In Learning and Behavior Therapy (1998). William O'Donohue (Ed.). Boston, MA: Allyn and Bacon.

CHAPTER 24

DETECTING CAUSAL RELATIONS

Helena Matute
Ralph R. Miller

Sometimes humans perceive causal relations that do not really exist (as in superstitious behavior, illusions of control, and various clinical problems); sometimes we do *not* perceive causal relations that do exist (as in cases in which our belief that minority people cause crime *blocks* our identifying poverty amidst wealth as a major cause of crime); yet at other times we show extreme accuracy in detecting causal relationships. The question of why (and how) we attribute causal roles to some events but not to others has interested psychologists and philosophers for many years (e.g., Hume 1739/1964). Within psychology, problems related to causal attribution have been studied in areas as diverse as social (Heider, 1958), clinical (e.g., Alloy & Abramson, 1979), developmental (Piaget & Inhelder, 1951), perception (Michotte, 1963), and information processing (Kahneman, Slovic, & Tversky, 1982).

In this chapter, we will focus on causal judgment in the framework of associative learning. This perspective bridges animal and human research in its study of how organisms process causal information. The basic idea underlying this approach is that conditioned responding (CR) occurs in animals under conditions analogous to those that would lead a human subject to conclude that a causal relationship exists between the conditioned stimulus (CS), or cause, and the unconditioned stimulus (US), or effect (see, e.g., Dickinson, 1980; Rescorla, 1988). That is, in contrast to more traditional theories of conditioning (e.g., Pavlov, 1927) that focused on the CS as directly eliciting or causing the CR, in current conditioning research the focus is generally on the subject's learning that the CS (e.g., light, tone, etc.) is a cause or a predictor of the US or effect (usually food, foot shock, or sex object), and hence, the occurrence of a CR is usually seen merely as evidence that the animal has learned that the CS signals the US. Similarly, in instrumental conditioning, the subject's response is frequently viewed as the cause of the outcome (e.g., the delivery of reinforcement, see e.g., Dickinson & Shanks, 1995).

The straightforward prediction following from this approach is that conditioning phenomena are replicable in human causal learning situations, both in laboratory experiments and in real life. More specifically, the conditions that affect animal conditioning, and that have been extensively studied over many years, should also affect causal attribution in human judgment. Recent reviews by Allan (1993), Shanks (1993), Wasserman (1990a, 1993), and Young (1995) provide summaries of the basic findings in this area from the theoretical and experimental perspectives (see also the recent volume edited by Shanks, Holyoak, & Medin, 1996). In this chapter we discuss the conditions that favor the detection (and nondetection) of causal and noncausal relations and show how these findings can be applied for the better understanding and treatment of clinical problems.

BASIC RESEARCH ON CAUSAL LEARNING AS A BRIDGE BETWEEN CONDITIONING EXPERIMENTS AND REAL-WORLD APPLICATIONS WITH HUMANS

Behavior therapists have known for many years that the results of conditioning research with animals can be applied to the understanding of human behavior. However, our point of view, as well as that of many other researchers currently studying basic associative learning in humans, is that the gap between animal conditioning experiments and human applications is enormous and needs to be filled with intermediate steps at the level of basic human research. More specifically, we need experiments showing exactly how human subjects (and not only animals) process and respond to causal information.

As an example of the gap between animal research and human applications, consider learned helplessness, a phenomenon that was developed in the field of animal conditioning (Overmier & Seligman, 1967; Seligman & Maier, 1967) and has been applied to almost all areas of human problems (see Peterson, Maier, & Seligman, 1993 for a recent review). According to learned helplessness theory (Abramson, Seligman, & Teasdale, 1978), when subjects are exposed to response-independent outcomes, they learn that outcomes are uncontrollable. This, in turn, may lead them (if certain attributions are made), to expect that outcomes will remain uncontrollable in the future. If this expectation takes place, subjects become helpless. This may manifest as depression, anxiety, and several other symptoms.

