

On the origin of personal causal theories

MICHAEL E. YOUNG

University of Minnesota, Minneapolis, Minnesota

Detecting the causal relations among environmental events is an important facet of learning. Certain variables have been identified which influence both human causal attribution and animal learning: temporal priority, temporal and spatial contiguity, covariation and contingency, and prior experience. Recent research has continued to find distinct commonalities between the influence these variables have in the two domains, supporting a neo-Humean analysis of the origins of personal causal theories. The cues to causality determine which event relationships will be judged as causal; personal causal theories emerge as a result of these judgments and in turn affect future attributions. An examination of animal learning research motivates further extensions of the analogy. Researchers are encouraged to study real-time causal attributions, to study additional methodological analogies to conditioning paradigms, and to develop rich learning accounts of the acquisition of causal theories.

It has been noted on many occasions that the concept of causality underlies much of human cognition and represents one of the major factors used in forward and backward inference. In social realms, inferential skill allows us to predict the behavior of others and to determine *why* they behave as they do. In science, we are engaged in an enterprise in which we seek to determine the causal relationships among observables, such as lesion locations and their behavioral effects, neural network architectures and their performance, and biological processes and cancer. In everyday life we strive to discover causes of events (Why is there excessive paint peeling on my house? Why is my child crying?) and to predict the effects of actions (Will she call me after what I said last night? What if I take the job?).

Causal attribution affords two primary advantages: prediction and control. If an organism can discover the causal nature of events, it can predict what events will follow others, and thus it will be able to prepare for the arrival of events that are important to its survival, its enjoyment, and so forth. Control over events can be established if the causing event is one that the organism can produce or prevent. This dichotomy, prediction and control, is captured in the traditional animal learning literature by the processes of classical conditioning and instrumental training. The parallels between causal attribution and animal learning are extensive (Killeen, 1981; Mackintosh, 1977; Shanks & Dickinson, 1987; Wasserman, 1990a).

The work of Shanks and Dickinson (1987) and Wasserman (1990b) is the basis of a significant trend in the understanding of causal attributions from an animal learning perspective (see, e.g., Baker, Mercier, Vallee-Tourangeau, Frank, & Pan, 1993; Chapman, 1991; Chapman & Robbins, 1990; Reed, 1992; Shanks, 1989; Van Hamme, Kao, & Wasserman, 1993; Wasserman, Elek, Chatlosh, & Baker, 1993; Young & DeBauche, 1993). In this manuscript, I will reconsider this framework in order to (1) determine how well recent research can be integrated into it, (2) address common criticisms of the approach, and (3) enlarge its boundaries by considering the validity of less known aspects of animal learning and their possible parallels in causality.

Recent reviews of the literature have focused on the contingency aspect of causality. Cheng (1993) compared the role of contingency with the role of prior knowledge. Allan (1993) surveyed the relative explanatory power of rule-based models and associative models of human contingency judgment. Shanks (1993b) presented a very detailed examination of the strengths and weaknesses of a relative contiguity model, of rule-based models, and of associative models as mechanisms for human instrumental learning. While I will discuss the relative role of contingency information as a cue to causality, my principal goal is to explore the limits of the animal learning analogy by considering a wider range of conditioning data and theories than that considered in these earlier reviews.

This exploration of causality centers on the origins of our *personal* causal theories. We learn about causal relationships by directly observing them in operation or by being informed about them. The former are *personal* theories, whereas the latter are *public* theories. I do not accept that humans are capable of viewing causality directly (Harré & Madden, 1975), but rather that causal relationships in the environment are learned about as a result of the interaction of two processes, one data driven and the other theory driven (cf. Nisbett & Ross, 1980).

Preparation of this article was supported by the University of Minnesota Center for Research on Learning, Perception, and Cognition, and the National Institute of Child Health and Human Development (HD-07151). I would like to thank Randy Fletcher, Bruce Overmier, Paul van den Broek, and four anonymous reviewers for comments on earlier versions of this manuscript. Correspondence concerning this article should be addressed to M. E. Young, Department of Psychology, University of Minnesota, 317 Elliott Hall, Minneapolis, MN 55455 (e-mail: young@cs.umn.edu).

Personal causal theories are acquired (data driven), and as these theories develop, they influence future perceptions via top-down processes (theory driven). In contrast, *public* theories are learned about via the process of cultural transmission. Some of these public theories are the result of the promulgation of a personal causal theory, and others are the result of analytical processes (Figure 1). The origin of personal theories thus has a significant impact on our causal knowledge; personal theories influence future attributions, and they can become public theories.

I will accept the basic tenets of a Humean view of causality as a working hypothesis: *Causality* is a psychological construct, a term that is used to describe a particular class of event interrelationships. Many environmental cues are important to data-driven attributions, with Hume's "cues-to-causality" representing one potentially important set. As data are gathered about event interrelationships, a generalization of these data can emerge in the form of a *theory*. This theory can then be applied to make predictions in novel situations to the degree that these situations are similar to the ones that underlie the theory. Prior knowledge in the form of these data-driven theories inevitably affects the integration of

new perceptions into the current body of causal knowledge. Few readers should find much difficulty accepting such a proposal. Since the focus is on the origins of theories of causation, it is necessary to determine *which* data drive attributions. Causality as a data-driven psychological construct should not imply that theories play no role. In fact, I will demonstrate that focusing on the role of learning in causal attributions *requires* that prior experience be considered.

Because I am taking the Humean stance, I must examine some basic tenets of Hume's theory and the evidence for and against it. I will adopt the term *neo-Humean* to describe the notions of the modern psychologist or philosopher who has indicated acceptance of Hume's framework but has amended it in various ways. This will help to defuse some of the most common objections raised against the specifics of Hume's 1739 *Treatise*, an endeavor that risks throwing the baby out with the bath water. It is important to remember throughout this discussion that I am considering the causal attributions that are readily made in everyday life (common notion causality). This is not an attempt to get at metaphysical or scientific causality.

Conditioning and Causal Attribution

Shanks and Dickinson (1987), Killeen (1981), and Wasserman (1990b) expounded upon a notion that has been informally acknowledged by learning theorists for decades: when an animal is being conditioned, it is learning the causal relations among environmental and internal events. They argue that causal attribution may be reduced to basic learning principles. The conditioning account of causal attribution will also be referred to here as an *associative* account. This does not suggest a definition of associationism that reduces learning to simple input-output relations, but a definition that stresses the relation to animal learning. In contemporary, real-time learning theories (Desmond & Moore, 1988; Grossberg & Schmajuk, 1989; Young, 1992), stimuli are processed within the context of ongoing neural activity that is itself a product of prior experience. Responding is thus a function of current stimuli *and* the mental state of the organism.

When an animal is undergoing classical or Pavlovian conditioning, its hypothetical goal is to predict the occurrence of the unconditioned stimulus (US), such as food or shock. The experimenter provides cues to assist in this predictive task by including conditioned stimuli (CSs). Prediction of a US is frequently beneficial to the animal and may permit it to prepare for its arrival. For example, if the US is an air puff to the eye, the animal can blink to protect the eye. In an analogous manner, instrumental training permits the subject to control the delivery of important events. In both cases, causal attribution can be ascribed to an animal. For classical conditioning, the CS is a candidate cause and the US is the effect; for instrumental training, the response is the cause and the outcome is the effect.

A difficult question arises when we posit causal processing in a lower species: How do we determine whether




Figure 1. The relationship between private and public causal theories, and the factors that influence their development. Elements in bold are the focus of this paper.

the subject judges two events as causally related? Conditioned responding in classical conditioning only tells us that the animal is able to predict the occurrence of the US. Prediction does not imply causation. Arguments like these focus solely on the temporal invariance of causality and ignore aspects such as spatial contiguity and prior knowledge.

During instrumental or Thorndikian training, the subject can be said to have learned to *produce* the outcome, which is not much different from saying that it has learned to *cause* the outcome. However, some researchers (e.g., Guthrie, 1952, and Hull, 1943, as cited in Rescorla, 1987) have argued that the chances of the desired response's occurring increases not because of knowledge of the outcome that will follow, but because the connection between environmental stimuli and the response behavior has been strengthened (i.e., the S-R association of the S[timulus]-R[esponse]-O[utcome] sequence has been increased, not the R-O association). Others (e.g., Rescorla & Solomon, 1967; Trapold & Overmier, 1972) have suggested that the outcome is learned about, but through being associated with environmental and discriminative stimuli (so both S-R and S-O associations are formed). Rescorla (1987) has recently proposed that three associations are acquired during instrumental training: a weak S-R, an R-O, and an S-(R-O). The latter signifies that the animal has learned that "when stimulus S is present, making response R will result in the delivery of outcome O." Only when R-O associations are being formed can we suggest that an animal subject is attributing the presentation of the outcome to its own behavior.

Killeen (1981; Killeen & Smith, 1984) took a more direct approach: he *asked* his subjects to tell him whether they caused an event or not. Pigeons were presented with three keylights and a food hopper. A peck on the center keylight occasionally turned it off, and it was sometimes turned off independently of a keypeck. When the center key went dark, two side keys were lit, one corresponding to "I caused it" and the other corresponding to "the computer caused it." The pigeon was trained to peck the "I caused it" key when the darkening of the center key was caused by their peck and to peck the "computer caused it" key when darkening was independent of their behavior; a correct response resulted in food delivery. Pigeons mastered the task within several days and reached 80% to 90% accuracy within a couple of weeks. It may still be argued that the pigeons developed some behavioral strategies that allowed them to solve this task without "inferring causality." However, human perception of causality may invoke the same strategies.

Thus far, I have presented no *definitive* evidence (nor will I) that an animal's behavior is a demonstration of causal attribution, despite the similar problems being faced by animals being conditioned and humans learning about their environment. As is common in science, further evidence bearing on this issue must be examined to help determine the weight of support for this hypothesis. In the following examination of the attribution-

Table 1
Definition of the Modern Neo-Humean Cues

Neo-Humean Cue to Causality	Description
Contiguity	
Temporal	An effect should quickly follow its causes.
Spatial	Causes should be in close spatial proximity to their effects.
Temporal priority	A cause always precedes its effects.
Covariation/contingency	The occurrence of an effect should depend on the occurrence of its causes.
Cue interaction/multiple causality	There are many candidate causes for an effect: these will frequently not operate independently of one another.
Prior experience	A causal attribution will depend not only on bottom-up data but also on an organism's expectations, given its prior experience.

conditioning analogy, I will systematically progress through a set of neo-Humean cues (see Table 1). Illustrative examples of the human use of these cues are followed by discussions of the status of each cue in animal learning. The manuscript closes with a proposal for a research strategy for increasing our understanding of causality by expanding the limits of the learning framework.

Temporal Contiguity

By definition, contiguity in time requires that two events, a cause and effect, occur one after the other with no delay. Hence, one aspect of a causal relation is that a causal event be immediately followed by its effect. However, there is a lack of clarity in the definition of required temporal relations. Hume's eighth principle suggested that a cause should not be operative "for any time" or it will no longer be considered a sufficient cause (events or objects that are operative for longer periods but satisfy other requirements are called "conditions"). In recent research, the requirement of absolute temporal contiguity (effect must immediately follow) has been relaxed. It is common to manipulate the delay between cause and effect as an independent variable, examining the effect that various degrees of contiguity have on causal judgment.

As a general principle, when the cause and effect are unfamiliar and minimally similar to other known causally related events, the closer in time the two events are, the more likely the events will be judged as causally related, assuming adequate support from the other cues to causality (Schlottmann & Shanks, 1992; Shanks, Pearson, & Dickinson, 1989; Wasserman & Neunaber, 1986). One method of bridging this time gap, thus increasing the likelihood of judging a temporally distant event as the cause, is to fill the gap with other events that are temporally contiguous with the preceding and following events (see Einhorn & Hogarth, 1986; Gruber, Fink, & Damm, 1957; Reed, 1992; Young & DeBauche, 1993). This establishes a causal chain of events, each of which is a cause of the next event in the chain. The first event

in the chain is the instigator and is judged as the ultimate cause.

Temporal contiguity between cause and effect is important for young children's causal judgments. Siegler and Liebert (1974) found that children of 5–9 years of age selected a contiguous event as cause more often than an event that preceded the effect by 5 sec. Mendelson and Shultz (1976) reached a similar conclusion using 4- to 8-year-olds. In both of these studies, temporal contiguity assumed greater importance than covariation for the youngest of children. This is surprising, given the assumed superiority of covariation in much of the literature. Any conclusion regarding the observed superiority is tempered by the observation that the importance of the cue was situation dependent; the Mendelson and Shultz children chose a cause with greater covariation and lower contiguity when there was an appropriate rationale for the delay (in this case, the cause, a moving ball, purportedly had to travel some distance before reaching the effect, a bell). This "rationale" functions as an event with temporal extent which bridges the delay.

Previous work has demonstrated that too short of a period of time between the cause and effect can lower causal judgments just as can too long of a period can (Michotte, 1963; Schlottmann & Shanks, 1992). The optimal temporal delays vary according to the nature of the cause and the effect. When judging which of various events has caused gastric upset, a subject is more likely to choose food ingestion that occurred some time ago as opposed to ingestion that immediately preceded the effect. This behavior by humans is usually explained by appeal to knowledge of the time it takes for a selected cause to bring about the effect. This knowledge may have been acquired from prior experience or cultural transmission.

