 to produce a 30-day treatment program. Subjects in the "same" condition were told that both potential cures cost \$20 to produce. Subjects then encoded the contingency information on a trial-by-trial basis and made causal ratings as in Experiment 2A.

Results

Causal Ratings of Target Cure

As depicted in Figure 6, the presence of nonnormative discounting in the final ratings of the target varied with the expense of the alternative. The ANOVA with contingency condition (HA-Ind, LA-Ind) and expense of alternative (expensive, same) as factors yielded no effect of contingency condition [$F(1,108) = 3.32, MS_e = 999, p = .071$] and no other effects ($F_s < 1$). Planned comparisons within each level of expense revealed that when the alternative was expensive relative to the target, subjects did not significantly discount; that is, they rated the target similarly in the HA-Ind and LA-Ind conditions [$t(49) = -.52, p = .604$]. However, when the alternative and target cost the same, subjects discounted. They rated the target lower in the HA-Ind condition than in the LA-Ind condition [$t(59) = -2.21, p = .031$]. This same pattern did not emerge in the initial trial-by-trial ratings of the target; in both expensive and same conditions, subjects did not discount [$t(49) = -1.43, p = .159$, and $t(59) = -.93, p = .357$, respectively].

Causal Ratings of Alternative Cure

An ANOVA on the causal ratings of the alternative revealed an effect of contingency condition [$F(1,108) = 143.15, MS_e = 1,526, p < .001$] and no other effects

Figure 6. Causal ratings for the target in Experiment 4 (expense of alternative manipulation). HA-Ind, high-alternative independent condition; LA-Ind, low-alternative independent condition. In both LA-Ind and HA-Ind conditions, the conditional contingency for the target was .33. Error bars represent ± 1 SE.

($F_s < 1$). Subjects appropriately gave the alternative higher causal ratings in the HA-Ind conditions ($M = 53.6$, $SE = 3.8$) than in the LA-Ind conditions ($M = -35.2$, $SE = 5.9$). Subjects' initial trial-by-trial ratings showed this same sensitivity to the contingency of the alternative [$t(110) = 12.42$, $p < .001$].

Summary

Experiment 4 demonstrates that when given a reason to more accurately assess the moderately contingent cue, subjects refrain from discounting. Whereas varying the expense of the alternative affected the causal ratings of the target, it did not affect the causal ratings of the alternative itself. These results suggest that the lack of discounting when the alternative was expensive was not the result of ignoring or discounting that expensive alternative. Nonnormative discounting effects are therefore not the necessary product of a competitive learning process.

GENERAL DISCUSSION

In the present set of experiments, we distinguished two types of cue interaction in simultaneous blocking paradigms: the process of controlling for alternative causes (i.e., conditionalization) and nonnormative discounting. We demonstrated that when subjects recognize the presence of a strongly contingent alternative cause, they exhibit nonnormative discounting. This nonnormative discounting is, however, sensitive to the structure of the causal cover story. Finally, we demonstrated that nonnormative discounting is not the necessary product of a contingency learning or causal judgment process,

because subjects refrain from discounting when motivated to more accurately evaluate the target cause.

This nonnormative discounting effect is a reliable one. We expected to see discounting in the Experiment 1 summary and trial-by-trial conditions; Experiment 2A summary and trial-by-trial conditions; Experiment 3; and Experiment 4 cheap condition. In five of these six theoretically driven comparisons in which we expected to find discounting, we did indeed find that subjects rated the target significantly lower in the HA-Ind condition than in the LA-Ind condition (Table 1).

Our nonnormative discounting results are consistent with a number of other findings in the literature in which subjects demonstrate cue interaction effects in situations in which it is not statistically normative (Baker et al., 1993; Busemeyer et al., 1993; Price & Yates, 1993; Vallée-Tourangeau et al., 1994). Existing theories of causal reasoning do not predict this kind of cue interaction; in fact, both the traditional Rescorla-Wagner associative account and the conditional contingency account predict that subjects' causal ratings should follow the pattern of conditional contingencies (i.e., they should not exhibit nonnormative discounting).⁶

Our results suggest that nonnormative discounting is not the result of a competitive learning process, since subjects do not necessarily discount. Likewise, although we make no claim about the algorithm that might compute the discounted contingencies themselves, because it is seen in both trial-by-trial and summary information encoding tasks, the nonnormative discounting effect is probably not the result of an associative learning process.

