

Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships

MARTIN J. DENNIS

Yale University, New Haven, Connecticut

and

WOO-KYOUNG AHN

Vanderbilt University, Nashville, Tennessee

The order in which people receive information has a substantial effect on subsequent judgment and inference. Our focus is on the order of covariation evidence in causal learning. The first experiment shows that the initial presentation of evidence suggesting a generative causal relationship (the joint presence or joint absence of cause and effect) leads to higher judged causal strength than does the initial presentation of evidence suggesting an inhibitory relationship (the presence of cause or effect in the absence of the other). Additional studies show that this primacy effect is unlikely to be due to fatigue or to an insufficient number of learning trials. These results are not readily explained by current contingency-based or associative theories of causal induction.

Learning causal relations between two events is a fundamental cognitive activity. For instance, explorers in a new country might observe that whenever they eat an exotic plant, they get sick. Subsequently, on the basis of the observation of the covariation, the reasoners may draw the causal conclusion that eating the exotic plant caused them to get sick. Although there are a number of ways of learning causal relations (e.g., Ahn & Kalish, 2000), our present focus is on causal induction based on the observation of covariation between two events.

We can imagine a simple case of covariation wherein a possible cause is either present or absent, and the target effect to be explained is either present or absent, as shown in Figure 1. One way to define covariation between two factors is to calculate an index $\Delta P = P(E|C) - P(E|\neg C)$: namely, the difference between the probability that the effect occurs given that the cause is present and the probability that the effect occurs given that the cause is absent (Jenkins & Ward, 1965; Ward & Jenkins, 1965). Numerous studies have demonstrated positive correlations between objective ΔP and the perceived covariation between events (Allan & Jenkins, 1980; Alloy & Abramson,

1979; Wasserman, Chatlosh, & Neunaber, 1983; cf. Shanks & Dickinson, 1987).

In everyday life, we rarely encounter events in a simultaneous manner as shown in Figure 1. Often, each event is experienced sequentially. Likewise, in most studies of causal induction, covariation information has been presented sequentially (e.g., Chapman, 1991; López, Shanks, Almaraz, & Fernández, 1998; Van Hamme & Wasserman, 1994). However, these studies examined only limited cases of real-life causal induction, in that the different types of evidence were evenly distributed throughout all of the learning trials, whereas the chances that the sequential information is perfectly evenly distributed in real-life situations are low.¹ Thus, an important question arises: Does the presentation order of the four types of evidence that determine covariation affect judgments of causal strength?

Order Effect In Judgments of Covariation and Causal Relations

It has long been known that the presentation order of stimuli, such as objects, persons, or events, can affect judgments about those items. For example, the presentation order of personality trait adjectives will bias one's impression of a person (Anderson & Jacobson, 1965; Asch, 1946; cf. Anderson, 1981, chap. 3). The first adjectives most affect our judgments. For instance, one will give more positive ratings of a target (e.g., a person) if the target is described with positive adjectives first (e.g., "self-disciplined, logical, intelligent, gloomy, cynical, moody") than if negative adjectives are presented first (Hendrick & Costantini, 1970). These findings exemplify *primacy effects*, in that initial information has the greatest impact on later judgments.

Support for this research was provided in part by National Institute of Health Grant NIH R01-MH57737 to W.-K.A. The experiments reported here were completed as part of the first author's dissertation research. The authors gratefully acknowledge the assistance of Yani Indrajana, Nancy Kim, Jennifer Amsterlaw, Heidi Wenk, and Becca Shaffer in the collection of data. The authors also thank Heidi Wenk and Paul Price for comments on a draft of this paper. Correspondence should be addressed to M. J. Dennis, Madsen Center, Augustana College, 2001 S. Summit Ave., Sioux Falls, SD 57197 (e-mail: dennis@inst.augie.edu).

		Effect	
		Present	Absent
Cause	Present	A	B
	Absent	C	D

Figure 1. Example contingency between a cause and an effect. Letters represent the names of the different types of evidence.

Although it seems reasonable to conjecture that primacy effects should hold in judgments of the causal relations between two events, only a few studies have systematically examined these effects in causal induction (Chapman, 1991; López et al., 1998; Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996). Furthermore, the available studies have produced conflicting results, including recency effects (López et al., 1998), primacy effects (Yates & Curley, 1986), and no effects (Wasserman et al., 1996).

For instance, López et al. (1998) used learning sequences in which participants received information about an outcome (in this case, a disease, X) and several cues (or possible symptoms). In one half of the learning sequence (the contingent block), one of these symptoms (A) was always paired with another (B). In general, when the compound AB was presented, the patient had the disease, but when symptom B occurred alone, it usually was not accompanied by the disease. This pairing suggests A was a better predictor of the disease than was B. In the other half of the sequence (the noncontingent block), symptom A was paired with a new cue (C). In this block, however, the disease mainly occurred with C alone and not with the compound AC, suggesting that symptom A was a worse predictor than was C. The order of these two blocks was manipulated so that, in one condition, the contingent block was presented first and, in another condition, the noncontingent block was first. López et al. found that ratings of the relationship between symptom A and the disease were higher when the contingent block was last (hence, a recency effect).

In contrast, Yates and Curley (1986) found a primacy effect in a covariation judgment task. Participants in this study were asked to consider a covariation (rather than a causal relation) between a hypothetical plant's color (light or dark) and the region it came from (A or B). They found that presentation of evidence suggesting positive covariation at the beginning of a sequence led to significantly higher covariation judgments than did presen-

tation of evidence suggesting negative covariation at the beginning of a sequence. Given previous studies demonstrating positive correlations between objective ΔP and the perceived strength of a causal relationship (e.g., Wasserman et al., 1983), it seems reasonable to expect that the same kind of primacy effect would be obtained in judging causal strengths. If so, Yates and Curley's (1986) results would be at variance with the recency effect found in López et al. (1998).

One obvious difference between the design of the experiments of López et al. (1998) and those of Yates and Curley (1986) is the use of multiple cues in the latter and the use of a single cue in the former, although it is not clear why this design difference explains the difference in results. More generally, there are several difficulties in directly generalizing Yates and Curley's results to causal induction. The first problem is that Yates and Curley's task involved covariation, not causal, judgment. In addition, the primacy effect was obtained from an indirect measure. That is, participants were asked to give estimates of the conditional probabilities, $P(\text{Light} | \text{Region A})$ and $P(\text{Light} | \text{Region B})$, and ΔP s were inferred by subtracting the second probability from the first one. Thus, the results might not be comparable to other findings from causal induction tasks, which tend to use a more direct assessment of causal strength (e.g., the degree to which one event causes another; Matute, Arcediano, & Miller, 1996).

Wasserman et al. (1996) used a direct causal strength judgment to examine the effect of presenting positive covariation first versus presenting negative covariation first. Unlike Yates and Curley (1986), however, Wasserman et al. (1996) found no primacy effect across two experiments, although the direction of the means was consistent with the primacy effect found by Yates and Curley. López et al. (1998) suggested that the lack of order effect in Wasserman et al. might have been due to the fact that they did not use enough trials to find the effect (i.e., total of 24 trials in Experiments 2 and 3 in Wasserman et al., 1996).

To summarize, although the recency effects found by López et al. (1998) seem to contradict the primacy effect found by Yates and Curley (1986), the two studies are not directly comparable due to differences in design and task. The first order of business, therefore, is to show that a primacy effect can indeed reliably occur in a causal induction task. Thus, the main purpose of the present study was to investigate order effects, using Yates and Curley's single-cue task, in which covariation between just two events is presented.