The effect of temporal contiguity on animals. During the early years of the study of animal learning, it was generally accepted that temporal contiguity between two events (the CS and US, or response and reinforcer) was the critical factor in learning. Pavlov (1927, p. 89) noted that at longer CS–US intervals, the conditioned response was "difficult if not impossible" to establish. Grice's (1948) study of the effect of delay of reinforcement in instrumental training established further that contiguity was a variable of import for animal learning.

An early difficulty with the contiguity requirement in classical conditioning centered on data demonstrating trace conditioning. In *delay* conditioning, the CS onset precedes the US onset (thus providing predictive value) and the CS is still present at US onset (see Figure 2). In *trace* conditioning, the CS offset precedes US onset, thus creating a trace interval between the two events (hence no direct temporal contiguity). Pavlov (1927) and others appealed to the presence of a *mental trace* of the CS that extends beyond its actual presence in the environment, and it is this trace which is contiguous with the onset of the US. This idea is prominent in many contemporary theories (see, e.g., Desmond & Moore, 1988; Durlach, 1989; Sutton & Barto, 1981).

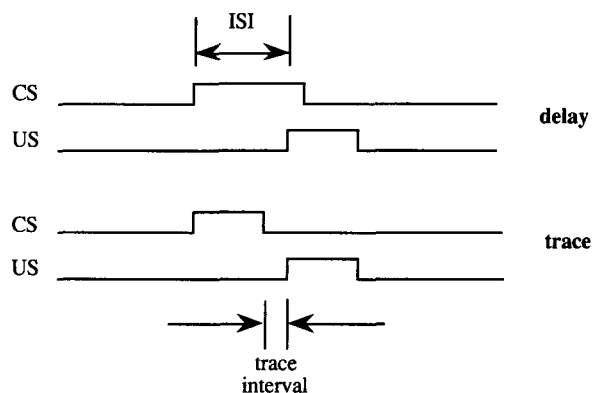


Figure 2. Delay and trace conditioning paradigms in classical conditioning. The interstimulus interval (ISI) is identical, whereas the trace interval for delay conditioning is zero.

A decrease in the relative importance of contiguity to animal learning came about in the wake of an early study by Rescorla (1967); his paper instead stressed the importance of a contingent relationship between a CS and US. If two animals receive equal numbers of contiguous pairings of a CS and US but the second animal also receives a large number of noncontiguous pairings, the first animal demonstrates much better acquisition of a conditioned response to the CS: the CS has greater predictive power for the first subject. Rescorla and Wagner (1972) tempered this emphasis on contingency by presenting a model that relied solely on contiguity (including contiguity of a US with contextual stimuli other than the CSs) which accounted for the findings of Rescorla (1967). However, the damage to contiguity had been done, and most modern animal learning texts and experiments relegate contiguity to a necessary but not sufficient variable for the establishment of a conditioned response. Its necessity is important, in that a CS must typically precede the US by a relatively short period of time for learning to take place (see, however, Garcia & Koelling, 1966); but in the absence of a contingent relationship, temporal contiguity has little or no impact on acquisition.

This is in accord with the role of contiguity in human causal attribution. It was noted that decreases in contiguity between the cause and effect result in corresponding decreases in the judgment of a causal relationship between the two events (Reed, 1992; Schlottmann & Shanks, 1992; Shanks et al., 1989; Siegler & Liebert, 1974; Wasserman & Neunaber, 1986). This generally holds true in conditioning as well. Cause–effect intervals that are too short can also result in weaker performance. Smith, Coleman, and Gormezano (1969) presented one of the most commonly cited studies demonstrating this, in which they used a rabbit eye blink conditioning preparation: acquisition of the CR was much better with a 200-msec CS–US delay than with a 50- or 100-msec delay. However, I am uncomfortable with making a strong argument for a causality conditioning homology here,

given the difference in time scales. Human adults will frequently rely on their previously established causal theories to eliminate as causes those that precede the effect by too short a period of time. To the extent that these theories are rooted in an individual's personal experience, we may effectively argue that these processes are similar to those of a nonhuman subject. Animals do encode the expected temporal interval between CS onset/offset and US onset (see, e.g., Kehoe, Horne, Macrae, & Horne, 1993; Kehoe & Napier, 1991).

Sufficiency of contiguity. The Rescorla-Wagner model (Rescorla & Wagner, 1972) provides a mechanism by which contingency-like performance can be exhibited by a system relying solely on contiguity. The comparator theories (Gibbon & Balsam, 1981; Miller & Schactman, 1985) also suggest that CS-US contiguity will produce learning regardless of the presence of other predictors of the US; the reason that there is an apparent reliance on contingency is a function of performance, not learning. The Rescorla-Wagner model and Gibbon and Balsam's (1981) comparator theory are both formulated with enough precision to warrant closer examination. The Rescorla-Wagner model has also been used to make predictions in human causality judgment (e.g., Baker et al., 1993; Shanks, 1989; Shanks & Dickinson, 1987), whereas comparator theory has not.

Rescorla and Wagner (1972) presented a framework in which a stimulus, the US, can support a limited amount of associative strength, λ . This quantity represents the importance of predicting a particular US. The degree to which a US is not being predicted—that is, the degree to which it is surprising (Kamin, 1969)—represents the amount of learning that is still necessary to reduce the surprisingness of the US. This is captured in the model as the difference in λ and V_T , where V_T is the sum of the associative strengths between all CSs present on a trial and the US—that is, the degree to which it is currently being predicted. Learning during a trial is a function of this difference:

$$\Delta V_x = \alpha_x \beta (\lambda - V_T), \quad (1)$$

where α_x is a parameter representing the salience of an individual CS (X in this equation) and β a learning rate particular to the US; $\lambda = 0$ when the US is absent. The associative strength between every CS present on a given trial and the US is modified by using Equation 1. The Rescorla-Wagner model accounts for contingency-like learning by treating the context as just another stimulus. When a US appears in the absence of a CS, thus decreasing contingency, the context gains associative strength. When the US appears in the presence of CSs, the associative values of the context and the CSs will sum to determine V_T . If the context acquires sufficient strength, it will interfere with the acquisition of strength by CSs during trials in which they occur, because the context is omnipresent. Thus, according to the Rescorla-Wagner model, a decrease in "contingency" is accompanied by an increase in context-US associative strength through a contiguity mechanism, ultimately resulting in

a decrease in the strength of the relationship between CSs and the US.

The Rescorla-Wagner model has been relatively successful at predicting the ordinal strength of causes and effects in human causal attribution (see, e.g., Baker et al., 1993; Shanks, 1989; Shanks & Dickinson, 1987). It provides more than simply asymptotic formulas, it is a reasonable description of the acquisition of causality judgments (Shanks & Dickinson, 1987), and it successfully predicts how CSs or causal candidates interact when more than one are present. In general, its ordinal predictions of the effect of various contingency schedules are correct, but there are some exceptions (see Shanks & Dickinson, 1987).

The comparator theories (Gibbon & Balsam, 1981; Miller & Schactman, 1985) represent an alternative class of contiguity models. Comparator theories suggest that subjects learn about every CS-US relation that they experience. Performance, however, is not a direct function of learning. In contrast with the Rescorla-Wagner model, CS-US contiguity is posited to produce learning regardless of the presence of other predictors of the US; but unless the CS is presented under conditions favorable for behavioral expression of that learning, knowledge of the CS-US relation will not be observed. Conditioned performance is determined by the associative value of the CS *relative to the value of its context*; hence it is the result of a comparison. When the value of the CS is greater than the value of the context, the CS will elicit a conditioned response; when it is lower in value, it will inhibit a conditioned response.

In Gibbon and Balsam's (1981) scalar expectancy theory (SET), the context is literally the experimental context; Miller and Schactman (1985) suggest that context may also be another punctate stimulus. There are other differences as well; Durlach (1989) does a nice job of comparing their relative predictions as well as predictions of the Rescorla-Wagner model. The principle difference between comparator theories and many associative theories is a difference in the relation between learning and performance. Each class of theories accounts for some data which the other does not. SET was originally developed to address a shortcoming in the Rescorla-Wagner model with respect to the parceling of an experimental session into discrete trials. Given the effect that the intertrial interval (ITI) has on learning, Gibbon and Balsam (1981) suggested that the associative connection between a CS and US is a function of the time between USs (the cycle duration, or C) and the length of time that a CS is present during a cycle (T). Learning is a function of the ratio C/T ; for example, as the cycle duration increases, learning rate increases (the classic advantage of spaced over massed trials).

Research in causal attribution has addressed in a very limited way the issue of learning a limited set of relations versus learning all of them but exhibiting only some. In certain cases this limitation is the result of an assumption that causal relations are not learned (Harré & Madden, 1975), but it is generally the result of the lack

of a methodology to resolve the issue. Researchers advocating an associative view of causality (Chapman & Robbins, 1990; Shanks & Dickinson, 1987; Wasserman, 1990b) commonly assume that only some relations are learned, whereas those advocating a more normative, contingency-based view (Cheng & Novick, 1992; Shaklee & Elek, 1988; Waldmann & Holyoak, 1992) suggest that most relations are encoded but only some are exhibited. Cheng and Novick (1992) proposed a view of causal attribution which bears a resemblance to that of the comparator theories. The production of a causal inference is dependent on the context in which the inference is required (they call this context a *focal set*). Unfortunately the choice of a focal set is relatively underdetermined, making it difficult to test the theory against alternative formulations (Shanks, 1993a).

The sufficiency of contiguity has been questioned in the causality literature as it has been in animal learning. When causal mechanisms are known, they facilitate a causal judgment in the absence of strict contiguity (see, e.g., Shultz & Kestenbaum, 1985). Einhorn and Hogarth (1986) apparently de-emphasized contiguity by excluding it from their theoretical equation of attribution. However, it is included as part of their requirement of a causal link and is very important in the construction of causal chains. Recall that filling a gap was shown to be important in alleviating the detrimental impact of a lack of contiguity between two events (Gruber et al., 1957; Reed, 1992). This work was largely motivated by the effect that this manipulation has on conditioning. Filling a trace interval is widely recognized as significantly alleviating the decrease in the rate of acquisition in trace conditioning as compared to delay conditioning (Bolles, Collier, Bouton, & Marlin, 1978; Egger & Miller, 1962; Kaplan & Hearst, 1982; Kehoe, Gibbs, Garcia, & Gormezano, 1979; Rescorla, 1982). *Serial conditioning* leverages this facilitatory effect: two or more sequential CSs are co-contiguous, with the last CS immediately preceding the US. For example, $A \rightarrow B \rightarrow C$ represents the initial occurrence of stimulus A , followed by B and finally by C (the US). The presence of B fills the temporal gap between A and C and facilitates the learning of the $A \rightarrow C$ relationship. Analogous results have been found in instrumental training of animals (Reed & Reilly, 1990).

Rescorla (1982) sought to determine *how* an intervening stimulus improves the acquisition of the relationship during classical conditioning. In Experiment 1, he used a between-subjects design involving four different manipulations of the trace interval. All, none, the first half, or the second half of a 10-sec trace interval was filled with another stimulus. As expected, filling the interval completely facilitated responding to the distant CS. Filling the first half was beneficial early in training, whereas filling the second half had only a small effect. Experiments 3 and 4 largely ruled out second-order conditioning as an explanation of these results whereby the intervening stimulus (IS) functioned as a reinforcer of the CS through conditioned reinforcement properties. Stimulus generalization was also eliminated as an ex-

planation via extinction of the IS following training—enhanced responding to the CS was still observed. Rescorla's conclusion was that the IS facilitated the association between the CS and the US, appealing to the rather vague notion of a Gestaltist perceptual continuity.

An alternative account of the facilitatory effect of filling the gap is that doing so differentiates the ITI from the trace interval (Mowrer & Lamoreaux, 1951, as cited in Kaplan & Hearst, 1982). This explanation is bolstered by results showing that filling the ITI rather than the trace interval also facilitates CS-US association (Kaplan & Hearst, 1982). From a causality perspective, this suggests that the absence of an event can function as a cause. One aspect of the causal explanation of serial conditioning is that the intervening stimulus is functioning as part of a causal chain (Young & DeBauche, 1993). Hence, the relationship between the first event, the CS, and the second, the IS, is important.

Kaplan and Hearst (1982, Experiment 4) manipulated the temporal relationship between the event filling the ITI, a tone, and a keylight CS. They describe their groups in terms of the degree of overlap between the tone and CS, but I will describe them in terms of the relationship between the CS and IS (absence of tone) to facilitate the causal chain perspective on this research. The relationship among stimuli in the four experimental groups is illustrated in Figure 3. They observed good responding to the CS in the complete overlap group, fair in the partial overlap group, and little or no responding in the no-overlap and trace compound groups. Focusing on the CS-IS relationship, we have delay conditioning for the complete overlap and partial overlap groups, and simultaneous conditioning for the no-overlap group. Therefore, it should not be surprising that the no-overlap group results in poor CS-US conditioning; the CS-IS relationship will be very weak, preventing the IS from playing a role in the CS-IS-US causal chain. The discrimination hypothesis offered by Mowrer and Lamore-




Figure 3. Relationship among stimuli in Experiment 4 groups of Kaplan and Hearst (1982).

aux cannot account for these data. The analysis I have offered is analogous to that of Kaplan and Hearst (1982), who note further that the Rescorla-Wagner model can be extended to encapsulate these data.