Confounds in Contingency Construction

Constructing contingencies is complicated in that many factors one might like to vary, or hold constant, are not independent of each other. We see two minor confounds relevant to our design and conclusions. First, a confound arises from differences in the number of cue–outcome pairings in the low- and high-alternative conditions. In the LA-Ind condition, the outcome never occurs in the absence of the target. In the HA-Ind condition, however, the outcome occurs in the absence of the target one third of the time. Such a difference may affect subjects' perceptions of the causal role of the target. In particular, they may be more confident of its causal role in the LA-Ind condition than in the HA-Ind condition, and this may translate into higher causal ratings in the LA-Ind condition. This difference, however, cannot explain all nonnormative discounting effects. For example, Baker et al. (1993, Experiment 1) found nonnormative discounting even though the outcome took place in the absence of the target in both low- and high-alternative conditions.

A second confound arises from the difference in the outcome density (i.e., the probability of the occurrence of the outcome) of our low- and high-alternative independent conditions. The probability of the outcome occurring in the LA-Ind condition is 0.17 and in the HA-Ind condition it is .50. Others have shown human causal

judgments to be sensitive to outcome density (Baker, Berbrrier, & Vallée-Tourangeau, 1989; Shanks, 1985). However, differences in outcome density typically produce patterns in causal judgments opposite that which we observed (i.e., higher outcome densities are associated with an increase in subjects' causal judgments). Therefore, although there are differences in the outcome density between the two independent conditions, the nonnormative discounting results cannot be explained by such differences.

Reasons for Nonnormative Discounting

We acknowledge two possibilities for why subjects demonstrate nonnormative discounting. One possibility is that nonnormative discounting is the result of some sort of general cognitive comparison process. For example, one could imagine that merely by virtue of being in the presence of something large, a moderately sized item might suddenly look smaller. Such context effects are seen in judgments about the size of animals (Cech & Shoben, 1985). Discounting may thus be a general cognitive phenomenon, one not unique to causal reasoning tasks. Nonnormative discounting has been demonstrated in a covariation-detection task that did not have an apparent causal cover story. Subjects judged whether a triangle, a circle, and a square appeared together. A strong relationship between the appearance of the triangle and square reduced subjects' estimates that the circle and square appeared together. This effect, however, was small in contrast to those seen with causal cover stories (Vallée-Tourangeau et al., 1994). Additionally, if nonnormative discounting were a general cognitive comparison or perceptual phenomenon, one would expect it to be apparent in frequency estimates; however, such effects do not occur (Price & Yates, 1995). Evidence as to whether discounting is the product of a more general cognitive comparison phenomenon is fairly ambiguous at this point.

A second possible explanation for nonnormative discounting is that it is the result of some sort of causal metabelief (a term introduced by Shanks, 1991) similar to the deductive reasoning account of blocking suggested by De Houwer and Beckers (2003). Causal inference is likely the result of an interaction between subjects' beliefs about causation and the actual covariation between the target cue and effect (Alloy & Tabachnik, 1984). This notion is consistent with the causal-model theory of Waldmann and his colleagues. According to causal-model theory, causal knowledge is acquired through an interaction of top-down and bottom-up processes. Waldmann and his colleagues have shown how in a variety of situations, preexisting knowledge about the particular causal situation, as well as about causal situations in general, determines whether or not subjects infer causation from covariation information (Waldmann, 1996; Waldmann & Hagmayer, 1995, 2001; Waldmann & Holyoak, 1990, 1992).

Causal metabeliefs, as well as domain-specific knowledge and motivation, appear to determine how humans process covariation information. These causal beliefs will determine over what alternative causes subjects condition-

alize. For example, as demonstrated in various "Simpson's paradox" experiments, subjects will conditionalize only on events believed to be causally relevant (Spellman, Price, & Logan, 2001; Waldmann & Hagmayer, 1995, 2001). Subjects are also more likely to conditionalize when motivated to do so (Schaller, 1992). The nonnormative discounting results in the present experiments may be due to another type of causal metaknowledge—that is, a belief that when there is a strong cause present, there are not likely to be other causally relevant factors. However, when motivated to look, people will do so.