Overview of Methods and Hypotheses

As in Yates and Curley's (1986) study, the present experiments manipulated the order in which evidence supporting positive covariation and evidence supporting

negative covariation are presented. Unlike in Yates and Curley's study, our participants were asked to make judgments of causal strengths between two events.

The specific manipulation of order in the present experiments can best be explained by referring to the four cells in Figure 1. Cases in cells A and D serve to confirm that a generative causal relationship exists between two factors—that is, that one event gives rise to the other. If, of course, an effect occurs without a cause (cell C) or an effect does not occur given a cause (cell B), then that evidence helps to disconfirm the generative causal relationship. Conversely, the frequencies in cells B and C help to confirm an inhibitory relationship, whereas the frequencies in cells A and D disconfirm such a relationship. Of particular interest in the present study is the effect of the order in which people observe cases that confirm a generative causal relationship (cells A and D in Figure 1) versus those that confirm an inhibitory relationship (cells B and C in Figure 1).

In the present experiments, all participants observed identical covariation between two events. In one condition, participants observed the bulk of A and D frequencies first, followed by the bulk of B and C frequencies (*AD-first condition*, henceforth). In the other condition, participants observed the bulk of B and C frequencies, followed by the bulk of A and D frequencies (*BC-first condition*, henceforth). We define the effect of order as being either primacy, in which the *AD-first condition* leads to judgments for a more positive causal relation than the *BC-first condition*, or recency, in which the *BC-first condition* leads to judgments for a more positive causal relation than the *AD-first condition*.

We propose that causal learning occurs through a process of belief formation and updating. In this view, the information that a person receives at the beginning is used to construct a model about possible causal relationships. This initial belief then helps to determine the final judgment, for example, through the well-known "confirmation bias" (Einhorn & Hogarth, 1978; Klayman & Ha, 1987; Snyder & Swann, 1978; Wason, 1960), or by providing an anchor point for future adjustments (Hogarth & Einhorn, 1992).

Thus, we hypothesize that if the initial evidence suggests a generative causal relationship, then people will prefer to focus on the evidence that confirms that relationship: cells A and D in Figure 1. In contrast, if the evidence initially suggests an inhibitory relationship, then people will tend to focus on evidence that confirms that relationship: cells B and C.² Alternatively, the different beliefs could simply anchor people's judgments on different sides of a bipolar scale of causal strength—that is, an initial hypothesized generative relationship may anchor judgments on the positive end of the scale, whereas an initial belief in an inhibitory relationship may anchor judgments on the negative end. In either case, the first pieces of evidence—those that are used to form the initial belief—should have a long-lasting impact on later judgments about a causal relationship. We thus generally predict primacy effects to occur in causal learning. After

presenting empirical evidence for our hypothesis, we will describe, in the General Discussion, alternative predictions of existing theories of causal learning (i.e., contingency-based theories and associative theories).

EXPERIMENT 1

Method

Participants. The participants were 17 Yale undergraduates. The participants either were paid \$7 for their participation in a group of unrelated experiments or received credit in partial fulfillment of an introductory psychology course requirement.

Procedure and Materials. In general, there were three phases in the experiment: instructions, a learning phase, and a test phase. In the learning phase, the participants observed a series of trials providing covariation information about two events. The test phase required that the participants make judgments about the causal relationship between the two events presented during the learning phase. The main experimental manipulations of order were made during the learning phases. Each phase is explained below. The entire experiment was presented on Power Macintosh 5260/100 computers, using the program *MacProbe* (Hunt, 1994). The presentation program also collected the participants' responses during the test phase.

Instruction phase. Each session began with a series of instruction screens on the type of events that the participants would observe during the session (see Appendix A for the actual text of the instructions). These instructions introduced the general classes of events used as content for the rest of the experiment—namely, the ingestion of an exotic plant and the exhibition of a physical reaction. One screen informed the participants of the three possible relationships between events: generative, inhibitory, or null relationships. The next screen then listed the four types of cases, corresponding to the four cells in Figure 1, that may be observed.

Then, the participants saw a series of eight practice trials in which equal numbers of each of the four different types of trials (A, B, C, D, in Figure 1) were presented in the order ABCDDBAC. The causal event in these practice trials was called "New Plant," and the outcome event was called "New Reaction." After going through these practice trials, the participants then read instructions for providing estimates of causal strength. These instructions first explained the scale, from -100 to 100 , that would be used to provide estimates. A score of -100 was described to mean an inhibitory relationship, a score of 100 a generative relationship, and a score of zero a null relationship. Another screen provided example points (-100 , -70 , -30 , 0 , 30 , 70 , and 100) with real-world examples that reinforced the probabilistic nature of the causal strength estimate.

Learning phase. After receiving the general instructions, the participants proceeded with the learning phases. During each learning phase, the participants observed a series of trials providing information about whether or not that patient had ingested an "exotic" fictional plant and whether or not that patient had exhibited a fictional physical reaction. In each trial, the question "Did Patient XX ingest *Plant*?" was followed by the word *yes* or *no*. The participants were then prompted to "Press the *SPACE* bar to continue." When the participant pressed the space bar, the question "Did Patient XX exhibit *Reaction*?" appeared below the previous question, followed by the word *yes* or *no*. The participants were again prompted to "Press the *SPACE* bar to continue."

The two fictional plant names were "Lanya" and "Hyaeth"; the fictional reaction names were "aliamenisia" and "burlosis." These names were randomly paired for each participant.³ The initials for each patient (represented as XX in the example above) were randomly drawn, and no combination of initials was repeated throughout the experiment. The total number of trials in each condition was 40.

Test phase. After observing the entire sequence of trials in a learning phase, the participants provided causal strength ratings for

the effect of the plant on a physical reaction. Following Wasserman, Elek, Chatlosh, and Baker (1993), the participants were asked, "Based on the cases you just observed, please answer the following question. To what extent does *Plant* cause *Reaction*?" where *Plant* and *Reaction* were the names of the plant and reaction presented during that session. Below this prompt, a slider appeared on screen. A number scale above the slider ran from -100 (labeled "*Plant* strongly inhibits *Reaction*") through 0 (labeled "*Plant* has no relationship with *Reaction*") to 100 (labeled "*Plant* strongly causes *Reaction*"). The participants controlled the position of the slider with the mouse; once the slider was in position, they could press a key to accept their estimate or another key to reestimate.

Design. Two experimental conditions were defined by the order in which covariation information was presented during the learning phase. In order to construct the experimental sequences, two different blocks of 20 individual trials were created as shown in Figure 2.⁴ In the AD block, the bulk of trials (90%) were from cells A (e.g., a patient took Lanya and exhibited aliamenisia) and D (e.g., a patient did not take Lanya and did not exhibit aliamenisia); in the BC block, the majority (90%) were from cells B (e.g., a patient took Lanya and did not exhibit aliamenisia) and C (e.g., a patient did not take Lanya and exhibited aliamenisia). Within each block (AD or BC), the trials were randomly ordered; this random order was fixed across participants. The two different experimental conditions were constructed by manipulating the order of these two blocks, so that, in the AD-first order condition, the AD block came before the BC block (this pattern is shown in Figure 2, both as a summary table and in the actual sequence of trials used), and in the BC-first order condition, the BC block came before the AD block (this pattern is illustrated if the blocks shown in Figure 2 are reversed). Note that both conditions have the same overall contingency, $\Delta P = 0$. Each participant went through both experimental conditions; the order of conditions was counterbalanced across participants.