Spatial Contiguity

Spatial contiguity necessitates a treatment analogous to that of temporal contiguity. Taking Hume literally, two causally related objects must be in spatial contact with each other, although we do find action at a distance (e.g., with magnetism). This dimension has also been investigated empirically (Gruber et al., 1957; Leslie & Keeble, 1987; Michotte, 1963; Schlottmann & Shanks, 1992), and the results suggest that spatial contiguity is a necessity for older children and adults, whereas younger children's use of the cue is inconsistent. A method parallel to that used for bridging temporal gaps is applicable to bridging spatial gaps (Gruber et al., 1957). The presence of a third object (or series of objects) that is spatially contiguous with both a cause and its effect connects the two. It is well known by billiard players that it is possible to cause a ball to move by striking another that is currently in contact with it. If the situation is just right, the intermediate ball will not move. Bridging spatial gaps is also evidenced when a person uses a long pole to push a box. In the absence of the pole, it would be unlikely for an observer to attribute the movement of the box to the person standing a few feet away.

At this point I must be careful to make a distinction between the roles I suggest for such bridgers of temporal and spatial gaps and the role suggested by the "generative theorists" (e.g., Shultz, 1982; White, 1989). I argue that these bridges facilitate attribution via basic learning mechanisms rooted in conditioning. This is suggested by the facilitatory role played by such bridges in nonhuman species. For the generative theorists, the most important cue to causality is knowledge of the mechanism by which the effect can be generated by a cause; temporal and spatial bridges serve as mediums for energy transfer. One can argue that at a conscious level, adults do at times rely on just such a theoretical argument for their causal judgments, but I suggest that this use of bridges occurs only after a personal causal theory has emerged from prior experiences. Adults frequently attribute causality and only later create a verbal explanation to support their (already made) conclusion, and young children have little, if any, concept of energy transfer at ages when they are inferring causality.

In classical conditioning experiments, the CS and US locations have usually been relatively close to one another given the confines of the typical apparatus. The same is true in instrumental training wherein the stimulus and response locations are kept close to the site of reinforcement. When general spatial contiguity is weakened, conditioning can become very difficult (see, e.g., Breland & Breland, 1961; Killeen, 1981; Rescorla & Cunningham, 1979; Testa, 1975). Rescorla and Cunningham (1979) per-

formed one of the very few animal learning studies addressing the relative importance of spatial contiguity to conditioning. Spatial contiguity was observed to facilitate the development of conditioned responding when the temporal relation between two events was held constant. However, the power of spatial contiguity as a cue is limited to a facilitatory role. If we observe changes in performance due to spatial contiguity in the absence of temporal contiguity, it is more natural to appeal to processes such as stimulus generalization or pseudoconditioning rather than associative learning. Temporal and spatial contiguity are not treated analogously in theories of conditioning.

Shultz and Kestenbaum (1985) suggested that spatial contiguity was more fundamental than temporal contiguity and was second only to temporal priority. The importance placed on spatial contiguity as a cue to causality in human research is not paralleled in the animal literature. The few studies that have been done relegate it to a secondary role. In order to rectify these positions, it is necessary to examine the tasks being used in the two types of research. In the human causal attribution literature, spatial contiguity is typically manipulated in launching tasks and other collision tasks (e.g., using a stick to push an object, or a ball striking an unseen object causing a Jack-in-the-box to pop up). In contrast, Testa (1975) and Rescorla and Cunningham (1979) used procedures with the CS and US in similar locations. These two classes of tasks differ with respect to the specificity of prior experience that the subjects can be said to have. Collisions involving one object striking or pushing another are very common in the experience of human (and probably animal) subjects. Although even more common, similarity in location between a cause and effect is more abstract than these specific collision experiences; collisions may be more likely to retrieve information on other collisions, owing to cue specificity.

Temporal Priority

Hume's second rule stated that the cause must be prior to the effect. There has been considerable debate in the developmental literature over the necessity of priority. Early research (e.g., Kuhn & Phelps, 1976; Shultz & Mendelson, 1975) suggested that young children do not rely on temporal priority as a causal cue. However, there was a source of difficulty for children in these studies: they were required to verbalize their knowledge (Bullock, 1984). It is possible that a child may understand the causal relations observed but have difficulty in correctly verbalizing this knowledge. Bullock and Gelman (1979) opted for a nonverbal task: to indicate their inference, subjects simply pointed at one of two alternative causes of an event. The majority of 3-year-olds (an average of 13 of 16) chose the temporally prior event, and almost all of the 4- and 5-year-olds chose this event. When asked to provide verbal explanations during the experiment, only 12% of 3-year-olds used an order argument,

whereas 60% of 5-year-olds did. This supports their assumption that verbalization is a poor assessment of the causal knowledge of very young children.

Although it is recognized by young children as a cue to causality, when contrasted with other causal cues, temporal priority may not assume a high status until later in childhood (Shultz, Altmann, & Asselin, 1986; Sophian & Huber, 1984). Sophian and Huber found that 3-year-olds did not emphasize temporal priority in the presence of other cues when the cause and effect were separated by 200 msec. In contrast, over 50% of 5-year-olds chose the temporally prior event despite contrary indications involving other causal information. Shultz et al. demonstrated that young children (3–4 years) relied on their prior experience rather than on temporal priority in deciding which of two objects was the cause and which the effect. Older children (7–8 years and 11–12 years) were significantly better at using temporal priority as a cue.

Temporal priority has been studied in animal learning by contrasting the learning of cues, or CSs, that precede the US with those that co-occur with or follow the US. It is commonly accepted that forward conditioning is effective (the CS precedes the US), whereas backward (the CS follows the US) and simultaneous conditioning are generally ineffective (see, e.g., Smith et al., 1969). Mackintosh (1974) concluded that "there is little reason to accept the reality of backward conditioning" (p. 60). As a result, it has been assumed that the CS must provide predictive value for its presence to be noted and remembered by the subject. In studies of learning following backward conditioning, it is generally found that a CS which follows the US becomes a conditioned inhibitor, ostensibly because it signals the termination of the US (e.g., Kamin, 1963, and Moscovitch & LoLordo, 1968, both as cited in Mackintosh, 1974).

Although this view is still prevalent, Spetch, Wilkie, and Pinel (1981) surveyed the empirical research and found good evidence for backward excitatory conditioning. The conditions under which it is observed are limited, however. Excitatory conditioning occurs after a few initial pairings followed by conditioned inhibition as training progresses. The excitatory phase is relatively short and only appears to occur with aversive stimuli. Wagner and Terry (1975) proposed that associations between the CS and US depend on the simultaneous rehearsal of the two events in short-term memory. This offers a method for rectifying the existence of backward conditioning with contemporary models that emphasize prediction: if the US trace persists beyond its termination, then a CS that follows it can be used as a predictor of this trace. The duration of the trace may be considerably longer when the US is aversive. To explain the transition from excitatory to inhibitory conditioning, Wagner and Terry proposed that US persistence is also dependent on its surprisingness, which decreases during learning.

Recent evidence suggests a slightly different theoretical perspective on the necessity of temporal priority during conditioning. Matzel, Held, and Miller (1988) used an alternative method for determining whether backward

conditioning could result in the formation of backward associations. In previous studies, experimenters measured subject response to the CS that followed the US. In contrast, Matzel et al. measured responding to a second CS that predicted the CS that followed the US. In the prototypical design, they first preconditioned a group of rats to a 5-sec click immediately followed by a 5-sec tone. The tone was subsequently paired in the next phase with a 5-sec shock in a forward relationship (tone→shock), a backward relationship (shock→tone), or simultaneously. As would be expected from previous research, the tone only elicited a significant CR following forward conditioning. However, the preconditioned click elicited a CR after all three types of training. They suggest that the rat encoded the following relationships depending on the training received: forward training, click→tone→shock; backward training, click/shock→tone; simultaneous training, click→tone/shock. In all these cases, the click is more predictive of the shock arrival than the tone is, and subsequently the click is capable of eliciting an avoidance response. Their conclusion was that temporal relations were encoded, including backward associations, but that this information was only utilized when it was of functional, predictive value.

The work of Matzel et al. (1988) suggests that a conditioning account does not preclude learning of backward associations. The observer can encode the temporal relations among environmental events and utilize this information in either a forward or a backward fashion. More generally, comparator theories (Gibbon & Balsam, 1981; Miller & Schactman, 1985) adopt the position that all CS–US relations are learned by a subject, but that only a subset of those are expressed as befits a given situation. This suggests that temporal priority is a performance, not a learning, variable. Although it is unlikely that *all* stimuli relations are learned, it is quite probable that only some of what is learned is being observed in a given experiment.

Temporal priority is thus important to animals as well as humans in making a predictive judgment. However, it is clear that the importance of this causal cue relative to the others changes considerably during human development. It is difficult to assess the presence of a similar trend in nonhuman animals, since the experiential base of laboratory animals is relatively limited.

Covariation

A very large body of research has been dedicated to the understanding of human judgments of covariation and contingency (e.g., Allan & Jenkins, 1983; Arkes & Harkness, 1983; Benedict, 1991; Chapman & Robbins, 1990; Cheng & Novick, 1992; Dickinson, Shanks, & Evenden, 1984; Mendelson & Shultz, 1976; Shaklee & Tucker, 1980; Shultz & Mendelson, 1975; Ward & Jenkins, 1965; Wasserman & Shaklee, 1984; Yates & Curley, 1986). This extensive literature has been recently reviewed from various perspectives (Allan, 1993; Cheng, 1993; Shanks, 1993b). A limited presentation of this history is necessary for framing the discussion, but I will orient the dis-

discussion around the issues that are important to an animal learning analysis.

Hume has often been criticized for his insistence on constant conjunction as the primary cue to causality. A number of publications have been dedicated to debunking the primacy of covariation (the modern embodiment of "constant conjunction"). Koslowski, Okagaki, Lorenz, and Umbach (1989) suggest that "covariation is not enough" in the title of their article. Their data do support this argument but simultaneously demonstrate the primary importance of covariation for subjects in their experiments. Shultz (e.g., 1982) has long argued that knowledge of a possible mediating causal mechanism is more important than covariation to children and adults, but his own earlier work (e.g., Mendelson & Shultz, 1976; Shultz & Mendelson, 1975) has illustrated the importance, if not the primacy, of covariation. In a recent article, White (1992) suggested that covariation does play an important role in causal attribution but proposed that it is used only when certain other types of information are not available and when the experimental question posed specifically refers to multiple observations. Although I believe that he is right in bringing to our attention the demand characteristics of the type of causality question asked, I disagree with his treatment of covariation as a secondary principle.

This disagreement is based on a bias exemplified by White, Shultz, and others. When it is observed that causal attribution can take place after a single observation, they argue that covariation is obviously not necessary. It is frequently disregarded that no observation takes place in a vacuum of experience. When a person observes one object striking another on a computer screen, he/she is necessarily going to implicitly process this experience in the context of all of his/her prior experiences with object collisions (which is quite extensive). It is arguably impossible for adults to observe any succession of events without tapping into similar events gleaned from their past, even when these historical events are only superficially similar.

A common response from the anti-Humeans is to cite studies which suggest that causal perception of launching event sequences occurs in infants as young as 24 weeks of age (e.g., Leslie, 1986; Leslie & Keeble, 1987). Assuming that these infants are demonstrating causal perception by around 6 months, that still provides for considerable experience with collision events—situations in which one object moves another—which will influence later attributions. White (1992) suggested that this is unlikely, because infants "presumably possess limited powers of analogical reasoning" (p. 163); but application of this prior experience requires nothing more than generalization, a skill possessed by many non-human animals. This "analogizing" also should not presuppose the conscious processing that is suggested by White's use of the term "reasoning." Importantly, the entire basis for this argument has been called into question by an attempt to replicate Leslie's experiments. Oakes and Cohen (1990) used real objects with 6- and 10-month-

old infants; Leslie used a cinematic display of moving red and green blocks. Oakes and Cohen's results suggest that by 10 months, but not by 6 months, infants discriminate between causal and noncausal events. The younger infants appeared to respond to individual object characteristics but not to the relationship between the objects.

Shultz (1982) has been one of the strongest advocates of a generative view of common notion causality. He considered covariation as a cue to causality secondary to a knowledge of causal mechanism. Shultz's view does not appear to be irreconcilable with a neo-Humean view, however. First, it should be noted that Shultz found it necessary to provide his young subjects with experience of the mechanisms in action before testing. Shultz (1982) stated that

it is, in short, unreasonable to suppose that experience (whether in the form of explicit pretraining or probable prior exposure) would not have influenced the present results. Some experience with particular transmissions may indeed be critical to correct application of the generative transmission rule. The issue here concerns the manner in which such experience is processed and interpreted by the child. (p. 17)

One method of rectifying Shultz's generative transmission rule (which actually appears to consist of a number of domain-specific rules) with a neo-Humean emphasis on contingency is to propose that the neo-Humean cues account for the development of domain-specific rules or theories which later interact with the cues (see Figure 1). Since in this account a theory is a distillation of a number of prior observations, the theory *should* play a major role in attributions, perhaps outweighing the current "data." Thus there appears to be no suitable objection to the argument put forth by Hume that covariation is an important cue to causality, since it also underlies the causal theories with which it later competes.