Although the use of a discounting metabelief may look nonnormative in the present procedure, given that this belief is derived from causal situations in general or judgment situations in general, it may in fact be a rational inference (Macrae, Milne, & Bodenhausen, 1994). For the case of causal attribution, Morris and Larrick (1995) have shown how discounting is a normative inference from most situations in which there are multiple possible causes. Additionally, in many situations it may be perfectly satisfactory to find only one cause. To return to the allergy example in Figure 1, if a physician prescribed both an antihistamine and a nasal steroid for your allergies, you would need to find only one causal agent in order to get allergy relief. Nonnormative discounting, therefore, may reflect a *rational* inference given the pragmatics of everyday causal reasoning.

Conclusion

Whereas much recent research has focused on conditionalization as a normative explanation for cue interaction effects, in the present paper we demonstrate that cue interaction may exist even when it is not statistically normative and that neither normative statistical models nor traditional associative models readily account for this nonnormative discounting. Nonnormative discounting exists beyond conditionalization; it occurs for both summary and trial-by-trial information; its use depends on the wording of causal cover story; and it may be flexibly applied (depending on motivation). Thus, our results suggest that people do control for alternative causes; that is, they act like intuitive scientists, but people also act rationally "irrational" by discounting moderately contingent causes in the presence of highly contingent alternatives.

REFERENCES

- AHN, W., KALISH, C. W., MEDIN, D. L., & GELMAN, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299-352.
- 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.
- ALLOY, L. B., & TABACHNIK, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, *91*, 112-149.
- BAKER, A. G., BERBRIER, M. W., & VALLÉE-TOURANGEAU, F. (1989). Judgements of a 2×2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology*, *41B*, 65-97.
- BAKER, A. G., MERCIER, P., VALLÉE-TOURANGEAU, F., FRANK, R., & PAN, M. (1993). Selective associations and causality judgments:

- Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 414-432.
- BUSEMEYER, J. R., MYUNG, J., & MCDANIEL, M. A. (1993). Cue competition effects: Empirical tests of adaptive network learning models. *Psychological Science*, **4**, 190-195.
- CECH, C. G., & SHOBEN, E. J. (1985). Context effects in symbolic magnitude comparisons. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **11**, 299-315.
- CHAPMAN, G. B., & ROBBINS, S. I. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, **18**, 537-545.
- CHENG, P. W. (1993). Separating causal laws from causal facts: Pressing the limits of statistical relevance. In D. Medin (Ed.), *The psychology of learning and motivation: Vol. 30* (pp. 215-264). San Diego: Academic Press.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, **104**, 367-405.
- CHENG, P. W., & NOVICK, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality & Social Psychology*, **58**, 545-567.
- CHENG, P. W., PARK, J., YARLAS, A. S., & HOLYOAK, K. J. (1996). A causal-power theory of focal sets. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 313-355). San Diego: Academic Press.
- DE HOUWER, J., & BECKERS, T. (2002). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology*, **53B**, 289-310.
- DE HOUWER, J., & BECKERS, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology*, **56B**, 345-357.
- DENNISTON, J. C., SAVASTANO, H. I., & MILLER, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65-117). Mahwah, NJ: Erlbaum.
- FUGELSANG, J. A., & THOMPSON, V. A. (2001). Belief-based and covariation-based cues affect causal discounting. *Canadian Journal of Experimental Psychology*, **55**, 70-76.
- HANSEN, R. D., & HALL, C. A. (1985). Discounting and augmenting facilitative and inhibitory forces: The winner takes almost all. *Journal of Personality & Social Psychology*, **49**, 1482-1493.
- KAO, S., & WASSERMAN, E. A. (1993). Assessment of an information integration account of contingency judgment with examination of subjective cell importance and method of information presentation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 1363-1386.
- KELLEY, H. H. (1972a). Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 1-26). Morristown, NJ: General Learning Press.
- KELLEY, H. H. (1972b). Causal schemata and the attribution process. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 151-174). Morristown, NJ: General Learning Press.
- KEPPEL, G. (1991). *Design and analysis* (3rd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- MACRAE, C. N., MILNE, A. B., & BODENHAUSEN, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality & Social Psychology*, **66**, 37-47.
- MANDEL, D. R., & LEHMAN, D. R. (1998). Integration of contingency information in judgments of cause, covariation, and probability. *Journal of Experimental Psychology: General*, **127**, 269-285.
- MORRIS, M. W., & LARRICK, R. P. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review*, **102**, 331-355.
- PRICE, P. C., & YATES, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition*, **21**, 561-572.
- PRICE, P. C., & YATES, J. F. (1995). Associative and rule-based accounts of cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 1639-1655.
- 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 research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- SCHALLER, M. (1992). In-group favoritism and statistical reasoning in social inference: Implications for formation and maintenance of group stereotypes. *Journal of Personality & Social Psychology*, **63**, 61-74.
- SHANKS, D. R. (1985). Continuous monitoring of human contingency judgments across trials. *Memory & Cognition*, **13**, 158-167.
- SHANKS, D. R. (1991). On similarities between causal judgments in experienced and described situations. *Psychological Science*, **2**, 341-350.
- SHANKS, D. R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology*, **48A**, 257-279.
- SHANKS, D. R., HOLYOAK, K. J., & MEDIN, D. L. (Eds.) (1996). *The psychology of learning and motivation. Vol. 34: Causal learning*. San Diego: Academic Press.
- SPELLMAN, B. A. (1996a). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, **7**, 337-342.
- SPELLMAN, B. A. (1996b). 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.
- SPELLMAN, B. A., PRICE, C. M., & LOGAN, J. (2001). How two causes are different from one: The use of (un)conditional information in Simpson's paradox. *Memory & Cognition*, **29**, 193-208.
- TANGEN, J. M., & ALLAN, L. G. (2003). The relative effect of cue interaction. *Quarterly Journal of Experimental Psychology*, **56B**, 279-300.
- VALLÉE-TOURANGEAU, F., BAKER, A. G., & MERCIER, P. (1994). Discounting in causality and covariation judgments. *Quarterly Journal of Experimental Psychology*, **47B**, 151-171.
- WALDMANN, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 47-88). San Diego: Academic Press.
- WALDMANN, M. R., & HAGMAYER, Y. (1995). When a cause simultaneously produces and prevents an effect. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 425-430). Hillsdale, NJ: Erlbaum.
- WALDMANN, M. R., & HAGMAYER, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition*, **82**, 27-58.
- WALDMANN, M. R., & HOLYOAK, K. J. (1990). Can causal induction be reduced to associative learning? In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 190-197). Hillsdale, NJ: Erlbaum.
- WALDMANN, M. R., & HOLYOAK, K. J. (1992). Predictive and diagnostic learning in cue competition. *Journal of Experimental Psychology: General*, **121**, 222-236.