However, if we analyze the available data on learned helplessness, we can see that for several reasons, some additional experiments on human causal learning would be necessary in order to draw such conclusions. First, the theory was initially developed from animal conditioning experiments. Thus, the central assumption that subjects detect that outcomes were response-independent could only be inferred from the animal's behavior. Second, most applied settings in which this theory has been used to explain subjects' problems (e.g., depression) represent conditions in which the symptom is already there by the time the researcher begins the investigation. Thus, the assumption that it was caused by the subject's detection of the absence of a causal relation between responses and outcomes is again an inferred one.[1] Third, laboratory experiments on learned helplessness in humans have generally focused on the attributional component that, according to the theory, takes place *once subjects have already detected* that outcomes are response-independent (Abramson et al., 1978). Consequently, learned helplessness researchers frequently inform subjects that they are not controlling the outcome (see e.g., Abramson & Alloy, 1981), and by giving this information, researchers are able to investigate the subsequent attributions and potential deficits. However, by

---

[1] Note that the difficulty of detecting the etiology of a real-world problem is not constrained to the area of learned helplessness, but is a well-known problem in clinical psychology.

using this strategy, the assumption that subjects can detect response-outcome independence remains untested. As noted by Abramson and Alloy (1981, p. 438), "with uncontrollability so transparent, it is no wonder that [nondepressives] detect noncontingency accurately." But what if subjects were unable to detect noncontingency by themselves (Schwartz, 1981a, 1981b)? Obviously, subjects would make no attributions concerning uncontrollability if they were not able to detect it, nor would they develop any of the subsequent potential deficits. Experiments that actually test whether subjects exposed to uncontrollable outcomes are able to detect the absence of a causal relation between their responses and the occurrence of the outcomes (and if so, how and in which conditions) are a necessary intermediate step between animal learned helplessness experiments and applications of learned helplessness theory to human real-world domains. Below we consider this and other related areas of basic research in human causal learning, and analyze in which cases and to which extent are applications justified.

DIFFICULTY IN DETECTING CAUSALITY WHEN CONTIGUITY IS WEAK

One of the best-known facts from animal and human conditioning is that the CS (cause) and the US (effect) must occur in temporal and spatial proximity for the animal to learn a predictive relationship between them (i.e., for conditioning to take place). In general, as the CS-US delay increases, conditioned responding becomes weaker (Pavlov, 1927). The same applies to the response-outcome contiguity in instrumental learning.

This impact of contiguity upon behavior (human and animal) is well known in behavior therapy (e.g., the problem of the long delay between going on a diet and losing weight). The interesting thing for our present purposes is that exactly the same general rule applies to human causal learning. For example, Shanks, Pearson, and Dickinson (1989) reported several causal judgment experiments with humans in which they manipulated the temporal delay between the potential cause (the subject's response [R] in a computer keyboard) and the effect (the appearance of a visual outcome [O] on the computer screen). Subjects experienced this relationship during several learning trials and were asked to rate the degree to which their own response was the cause of the appearance of the outcome. In all conditions, increasing the R-O delay produced lower judgments of causality, all other things being equal. Similarly, Wasserman and Neunaber (1986) showed that subjects tended to perceive a causal R-O relation when the response advanced the occurrence of the outcome but not when the response postponed the occurrence of the outcome.

Thus, contiguity affects not only behavior, but also the causal attributions that subjects make. If contiguity is weak, subjects may fail to detect a causal relation that does exist. As an example, consider a patient's difficulty in recognizing the beneficial effects of a therapy with a necessary long delay before a beneficial outcome occurs. Not surprisingly, many patients may prefer pharmacological over behavioral therapy, due to the differential response-outcome intervals.

Thus, if we aim to change a patient's attribution of causality, one of the aspects we ought to consider manipulating is contiguity and the factors that affect it. For example, we might think of "filling the gap" between the response and the outcome by introducing a signaling stimulus when contiguity is weak (e.g., Kaplan & Hearst, 1982). According to some authors (e.g., Young, 1995), this establishes a causal chain that enhances the likelihood that the response be perceived as the ultimate cause for the distant outcome (see Rescorla, 1982, for a different interpretation). This result has been observed in several human causal learning experiments that used weak R-O contiguity (e.g., Reed, 1992; Shanks, 1989, Experiment 2): The insertion of a CS during the delay period enhanced the perception of a causal relation between the response and the delayed outcome. Applications may include recognizing the several steps (mood changes, etc.) the patient may experience during the delay period as part of the causal chain leading to the distant outcome.

So far, we have focused our discussion on the difficulty in detecting a causal relation when contiguity is weak. The symmetrical problem is that high R-O contiguity favors the perception of a causal relation even if the relationship occurs by pure chance. We discuss this prediction in the next section.

EXPOSURE TO UNCONTROLLABLE OUTCOMES: HELPLESSNESS OR SUPERSTITION?