Cheng (1993) presented a similar argument for the formation of causal mechanism knowledge through a data-driven process. She noted the importance of contingency data for the grounding of causal laws. Her analysis centered on contingency as a grounding principle, but there are limits to the explanatory power of this cue, some of which can be addressed through consideration of additional cues. However, a persistent limitation of the neo-Humean framework is reflected in Figure 1 by the additional effect of public causal theories. An animal learning perspective has no obvious explanatory power for learning due to linguistic instruction. For example, Rosenfarb, Newland, Brannon, and Howey (1992) exposed subjects to an action-outcome contingency and found that their performance was a reflection of their verbal knowledge, even when that knowledge was manipulated by the experimenters. Although contingency cues do not operate independently of other processes, they are still central to the development of personal causal theories.

It should not be concluded that Hume had the last word on the notion of constant conjunction. This is one cue to

causality that has undergone considerable revision. Hume used the term *constant* conjunction, but elsewhere in his *Treatise*, he gave consideration to the notion of frequent conjunction (Hume, 1739/1969). There are many conditions under which an effect may not follow its usual cause, yet such observations do not seriously dampen our causal judgments. When striking a match fails to light it, we simply try again, not knowing what conditions may have prevented its lighting the first time.

Suppes (1970, 1984) argues that not only are causal relations apparently probabilistic to the human observer, but the basic metaphysics are probabilistic. He argues that Kant and Hume were products of their time: science was dominated by Newtonian mechanics and the deterministic nature of events. Any relationship that was apparently probabilistic was judged so only because of a lack of knowledge of all the factors impinging on a given event's occurrence. Once all the factors were known, which may be infinite, it would be theoretically possible to predict the future occurrence of any event. This notion still carries great appeal today. Even if the universe may be deterministic, pragmatism requires us to treat causal relations as probabilistic. Attempts to understand the complexity of everyday phenomena have given rise to modern chaos theory. Researchers in chaos have accepted that although there may be an underlying order, this does not imply that events are predictable.

Temporal invariance as sufficient. I find it necessary at this stage to digress in order to distinguish between a causal relationship and "mere" temporal invariance. Examples of the latter are frequently cited to discount a Humean perspective. It is adaptive to be able to predict future events, but these predictions can be based on temporal invariances as well as causal relationships. For example, when we observe someone taking a step, we can judge with reasonable certainty that another step will follow. Simply that we can confidently predict an event does not imply that a causal relationship is necessarily present.

Consider the prototypical example of night following day (and vice versa). Here we have a temporal relationship that is invariant but is not described as causal: day does not cause night. We must first note that the temporal parameters in operation for day–night are much longer than is represented in everyday causal judgments. Days are long and thus even by Humean principles cannot be causes, since events lasting for "any time" cannot be causes. However, appeal to the duration of these events does not account for circumstances in which the events are long but inferences are still made. For example, the orbit of the moon causes the tides, short days cause deciduous trees to lose their leaves, and the oxygen in water causes metal to rust. These examples apparently discount one of Hume's principles. What differentiates the day→night scenario from the moon→tide and water→rust scenarios? It is the fact that day and night are not distinct from one other: each represents segments on a continuum of lamination.

Although events of great duration are candidates as causes, it is likely that such attributions are not readily made. These attributions are likely the result either of a pseudoscientific method whereby conditions are carefully observed to ensure that distant events remain proximal in an observer's analysis, or of noting the difference between an observed situation and other similar situations in the observer's experience. When an individual sees rust on metal in a damp room, the unusual humidity level of the room suggests it as a candidate cause. The study of attributions involving long events is infeasible in the laboratory unless the events are verbally summarized. It is thus difficult to establish or discount the operation of Hume's eighth rule in common notion causal attributions.

Consider next an invariant sequence which consists of shorter duration events: one clap does not cause another. If we could perform an experiment to test the contingency of the two events, we would conclude that they are not causal, because there is no contingent relationship. In making everyday judgments, we frequently do not have the luxury of manipulating our environment to determine contingency. So why are series of events like clapping or walking not perceived as causally interrelated? When very similar events (one clap is similar to another) occur at close, regular intervals, they may be perceived as a unique gestalt event (H. J. Hilton, personal communication, February 23, 1993). Thus we describe a sequence of claps as *clapping*, a sequence of steps as *walking*, a sequence of up and down head motions as *nodding*, and a series of short, staccato sounds as *tapping*. By treating the sequence as a gestalt event, we can determine the cause of the event *tapping* and in turn treat it as a potential cause—for example, of a headache.

Another oft-mentioned refutation of Hume's cues is the two-clocks example (R. Giere, personal communication, March 3, 1994). Two different clocks are sitting next to each other (spatial contiguity), with one set to run 1 sec slower than the other. If they both strike the hour with a single chime, then every hour (covariation) the chime of one will be followed 1 sec later by the chime of the other (temporal priority and temporal contiguity). Hume's cues to causality are apparently satisfied, but this is a temporal invariance which is not judged as causal. However, given our prior knowledge about clocks (see the next section on prior knowledge), our extensive experience suggests that clocks do not cause other clocks to chime. An observer without this experience might very well judge the events as causal. In fact, we could rig the experiment so that the second clock echoed the first clock's chime through an electronic sound detector, so that a person within our culture would erroneously judge the events as noncausal.

Cultural rituals represent another collection of temporally invariant events which are not described as causally interrelated. These examples carry little weight, because most rituals either are not truly invariant because they have lots of variations on a theme, or are not

similar to other invariant sequences in the person's experience, thus lacking evidence for a causal judgment.

Although I cannot pretend to have developed an exhaustive list of the temporal invariances typically found in our world, I have surveyed some that are commonly cited as presenting difficulties for the Humean view. Temporal invariances that involve long events (such as day and night) will not naturally be perceived as being directly causally related. When the events are of short duration, and when they are similar and follow each other in relatively quick succession, the series of events is perceived as a longer, unique event. As the durations of the component events lengthen and their similarity decreases, I predict that causal attributions will become more common. Other examples of invariant sequences (e.g., cultural rituals) either lack the unvarying relation typified by causal chains of events or are unique, with no similar sequences that are also experienced as invariant. Remember that causal attribution is not a function of one cue to causality or strictly of the current perceptual evidence. Evidence for an attribution is taken from multiple cues and is influenced by prior experiences (in the form of personal causal theories) of these and similar sequences.

Covariation versus contingency. Covariation and contingency are often used interchangeably in the literature on human judgment of event relationships. This is a loose use of terminology, one that I have adopted in the discussion so far. It is important to note, however, that covariation (as measured by the various correlation statistics) does not imply a direction of dependence; there is an assumed *two-way* relationship between the correlated or covarying variables. Statistics such as Pearson's *r* for continuous variables and Pearson's χ^2 for binary variables represent bidirectional dependencies. As Allan (1980) demonstrated, Pearson's χ^2 is proportional to the products of two ΔP contingency measures, $P(A|B) - P(A|\text{not } B)$ and $P(B|A) - P(B|\text{not } A)$. However, in the majority of experiments in which subjects have been asked to make a "covariation" judgment (e.g., Alloy & Tabachnik, 1984; Shaklee, 1983; Shaklee & Elek, 1988; Shaklee & Mims, 1981; Shultz & Mendelson, 1975), the instructions have in reality asked for a unidirectional judgment. This is not surprising, given that *causality is directional*; correlation does not imply causality.

The problems arise when researchers examine "rules" that their subjects might be using in making their covariation (actually contingency) judgments. Consider the 2×2 contingency table presented as Table 2 (cell entries are frequencies).

Shaklee (1983) compared the degree to which subject performance appeared consistent with a "cell *a*" rule

(considering only the frequency of co-occurrence of the two events), an "*a* versus *b*" rule (the difference between frequency of co-occurrence and frequency of occurrence of the effect when the cause was absent), a "sum of diagonals" rule [$(a+d) - (b+c)$], and a "conditional probability" rule [$P(E|C) - P(E|\text{not } C)$, which is equivalent to $a/(a+c) - b/(b+d)$]. The "cell *a*" rule and "sum of diagonals" rule both lack directionality from the cause to the effect. If subjects being asked to make a contingency judgment are using covariation data (e.g., Shaklee, 1983), we should be curious why. Shaklee and Mims (1981) presented fourth grade through college age subjects a series of unidirectional problems (e.g., plants pictured as healthy or sick as a function of the presence or absence of plant food, bug spray, or special medicine). Across all age groups, a significant number of subjects responded in a fashion consistent with the sum of diagonals rule, a bidirectional assessment. The problem domains and question form do not encourage such a judgment, but the presentation of the data in a 2×2 contingency table may be sufficiently symmetrical to mislead or otherwise confuse subjects.

There have been a handful of studies investigating the effect of presentation form involving a contingency table on subject performance (e.g., Baker, Berbrier, & Vallee-Tourangeau, 1989; Van Hamme & Wasserman, 1993; Ward & Jenkins, 1965; Wasserman, 1990a; Wasserman & Shaklee, 1984). Although in many of these studies it was not possible to differentiate unidirectional from bidirectional rule use, since the "*a* versus *b*," "sum of diagonals," and "conditional probability" rules would all have come to the same conclusions, a recent study by Kao and Wasserman (1993) permits a detailed assessment of the effect of presentation form. In Experiment 2, they presented subjects with a series of 8 contingent problems and 13 noncontingent problems using either a summary table or a trial-by-trial serial presentation. Given the rule classification criteria of Shaklee, the summary table presentation resulted in 21% of the subjects conforming to "*a* versus *b*" or "*a* versus *c*" rule use (unidirectional) and 17% conforming to use of sum of diagonals (bidirectional). Trial-by-trial presentation resulted in 24% unidirectional rule use and 5% bidirectional. Further evidence of the bidirectional bias created by summary table presentation is evidenced by comparing subject performance on the individual problems the authors presented: if the *b* and *c* frequency counts are reversed, subjects using unidirectional rules should be influenced more than those biased toward bidirectional use. Of the 21 problems used, 8 pairs permitted such a comparison. The average change in rating when *b* and *c* counts were swapped was .6 using summary tables and 1.13 using trial-by-trial presentation.¹ The mode of presentation has a systematic influence on judgments and must be considered in all studies of human causal attribution.

It is generally recognized in quantitative psychology that measures of association between two variables include symmetric and asymmetric metrics. Pearson's χ^2 and related measures (e.g., Cramer's *V*, Pearson's ϕ ,

Table 2
2 × 2 Contingency Table

Effect	Cause	
	C	not C
E	<i>a</i>	<i>b</i>
not E	<i>c</i>	<i>d</i>

Kendall and Stuart's contingency coefficient; Liebetrau, 1983) are symmetric measures of association in that they are insensitive to an interchange of the two variables being analyzed. Asymmetric measures (e.g., ΔP , Goodman & Kruskal's λ and τ , Somers's D ; Liebetrau, 1983) are directional: A 's dependence on B may be different from B 's dependence on A . A normative judger of event relations should use a symmetric measure of correlation when making a covariation (bidirectional) judgment and an asymmetric measure for unidirectional contingency judgments. This distinction has not been honored in much of the judgment of covariation literature (see Bindra, Clarke, & Shultz, 1980, and Einhorn & Hogarth, 1986, p. 8, for related discussions).

But do we engage in complex contingency analysis when we attribute causality? The question is: what do we do and how do we do it? It is doubtful that we *consciously* calculate the probabilistic relations among events or perform set operations on our episodic memories. Wasserman et al. (1993) rejected a statistical model of causal judgment when they found that their subjects' estimates of separate conditional probabilities did not map well onto their estimates of the strength of a causal relationship. An alternative is that contingency and/or covariation data are acquired implicitly—that is, unconsciously. Lewicki (1986) and Reber (1989) demonstrated that not only can covariation information be detected implicitly, but our implicit methods are superior to our explicit ones (cf. Brody, 1989). In addition, the ability of animals to encode contingencies also suggests the possibility that implicit mechanisms may provide good estimates of measures of association.

Animal learning, covariation, and contingency. Confusing covariation and contingency has not typically been a problem in animal learning and associationist accounts of causality. The principle reason is the requirement that for learning to occur, a CS must precede the US and a response must precede the outcome, thus stressing the directional nature of the relationship. Until recently, it was assumed that animals required a forward, predictive relationship between stimuli, thus implying a strong unidirectionality. Animals are described as making a contingency judgment; their goal is to determine on which cues and/or responses the arrival of a biologically relevant stimulus is contingent. For example, in Pavlovian or classical conditioning, when $\Delta P = P(\text{US}|\text{CS}) - P(\text{US}|\text{no CS})$ is positive, responding to a CS will be higher; when ΔP is negative, responding will be lower following the CS; and when ΔP is 0, the animal learns that the CS is irrelevant (see Mackintosh, 1974, for a discussion of learned irrelevance). Parallels hold in instrumental training for the controllability of an outcome, $P(\text{O}|\text{R}) - P(\text{O}|\text{no R})$, and for the elicibility of a discriminative stimulus, $P(\text{R}|\text{S}^D) - P(\text{R}|\text{no S}^D)$. This does not imply that covariation is never mistaken for contingency in the discussion of animals, although interestingly the confusion of terms is invariably found in pa-

pers drawing parallels between animal learning and human contingency judgments (e.g., Alloy & Tabachnik, 1984; Benedict, 1991).