To summarize, after receiving general instructions about the experiment, each participant observed a series of trials in either the AD-first order condition or the BC-first order condition and then made a causal strength judgment. After an optional short break, the participants observed another series of trials about two new events in the opposite order and then made a second causal strength judgment.

Results and Discussion

The counterbalanced order in which the conditions (AD-first and BC-first) were presented had no significant

effect on its own [$F(1, 15) = 0.08$] or in interaction with the experimental conditions [$F(1, 15) = 3.11$]. Hence, the results reported here are for data collapsed over order of conditions.

The participants gave significantly higher causal strength estimates in the AD-first condition ($M = 17.05$, $SD = 22.15$) than in the BC-first condition ($M = -3.04$, $SD = 20.68$) [$t(16) = 2.90$, $p = .01$, Cohen's $d = 0.75$]. The mean ratings in the AD-first condition were significantly different from zero [$t(16) = 3.17$, $p = .006$, Cohen's $d = 1.09$], whereas the mean ratings in the BC-first condition were not [$t(16) = 0.61$, Cohen's $d = 0.21$].

In line with the results of Yates and Curley (1986), initial presentation of the evidence supportive of positive relationships (AD) led to higher causal strength judgments than presentation of evidence supportive of negative relationships (BC). These results provide the first clear evidence that there is an effect of the order of evidence on the perceived strength of causal relationships between events. This effect takes the form of primacy—that is, the positive evidence of a causal relationship had more impact when it was presented first. Experiments 2 and 3 expanded the present finding by ruling out alternative explanations for this primacy effect.

EXPERIMENT 2

One explanation for the primacy effect found in Experiment 1 is that the participants may have paid less attention to later trials in each learning sequence. This "attention decrement" hypothesis has been proposed to explain primacy effects in impression formation (Anderson, 1981) and contingency judgment (Yates & Curley, 1986). It is important, however, to distinguish between different processes that might be covered by the rubric of attention decrement. One interpretation is that evidence presented early in a sequence is better attended because it serves as a basis for forming an initial hypothesis. That

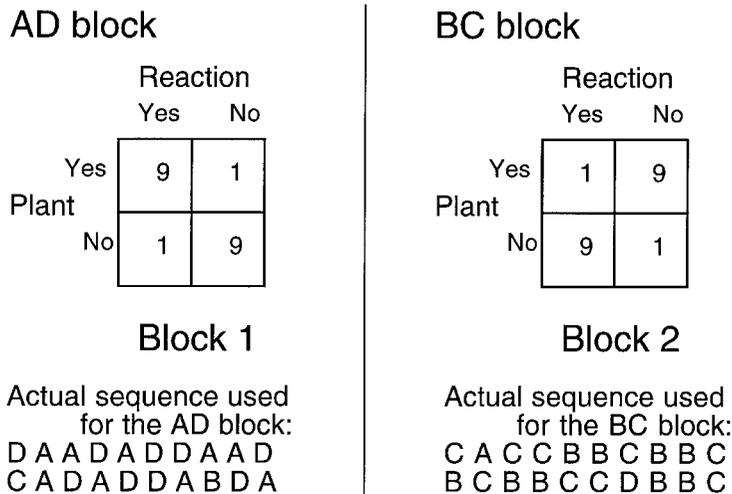


Figure 2. Distribution of trial types used in Experiment 1 for Block 1 and Block 2 of the AD-first condition, with respective trial sequences.

is, the variation in attention paid to pieces of evidence throughout a learning sequence may simply be an epiphenomenon of the belief-updating process that we are proposing. Another, less theoretically interesting, interpretation of attention decrement is that it is simply a fatigue or boredom effect. In this light, primacy effects would not occur because of an underlying belief-updating process but instead would arise because the participants became tired of observing a long sequence of trials.

Previous evidence suggests that a lack of attention due to fatigue is not the sole cause of order effects. Chapman (1991), for example, ensured that participants would attend to every trial by requiring that they type in the first letter of every item in a list of symptoms (out of six possible) that was presented on each trial. Order effects still occurred in Chapman's experiments, even with participants' attention presumably held stable throughout the learning sequence. Since Chapman's experiments concerned a different type of order effect, an additional experiment is needed to rule out the possibility that the primacy effect in Experiment 1 was a by-product of fatigue.

Chapman's (1991) method of ensuring attention was not appropriate for the present experiment because our task involved the presence or absence of just two events; therefore, we used another manipulation—namely, that of rewarding accurate judgments at the end of learning. Because we used noncontingent sequences, accuracy was easily defined as causal strength estimates around zero. Although offering a reward at the end of the experiment does not guarantee that participants read every trial during learning, it does suggest that they would pay closer attention to the trials than would participants who had no opportunity to receive a reward. If fatigue is driving the primacy effect, then we would predict at the very least a weakened effect of order in a reward condition, relative to a no-reward condition.

Method

Participants. The participants were 62 Yale undergraduates. The participants either were paid \$7 for their participation in a group of unrelated experiments or received credit in partial fulfillment of an introductory psychology course requirement. Many more participants were run in this experiment than in Experiment 1, because (1) we included a between-subjects factor (reward/no reward) in our design and (2) we predicted little or no difference between the amount of primacy in the two groups.

Procedure and Materials. Each participant was randomly assigned to one of two groups: reward or no reward. The general method of the experiment was identical to that of Experiment 1, except for the following changes to the instructions at the beginning of the experiment.

Immediately before seeing the example cases in the instructions, the participants in both groups were told, "Do not use paper-and-pencil, fingers-and-toes, etc., to tally the occurrences of the cases. Simply pay attention while observing the cases." This instruction is similar to those used in previous studies (e.g., Ward & Jenkins, 1965; Yates & Curley, 1986) that warned participants against keeping written records. It was included to discourage the participants from using external devices to count cases, especially those in the reward group, for whom the motivation to be numerically accurate would be greater.

The main experimental manipulation came immediately before the starting of the actual learning trials. The participants in the reward group saw the following instructions on the computer screen:

Please try to be as accurate in your estimates as possible. If your estimates are accurate within a reasonable margin of error, we will pay you a \$5 reward at the end of the experiment. Note that you must give correct estimates for all of the blocks to earn the reward. We will explain the correct answers for each block at the end of the experiment.

The no-reward group of participants received the same passage minus the two middle sentences that told of the reward. Because there was never a contingency presented in the set of trials, the participants in the reward condition received \$5 if they gave estimates within 7 points of zero. This was explained to these participants at the end of the experiment.

Results and Discussion

The counterbalanced order in which the conditions (AD-first and BC-first) were presented had no effect on its own [$F(1,58) = 0.07$] or in interaction with the experimental conditions [$F(1,58) = 0.34$]. Hence, the results reported here are for data collapsed over the order of conditions.

Even with the prospect of receiving reward, less than one third of the participants in the reward condition (31%) made "correct" responses (0 ± 7) in both the AD-first condition and the BC-first condition (and hence received the reward). Indeed, an analysis of the causal strength ratings indicated that there were no differences between the reward and no-reward groups, as shown in Figure 3. A 2 (group) \times 2 (trial order) analysis of variance (ANOVA) shows that the only significant effect was that of trial order [$F(1,60) = 21.69, p = .0001, \eta^2 = .27$]. That is, the ratings in the AD-first condition ($M = 12.21, SD = 24.61$) were higher than those in the BC-first condition ($M = -6.75, SD = 24.85$). There was no effect of whether the participants were in the reward group or the no-reward group [$F(1,60) = 0.05, \eta^2 < .01$], and group did not interact with trial order [$F(1,60) = 0.36, \eta^2 < .01$]. Given that an η^2 of about .01 is defined by Cohen (1988) as a small effect, and given that the observed interaction effect was even smaller than this value, it seems that there is little reduction in the primacy effect due to counteracting fatigue. We hypothesized that a reward would counteract the influence of fatigue. That is, if fatigue were responsible for the primacy effect, we would see a reduction of that effect in the no-reward condition, relative to the reward condition. This was clearly not the case; thus, we suggest that the primacy effect in causal induction is not a by-product of fatigue.