Consider a patient who has been diagnosed as having a serious illness and has decided to follow some "newly discovered" therapeutic exercises that are claimed to produce incredible improvements. The patient is indeed experiencing some recovery. Should this recovery be attributed to the exercises? Of course, in order to know whether or not a relationship is causal, the subject needs to know, in some way or another, two things. First, what is the probability of the outcome's occurring in the *presence* of responding, p(O|R), and second, what is the probability of the outcome's occurring in the *absence* of responding, p(O|noR) (see Allan, 1993; Shanks, 1993 for discussions on how this learning may take place). If p(O|R) is greater than p(O|noR), then the subject may conclude that performing the response increases the likelihood of the outcome occurring. By contrast, if p(O|noR) is greater, then the subject should refrain from responding in order to obtain the outcome. These two cases represent contingent conditions in which the response (or its absence) affects the occurrence of the outcome. Finally, if p(O|R) equals p(O|noR), then the response does not influence the outcome; hence, the outcome is said to be response-independent because it occurs with the same probability whether or not the subject responds.

Many researchers have examined whether subjects are actually sensitive to these event contingencies and can discriminate contingency (conditions in which the response affects the outcome, that is, p[O|R] < or > p[O|noR]) from adventitious contiguity (p[O|R] = p[O|noR]). The results of these experiments are somewhat mixed. Whereas some researchers have reported that animal and human subjects are sensitive to response-outcome independence and can learn that they do not have control over the outcomes (e.g., Seligman, 1975), others have shown that subjects exposed to response-independent outcomes tend to think they are controlling the outcome (i.e., illusion of control) and to behave superstitiously (e.g., Alloy & Abramson, 1979; Langer, 1975; Matute, 1994; Skinner, 1948; Wright, 1962). What evidence supports each of these discrepant views?

The seminal paper on the exposure to response-independent outcomes was published by Skinner in 1948. Skinner exposed a group of pigeons to uncontrollable (free) food on a fixed time schedule and observed that pigeons tended to develop repetitive patterns of behavior as if they "thought" their behavior was causing the reinforcer to occur. This was proposed as a model for human superstitious behavior, which according to Skinner would take place in much the same way as superstitious behavior in the pigeon. That is, even if the response is not the cause of the outcome, the response is adventitiously reinforced through R-O contiguity; thus, the response will tend to occur more frequently in the future. Several researchers replicated Skinner's findings in human subjects (e.g., Catania & Cutts, 1963; Wright, 1962). However, these views were neglected, in part because of a very convincing article by Staddon & Simmelhag (1971) that explained "superstition in the pigeon" in terms of reinforcers predisposing subjects toward certain temporally and reinforcer-specific responses, rather than randomly selecting ongoing behavior as suggested by Skinner. Additionally, Skinner's ideas were contrary to the intuitive view that organisms (especially humans) are sensitive to causal relations and can distinguish them from spurious correlations.

Learned helplessness theory (Abramson et al., 1978; see Overmier & LoLordo in this volume for a current review) was then developed, based on the view that animal and human subjects exposed to noncontingency between responses and outcomes were able to learn that their behavior did not control (i.e., cause) reinforcement. According to Abramson et al., learning that outcomes were uncontrollable could lead, if certain attributions were made, to several deficits that are collectively termed the *learned helplessness* effect (Overmier & Seligman, 1967).

The proposal that subjects were able to learn response-outcome independence was contrary to traditional views (e.g., Skinner, 1948) that subjects learn by simple contiguity, regardless of whether reinforcement was controllable or not. This difference was not only relevant at the theoretical level, but it had important implications for clinical problems as well. For instance, according to some traditional models of behavior therapy, depression was caused by an overall absence of reinforcers. In contrast, according to learned helplessness theory, it was not the absence of reinforcers that caused depression, but the subject's perception (and later expectations) that reinforcers were occurring independently of her or his behavior (see Maier, 1989; Overmier & LoLordo, this volume; Seligman, 1975).

However, most of the evidence seemingly supporting learned helplessness theory has been obtained under conditions in which subjects are exposed to "failure feedback," and consequently is open to alternative interpretations. That is, in a typical learned helplessness experiment (e.g., Hiroto & Seligman, 1975), subjects are either (a) told that they gave the wrong answer each time they finish a cognitive problem, or (b) exposed to uncontrollable reinforcement (noise termination) but told, after each trial, that they failed to stop the noise (see Abramson & Alloy, 1981; Hiroto & Seligman, 1975, for procedural details commonly used in human learned helplessness research; Matute, 1994, for further discussion of methodological issues). Consequently, alternative explanations such as failure exposure (e.g., Buchwald, Coyne, & Cole, 1978), egotism (Frankel & Snyder, 1978), or even extinction due to the continuous failure to get the desired event may be more relevant to these results than learned helplessness theory. As acknowledged by Abramson and Alloy (1981), these experiments do not test whether the subjects detect response-reinforcer independence (see also Peterson et al., 1993, for a current discussion of alternative theories).