As mentioned earlier, there is some disagreement about the roles of contiguity and contingency in animal learning. Most of the recent debate has centered on the necessity for contingency in classical conditioning experiments (see, e.g., Papini & Bitterman, 1990). The use of contingency as an independent variable is very common, and its use is based on the assumption that contingency is a driving variable in conditioning experiments. It is not necessary, however, to require that subjects are calculating conditional probabilities as they learn. The Rescorla-Wagner model represents one method for accounting through contiguity for many results typically attributed to contingency. If all things were equal, it would not be important in the short run to determine whether contiguity or contingency was the controlling variable in learning experiments. However, Papini and Bitterman argue that there is sufficient evidence to suggest that all things are *not* equal and that contingency represents little more than a convenient explanatory construct that fails to account for many experimental results.

This debate regarding the relation between contiguity and contingency has spread to the study of causality. Wasserman and Neunaber (1986) demonstrated a decrease in performance when the contingency between events was held constant while contiguity was decreased. They concluded that the findings implicated contiguity as the determining variable rather than contingency. Shanks (1993b) discusses some of the specific limitations of the *relative contiguity* mechanism proposed by Wasserman and Neunaber (1986), including its difficulty with accounting for the effect of temporal bridges of the sort discussed earlier. Instead, Shanks has proposed that contiguity and contingency have independent effects: contingency information is captured with a mechanism similar to the Rescorla-Wagner model, while the effect of contiguity is encapsulated by postulating that event representations decay over time. It has yet to be seen whether the relative contiguity approach can be extended to encompass the breadth of Shanks's proposal. However, contingency formulations have their own set of shortcomings.

According to contingency theories, a subject's behavior should be the result of a retrospective evaluation of the events that have occurred within the experimental context and the computation of the relative contingencies of one set of stimuli (e.g., the CSs) with another set (the USs). This suggests that trial order should have no effect as long as the overall contingency is the same. However, prior research (Benedict, 1991; Benedict & Ayres, 1972; Yates & Curley, 1986) has clearly indicated that early contingencies have a significant impact on later performance in humans and nonhuman animals. Under noncontingent conditions, an early positive contingency resulted in conditioning in rats and positive ratings by humans of a CS-US relationship and an early

negative contingency had the opposite effect. A contingency theory would need to weigh experience differentially in the calculation of probabilities.

These early contingency experiments are one demonstration that it is not strictly the CS-US contingency that is important but the particular series of events. The Rescorla-Wagner contiguity model predicts the effect of early pairings, because it is a trial-by-trial learning model. Contingency models must compute contingencies from a representation analogous to the cells in a 2×2 contingency table (see Hammond & Paynter, 1983). Computation may take place using episodic memories of all prior trials or using summary information (e.g., the 2×2 table). Although it is unlikely that all trial episodes are available as distinct memories, it is possible. It is more likely that some episodes are stored quite distinctly (e.g., those that are particularly surprising), and that those memories interact with summarized or extracted information stored in the brain. This proposal, however, would complicate contingency calculation, and it is more generally posited that the animal is logging event frequencies in a 2×2 contingency table, thus permitting local learning and allowing easy computation of contingency using whatever mechanism the brain of that species has developed (e.g., ΔP , Goodman-Kruskal λ , Pearson's ϕ , χ^2). It is interesting to note that humans are not particularly good judges of event frequencies; these judgments are affected by memory trace strength, retrievability of the memory (the availability heuristic), and similarity to other memories (representativeness heuristic) (Eysenck & Keane, 1990). Hence, the data available to a contingency computational mechanism will be imperfect.

Gibbon and Balsam's (1981) SET theory is similar to contingency theories in that it requires the computation of a ratio between overall cycle time and time that a positive trial is present between reinforcement occasions (the C/T ratio). To compute C and T, the subject is faced with the same tradeoff between storing episodic memories of all trials and storing running totals of C and T. Their model also fails to predict the effect that early contingencies will have on performance.

Studies that involve the addition of un signaled USs intermixed with USs signaled by the CS are usually considered to support the role of contingency in conditioning. As un signaled USs are added, contingency decreases, resulting in weaker responding to the CS (see, e.g., Durlach, 1983; Hammond & Weinberg, 1984; Rescorla, 1967). However, signaling these additional USs goes a long way toward alleviating the impact they have on CS responding (Durlach, 1983; Hammond & Weinberg, 1984). Contingency theories predict that signaling USs should not affect responding, because DP has not changed with respect to the CS under study.

The differential effect of signaled versus un signaled USs is predicted by associative theories like the Rescorla-Wagner model. Recall that CSs must compete as predictors of the US. When un signaled USs are added, the con-

text gains in associative strength and directly competes with the CS when both are present. Signaling the additional USs results in a second CS, the signal, gaining associative strength, decreasing context-US association through this same competitive mechanism. From a causality perspective, this is not surprising. When a single event is not a particularly good predictor of an effect, our judgments of that event as causing the effect are lower. When we observe the effect occurring sometimes following one event and sometimes following another, it is easy to conclude that we are simply observing two different causes of the effect, each sufficient and neither individually necessary. This behavior has been observed in human causality judgments in the face of un signaled versus signaled effects (Reed, 1992; Shanks, 1989).

Hammond and Paynter (1983) evaluated six different contingency theories (including ΔP , Gibbon & Balsam's SET, and ϕ , a relative of χ^2) and their ability to account for some important empirical variables in the animal contingency literature. They concluded that all of the contingency theories reviewed failed to predict the effects of one or more of the variables on animal conditioning. In addition, contingency formulas do not confront the issue of deciding in which cell of a contingency table a particular event should be placed. It is necessary to specify a unit of time analysis for making this decision. For example, when a 5-sec CS is followed by a 1-sec US, the trial may be classified as (1) a single occurrence in the "a" cell, (2) one occurrence of "c" (CS without US) and one of "a," or (3) 5 of "c" and one of "b" (US without CS). In the first case, the unit of analysis is 6 sec; in the second, 3 sec; and in the third, 1 sec. This issue is especially relevant when one is attempting to analyze conditioning in the presence of trace intervals. Under what conditions is a CS considered to be "close enough" to the US to warrant placement in the "a" cell and when is it temporally too far away, thus resulting in its classification as a CS in the absence of a US and the US's classification as an un signaled US? The issue of time is not managed well by present formulations. Hallam, Grahame, and Miller (1992) propose that each US event should contribute *fractionally* to the appropriate cells as a function of the interval between CS and US events. For example, the shorter the interstimulus interval (ISI), the more the US counts as an "a" cell entry; the greater the ISI, the more it counts as a "b" cell entry. Unfortunately, this proposal lacks a formal specification (for a related proposal in the causality literature see Shanks, 1993b; Shanks et al., 1989).

Hammond and Paynter (1983) noted that despite its drawbacks, the Rescorla-Wagner theory is superior to contingency formulations in its theoretical scope and prediction of the effects of stimulus competition. They conclude:

it is our guess that the problems of the existing contingency theories reviewed here will not be solved by concatenating another probabilistic contingency formula. . . . we lean toward more complex approaches—complex theories to deal with complex problems. (p. 547)

Given the evidence, I find little reason to advocate for any of the current contingency accounts of either conditioning or causality. The shortcomings mentioned above have not been adequately explained by existing contingency theories. An analysis relying on contiguity appears better suited to the problems at hand, although contiguity has yet to be clearly demonstrated as sufficient.

Multiple causes and cue interaction. Having progressed from Hume's concept of constant conjunction to covariation and contingency, we have yet another step to take. Each of these concepts presupposes that humans extract conjunction data by analyzing pairs of events, a candidate cause and the effect. This assumption is unfounded. Interaction between cues is important to human (and animal) judgments of causality. The presence of alternate causes affects current attributions (see, e.g., Chapman & Robbins, 1990; Einhorn & Hogarth, 1986; Price & Yates, 1993; Shanks & Dickinson, 1987; Shustack & Sternberg, 1981). For example, Chapman and Robbins (1990) demonstrated that the judged predictive strength of a cue is influenced by the predictive strengths of contemporary cues. In their second experiment, a cue gained negative strength when paired with the absence of an effect (i.e., it was perceived as preventing rather than causing the effect) when consistently paired with a cue that had previously been judged as a cause of the same effect. This was contrasted with a cue that was paired the same number of times with the absence of the effect but failed to gain significant negative strength. The only difference between the two cues was whether they were paired during this phase of the experiment with a cue previously trained to predict the effect. These results and others have led many authors (e.g., Allan, 1993; Chapman & Robbins, 1990; Wasserman et al., 1993) to suggest that approaches based on animal learning theories (e.g., Rescorla & Wagner, 1972) are better suited to capture the complexities of human judgments of contingency.

Multiple CSs and CS interaction. A significant reason for the continued popularity of Rescorla and Wagner's (1972) model is its elegant method of addressing experimental paradigms where multiple CSs are present during a trial. In Equation 1, the amount of associative strength remaining during a trial is a function of the strengths of other CSs present; for example, if the US is being fully predicted by a subset of the contemporary CSs, no change in associative values will occur. Blocking is a commonly studied paradigm in which CS interaction is observed (e.g., Kamin, 1969). In *forward blocking*, a cue, B , is paired with another predictor, A , that has received pretraining as a predictor of a third event, the US ($A+$, $AB+$ training; the "+" indicates reinforcement, the presentation of the US following the CSs). B elicits significantly less conditioned responding than does a control ($AB+$ without prior $A+$ training), a result accounted for by competitive cue models like Rescorla and Wagner's. As an example, assume that at the end of Phase 1, A has an associative value, V_A , of 80, λ is 100,

α is .10, and β is 1. At asymptote, V_A will be 90 and V_B will be 10, in contrast to asymptotic values of 50 and 50 with no prior training of A . It is generally the case that the order of phases is important for the emergence of blocking. When the $A+$ trials are interspersed with the $AB+$ trials, A can still block B (Wagner, 1969), but not over a relatively short number of trials. The results of an analysis using the Rescorla-Wagner model suggest that the asymptotic strength of B will be zero, but that the strengths will pass through a stage of higher values. This results from the increasing competition from A as learning progresses.

When the order of phases is completely reversed ($AB+$, $A+$), there is very little evidence of a *backward blocking* effect during animal conditioning. In one study by Kamin (1969), rats demonstrated backward blocking. He found that 8 LN+ trials (light-noise) resulted in less fear to L if training was followed by 16 N+ trials before testing of L. In general, though, attempts to obtain backward blocking in animals have been largely unsuccessful, and the observed effects have been small (see Durlach, 1989).

Forward blocking of causes has been observed in a number of human causal and contingency judgment studies (Dickinson et al., 1984; Price & Yates, 1993; Shanks, 1993a; Van Hamme et al., 1993; Waldmann & Holyoak, 1992). Some studies (e.g., Price & Yates, 1993) included interspersing of the relevant trials with results similar to those of Wagner (1969). However, evidence for backward blocking is not as clear. Backward blocking has been observed in some human judgment experiments (see, e.g., Chapman, 1991; Shanks, 1985, 1991), but not in others (see, e.g., Van Hamme et al., 1993; Waldmann & Holyoak, 1992). Chapman (1991) and Shanks (1985, 1991) have observed backward blocking, although Chapman found a much smaller effect in backward than in forward blocking.

Williams, Sagness, and McPhee (1994) used pre-training manipulations to modify the strategy of subjects engaging in a predictive task. The groups that were trained to use an elemental strategy (in which the effects of candidate causes are additive) exhibited forward as well as strong backward blocking. The groups that used a configural strategy (treating multiple cues as unique predictive compounds) failed to demonstrate backward or forward blocking. Only when Williams et al. used an instructional manipulation similar to that used by Chapman and Robbins (1990) did they observe forward blocking and no backward blocking.² They suggest that configural processes may have a more detrimental effect on backward than on forward blocking. Allan (1993) put forth a similar argument, although she anticipated that there would be difficulties in accounting for other data when this perspective is taken. The difference between the two forms of blocking may depend on the interaction of elemental and configural learning processes, but a clear specification has yet to emerge.

Shanks and Dickinson (1987) and Van Hamme and Wasserman (1994) presented reformulations of the

Rescorla-Wagner model in an attempt to account for backward blocking. The Shanks and Dickinson version demonstrated backward blocking but had other shortcomings (Shanks, 1993b). Van Hamme and Wasserman (1994) extended the Rescorla-Wagner equation by including strength modification for cues not present on a trial. This extension is sufficient to account for backward blocking since it permits the modification of B 's strength during the second phase of the experiment when it is not occurring. Unfortunately, Van Hamme and Wasserman's modified Rescorla-Wagner model either must learn about *all* the cues absent on a trial or must decide which of the absent cues are relevant to the context. It is unlikely that an organism would be so poorly designed as to spend the amount of effort necessary to learn the relation between every event, present or not, and the outcome of interest. Van Hamme and Wasserman (1994) thus chose the second option, permitting application of the learning equations to those cues *relevant* to the prediction of the outcome. A cue is judged relevant after acquiring some positive or negative associative strength, either via prior experience of the cue in the experimental context or via verbal instructions. Their revised model holds promise; but further testing of its predictive power is required, and it will retain many of the known limitations of the Rescorla-Wagner model (Miller, Barnet, & Grahame, in press).