As in Experiment 1, the mean ratings in the AD-first condition were significantly greater than zero [$t(61) = 3.91, p = .0002$, Cohen's $d = 0.70$]; in Experiment 2, however, the mean ratings in the BC-first condition were significantly less than zero [$t(61) = -2.14, p = .04$, Cohen's $d = 0.38$].

To summarize, counteracting the effect of fatigue did not eradicate the effect of order. Assuming that the prospect of a reward would keep participants more motivated throughout the learning trials than would no prospect of a reward, these results indicate that a primacy effect is

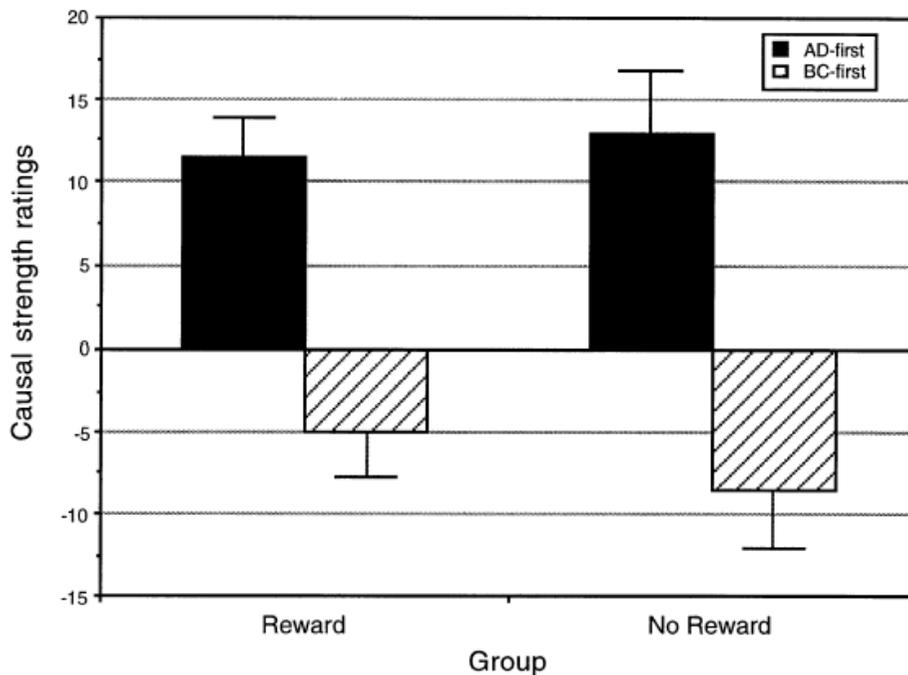


Figure 3. Mean causal strength ratings for the reward and no-reward groups. Error bars represent standard error.

unlikely to be due solely to fatigue. This further suggests that the primacy effect found in Experiment 1 was likely to have been due to a process of belief formation and updating.

EXPERIMENT 3

Because the results of Experiment 2 suggest that the primacy effect is not merely due to fatigue, we may now consider another possible determinant of the primacy effect: the length of the learning sequence.

One major point of difference between López et al.'s (1998) study, which showed recency, and Yates and Curley's (1986) study, which showed primacy, was in the length of sequence. In López et al.'s Experiment 1, for example, the participants saw 160 trials in the learning sequence. In contrast, Yates and Curley's experiment used just 28 trials. Of course, the studies differed on many other procedural points as well, such that it is impossible to tell whether the length of the trial sequence was the sole factor responsible for the conflicting results. However, in order to demonstrate the generality of the primacy effect under trial sequences of different lengths, we doubled the length of the sequence used in Experiments 1 and 2, from 40 to 80 trials.⁵

In addition, Experiment 3 included an *even* condition, in which evidence supporting a generative relationship (AD) and evidence supporting an inhibitory relationship (BC) were evenly distributed. The even condition was included in order to ensure that the sequence was long

enough and that any order effect was not a by-product of preasymptotic or unstabilized learning. That is, it may have been the case that the primacy effect was obtained only when the participants had not sufficiently learned the contingency between two factors and, therefore, were not certain about their responses (see Shanks & Dickinson, 1987, for an example of a transitory bias early in zero contingency learning sequences). If participants, when given evenly distributed trials, can accurately detect a zero contingency in the experiment, it seems safe to conclude that the number of trials was long enough to allow accurate learning, and, therefore, it cannot be the source of any order effect.

Method

Participants. The participants were 63 Yale undergraduates. The participants either were paid \$7 for their participation in a group of unrelated experiments or received credit in partial fulfillment of an introductory psychology course requirement.

Procedure and Materials. The procedure and materials were the same as those used in Experiment 1, except for the following. In addition to the AD-first and BC-first conditions, there was an even condition in which all of the trials were randomly distributed.⁶ The sequence of trials in the even condition was constructed so that there were equal numbers of each type of trial (i.e., AD and BC) in the first and second halves of the condition. In all three conditions, the order within each half of the condition was randomized once and then fixed, so that the entire sequence in each condition was the same for all the participants.

Forty-two participants were run through a within-subjects design in which the order of conditions was counterbalanced across participants, so that all conditions were observed first equally fre-

quently. The additional 21 participants were run through a between-subjects design in which each participant was run through just one of the three conditions. The larger number of participants in the between-subjects design was included to rule out the possibility that the trial order effect in Experiments 1 and 2 was not the result of a bias that arose from the repeated observation and estimation of causal relationships in those experiments' within-subject designs.

As Figure 4 shows, the other difference between Experiment 3 and Experiments 1 and 2 was the length of the learning sequence. Whereas Experiments 1 and 2 had 40 trials apiece, Experiment 3 had double that number, although the proportions of trials within each block were the same across the experiments. The contingency of all of the conditions was also identical to the previous contingencies ($\Delta P = 0$).

Results and Discussion

In reporting the results, we turn first to an analysis including just the participants from the within-subjects design. We then look at the results from the between-subjects design to show that the trial order effect occurs independently of any possible condition order effect.

For the within-subjects design, the order in which the conditions was presented had no significant effect on its own [$F(2,39) = 0.82$] or in interaction with the experimental conditions [$F(4,78) = 0.57$]. Hence, the results reported here are for data collapsed over the order of conditions.

One important point to test is the ratings for the even condition. These ratings should be zero if training was complete. Indeed, the ratings of the even condition ($M = 2.61, SD = 19.35$) were not different from zero [$t(41) = 0.87, Cohen's d = 0.19$]. Further tests on the other conditions showed that the AD-first condition ($M = 17.67, SD = 25.66$) was significantly greater than zero [$t(41) = 4.46, p = .0001, Cohen's d = 0.97$], whereas the causal strength ratings in the BC-first condition ($M = -5.50, SD = 22.27$) were not different from zero [$t(41) = -1.60, Cohen's d = 0.35$].