One solution to this problem is to eliminate the failure feedback used in the typical manipulation. That is, subjects can be exposed to uncontrollable reinforcement (noise termination) in the absence of failure feedback. However, we found not one reference reporting a true learned helplessness effect in the absence of failure feedback. For example, among the very few human experiments that have not used failure feedback, Thornton & Jacobs (1971) found that the behavior of the "helpless" group in a subsequent test phase did not differ from that of the control group. Kofta and Sedek (1989) were able to obtain an effect, but this effect was interpretable by an alternative mechanism (Sedek & Kofta, 1990). Moreover, when failure feedback is eliminated from the typical learned helplessness procedure, subjects do not detect response-outcome independence. On the contrary, they ordinarily develop an illusion of control and superstitious behavior (Matute, 1994, 1995), similar to what was previously described by Alloy and Abramson (1979), Langer (1975), Skinner (1948), Wright (1962), and many others. Indeed, many subjects (as well as some patients) adopt a suspicious attitude if they are subsequently informed that their behavior was not causing the outcome.

Even though "true" learned helplessness (i.e., deficits unambiguously attributable to the perception and expectations of uncontrollability) has not yet been demonstrated in human subjects, and even though "pure" uncontrollability (i.e., in the absence of failure feedback) seems to lead to the opposite outcome (i.e., illusion of control and superstitious behavior), learned helplessness and superstitious behavior could simply represent opposite ends of the same continuum. Consider, for example, the patient who has been diagnosed as having an incurable illness. This patient may take either of two courses of action.

First, she or he may realize that outcomes (e.g., transient improvements) are response-independent (i.e., therapy-independent), that nothing can really change the outcome, and consequently may become depressed. Second, this patient may try all types of therapies, including the most superstitious ones (e.g., witchcraft), and associate transient recoveries with some of those "therapies." This behavior may be called superstitious, but the pejorative name notwithstanding, it may be preventing depression.

It is not unusual to observe those two opposite patterns of responding to uncontrollable outcomes in different individuals in real life. Under desperate situations people may tend to either superstition or depression. The finding by Alloy and Abramson (1979) that nondepressed subjects are more prone to illusions of control than are depressed subjects is consistent with this hypothesis. A clear implication (see Alloy & Abramson, 1988; Alloy & Clements, 1992) is that, in cases in which the outcome is uncontrollable, letting subjects keep their illusions and superstitions, or even favoring the development of new ones, can have prophylactic effects against depression and other problems.

ACCURATE DETECTION OF CAUSAL AND NON-CAUSAL RELATIONS

We have argued that in many uncontrollable situations, subjects do not realize that they lack control over the outcome. But is this always true? Thus, an important remaining problem is to identify the conditions that lead to the perception of response-outcome independence because these conditions will ultimately determine, among other things, whether learned helplessness may occur (Maier & Seligman, 1976).

Even though these conditions have not been well identified in the learned helplessness literature, there are several experiments conducted on the judgment of causality as well as in animal learning that indicate under which conditions response-outcome dependence and independence will be perceived. In general, results have been quite consistent. In contingent conditions—that is, conditions in which outcomes do depend on the subject's active or passive responding—animal subjects learn to respond differentially to different event contingencies (e.g., Hallam, Grahame, & Miller, 1992; Rescorla, 1968) and the causal ratings of human subjects reflect that they too can accurately detect the actual event contingencies (see Shanks & Dickinson, 1987; Wasserman, 1990a for reviews). The results are also consistent with respect to detection of noncontingent conditions (those in which the effect is independent of the cause), provided that the potential cause is an exogenous event (i.e., not a subject's response, e.g., Baker & Mackintosh, 1977; Rescorla, 1968).

Moreover, within the laboratory, causal judgment data indicate that humans are able to perceive, *at least under ideal conditions*, when their response does not cause an outcome to occur (see Shanks & Dickinson, 1987; Wasserman, 1990a for reviews). It should be noted, however, that these experiments are more theoretically oriented than the experiments on superstition and learned helplessness. That is, subjects in these causal judgment experiments are informed that their goal is to find out how much control they have over a certain outcome (e.g., a triangle flashing on the computer's monitor). They are also often told that the best strategy they can use in order to find out how much control is possible is to respond on about 50% of the trials so that they can be equally exposed to both p(O|R) and p(O|noR). Thus, what these experiments show is that, at least under these ideal laboratory conditions, human subjects do have the capacity to learn that outcomes are response-independent. The problem is that this "scientific behavior" (i.e., responding on 50% of the trials to test both p[O|R] and p[O|noR]) does not seem to occur spontaneously in more naturalistic conditions (e.g., learned helplessness experiments) in which subjects are exposed to aversive outcomes and are trying to escape. For example, Matute (1996) observed that in their attempts to escape, most subjects exposed to uncontrollable aversive conditions responded at almost every opportunity.