Blocking represents only one way in which CSs interact in the formation of CS-US relationships during classical conditioning. Others include conditioned inhibition, compound conditioning (e.g., negative patterning), occasion setting, and serial conditioning. Occasion setting and serial conditioning represent two particular areas of interest because of some interesting parallels to causal attribution. In occasion setting (P. C. Holland, 1986, 1989a, 1989b), subjects are exposed to a Pavlovian feature-positive discrimination ($XA+$, $A-$) paradigm. The feature, X , typically precedes the common element, A , by a short period of time. As training progresses, subjects respond more to the common element when it follows the feature than when it does not (assuming reasonable temporal and salience characteristics of the cues involved). In Holland's terms, the feature, X , "sets the occasion" for the common element to elicit a CR. The role of an occasion setter is reminiscent of the role of a "condition" in causal attributions. A condition is an event or cue that is necessary for the occurrence of an effect and is present for some time preceding it (hence not the triggering event). The paradigmatic example is the presence of oxygen for a fire. The striking of a match or other event is usually judged the cause, but oxygen is one of the conditions for the occurrence of a fire. When oxygen is absent (e.g., in space), an observer would not predict the occurrence of fire following the striking of a match. Hence, the presence of oxygen sets the occasion for the operativity of the match strike. This parallel provides another avenue of exploration of the conditioning/causal attribution analogy with methodologies used in occasion-setting research. There is one difference,

however, between the way in which a condition is discussed in the causality literature and the role of occasion setters. Occasion setters are stimuli that exist for relatively short periods of time (at most, minutes), but conditions typically exist for much longer periods, as with oxygen as a condition for combustion. In the latter case, the vast conditioning literature on the role of context is more relevant.

In serial conditioning two predictive cues are also offset in time, but both are present on all trials ($A \rightarrow B+$, where \rightarrow indicates "is followed by"). This paradigm has been most vigorously pursued by Kehoe and his colleagues (Kehoe et al., 1993; Kehoe, Marshall-Goodell, & Gormezano, 1987; Kehoe & Napier, 1991; Kehoe, Schreurs, & Graham, 1987). The use of two sequential predictors is analogous to the presence of a simple causal chain, in which each event causes the next. In serial conditioning, the intervening event, B , facilitates the development of the A -US relationship, often to the detriment of its own B -US strength. In fact, the initial event of a longer chain usually obtains the most US associative strength (Kehoe & Napier, 1991). This suggests that intervening events may be serving a role analogous to the apparent temporal bridging role of causal mechanisms (Young, 1992; Young & DeBauche, 1993). Causal mechanisms provide an event structure to bridge a temporal gap between an initiating event (the candidate cause) and the ultimate effect. If playing a role analogous to the B event in serial conditioning, a mechanism should be serving a modulatory role and not be acquiring much of a predictive relationship between itself and the effect. These and other features of serially conditioned CSs represent empirical questions to be pursued in the exploration of the scope of the attribution/conditioning analogy.

Prior Experience and Knowledge

Inherent in all the discussions of neo-Humean cues is the importance of experience. Hume was an empiricist who deemed experience the cornerstone of causal attribution. However, as language users, we have methods other than direct experience for determining whether events might be causally related. I have mentioned "cultural transmission" in passing, but this avenue likely plays a major role in some adult attributions. Through formal education, we learn that static electricity can cause hair to stand on end, that viruses cause colds, and that demand increases as supply dwindles. These *public* causal theories are relied on in many of our judgments. It is this notion that Bullock, Gelman, and Baillargeon (1982) and Shultz (1982) are building on when they stress the importance of causal mechanisms. Mechanisms are embodiments of cultural expertise on causality passed on through the generations. They collect the experiences of others, thus freeing us from having to directly discover all causal relationships ourselves. In this sense, mechanisms play a role different from the one that I proposed earlier, and I have little doubt that they are rightly utilized by adult observers everywhere.

Unfortunately, humans frequently ignore these cultural nuggets and leave them locked within some deep, dark chamber of their minds. The extensive research on folk or naive physics is a depressing example of the difficulties of overriding our prior belief systems (Champagne, Klopfer, & Anderson, 1980; Clement, 1982; Kaiser, Proffitt, & McCloskey, 1985). J. H. Holland, Holyoak, Nisbett, and Thagard (1986) summarize the current research as substantial evidence that the new rules obtained via education do not merely compete with the old rules (our prior belief systems) but are forced to modify the old ones, resulting in slightly updated versions. Any field of expertise in which the learner has extensive prior knowledge or experience of the domain will suffer from these problems. Highly educated people still believe in the predictive power of astrology and tarot cards, despite their schooling. Coaches and athletes continue to wear "lucky" clothes or adopt routines, despite ignorance of a mechanism that could bring about a win.

This is not to suggest that we must experience everything directly. The problems in applying cultural knowledge arise when our prior experience or knowledge (which may have been culturally obtained at a younger age) is quite extensive in a domain. This is true with common notion physics, psychology, statistics, and biology, *inter alia*. Prior experience weighs heavily on our current judgments in the form of personal causal theories. The research showing cue interaction in causal judgment (e.g., Chapman & Robbins, 1990; Waldmann & Holyoak, 1992) illustrates the significant impact that prior experience and expectations have on current attributions.

Although not usually treated as a cue to causality, the amount of prior experience of similar events is critical in determining whether causality is inferred. Prior biases and expectations can override many of the other cues, in both animals and humans (for a review, see Alloy & Tabachnik, 1984). Ausubel and Schiff (1954) investigated the effects of prior knowledge in kindergarten, third grade, and sixth grade subjects who were required, after observing a sequence of trials, to determine what caused one side of a teeter-totter to drop. The youngest subjects showed no difference in inferring a relationship between object color and teeter-totter side to drop and between side length and side to drop. In contrast, older children had difficulty learning the color relationship.

The presence of alternative causes also affects the extent to which an event is judged as the cause of a particular effect. It is well accepted that the presence of plausible alternatives decreases judgments regarding a cause in question (Einhorn & Hogarth, 1986; Kelley, 1973; Koslowski & Okagaki, 1986; Mackie, 1974; Nisbett & Ross, 1980). This is driven from two types of processes: competition among candidate causes, and prior expectations. When there are multiple plausible candidates, the surety of one of them being *the* cause is decreased. The observer is left to choose one of the candidates, ostensibly relying upon other cues to causality, or to infer that multiple causation is germane to the problem at hand. When the subject believes, on the basis of prior experi-

ence, that one of the candidates is more likely than another, this will strongly affect his/her judgment (the Ausubel & Schiff, 1954, study is a good example). These biases frequently result in misconceptions regarding "true" causal relationships (see Nisbett & Ross, 1980, for an excellent discussion of the many effects that biases have on judgments of covariation and causality).

This leads me to suggest that the generative theorists and the neo-Humeans are not as far apart as originally surmised. (Nor are they as distant as is typically portrayed in the literature—e.g., Shultz, 1982; White, 1989; but see Cheng, 1993, for an excellent integration of the two views.) They both argue that prior experience and knowledge play a significant role in causal attribution. They differ in the suggested form that this knowledge takes and, potentially, its origins. A generative theorist like Shultz suggests that it is stored in the form of rules (Shultz & Kestenbaum, 1985), whereas many neo-Humeans make no claims about representation except to suggest that current experiences are influenced by prior events. A connectionist interpretation might best rectify the two positions. Experience is encoded within a connectionist network over time, leaving traces of those experiences. This learning not only affects future behavior but also alters the processing and storage of new stimuli within the network. The effects of this prior learning may appear rule-like to an outside observer (like Shultz or White), but this does not imply that the underlying representation is in the form of symbolic rules.

Einhorn and Hogarth (1986) attempted to integrate the views of those both for and against the priority of the cues to causality. Shultz has emphasized the importance of generative transmission over the traditional Humean cues. In contrast, Einhorn and Hogarth do not de-emphasize the cues to causality. They stress that these cues are not *rules* but *heuristics*. As heuristics, they are fallible and thus are only probabilistic indicators of the likelihood of a causal attribution. In addition, they note that the cues "play a crucial role in delimiting the particular models of generative transmission that people use in causal reasoning, . . . [so] it cannot be said that the cues *per se* are disregarded in favor of [generative transmission]" (Einhorn & Hogarth, 1986, p. 13). They are emphasizing the role of neo-Humean cues to causality in the origins of causal theories, an argument that I advocate and one that has not generally been attributed to Einhorn and Hogarth's theoretical framework.

In an attempt to bring the two views closer together, I find it necessary to point out that certain points cannot be emended. The occasional generative theorist who still insists that some causal relations can be directly observed (e.g., Harré & Madden, 1975; White, 1989), without reducing this observation to something like the cues to causality, will always be at odds with a Humean position.

The matter of preparedness. Experience is central to the empirical and theoretical approach historically pursued in the field of animal learning. Hence it comes as no surprise that prior experience (i.e., knowledge) has an effect on ongoing attributions in conditioning experi-

ments, as is illustrated in the blocking design. What is more interesting for the discussion at hand is the richness of the effects that this experience has on animals. Some of these effects are cited in support of the inadequacy of associationist views of causality, but I suggest that these arguments are more often evidence of the narrow fashion in which some cognitive psychologists still view animal learning theories. The most commonly cited argument against associative views is the evidence of stimulus relevance or belongingness.

Relevance and belongingness. Early work in animal learning (e.g., by Pavlov, Skinner) supported the idea that any stimulus could serve as a CS for any US. This was, in a sense, necessary if one was to retain the proposed goal that the laws of learning were universal: any CS-US or S-R relationship could be acquired. However, some observations did not easily accord with this view. Thorndike (1932) was unable to train cats to yawn or to scratch to receive a reinforcer. Breland and Breland (1961) had difficulties training a raccoon to pick up a coin and deposit it in a coin bank, and pigs also failed to learn the same task. Garcia and Koelling (1966) provided the most vexing data of this sort for researchers in Pavlovian conditioning. Four groups of rats were given different CS-US relationships to learn. For one group, the taste of a water solution (sweet) predicted the onset of gastric upset (induced by lithium chloride or radiation). For another, the taste predicted the onset of shock. The last two received an audiovisual stimulus (flashing light plus click) predicting either shock, for one group, or gastric upset, in the other. The rats were readily able to learn that taste predicted gastric upset and that an audiovisual stimulus predicted shock, but they had severe difficulties in learning the other relationships. There have been numerous demonstrations of such relevance, although results as dramatic as those observed in taste aversion are not common. Garcia and Koelling suggested that some CSs were more "relevant" in the prediction of certain USs and that this relevance is inherent in a species, the result of natural selection. In fact it is relatively easy to develop post hoc a story of why such a CS-US belongingness is adaptive for a particular species (fraught with the usual difficulties with this type of inference).

The other notable characteristic of such relevant relationships is the rapidity with which they are learned. Taste aversion is often acquired after a single exposure. This represents a nice parallel for causal attribution researchers of the inadequacy of associative approaches. These investigators (Alloy & Tabachnik, 1984; Nisbett & Ross, 1980; Shultz & Kestenbaum, 1985; White, 1989) note that attributions are often made by humans after a single observation of a cause-effect sequence. This is analogous, in their opinion, to the processes observed in animals in taste aversion studies and demonstrates that covariation is unnecessary as a cue to causality.

There are two problems with this view. First, the variety of occasions under which humans exhibit single trial learning is much greater than that observed in animals.

Although it may be adaptive for taste-gastric upset relationships to be rapidly learned (as it is even in day-old rats), it is much more difficult to contend that evolution has selected for the relevance of the pushing of a button for the turning on of an electric appliance. I find it more parsimonious to suggest that prior experience (in the form of causal theories), not innate knowledge, is being leveraged under these conditions. When observing a "novel" event pairing that is similar to experienced event pairs in which the prior event predicts the latter, the process of generalization will operate to "infer" a causal relationship between the novel events—after a single trial. This position has also been advocated in the conditioning literature to explain some, but not all, of the observations of CS-US belongingness (Mackintosh, 1973). It is difficult to prove that an organism has had no prior experience relevant to a particular pair of event relationships; this experience may merely be similar and not identical to the current events, or it may be obtained vicariously.

This is not to suggest that humans, or animals for that matter, *always* require some background in related events when they exhibit single-trial attributions. There are certainly cases in which biological preparedness can be effectively argued (e.g., taste aversion) and, in humans at least, circumstances under which higher cognitive mechanisms are being brought to bear. In my opinion, however, this leaves a great number of occasions wherein inferences are made quickly, spontaneously, and without prolonged thought. One must argue that prior experience is quickly being leveraged, that the organism has innate knowledge of the relationship to be learned, or that the other cues to causality provide sufficient evidence to support an inference.

Latent inhibition and learned irrelevance. Latent inhibition and learned irrelevance represent two important ways in which prior experience can affect current attributions without the effect being readily apparent to the casual observer. In latent inhibition (Lubow, 1973)—that is, the CS preexposure effect—a subject is slower than a control in learning a CS-US relationship when the subject has received numerous prior presentations of the CS alone. Latent inhibition provides a mechanism for addressing the set of objections to associative accounts of causality exemplified by the following: "The color of the walls in my bedroom are always paired with my going to sleep, but other rooms of that color do not cause me to anticipate sleeping." Through numerous prior experiences in one's lifetime, wall color has been shown to be an ineffective predictor of other events. An analogous situation was observed in the older subjects in the Ausubel and Schiff (1954) experiment, who had difficulty learning that the color of the blocks predicted which side of a teeter-totter would fall. Their prior experience suggested that object color was a poor predictor of effects; this is latent inhibition at work.