NOTES

1. The term *discounting* was chosen because the nonnormative cue interaction we describe in this paper fits the use of this term in the causal attribution literature. Kelley (1972a, 1972b) introduced two different uses of the term, defining the discounting principle as follows:

The role of a given cause in producing a given effect is discounted if other plausible causes are also present. This discounting is reflected in various ways. The attributor is less confident that the observed effect reflects the given cause. He is less willing to infer that the magnitude of the given cause is as great as might otherwise be indicated by the magnitude of the observed effect. (1972a, p. 8)

Thus, one use of the term *discounting* is in judging whether a target cause was *present* given the existence of alternative causes (Kelley, 1972b). Discounting occurs when one becomes less confident that Cause X was present when told that Cause Y was also present (see Morris & Larrick, 1995, for a review). The second use of the term *discounting* is in judging a target cause's *strength* given the existence of alternative causes (Kelley, 1972a). Discounting occurs when one judges Cause X as less causal when it is learned about in the presence of

Cause Y (e.g., Hansen & Hall, 1985). Our use of the term *discounting* is in accord with this second usage, which we believe also reflects its common English meaning.

2. One can also conditionalize on the presence of the alternative. To do so using the allergy example, one would evaluate ΔP for the target only in those cases in which the nasal steroid was present. Evaluating the left-hand column of the contingency table in Figure 1, we calculate $18/18 - 6/6 = 0$. In this instance, conditionalizing on the presence of the alternative yields the same value as conditionalizing on the absence of the alternative. This equality need not hold. However, it does so in all experiments reported in this paper.

3. One can also calculate derived causal ratings on the basis of the predictions subjects make on each trial (e.g., see Tangen & Allan, 2003), although such calculations require the assumption that subjects are using a probability-matching strategy. Results based on the derived ratings for Experiments 2–4 are available from K.M.G.