Therefore, no matter how uncontrollable reinforcement (termination of the aversive stimulus) was, it tended to occur in the presence, rather than in the absence, of responding. Thus, subjects could not learn that the outcome would have occurred with the same probability if they had not responded. This in turn caused subjects to believe that their response was controlling the outcome, which led to an illusion of control.

Going back to our previous example of the patient diagnosed as having a serious illness and who acted at a very high rate (trying witchcraft and all other possible types of therapy), most transient recoveries would probably be contiguous with his or her acting rather than being passive. Thus, it would not be surprising if this patient continued to use superstitious therapies rather than becoming helpless or depressed. Only if this patient stops responding (perhaps because of a lack of reinforcement during a prolonged interval) will some eventual reinforcement occur in the absence of responding thereby allowing the subject to learn that his or her behavior did not influence the temporary recovery.

In summary, we have seen that subjects do have the capacity to detect response-outcome independence, but they tend toward illusions of control and superstitious behavior in naturalistic conditions in which they are trying to obtain a valued outcome that occurs noncontingently. The subjects' probability of responding, p(R), appears to be an important modulating factor: subjects responding at high rates (e.g., subjects responding at every opportunity to escape an aversive situation) tend toward stronger illusions of control than subjects responding at lower rates (probably because a higher percentage of responses become reinforced—in a partial schedule—as the subject's p[R] approaches 1). According to this view, conditions that reduce p(R), such as fatigue, punishment, depression, extinction (lack of reinforcers), or any other source for a low (passive) response rate can lead to a detection of response-outcome independence by allowing subjects to be exposed to both p(O|R) and p(O|noR). Consequently, learned helplessness effects, rather than superstitious behavior, could potentially be observed under those conditions.

One more factor that could reduce the illusion of control (and perhaps favor depression in some cases) is the availability of an alternative cause (different from the subject's response) to which the outcome could be attributed. For example, early people's illusion that they caused rain by dancing probably disappeared when science provided an alternative, competing cause for rain. Of course, this process of discounting one potential causal factor on the basis of the availability of alternative causes is not constrained to the illusion-of-control problem and has received considerable attention on its own right. The rest of the chapter addresses these findings.

DISCOUNTING POTENTIAL CAUSES: THE INFLUENCE OF PRIOR (AND SUBSEQUENT) KNOWLEDGE

Let us consider a totally different example: people suffering from sleep disorders who are given both hypnotics *and* relaxation training sometimes attribute their improvement to the drug rather than to their new relaxation abilities (Davison, Tsujimoto, & Glaros, 1973). Why? We cannot claim the effects of contiguity (nor contingency) in this case. The patient falling asleep after having consumed a hypnotic *and* having performed a relaxation exercise is exposed to similar contiguity (and contingency) between each of the two potential causes and the reinforcer. We may claim, however, that prior knowledge (perhaps through earlier sharing of information by others, perhaps through the initial sessions in which the relaxation abilities were poor) tells the subject that the drug alone is sufficient to induce sleep. Consequently, the subject discounts relaxation as a therapeutic factor.

This effect is called *forward blocking* and is well documented in animal (e.g., Kamin, 1968) and human (e.g., Shanks, 1985) research. In a typical animal experiment, the subjects may be exposed, during Phase 1, to several trials in which a light CS is followed by a foot-shock US.

Then, in Phase 2, the light is presented simultaneously with a tone and followed again by a foot shock. That is, the experiment takes the form of A→E (A followed by E) in Phase 1, and AX→E in Phase 2. Then at test, X (the tone in our example) is presented alone to see if it produces fear of the foot shock US, which is equivalent to asking the subject whether X predicts that E will occur. In general, X is not perceived as a predictor of the foot shock (experimental subjects do not show fear of X). In contrast, control animals that lack the critical prior knowledge provided by the A→E pairings in Phase 1 perceive X as a predictor of E. The usual explanation is that in the experimental group, prior learning that A predicts E *blocks* the attribution of a causal role to X if X is trained in the presence of A. A similar result has been observed in many instrumental animal experiments in which a CS predicting reward caused a decrement in instrumental responding (e.g., Pearce & Hall, 1978; Williams, 1975). (As a real-world example, consider also the differential causal explanations given to a clinical case by therapists coming from different theoretical perspectives).