The most crucial comparison, of course, was to see whether the order effect still held, even given the increased number of trials. A one-way ANOVA across the three conditions showed that there was a significant effect of the order in which the trials were presented within each condition [$F(2,82) = 9.71, p = .0002, \eta^2 = .19$]. To find which groups were different from the others, paired *t* tests compared each of the condition means with every other. The Bonferroni-corrected significance level for the three paired comparisons was set to .0167. Paired comparisons showed that the AD-first condition received significantly higher ratings than the BC-first condition [$t(41) = 3.87, p = .0004, Cohen's d = 0.52$]. In addition, the even condition was rated significantly lower than the AD-first condition [$t(41) = 3.24, p = .002, Cohen's d = 0.54$], although it was not different from the BC-first condition [$t(41) = -1.53, Cohen's d = 0.20$].

The next analysis includes data from the participants in the between-subjects design and from the first condition only of the within-subjects design. A one-way ANOVA across the three conditions was marginally significant [$F(2,62) = 2.41, p = .10$]. However, the main comparison of interest, between the two skewed order conditions, showed that the AD-first condition ($M = 14.82, SD = 22.20$) was significantly higher than the BC-first condition ($M = 1.50, SD = 20.72$) [$t(39) = 1.98, p = .05$]. (The fact that the ratings from the even condition [$M = 5.68, SD = 16.61$] fell between the other conditions suggests that its inclusion in the full ANOVA served more to add error variance than to distinguish means.) Hence, the primacy effect appears to occur in the absence of learning across conditions.

The main results of Experiment 3 strongly suggest that the order effects found in Experiments 1 and 2 were not due simply to an insufficient number of trials. With

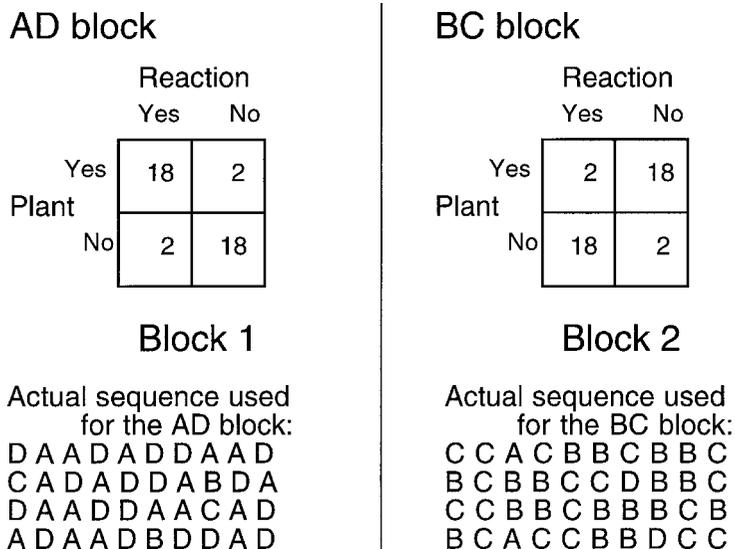


Figure 4. Distribution of trial types used in Experiment 3 for Block 1 and Block 2 of the AD-first condition, with respective trial sequences.

twice the number of trials as the previous experiments, the participants still tended to rate the conditions with the bulk of cell A and cell D trials presented early in learning higher than the conditions with the bulk of cell B and cell C trials first. The primacy effect still occurs reliably in a reasonably long sequence.

GENERAL DISCUSSION

The present experiments show that the order of covariational information has a significant impact on participants' judgments of the strength of a causal relationship between two events, yielding a primacy effect.⁷ Informal comparisons of the effect sizes for the differences between the AD-first and BC-first conditions across experiments (Cohen's *ds* of 0.75, 0.57, 0.71, and 0.52 for Experiment 1, Experiment 2 [reward], Experiment 2 [no reward], and Experiment 3, respectively) suggest that this primacy effect can reliably be obtained across manipulations of attention and sequence length. Although there is some reduction in the numerical size of the effect, they all fall between "medium" and "large" effect sizes, according to Cohen (1988). Thus, it seems safe to conclude that this primacy effect is robust enough to be obtained across various circumstances. Furthermore, the alternative explanations proposed here, fatigue and insufficient sequence length, are contradictory; that is, if the primacy effect is merely due to fatigue (so that reducing the effect of primacy also reduces the effect of order), then we would expect a larger primacy effect with a longer sequence, which does not seem to be the case in our data. This result especially casts doubt on the fatigue explanation of the primacy effect. Instead, we explain this primacy effect as being due to a process of belief updating, such that the initial information affects the hypothesis one draws; this hypothesis then impacts later judgments. What other explanations might there be for this order effect? We turn now to two models currently inspiring research into causal induction. We argue that these models are unable to explain the primacy effect found in the present experiments.

Associative Learning Model

One popular mechanism proposed to be at the core of causal induction is associative learning. In particular, the Rescorla–Wagner model (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972) has been cited as providing a good description of a variety of phenomena in causal induction, including the judgment of contingency (e.g., Wasserman et al., 1993) and competition between two possible causes of an event (Shanks & López, 1996). However, the Rescorla–Wagner model fails to account for the present results.

The nature of this prediction can perhaps be seen more clearly through an examination of the learning function itself. Formally, this function is expressed as

$$\Delta V_{\text{cue}} = \alpha\beta(\lambda - \sum V_{\text{total}}), \quad (1)$$

where V_{cue} represents the associative strength between the cue and the outcome (and hence the judged causal strength of the cue when the model was applied for causal induction tasks); ΔV_{cue} represents the change in associative strength after an individual trial; α and β represent learning rate parameters determined by aspects of the cue and the outcome (i.e., $\alpha = 0$ when a cue does not occur on a trial, and $\alpha > 0$ when a cue does occur; the value of β when the outcome is present is greater than when the outcome is absent); λ represents the maximum associative strength (i.e., the maximum amount of learning) that the outcome can support; and $\sum V_{\text{total}}$ represents the total of associative strength combined across all of the cues present on that trial (if there is only one explicit cue, the context is assumed to play the role of a second cue). The inclusion of the final term in Equation 1 means that cues compete for association with an outcome. As cues and outcome co-occur in repeated learning trials, the cue on which the outcome is most contingent will take the majority of the possible associative strength. Note that the strength of association between cue and outcome is assumed to be updated at each trial. Therefore, the Rescorla–Wagner model is sensitive to the sequence in which a series of learning trials is presented.

Wasserman et al. (1996) ran computer simulations of the Rescorla–Wagner model, using order manipulations similar to ours, across a range of learning parameters (see Wasserman et al., 1996, Figure 5) and concluded that "the Rescorla–Wagner model *never* predicts primacy in causal judgment. . . . Given this evidence, if primacy were to be a reliable effect in human contingency and causal judgments, then the utility of Rescorla–Wagner model would be called into question" (p. 231). We also ran simulations of the trial sequences from Experiments 1 and 3; details of the simulation methods and the results are provided in Appendix B. For the 40-trial sequence from Experiment 1, the Rescorla–Wagner model yields recency effects in almost all of the parameter combinations tested (16 of 18 cases; see Tables B1 and B2). When the primacy effect was predicted by the model, the effect was close to zero. Furthermore, these cases contradicted the usual assumption that the value of α_{context} is smaller than α_{cause} and so are themselves questionable predictions of the model. Use of the 80-trial sequence of Experiment 3 yielded clear recency effects for every combination of parameters included in the simulation as in the simulations of Wasserman et al. (see also López et al., 1998, and Shanks, López, Darby, & Dickinson, 1996, for more on this prediction).