4. Using the approach recommended by Keppel (1991, pp. 177–180), we perform a limited number of planned comparisons and control the per-comparison error rate rather than the experiment-wise error rate. We acknowledge that there are true theoretical divides concerning how to deal with multiple comparisons. The arbitrariness involved in both defining a family over which to control Type I error and the decision to control Type I error at the sacrifice of increasing Type II error have led us to not correct for multiple comparisons. Although our own conclu-

sions are based on the uncorrected alpha values, we report exact p values to 3 decimal places for all of our comparisons. As such, if a reader wished to correct for multiple comparisons and apply a Bonferroni correction, he/she need only divide alpha by the number of comparisons and compare our p values to this conservative criterion.

5. One might note that in the Baker et al. (1993) tank experiments in which nonnormative discounting was observed, an increase in the positive effectiveness of the interventions (camouflage and spotter plane) was associated with a maintenance in the status quo much like the original Tangen and Allan (2003) bacteria story. In the Baker experiments, interventions work when the intact tank successfully crosses the minefield (i.e., still intact). The important difference between the original bacteria story and the tank experiment is that there is a causal mechanism for tanks exploding in the absence of any treatment: mines. There is not, however, a causal mechanism for bacteria dying in the petri dish in the absence of any treatment. This difference may lead subjects to make assumptions about what would have happened in the absence of any treatment: exploding tanks (change) and living bacteria (status quo). A priori beliefs about causal mechanisms are known to influence causal judgments (e.g., Ahn, Kalish, Medin, & Gelman, 1995; Fugelsang & Thompson, 2001).

6. The predictions of these two accounts converge on the conditional contingency only when the two conditional contingencies (i.e., that based on the absence and that based on the presence of the alternative cause) are equal. Such was the case in the present experiments.

APPENDIX

Cover Stories for All Experiments

Experiment 1: Plant Cover Story

While looking through the garage of the house you have just rented, you find some very interesting-looking containers of liquid. Your landlady tells you that some of them are very expensive plant-treatment liquids and some of them are just colored water. Of the plant treatment liquids, she remembers that some of them are flower-growth stimulators (fertilizers) and some are flower-growth inhibitors and that the liquids came in various strengths—but she does not remember which liquid is which. She does want you to find out, however, and is willing to reduce your rent if you can prove to her that you can distinguish them. You can do so by pouring the liquids on various plants in various combinations and then accurately predicting whether or not the plant will produce a flower.

Experiment 2A: Disease Cover Story

The World Health Organization (WHO) has appointed you head of a project to track down a cure for a deadly virus that has recently cropped up. You are currently evaluating two potential cures for the virus nicknamed Clear and Cloudy for how they look in liquid-treatment form. You have the opportunity to determine the effectiveness of the potential treatments by monitoring the outcomes of patients treated with various combinations of the treatments.

Experiment 2B: Bacteria Cover Story (after Tangen & Allan, 2003)

Scientists have recently discovered a new strain of bacteria in the human digestive system. Scientists are interested in discovering whether certain chemicals affect the survival of the bacteria. As a member of the National Science Foundation (NSF), you have been charged with this investigation. Currently, you are interested in testing two different chemicals that you have nicknamed PURPLE and RED for how their molecules look when stained. You will have the opportunity to see the results of trials in which these chemicals have been added in various combinations to petri dishes containing the bacteria.

Experiment 3: Modified Bacteria Cover Story

Scientists have recently discovered a new strain of bacteria in the human digestive system. Scientists are interested in discovering whether certain chemicals are effective in eliminating (i.e., killing) the bacteria. As a member of the National Science Foundation (NSF), you have been charged with this investigation. Currently, you are interested in testing two different chemicals that you have nicknamed PURPLE and RED for how their molecules look when stained. You will have the opportunity to see the results of trials in which these chemicals have been added in various combinations to petri dishes containing the bacteria.

Experiment 4: Disease Cover Story With Cost Manipulation

The World Health Organization (WHO) has appointed you head of a project to track down a cure for a deadly virus that has recently cropped up. You are currently evaluating two potential cures for the virus nicknamed Clear and Cloudy for how they look in liquid-treatment form. You are under some pressure from the government to find a cure that could be produced very cheaply. The Clear liquid is very expensive to produce, costing well over \$3,000 (U.S.) for a 30-day treatment program. The Cloudy liquid is very inexpensive to produce, costing only \$20 for a 30-day treatment program. You have the opportunity to determine the effectiveness of the potential treatments by monitoring the outcomes of patients treated with various combinations of the treatments.