Competition between potential causes has also been demonstrated in the human laboratory. For example, Shanks (1985) and Hammerl (1993) have shown that humans reduce the causal role that they attribute to their own behavior if an additional potential cause (a CS) for the outcome is present.

Other human experiments on competition between causes (e.g., Wasserman, 1990b) have frequently used a medical diagnostic situation in which subjects are shown the records of fictitious patients who have consumed some allergens (which are analogous to potential causes or CSs) and then developed some allergies (analogous to potential effects [E] or USs). Subjects are subsequently asked about the degree of causal relation between the target allergen X and the allergic reaction. As in the animal experiments described above, human subjects tend to discount the potential causal role of allergen X if the competing allergen A has a stronger association to the allergic reaction. Whether this deficit in attributing a causal role to X is due to the subjects' not having learned the X→E relationship (e.g., Rescorla & Wagner, 1972) or to their subsequent discounting (even though they learned it) the influence of X on the basis of the stronger A→E relation (e.g., Miller & Matzel, 1988; Shanks & Dickinson, 1987) is a matter of current theoretical debate.

This debate, despite its theoretical nature, has important implications. For example, suppose a patient has been exposed to several A→E episodes in which A caused effect E (e.g., alcohol→sexual impotence). Then A (alcohol) begins to occur simultaneously with X (e.g., disturbing thoughts) and is also followed by E (i.e., AX→E). According to the results of many experimental reports with both animals and humans, this person will probably discount the potential causal role of X (the disturbing thoughts in our example). But what if this were also an important causal factor? The patient's behavior would be more adaptive if the X→E relation were learned, and in some way stored, even if X were discounted as a causal factor, than if the X→E relation had never been encoded. In other words, if the information about the relationship between X and E had not initially been acquired due to competition from the stronger A→E association (e.g., "Because I know that A is the cause of E, I do not need to pay attention to X"), then the attribution could never be reversed. In contrast, if the information about X had been processed and stored, though discounted as a causal factor, then this attribution would be available to be reversed if necessary.

On the one hand, traditional associative theories (e.g., Rescorla & Wagner, 1972) posit that competition between causes occurs during acquisition (if one association is very strong, the other one is not learned) and thus it is not subject to reevaluation. On the other hand, comparator theories of learning (e.g., Miller & Matzel, 1988; Shanks & Dickinson, 1987) posit that subjects acquire associations concerning all potential causes present on a given situation. According to this view, discounting effects take place at a postlearning (i.e., reasoning or attributional) stage, and the information is there to be reevaluated if necessary (see also Van Hamme & Wasserman, 1994, for an explanation of retrospective evaluation as learning, rather than postlearning, process; and Miller & Matute, 1996, for difficulties with such a model).

Many experiments have been conducted to investigate whether information is subject to retrospective reevaluation. For example, consider the following design: AX→E in Phase 1 followed by A→no E in Phase 2. That is, the subjects first learn that A and X, presented together, lead to E, and then they learn that A by itself does not lead to E (e.g., they first learn that the combination of two therapeutic techniques produces an improvement, and then they learn that technique A, when used by itself, does not produce the effect). Will learning that A does not lead to E enhance the causal attribution to X that occurred during Phase 1? The two approaches already mentioned address this question in different ways. Traditional associative theories (e.g., Rescorla & Wagner, 1972) predict that during Phase 1, subjects learn about both A and X, and then further learning during Phase 2 that A by itself does not cause E has no impact on what subjects had initially learned about X during Phase 1. Comparator theories (Miller & Matzel, 1988; Shanks & Dickinson, 1987), on the other hand, predict that learning in Phase 2 that A does not produce the effect will prompt subjects to reexamine their attribution and give a stronger causal role to X. That is, subjects would conclude that the therapeutic technique X was more effective than they had initially thought. This prediction has been confirmed in several experiments (e.g., Dickinson & Charnock, 1985; Kaufman & Bolles, 1981; Matzel, Schachtman, & Miller, 1985). Similarly, many other experiments have also demonstrated that other types of posttraining manipulations (variations of the Phase 2 training in this example) result in subjects' (animals and humans) reexamining their initial attribution and giving a second, modified, conditioned (or attributional) response (Chapman, 1991; Cole, Barnet, & Miller, 1995; Shanks, 1985; Van Hamme & Wasserman, 1994). Thus, we can conclude that people can reevaluate their initial attributions of causality, but of course a careful examination of the conditions under which retrospective processing takes place is important. These are discussed below.