The Rescorla–Wagner model predicts the recency effect because, according to the model, associative learning is determined by the extent to which a cue is "surprising."⁸ For instance, evidence suggesting a generative relationship would be more surprising and hence would have a larger impact on associative strength if it follows the bulk of evidence suggesting an inhibitory relationship (i.e., BC-first condition) than if it is presented in the absence of any prior association (i.e., AD-first condi-

tion). This pattern of changes in associative strength leading to recency is illustrated in Figure B1, in which the slope of the AD block is steeper in the BC-first condition (i.e., the second block of the BC-first condition) than in the AD-first condition (i.e., the first block of the AD-first condition). This more extreme rate of change in the second phase of the learning sequence makes it more likely that a recency effect will occur. Thus, the primacy effects found in our experiments are at variance with the model.

Contingency-Based Model

Another possible mechanism of causal learning relies on the subjective calculation of contingency indices such as ΔP . To date, the most complete version of contingency-based human causal induction is Cheng's (1997) power probabilistic contrast (PC) theory. According to this theory, the causal power of a causal candidate i , p_i , is

$$p_i = \frac{\Delta P_i}{1 - P(e|i)}. \quad (2)$$

In essence, Equation 2 shows that the judged causal strength of an event is based on the ΔP , weighted by the presence of the effect in the absence of the cause (see Spellman, 1996, for empirical support). According to the power PC theory, the causal power of the factors used in our experiments was zero. Note, however, that p_i did not vary as a function of order. That is, it was identical for both order conditions in Experiments 1 and 2 and for the even and skewed conditions in Experiment 3. Yet, in the present experiments, manipulation of trial order introduced significant changes in the perceived causal strength.

At this stage, any order effect is beyond the boundary conditions of the power PC theory because the causal strength of an event is calculated over all available trials all at once when enough observations are assumed to have been accumulated (see also Price & Yates, 1995, for a similar interpretation). This theory may be modified to allow calculations of two contingencies (one for each block) and to give more weight to the contingency of the primary block. However, this model is meant to be normative (Cheng, 1997; Glymour & Cheng, 1998), and it is not clear why giving more weight to the contingency from initial trials would be normative. For instance, Bayes's theorem, which is considered a good example of a normative model, predicts no order effects under any circumstances (Slovic & Lichtenstein, 1971; see López et al., 1998, for a similar interpretation). If anything, it seems easier to speculate on normative reasons for giving more weight on contingency from later trials (e.g., more recent data are more valid because they are more representative of the current state of the world).

Finally, is the primacy effect obtained in the present experiments in conflict with the recency effect obtained in López et al. (1998)? It is difficult at this point to tell how to reconcile the discrepancy because their experimental paradigm used multiple and compound cues, whereas ours used single cues. It might be that people

use a different hypothesis-testing strategy when multiple and compound cues are presented and when the interaction among cues should be considered. The present experiments, however, made two contributions with respect to this issue. First, we have shown that the length of learning sequence is unlikely to be the reason for the discrepancy. Second, we reliably demonstrated the primacy effect in a single-cue learning paradigm. Thus, the results of the present experiments call for future empirical work exploring the paradigm differences responsible for primacy and recency effects (see Hogarth & Einhorn, 1992, for potential factors involved in primacy vs. recency effects).

Conclusion

The experiments reported here show that primacy effects in causal learning do occur, and these effects are not due to mere fatigue or to the length of the trial sequence. In addition, the present results present clear problems for current models of causal learning. The effect of evidence order falls outside the scope of contingency-based models, emphasizing the restricted range of application of those models. The dominant associative model is sensitive to the effect of evidence order, but it makes predictions in the opposite direction of the present results. Clearly, the cognitive mechanisms by which people infer causal relationships between events are not fully understood. The present study makes a contribution by presenting another important factor that influences learning of new causal rules—namely, the order of evidence. We suggest that the primacy effects found here are most compatible with a belief-updating process.

REFERENCES

- AHN, W., & KALISH, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 199-225). Cambridge, MA: MIT Press.
- ALLAN, L. G., & JENKINS, H. M. (1980). The judgment of contingency and the nature of the response alternatives. *Canadian Journal of Psychology*, **34**, 1-11.
- ALLOY, L. B., & ABRAMSON, L. Y. (1979). Judgments of contingency in depressed and nondepressed college students. *Journal of Experimental Psychology: General*, **108**, 441-485.
- ANDERSON, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- ANDERSON, N. H., & JACOBSON, A. (1965). Effect of stimulus inconsistency and discounting instructions in personality impression formation. *Journal of Personality & Social Psychology*, **2**, 531-539.
- ASCH, S. E. (1946). Forming impressions of personality. *Journal of Abnormal & Social Psychology*, **41**, 258-290.
- CHAPMAN, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 837-854.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, **104**, 367-405.
- COHEN, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.
- EINHORN, H. J., & HOGARTH, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, **85**, 396-416.
- GLYMOUR, C., & CHENG, P. W. (1998). Causal mechanism and probability: A normative approach. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 295-313). Oxford: Oxford University Press.

- HENDRICK, C., & COSTANTINI, A. F. (1970). Effects of varying trait inconsistency and response requirements on the primacy effect in impression formation. *Journal of Personality & Social Psychology*, **15**, 158-164.
- HOGARTH, R. M., & EINHORN, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, **24**, 1-55.
- HUNT, S. M. J. (1994). MacProbe: A Macintosh-based experimenter's workstation for the cognitive sciences. *Behavior Research Methods, Instruments, & Computers*, **26**, 345-351.
- JENKINS, H. M., & WARD, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, **79**(1, Whole No. 594).
- KLAYMAN, J., & HA, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, **94**, 221-228.
- LÓPEZ, F. J., SHANKS, D. R., ALMARAZ, J., & FERNÁNDEZ, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 672-694.
- MATUTE, H., ARCEDIANO, F., & MILLER, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 182-196.
- MILLER, R. R., BARNET, R. C., & GRAHAME, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin*, **117**, 363-386.
- 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 theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.
- SCHUSTACK, M. W., & STERNBERG, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, **110**, 101-120.
- 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). San Diego, CA: Academic Press.
- SHANKS, D. R., & LÓPEZ, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, **24**, 511-522.
- SHANKS, D. R., LÓPEZ, F. J., DARBY, R. J., & DICKINSON, A. (1996). Distinguishing associative and probabilistic contrast theories of human contingency judgment. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *Causal learning: The psychology of learning and motivation* (Vol. 34, pp. 265-311). San Diego, CA: Academic Press.
- SLOVIC, P., & LICHTENSTEIN, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior & Human Performance*, **6**, 649-744.
- SNYDER, M., & SWANN, W. B., JR. (1978). Hypothesis-testing in social interaction. *Journal of Personality & Social Psychology*, **36**, 1202-1212.
- SPELLMAN, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, **7**, 337-342.
- VAN HAMME, L. J., & WASSERMAN, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning & Motivation*, **25**, 127-151.
- WAGNER, A. R., & RESCORLA, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning* (pp. 301-336). London: Academic Press.
- WARD, W. C., & JENKINS, H. M. (1965). The display of information and the judgment of contingency. *Canadian Journal of Psychology*, **19**, 231-241.
- WASON, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, **14**, 129-140.
- WASSERMAN, E. A., CHATLOSH, D. L., & NEUNABER, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning & Motivation*, **14**, 406-432.
- WASSERMAN, E. A., ELEK, S. M., CHATLOSH, D. L., & BAKER, A. G. (1993). Rating causal relations: Role of probability in judgments of response-outcome contingency. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 174-188.
- WASSERMAN, E. A., KAO, S.-F., VAN HAMME, L. J., KATAGIRI, M., & YOUNG, M. E. (1996). Causation and association. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *Causal learning: The psychology of learning and motivation* (Vol. 34, pp. 208-264). San Diego, CA: Academic Press.
- WILLIAMS, D. A. (1996). A comparative analysis of negative contingency learning in humans and nonhumans. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *Causal learning: The psychology of learning and motivation* (Vol. 34, pp. 89-131). San Diego, CA: Academic Press.
- YATES, J. F., & CURLEY, S. P. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, **62**, 293-302.