Learned irrelevance represents another way in which prior experience may inhibit certain attributions from being made (Baker & Mackintosh, 1977). A subject who

has prior experience of a CS that is a poor predictor of a US (i.e., there is a near-zero contingency between the two) is retarded in the later acquisition of the CS as a predictor of this particular US when it *does* consistently follow the CS. This differs from latent inhibition in two ways: (1) in latent inhibition, a subject is not required to have had any prior exposure to the US; and (2) learned irrelevance is particular to a CS-US pair, whereas latent inhibition affects all subsequent CS-US pairs involving the CS (however, note that latent inhibition has been shown to be sensitive to the context in which the CS pre-exposure has taken place; Hall, 1991). When someone suggests that the associative position is wrong because "his alarm clock doesn't cause the neighborhood children to go to school," he has failed to note the relevance of the vast amount of prior experience he has had wherein his alarm clock has been a poor predictor of *other* people's behavior (not to mention the problem of lack of spatial contiguity). Note again that this is not to suggest that this is all there is to the human story. Other factors can influence an attribution; in this case, it would be important to note that the children cannot *hear* his alarm clock. However, how do we know that this is a relevant fact? Perhaps our prior experience of sounds as predictors of others' and our own behavior are at the root of our assumption that hearing a sound is a necessary condition for it to function as a causal factor. I have found the scope of the associative approach to causality to be quite surprising.

Summary

The preceding survey of the human causal attribution literature has demonstrated the importance of Hume's principles. The neo-Humean cues affect human judgment, although the degree to which each does so is situation dependent. Through a rational Humean analysis of a variety of data, I hope to have shed some light on the strength of a learning perspective on the origin of per-

sonal causal theories. An examination of causal attribution from a learning stance (Shanks & Dickinson, 1987; Wasserman, 1990b) has helped refine the definition of the Humean cues and facilitated their proper application to the human literature.

The preceding survey of the animal learning literature demonstrates the strength of the conditioning-causal attribution analogy. The work of Shanks and Dickinson (1987) and Wasserman (1990b) provides a solid framework for the understanding of causality. Accepting and extending the analogy has provided a new perspective on understanding human cognition. A number of methodologies and theories that originate in the conditioning literature offer opportunities for enriching our understanding of the human judgment process. While promising, this view is not without its critics. In addition to the earlier mentioned arguments posed by generative theorists, proponents of mediation of attributions by higher order causal induction have questioned the sufficiency of an associative formulation (Cheng & Novick, 1992; Melz, Cheng, Holyoak, & Waldmann, 1993; Waldmann & Holyoak, 1992). Many of the concerns of these authors have been addressed by Allan (1993) and Shanks (1993a, 1993b).

Conclusions

In the prior discussion, I have presented evidence for the analogy between many human causal attributions and animal conditioning. The parallels are summarized in Table 3. This obviously does not function as proof of the relationship. In my opinion, however, the evidence has been compelling enough to warrant the continued investigation of this theoretical approach. The conditioning literature not only offers a framework within which to couch human attribution research, but also presents a wealth of data, empirical methods, and theories from which to draw.

Current investigations have only touched on the wealth of data, methodologies, and theories that the an-

Table 3
Summary of Effects of the Cues to Causality on Human Causal Attribution and Animal Conditioning

Cue to Causality	Human Causal Attribution	Animal Conditioning
Temporal contiguity	Very important: gaps can be bridged via real or hypothesized events.	Very important: gaps can be bridged via real events and if no intervening events interfere with recall.
Spatial contiguity	Very important for adults and older children as studied in launching paradigms.	Secondary role, but only studied as spatial proximity, not contiguity.
Temporal priority	Necessary for events to be at least contemporary.	Necessary for behavior change to be observed.
Covariation/contingency	Mixed results. Appears to depend on longer term experience as basis for contingency judgment.	Mixed results. May be encompassed by temporal contiguity.
Cue interaction	Very important as seen in human analogs of blocking and inhibition. Some evidence for backward blocking.	Very important, e.g., blocking, inhibition, configural learning. Weak evidence for backward blocking.
Prior experience	Major effect in the form of prior knowledge. Belongingness may be based on experience or possibly nature of event.	Important in entire notion of conditioning as a learning process. Belongingness <i>may</i> be based on innate mechanisms or prior experience.

imal learning perspective offers the study of human causality. As I have illustrated, there are a number of new avenues to consider. The present analysis suggests a series of research recommendations:

1. Temporal contiguity: Intervening stimuli may be perceived as causal mechanisms and mechanisms may only be functioning to bridge a contiguity gap. We need studies in which the temporal placement of the intervening stimulus is manipulated to assess the role of these stimuli (e.g., analogies to Rescorla, 1982).

2. There has been a trend toward discrete trial presentation in the study of causality. Given the importance of temporal cues, studies of real-time events should be emphasized in which temporal relationships are manipulated.

3. Covariation and contingency must be clearly differentiated in causal attribution research and the effect of presentation form considered in all studies.

4. Spatial contiguity may have a secondary effect when studied in domains other than the launching paradigm and may be superseded by temporal contiguity.

5. Blocking represents but one type of cue interaction. Studies should adapt methodologies from research on occasion setting, serial conditioning, latent inhibition, CS-US irrelevance, and similar areas.

6. Prior knowledge has a profound impact on attributions. Domains of inquiry must be carefully selected and analyzed in terms of the extent of prior experience of the events involved. This experience may be either positive (wherein the two events have been previously associated), or negative (wherein the events have previously been shown to be unrelated, or the proposed cause has held no predictive power over any environmental events).

7. An associative learning theory can explain the formation and leverage of personal causal theories, although current approaches have their limitations. Connectionist models embody the associative mechanisms required, but should be extended to include temporal mechanisms more flexible than those found in current theories.

Although researchers in animal learning have been studying the issue of event relations for a much longer period of time than have cognitive psychologists, they do not have the final answers. The variety of theories of Pavlovian conditioning is evidence of the failure to converge on a single explanation that can account for most of the existing data. It has become apparent that because of the attempt to remain parsimonious, existing theories have failed to cover the gamut of conditioning data. This is not surprising nor is it a great failing; existing theories have generated an abundance of excellent studies and have provided a suitable framework for dealing with the immeasurable amount of data. If one were constrained only to develop psychological theories that accounted for all known data, no theories would be advanced. There is growing recognition in the field of associative models of causal attribution, however, that connectionism may offer a framework with enough richness to capture more of the currently known data than do present theories. The development of recurrent network models of condi-

tioning and human causal attribution is a step in this direction (Young, 1992; Young & Bailey, 1994; Young & DeBauche, 1993).

REFERENCES

- ALLAN, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, **15**, 147-149.
- ALLAN, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, **114**, 435-448.
- ALLAN, L. G., & JENKINS, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning & Motivation*, **14**, 381-405.
- 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.
- ARKES, H. R., & HARKNESS, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General*, **112**, 117-135.
- AUSUBEL, D. P., & SCHIFF, H. M. (1954). The effect of incidental and experimentally induced experience in the learning of relevant and irrelevant causal relationships by children. *Journal of Genetic Psychology*, **84**, 109-123.
- BAKER, A. G., BERBRIER, M. W., & VALLEE-TOURANGEAU, F. (1989). Judgments of a 2x2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology*, **41B**, 65-97.
- BAKER, A. G., & MACKINTOSH, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning & Behavior*, **5**, 315-319.
- BAKER, A. G., MERCIER, P., VALLEE-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.
- BENEDICT, J. O. (1991). Judgment of covariation in classical and instrumental conditioning contexts. *Bulletin of the Psychonomic Society*, **29**, 457-459.
- BENEDICT, J. O., & AYRES, J. B. (1972). Factors affecting conditioning in the truly random control procedure in the rat. *Journal of Comparative & Physiological Psychology*, **78**, 323-330.
- BINDRA, D., CLARKE, K. A., & SHULTZ, T. R. (1980). Understanding predictive relations of necessity and sufficiency in formally equivalent "causal" and "logical" problems. *Journal of Experimental Psychology: General*, **109**, 422-443.
- BOLLES, R. C., COLLIER, A. C., BOUTON, M. E., & MARLIN, N. A. (1978). Some tricks for ameliorating the trace-conditioning deficit. *Bulletin of the Psychonomic Society*, **11**, 403-406.
- BRELAND, K., & BRELAND, M. (1961). The misbehavior of organisms. *American Psychologist*, **16**, 681-684.
- BRODY, N. (1989). Unconscious learning of rules: Comment on Reber's analysis of implicit learning. *Journal of Experimental Psychology: General*, **118**, 236-238.
- BULLOCK, M. (1984). Preschool children's understanding of causal connections. *British Journal of Developmental Psychology*, **2**, 139-148.
- BULLOCK, M., & GELMAN, R. (1979). Preschool children's assumptions about cause and effect: Temporal ordering. *Child Development*, **50**, 89-96.
- BULLOCK, M., GELMAN, R., & BAILLARGEON, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209-254). New York: Academic Press.
- CHAMPAGNE, A. B., KLOPPER, L. E., & ANDERSON, L. H. (1980). Factors influencing the learning of classical mechanics. *American Journal of Physics*, **48**, 1074-1079.
- CHAPMAN, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 837-854.
- CHAPMAN, G. B., & ROBBINS, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, **18**, 537-545.
- CHENG, P. W. (1993). Separating causal laws from causal facts: Press-

- ing 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., & NOVICK, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, **99**, 365-381.
- CLEMENT, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics*, **50**, 66-71.
- DESMOND, J. E., & MOORE, J. W. (1988). Adaptive timing in neural networks: The conditioned response. *Biological Cybernetics*, **58**, 405-415.
- DICKINSON, A., SHANKS, D. R., & EVENDEN, J. L. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, **36A**, 29-50.
- DURLACH, P. J. (1983). Effect of signaling intertrial unconditioned stimuli in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **9**, 374-389.
- DURLACH, P. J. (1989). Learning and performance in Pavlovian conditioning: Are failures of contiguity failures of learning or performance? In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 19-59). Hillsdale, NJ: Erlbaum.
- EGGER, M. D., & MILLER, N. E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology*, **64**, 97-104.
- EINHORN, H. J., & HOGARTH, R. M. (1986). Judging probable cause. *Psychological Bulletin*, **99**, 3-19.
- EYSENCK, M. W., & KEANE, M. T. (1990). *Cognitive psychology: A student's handbook*. Hove, U.K.: Erlbaum.
- GARCIA, J., & KOELLING, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, **4**, 123-124.
- GIBBON, J., & BALSAM, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219-253). New York: Academic Press.
- GRICE, G. R. (1948). The relation of secondary reinforcement to delayed reward in visual discrimination learning. *Journal of Experimental Psychology*, **38**, 1-16.
- GROSSBERG, S., & SCHMAJUK, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, **2**, 79-102.
- GRUBER, H. E., FINK, C. D., & DAMM, V. (1957). Effects of experience on perception of causality. *Journal of Experimental Psychology*, **53**, 89-93.
- HALL, G. (1991). *Perceptual and associative learning*. Oxford: Oxford University Press, Clarendon Press.
- HALLAM, S. C., GRAHAME, N. J., & MILLER, R. R. (1992). Exploring the edges of Pavlovian contingency space: An assessment of contingency theory and its various metrics. *Learning & Motivation*, **23**, 225-249.
- HAMMOND, L. J., & PAYNTER, W. E. (1983). Probabilistic contingency theories of animal conditioning: A critical analysis. *Learning & Motivation*, **14**, 527-550.
- HAMMOND, L. J., & WEINBERG, M. (1984). Signaling unearned reinforcers removes the suppression produced by a zero correlation in an operant paradigm. *Animal Learning & Behavior*, **12**, 371-377.
- HARRÉ, R., & MADDEN, E. H. (1975). *Causal powers: A theory of natural necessity*. Oxford: Blackwell.
- HOLLAND, J. H., HOLYOAK, K. J., NISBETT, R. E., & THAGARD, P. (1986). *Induction: Processes of inference, learning and discovery*. Cambridge, MA: MIT Press.
- HOLLAND, P. C. (1986). Temporal determinants of occasion setting in feature-positive discriminations. *Animal Learning & Behavior*, **14**, 111-120.
- HOLLAND, P. C. (1989a). Feature extinction enhances transfer of occasion setting. *Animal Learning & Behavior*, **17**, 269-279.
- HOLLAND, P. C. (1989b). Occasion setting with simultaneous compounds in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, **15**, 183-193.
- HUME, D. (1969). *A treatise of human nature*. New York: Penguin Books. (Original work published 1739)
- KAISER, M. K., PROFFITT, D. R., & MCCLOSKEY, M. (1985). The development of beliefs about falling objects. *Perception & Psychophysics*, **38**, 533-539.
- KAMIN, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42-64). Halifax, NS: Dalhousie University Press.
- 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.
- KAPLAN, P. S., & HEARST, E. (1982). Bridging temporal gaps between CS and US in autoshaping: Insertion of other stimuli before, during, and after CS. *Journal of Experimental Psychology: Animal Behavior Processes*, **8**, 187-203.
- KEHOE, E. J., GIBBS, C. M., GARCIA, E., & GORMEZANO, I. (1979). Associative transfer and stimulus selection in classical conditioning of the rabbit's nictitating membrane response to serial compound CSs. *Journal of Experimental Psychology: Animal Behavior Processes*, **5**, 1-18.
- KEHOE, E. J., HORNE, P. S., MACRAE, M., & HORNE, A. J. (1993). Real-time processing of serial stimuli in trace conditioning of the rabbit's nictitating membrane response. *Journal of Experimental Psychology: Animal Behavior Processes*, **19**, 265-283.
- KEHOE, E. J., MARSHALL-GOODSELL, B., & GORMEZANO, I. (1987). Differential conditioning of the rabbit's nictitating membrane response to serial compound stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **13**, 17-30.
- KEHOE, E. J., & NAPIER, R. M. (1991). Real-time factors in the rabbit's nictitating membrane response to pulsed and serial conditioned stimuli. *Animal Learning & Behavior*, **19**, 195-206.
- KEHOE, E. J., SCHREURS, B. G., & GRAHAM, P. (1987). Temporal primacy overrides prior training in serial compound conditioning of the rabbit's nictitating membrane response. *Animal Learning & Behavior*, **15**, 455-464.
- KELLEY, H. H. (1973). The processes of causal attribution. *American Psychologist*, **28**, 107-128.
- KILLEEN, P. R. (1981). Learning as causal inference. In M. L. Commons & J. L. Nevin (Eds.), *Quantitative analyses of behavior: Discriminative properties of reinforcement schedules* (pp. 89-112). Cambridge, MA: Ballinger.
- KILLEEN, P. R., & SMITH, J. P. (1984). Perception of contingency in conditioning: Scalar timing, response bias, and erasure of memory by reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 333-345.
- KOSLOWSKI, B., & OKAGAKI, L. (1986). Non-Humean indices of causation in problem-solving situations: Causal mechanism, analogous effects, and the status of rival alternative accounts. *Child Development*, **57**, 1100-1108.
- KOSLOWSKI, B., OKAGAKI, L., LORENZ, C., & UMBACH, D. (1989). When covariation is not enough: The role of causal mechanism, sampling method, and sample size in causal reasoning. *Child Development*, **60**, 1316-1327.
- KUHN, D., & PHELPS, H. (1976). The development of children's comprehension of causal direction. *Child Development*, **47**, 248-251.
- LESLIE, A. M. (1986). Getting development off the ground: Modularity and the infant's perception of causality. In P. van Geert (Ed.), *Theory building in development* (pp. 405-437). Amsterdam: North-Holland.
- LESLIE, A. M., & KEEBLE, S. (1987). Do six-month old infants perceive causality? *Cognition*, **25**, 265-288.
- LEWICKI, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **12**, 135-146.
- LIEBETRAU, A. M. (1983). *Measures of association*. Beverly Hills, CA: Sage.
- LUBOW, R. E. (1973). Latent inhibition. *Journal of Comparative & Physiological Psychology*, **79**, 398-407.
- MACKIE, J. L. (1974). *The cement of the universe: A study of causation*. Oxford: Oxford University Press, Clarendon Press.
- MACKINTOSH, N. J. (1973). Stimulus selection: Learning to ignore stimuli that predict no change in reinforcement. In R. A. Hinde & S. J. Hinde (Eds.), *Constraints on learning* (pp. 75-96). London: Academic Press.
- MACKINTOSH, N. J. (1974). *The psychology of animal learning*. New York: Academic Press.