STRONGLY VALENCED CAUSES ARE NOT DISCOUNTED

Consider now a patient who attributes her or his success at business meetings to two causes, (A) preparing for the meeting, and (X) performing a ritual of obsessive (superstitious?) behaviors. Thus, in Phase 1 we have AX→E. Then, you try to convince this patient that being prepared is sufficient to be successful (A→E) and that, therefore, the ritual (X) is superfluous. Will the patient finally accept the hypothesis that X was not a causal factor based exclusively on the A→E information? Or will some additional therapy (e.g., extinction of X) be necessary?.

This is called *backward blocking* because the traditional two phases of a forward blocking design (see above) are reversed (i.e., the compound phase is here presented as Phase 1 rather than Phase 2). Based on what we have already discussed, we would expect that A→E training in Phase 2 would result in subjects' discounting the potential causal role they may have attributed to X during Phase 1. This prediction has been confirmed in several human causal learning studies (e.g., Chapman, 1991; Shanks, 1985) but has resisted the attempts of animal researchers for many years (Miller, Hallam, & Grahame, 1990; Schweitzer & Green, 1982). Thus, this type of retrospective evaluation is a potential difference in information processing between animals and humans, suggesting that sometimes, even in simple laboratory tasks, humans perform differently than animals.

However, this is not necessarily so. Miller and Matute (1996) noted an important difference between the experiments that attempted to obtain this effect in animal subjects and those that had obtained it in human subjects. In animal experiments, the effect E was always an event of high biological significance (usually a foot-shock

US), whereas the effect E used in human experiments was always a neutral, fictitious event occurring on a computer screen (e.g., a fictitious allergy developed by a fictitious patient). Thus, Miller and Matute ran rats as subjects in a procedure as similar as possible to the one that had been used with human subjects. Specifically, they used neutral tones, rather than biologically relevant stimuli, as both causes and effects. (Only after completion of the two phases of backward blocking was the effect made biologically significant through pairing it with a foot-shock US, in order to provide the motivation needed to assess responding in a subsequent test phase.)

The results showed a retrospective processing effect in animals using the backward blocking design (i.e., AX→E; A→E; Test X). Rat subjects demonstrated that they, like humans, can reconsider the causal role attributed to X during Phase 1 (i.e., AX→E) if they later learn that A alone is sufficient to produce E (i.e., A→E). Compared with that of control rats that had not received the critical A→E training in Phase 2, the response of the experimental subjects to X was substantially weakened.

Thus, whereas traditional attempts to obtain backward blocking in animals using biological significant USs as effects had failed, studies using neutral stimuli that are known to yield backward blocking in humans produced backward blocking in animals. This suggests that, whereas organisms can discount the role of potential causes in neutral settings, this discounting process is less apt to occur in valenced settings.

The conditions under which backward blocking was obtained in this experiment as compared with previous attempts are of potential importance when considering applications to behavior therapy. For example, increasing the biological significance of the US (nausea) or the alcohol concentration and taste in aversion therapy with alcoholics should increase the efficacy of these techniques by reducing the likelihood that the patients discount alcohol as the cause of their sickness. Moreover, if the business success of our previously discussed patient is a biologically significant event (those that are important for the well-being of the organism) we should predict that this patient will not be able to discount the role of the obsessive ritual as a causal factor in producing success.

This prediction was further supported in the laboratory by Miller and Matute (1996), who observed that even forward blocking, which is commonly observed in animal subjects (e.g., Kamin 1968), failed to occur if the auditory stimuli used as potential causes were of high intensity (and thus, presumably, of higher biological significance than neutral stimuli commonly used as CSs in conditioning experiments). That is, when X was biologically significant (e.g., very intense), subjects did not discount its potential causal role as is usually the case. Moreover, Denniston, Miller, and Matute (1996) observed that the influence of biological relevance on backward blocking was not constrained to cases in which biological significance was a feature inherent to the stimulus (e.g., high intensity). Rather, neutral stimuli (e.g., moderate intensity sounds) can also become biologically significant, and thus protected from discounting processes, if they are previously associated with a significant event (e.g., foot-shock US). In this study, E→US (Phase 1), followed by AX→E (Phase 2), A→E (Phase 3), and E→US (Phase 4), did not produce backward blocking when X was presented alone during the subsequent test phase. (Rats showed fear of X, thus suggesting that they had not discounted X as a causal factor in producing E, and thus, the US foot-shock.) However, when Phase 1 (E→US) was eliminated (i.e., X was a neutral event not associated with a US), rats did discount X as a causal factor producing E (rats showed little fear of X). This suggests that an important factor in modulating the occurrence of these selective attribution effects is whether the potential causes are biologically significant (either because of their inherent properties or because they have been previously conditioned through pairings with a significant event.)