NOTES

1. For instance, assume that, for a sequence of 20 events, 10 are of a certain type randomly interspersed throughout the sequence, and the remaining 10 are of different types. If we divide the sequence in half, on average we would expect that 5 of those events would occur in the first half. According to the binomial distribution, however, there is over a 37% chance that a given sequence will have 6 or more of that type of event in the first 10 trials.

2. Actually, the formation of a hypothesis of an inhibitory relationship appears to be more complicated than that of a generative relationship (Williams, 1996). We can imagine that cell B and cell C trials might give rise to two possible models: one of an inhibitory relationship (if conditions are right), or one of a null relationship. In either case, terminal ratings are predicted to be lower than in conditions that give rise to models of generative relationships.

3. Although the present experiments did not assess the participants' initial estimates on causal strengths between a specific plant and a specific reaction, it is reasonable to assume that specific stimulus contents did not have any systematic effect because all names were randomly paired for each participant.

4. Also note that the phrase *block of trials* is here merely an expository convenience; the participants received no explicit notice that the condition was divided into separate sections.

5. We did not use 160 trials, in order to keep the risk of fatigue to a minimum.

6. The actual sequence for the even condition was ACBCDDA ACDBBCABACADBBDCDBBAACDDCCDBCAABAABCDAC DBDCADCBCBDCBBACDCBCAADABCDABDDAB. The actual sequences for the AD-first and BC-first conditions can be seen by combining the subsequences for the AD block and the BC block shown in Figure 4.

7. Primacy in our experiments seems to occur only when generative evidence is presented first (i.e., only the AD-first condition ratings are different from zero in the experiments and are different from the even condition ratings in Experiment 3). This result may stem from the well-known differential weighting of evidence, such that generative evidence is weighted more than inhibitory evidence (e.g., Schustack & Sternberg, 1981). That is, there may be a weaker primacy effect with inhibitory evidence in our experiments, which is simply not detected.

8. The model represents "surprise" in terms of the difference between the cue presented on the trial in question and the cue that is expected on the basis of the summed predictive value of all of the cues present on the trial (i.e., the absolute value of $\lambda - \Sigma V_{\text{total}}$; Miller, Barnet, & Grahame, 1995).

APPENDIX A

Instructions for Experiment 1

(Note: Specific changes to these instructions for Experiments 2 and 3 are detailed in the Method sections of those experiments.)

Imagine that you are a botanist. Your current project is to catalog the effects that ingesting certain exotic plants has on humans. When a person ingests a plant, three things may happen:

1. The plant may inhibit a physical reaction. For example, some plants help to suppress one's appetite.
2. The plant may cause a physical reaction. For example, some plants help to increase one's appetite.
3. The plant may cause no physical reaction. For example, some plants do nothing to one's appetite.

In particular, you are studying the extent to which two exotic plants cause any one of two physical reactions. To determine this, you will observe a set of cases for each plant. In each case, a person might or might not have ingested a particular plant. That person also might or might not exhibit a particular physical reaction. Therefore, four types of cases can be identified:

- A. The person ingested the plant and exhibited the physical reaction.
- B. The person ingested the plant and did not exhibit the physical reaction.
- C. The person did not ingest the plant and exhibited the physical reaction.
- D. The person did not ingest the plant and did not exhibit the physical reaction.

(The participants then saw eight practice cases. After those cases, the following instructions were presented.)

After you have seen all of the cases for that plant and reaction, you will be asked to give an estimate that indicates your belief about the ability of the plant to cause the reaction. You will give your estimate on a scale from -100 to 100 . A score of -100 means that you think the proposed factor may inhibit the physical reaction. A score of 0 means that you think the proposed factor may have no relationship to the reaction. A score of 100 means that you think the proposed factor may be a strong cause of the reaction. Scores between -100 and 100 indicate different degrees of ability to cause the reaction. Note that the sum of the two estimates does not have to equal any particular number.

For example, you might decide to assign these values to the following events:

- | | |
|--------|--|
| -100 | Very strong prevention, e.g., The degree to which rain prevents the ground from being dry. |
| -70 | Strong prevention, e.g., The degree to which helmets prevent serious rollerblading injury. |
| -30 | Weak prevention, e.g., The degree to which exercise prevents heart disease. |
| 0 | Not a cause, e.g., The degree to which one day of rain causes the stock market to rise. |
| 30 | Weak cause, e.g., The degree to which getting wet causes a cold. |
| 70 | Strong cause, e.g., The degree to which being exposed to a virus causes a cold. |
| 100 | Very strong cause, e.g., The degree to which rain causes the ground to be wet. |

APPENDIX B

The simulations of the Rescorla–Wagner model using the trial sequences from Experiments 1 and 3 were performed as follows as in Wasserman et al. (1996). First, two cues were defined: the putative cause and a context cue, which represented anything that was not the putative cause and which was assumed to be always present. The change in associative strength (ΔV_{cue}) for cause and context was defined for each of the cells in the contingency table (following Wasserman et al., 1996). For cell A (cause and effect present),

$$\Delta V_{\text{cause}} = \alpha_{\text{cause}} \beta_{\text{outcome}} (\lambda_{\text{outcome}} - \sum V_{\text{cause+context}})$$

and

$$\Delta V_{\text{context}} = \alpha_{\text{context}} \beta_{\text{outcome}} (\lambda_{\text{outcome}} - \sum V_{\text{cause+context}}).$$

For cell B (cause present but effect absent),

$$\Delta V_{\text{cause}} = \alpha_{\text{cause}} \beta_{\text{no outcome}} (\lambda_{\text{no outcome}} - \sum V_{\text{cause+context}})$$

and

$$\Delta V_{\text{context}} = \alpha_{\text{context}} \beta_{\text{no outcome}} (\lambda_{\text{no outcome}} - \sum V_{\text{cause+context}}).$$

For cells C and D, in which the cause is absent, ΔV_{cause} is equal to zero. Furthermore, the total associative strength subtracted from λ includes only the cues that are present on that trial, so that, for cell C (cause absent but effect present),

$$\Delta V_{\text{context}} = \alpha_{\text{context}} \beta_{\text{outcome}} (\lambda_{\text{outcome}} - \sum V_{\text{context}}),$$

and for cell D (cause and effect absent),

$$\Delta V_{\text{context}} = \alpha_{\text{context}} \beta_{\text{no outcome}} (\lambda_{\text{no outcome}} - \sum V_{\text{context}}).$$

The value of λ_{outcome} may be set arbitrarily to scale the associative strengths of the model; in our simulations, it was set to 1 for cell A and cell C trials (Wasserman et al., 1996). Two values of $\lambda_{\text{no outcome}}$ were used in the simulations, following previous uses of the model: 0 (Van Hamme & Wasserman, 1994) and -1 (Wasserman et al., 1996).