- MACKINTOSH, N. J. (1977). Conditioning as the perception of causal relations. In R. E. Butts & J. Hintikka (Eds.), *Foundational problems in the special sciences* (pp. 241-250). Dordrecht: D. Reidel.
- MATZEL, L. D., HELD, F. P., & MILLER, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning & Motivation*, **19**, 317-344.
- MELZ, E. R., CHENG, P. W., HOLYOAK, K. J., & WALDMANN, M. R. (1993). Cue competition in human categorization: Contingency or the Rescorla-Wagner learning rule? Comment on Shanks (1991). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 1398-1410.
- MENDELSON, R., & SHULTZ, T. R. (1976). Covariation and temporal contiguity as principles of causal inference in young children. *Journal of Experimental Child Psychology*, **22**, 408-412.
- MICHOTTE, A. (1963). *The perception of causality*. London: Methuen.
- MILLER, R. R., BARNET, R. C., & GRAHAME, N. J. (in press). Assessment of the Rescorla-Wagner model. *Psychological Bulletin*.
- MILLER, R. R., & SCHACTMAN, T. R. (1985). Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In P. D. Balsam & A. Tomie (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 51-88). Hillsdale, NJ: Erlbaum.
- NISBETT, R. E., & ROSS, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- OAKES, L. M., & COHEN, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, **5**, 193-207.
- PAPINI, M. R., & BITTERMAN, M. E. (1990). The role of contingency in classical conditioning. *Psychological Review*, **97**, 396-403.
- PAVLOV, I. P. (1927). *Conditioned reflexes* (G. V. Anrep, trans.). London: Oxford University Press.
- PRICE, P. C., & YATES, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition*, **21**, 561-572.
- REBER, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, **118**, 219-235.
- REED, P. (1992). Effect of a signalled delay between an action and outcome on human judgement of causality. *Quarterly Journal of Experimental Psychology*, **44B**, 81-100.
- REED, P., & REILLY, S. (1990). Context extinction following conditioning with delayed reward enhances subsequent instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes*, **16**, 48-55.
- RESCORLA, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review*, **74**, 71-80.
- RESCORLA, R. A. (1982). Effect of a stimulus intervening between CS and US in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **8**, 131-141.
- RESCORLA, R. A. (1987). A Pavlovian analysis of goal-directed behavior. *American Psychologist*, **42**, 119-129.
- RESCORLA, R. A., & CUNNINGHAM, C. L. (1979). Spatial contiguity facilitates Pavlovian second-order conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **5**, 152-161.
- RESCORLA, R. A., & SOLOMON, R. L. (1967). Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review*, **74**, 151-182.
- 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.
- ROSENFARB, I. S., NEWLAND, M. C., BRANNON, S. E., & HOWEY, D. S. (1992). Effects of self-generated rules on the development of schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior*, **58**, 107-121.
- SCHLOTTMANN, A., & SHANKS, D. R. (1992). Evidence for a distinction between judged and perceived causality. *Quarterly Journal of Experimental Psychology*, **44A**, 321-342.
- SHAKLEE, H. (1983). Human covariation judgment: Accuracy and strategy. *Learning & Motivation*, **14**, 433-448.
- SHAKLEE, H., & ELEK, S. (1988). Cause and covariate: Development of two related concepts. *Cognitive Development*, **3**, 1-13.
- SHAKLEE, H., & MIMS, M. (1981). Development of rule use in judgments of covariation between events. *Child Development*, **52**, 317-325.
- SHAKLEE, H., & TUCKER, D. (1980). A rule analysis of judgments of covariation between events. *Memory & Cognition*, **8**, 459-467.
- SHANKS, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, **37B**, 1-21.
- SHANKS, D. R. (1989). Selectional processes in causality judgment. *Memory & Cognition*, **17**, 27-34.
- SHANKS, D. R. (1991). Categorization by a connectionist network. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 433-443.
- SHANKS, D. R. (1993a). Associative versus contingency accounts of category learning: Reply to Melz, Cheng, Holyoak, and Waldmann (1993). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 1411-1423.
- SHANKS, D. R. (1993b). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, **84**, 319-354.
- 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: Academic Press.
- SHANKS, D. R., PEARSON, S. M., & DICKINSON, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *Quarterly Journal of Experimental Psychology*, **41B**, 139-159.
- SHULTZ, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, **47**(1, Serial No. 194).
- SHULTZ, T. R., ALTMANN, E., & ASSELIN, J. (1986). Judging causal priority. *British Journal of Developmental Psychology*, **4**, 67-74.
- SHULTZ, T. R., & KESTENBAUM, N. R. (1985). Causal reasoning in children. *Annals of Child Development*, **2**, 195-249.
- SHULTZ, T. R., & MENDELSON, R. (1975). The use of covariation as a principle of causal analysis. *Child Development*, **46**, 394-399.
- SHUSTACK, M. W., & STERNBERG, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, **110**, 101-120.
- SIEGLER, R. S., & LIEBERT, R. M. (1974). Effects of contiguity, regularity and age on children's causal inference. *Developmental Psychology*, **10**, 574-579.
- SMITH, M. C., COLEMAN, S. R., & GORMEZANO, I. (1969). Classical conditioning of the rabbit's nictitating membrane response at backward, simultaneous, and forward CS-US intervals. *Journal of Comparative & Physiological Psychology*, **69**, 226-231.
- SOPHIAN, C., & HUBER, A. (1984). Early developments in children's causal judgments. *Child Development*, **55**, 512-526.
- SPECH, M. L., WILKIE, D. M., & PINEL, J. P. J. (1981). Backward conditioning: A reevaluation of the empirical evidence. *Psychological Bulletin*, **89**, 163-175.
- SUPPES, P. (1970). *A probabilistic theory of causality*. Amsterdam: North-Holland.
- SUPPES, P. (1984). *Probabilistic metaphysics*. New York: Blackwell.
- SUTTON, R. S., & BARTO, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, **88**, 135-170.
- TESTA, T. J. (1975). Effects of similarity of location and temporal intensity pattern of conditioned and unconditioned stimuli on the acquisition of conditioned suppression in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, **104**, 114-121.
- THORNDIKE, E. L. (1932). *The fundamentals of learning*. New York: Macmillan.
- TRAPOLD, M. A., & OVERMIER, J. B. (1972). The second learning process in instrumental learning. In A. A. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 427-452). New York: Appleton-Century-Crofts.
- VAN HAMME, L. J., KAO, S.-F., & WASSERMAN, E. A. (1993). Judging interevent relations: From cause to effect and from effect to cause. *Memory & Cognition*, **21**, 802-808.
- VAN HAMME, L. J., & WASSERMAN, E. A. (1993). Cue competition in causality judgments: The role of manner of information presentation. *Bulletin of the Psychonomic Society*, **31**, 457-460.
- 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. (1969). Stimulus selection and a "modified continuity theory." In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 3, pp. 1-41). New York: Academic Press.
- WAGNER, A. R., & TERRY, W. S. (1975). Backward conditioning to a CS following an expected vs. a surprising UCS. *Animal Learning & Behavior*, **3**, 370-374.
- WALDMANN, M. R., & HOLYOAK, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, **121**, 222-236.
- WARD, W. C., & JENKINS, H. M. (1965). The display of information and the judgment of contingency. *Canadian Journal of Psychology*, **19**, 231-241.
- WASSERMAN, E. A. (1990a). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science*, **1**, 298-302.
- WASSERMAN, E. A. (1990b). Detecting response-outcome relations: Toward an understanding of the causal structure of the environment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 27-82). San Diego: Academic Press.
- 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., & NEUNABER, D. J. (1986). College students' responding to and rating of contingency relations: The role of temporal contiguity. *Journal of the Experimental Analysis of Behavior*, **46**, 15-35.
- WASSERMAN, E. A., & SHAKLEE, H. (1984). Judging response-outcome relations: The role of response-outcome contingency, outcome probability, and method of information presentation. *Memory & Cognition*, **12**, 270-286.
- WHITE, P. A. (1989). A theory of causal processing. *British Journal of Psychology*, **80**, 431-454.
- WHITE, P. A. (1992). Causal powers, causal questions, and the place of regularity in causal attribution. *British Journal of Psychology*, **83**, 161-188.
- WILLIAMS, D. A., SAGNESS, K. E., & MCPHEE, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 694-709.
- YATES, J. F., & CURLEY, S. P. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, **62**, 293-302.
- YOUNG, M. E. (1992). A simple recurrent network model of serial conditioning: Implications for temporal event representation. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 1164-1169). Hillsdale, NJ: Erlbaum.
- YOUNG, M. E., & BAILEY, T. M. (1994). Event prediction: Faster learning in a layered Hebbian network with memory. In M. C. Mozer, P. Smolensky, D. S. Touretzky, J. L. Elman, & A. S. Weigend (Eds.), *Proceedings of the 1993 Connectionist Models Summer School* (pp. 245-252). Hillsdale, NJ: Erlbaum.
- YOUNG, M. E., & DEBAUCHE, B. (1993). Causal mechanisms as temporal bridges in a connectionist model of causal attribution. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 1092-1097). Hillsdale, NJ: Erlbaum.

NOTES

1. The differences were derived from a summary table and not individual subject data. The absolute differences were computed by using the algebraic means reported in Kao and Wasserman's (1993) Table 5 comparing G13-SUM with G13-TBT and G9-SUM with G9-TBT. The following pairs of problems were used: 2-4, 3-5, 6-8, 7-9, 10-12, 11-13, 15-16, 19-20.

2. Subjects were being asked to predict the rise and fall of the stock market after observing changes in the performance of individual stocks. Chapman and Robbins (1990) had informed their subjects that the stocks would fall into three categories (positive, negative, and neutral) depending on their influence on the market as a whole.

(Manuscript received July 5, 1994;
revision accepted for publication November 15, 1994.)