Again, following Wasserman et al. (1996), nine combinations of learning rate parameters were created by crossing three sets of α parameters ($\alpha_{\text{cause}} < \alpha_{\text{context}}$, $\alpha_{\text{cause}} = \alpha_{\text{context}}$, and $\alpha_{\text{cause}} > \alpha_{\text{context}}$) with three sets of β parameters ($\beta_{\text{outcome}} < \beta_{\text{no outcome}}$, $\beta_{\text{outcome}} = \beta_{\text{no outcome}}$, and $\beta_{\text{outcome}} > \beta_{\text{no outcome}}$). Note that it is usually assumed that $\alpha_{\text{cause}} \geq \alpha_{\text{context}}$ (e.g., Wasserman et al., 1993) and $\beta_{\text{outcome}} \geq \beta_{\text{no outcome}}$ (e.g., Wasserman et al., 1996) in most situations.

APPENDIX B (Continued)

Table B1
Final Associative Strengths and Primacy Scores From Rescorla–Wagner Simulations of Experiments 1 and 3
Across All Combinations of Parameter Pairs, When $\lambda = 0$

$\beta_{\text{outcome}}, \beta_{\text{no outcome}}$	$\alpha_{\text{cause}}, \alpha_{\text{context}}$								
	0.1, 0.9			0.5, 0.5			0.9, 0.1		
	AD First	BC First	Primacy	AD First	BC First	Primacy	AD First	BC First	Primacy
Experiment 1 (40-Trial Sequence)									
0.9, 0.1	0.04	0.02	0.02	-0.06	0.13	-0.19	0.07	0.41	-0.34
0.5, 0.5	-0.02	0.08	-0.10	-0.35	0.38	-0.73	-0.27	0.45	-0.72
0.1, 0.9	0.001	0.03	-0.03	-0.06	0.13	-0.19	-0.05	0.11	-0.16
Experiment 3 (80-Trial Sequence)									
0.9, 0.1	-0.007	0.03	-0.04	-0.24	0.21	-0.45	-0.003	0.25	-0.25
0.5, 0.5	-0.12	0.15	-0.27	-0.62	0.55	-1.17	-0.38	0.42	-0.80
0.1, 0.9	-0.02	0.05	-0.07	-0.21	0.21	-0.42	-0.12	0.12	-0.24

Note—Primacy scores are equal to AD-first strength minus BC-first strength. Hence, positive scores indicate primacy, and negative scores indicate recency.

Table B2
Final Associative Strengths and Primacy Scores From Rescorla–Wagner Simulations of Experiments 1 and 3
Across All Combinations of Parameter Pairs, When $\lambda = -1$

$\beta_{\text{outcome}}, \beta_{\text{no outcome}}$	$\alpha_{\text{cause}}, \alpha_{\text{context}}$								
	0.1, 0.9			0.5, 0.5			0.9, 0.1		
	AD First	BC First	Primacy	AD First	BC First	Primacy	AD First	BC First	Primacy
Experiment 1 (40-Trial Sequence)									
0.9, 0.1	0.02	0.02	-0.007	-0.33	0.22	-0.55	-0.43	0.47	-0.90
0.5, 0.5	-0.07	0.14	-0.21	-0.76	0.71	-1.47	-0.93	0.53	-1.46
0.1, 0.9	-0.000	-0.003	0.003	-0.14	0.14	-0.28	-0.45	-0.20	-0.25
Experiment 3 (80-Trial Sequence)									
0.9, 0.1	-0.06	0.05	-0.11	-0.52	0.41	-0.93	-0.22	0.36	-0.58
0.5, 0.5	-0.25	0.30	-0.55	-1.24	1.09	-2.33	-0.92	0.69	-1.61
0.1, 0.9	-0.03	0.06	-0.09	-0.42	0.37	-0.79	-0.38	0.04	-0.42

Note—Primacy scores are equal to AD-first strength minus BC-first strength. Hence, positive scores indicate primacy, and negative scores indicate recency.

The unequal parameter sets were given extreme values (0.1 and 0.9) to sample the space of possible parameter combinations broadly. The equal parameter sets were given values of 0.5.

Figure B1 shows an example of the trial-by-trial changes in V_{cause} , the causal strength judgment, when $\lambda_{\text{no outcome}}$ is zero and all learning rate parameters are set at 0.5, for all three conditions (AD-first, BC-first, and even) in Experiment 3. The pattern of results shows a clear recency effect, in that the final value for the BC-first condition is higher than the final value for the AD-first condition. Remember that, in the BC-first condition, the second block contains mostly cell A and cell D tri-

als and that the last block of the AD-first condition contains mostly cell B and cell C trials. The even condition stays centered around zero, the actual contingency.

Tables B1 and B2 show the sets of AD-first and BC-first final values for V_{cause} in each of the parameter combinations, for both the 40-trial sequence from Experiment 1 and the 80-trial sequence from Experiment 3, when $\lambda_{\text{no outcome}} = 0$ (Table B1) and -1 (Table B2). To aid in inspection of the tables, primacy values are also reported. These scores are the differences of the AD-first and BC-first conditions. Positive scores indicate primacy, whereas negative scores indicate recency.

APPENDIX B (Continued)

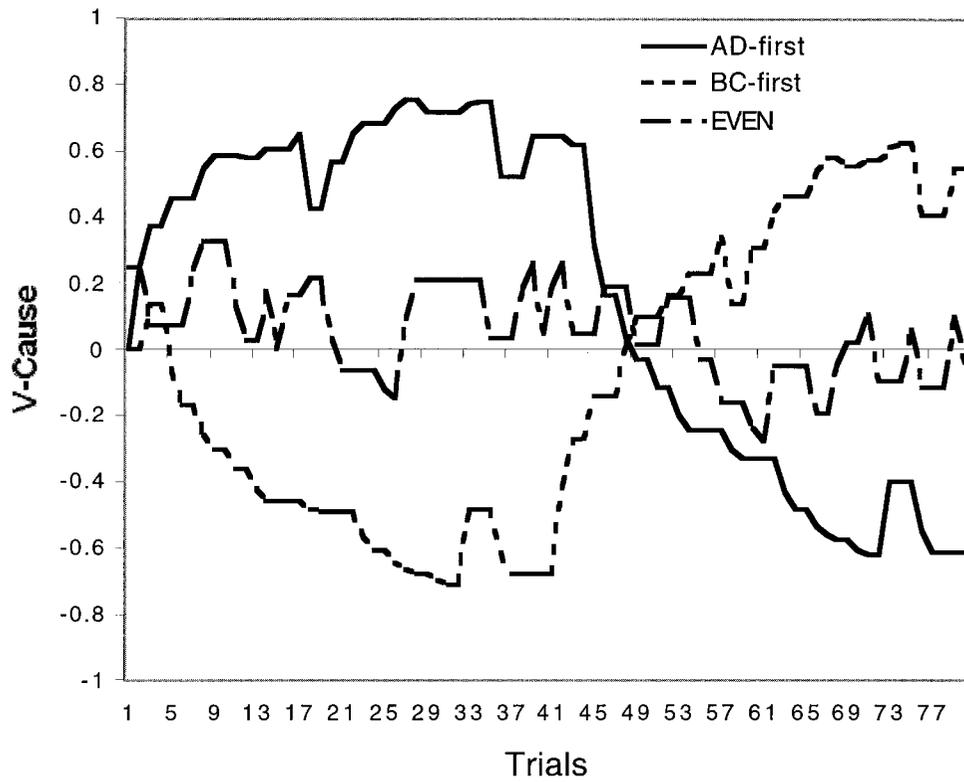


Figure B1. Trial-by-trial causal strength estimates from the Rescorla–Wagner simulation of the Experiment 3 learning sequence.

(Manuscript received October 28, 1998;
revision accepted for publication February 18, 2000.)