Most situations of interest for applied psychology involve valenced conditions that are important to the well-being of the subject, either because the effect is inherently significant or because it has been associated to significant events (USs) during the patient's history. Consequently, we anticipate that backward blocking will probably be difficult to obtain in applied situations. Thus, if patients have experienced causes A and X followed by a significant US effect in the past, they will probably be reluctant to discount X as the cause of the US, no matter how many times we show them that A is the cause of the US. According to this view, extinction of X, rather than (or in addition to) showing that A alone was the cause of the US, will probably be necessary (recall our obsessive business person above; or the difficulty in eliminating stereotypes).

PREDICTING EFFECTS AND DIAGNOSING CAUSES

We have seen that animals exposed to A→E during Phase 1 and AX→E during Phase 2 do not expect E (foot-shock) at test when X is presented alone; and that the patient who is exposed to hypnotics in Phase 1 and hypnotics plus relaxation in Phase 2 does not anticipate sleep when relaxation alone is used.

Now imagine that A usually precedes (causes) symptom E (A→E), but that later on you see new patients who have developed the symptoms E and X simultaneously after being exposed to A (i.e., A→EX). That is, here, instead of having two potentially competing causes we have two potentially competing effects. In the future, would you tend to view these two symptoms, E and X, as competing cues for diagnosing that cause A has occurred? That is, would you discount the potential diagnostic value of X as a result of the initial A→E pairings? Or in contrast, would you interpret those two syndromes as cumulative evidence that cause A has occurred?

Waldmann and Holyoak (1992) predicted this later possibility as the only possible outcome of such a scenario. That is, they posited that collaboration rather than competition between effects would occur in human subjects exposed to multiple potential effects of a common cause. Indeed, their results, which were replicated by Van Hamme, Kao, and Wasserman (1993), seemed to confirm this hypothesis. However, there is also considerable evidence indicating that if a certain effect or symptom has a strong diagnostic value, other effects tend to be ignored as potential diagnostic cues both in human (Chapman, 1991; Matute, Arcediano, & Miller, 1996; Price & Yates, 1993; Shanks & Lopez, 1996) and animal subjects (Esmoris-Arranz, Miller, & Matute, 1997). Apparently, effects compete in some situations as diagnostic cues and collaborate in other situations. Matute et al. (1996) investigated which conditions yielded competition between effects and which ones did not. They studied several conditions and concluded that the wording of the question used for assessment of the causal attribution plays an important role in determining whether effects compete or collaborate among themselves. That is, when subjects are asked whether each of the effects is an "effect of the cause," effects seem to collaborate with each other (subjects accept all potential effects of the cause in a noncompetitive manner); however, when subjects are asked about the "diagnostic value" of each of the effects, effects seem to compete among themselves (i.e., accepting effect A as a good diagnostic cue tends to lower the diagnostic value of effect X.) Thus, for example, a patient may view just one symptom as *the* diagnostic cue to recognize when an illness has come back (when is time to go back into therapy); but on the other hand, once the illness is present, the patient will probably predict all possible symptoms. As in the multiple-causes conditions described above, these attributions concerning multiple effects should also be subject to reevaluation under some conditions if the initial attributions prove wrong. However, research on competition and collaboration between effects has just started, and there is much to be done before the conditions under which reevaluation occurs in diagnostic settings can be identified.

Finally, note that this discussion of competition versus collaboration between effects also raises the issue of whether causes, which have traditionally been assumed to compete among themselves for predictive value (see previous section), could also collaborate among themselves in some cases. Indeed, collaboration between causes has been described in several different settings (e.g., Leddo, Abelson, & Gross, 1984). As an example, knowledge of a history of childhood abuse and of a history of mental illness in kin can collaborate, rather than compete, as predictors of mental illness in a patient. Based on available evidence, causes and effects do not appear to substantially differ in their susceptibility to competition or collaboration. Effects compete for diagnostic value just as causes compete for predictive value, but only when demand characteristics foster competition among the relevant stimuli.

CONCLUSIONS

We have tried to bring together research performed with animal and human subjects in the field of causal learning and to show some avenues through which this approach can benefit clinical psychology. In doing so, we have attempted (a) to call the attention of the reader to some points that, in our view, are not yet fully developed in basic research and to which we recommend caution if they are going to be applied, and (b) to suggest potential applications of those aspects that have been more thoroughly elaborated in basic research.

In summary, basic research on causal learning has revealed four important principles. First, causal judgments are sensitive to both contiguity and contingency between cause and effect. Thus, enhancing these two factors should prove beneficial when we want our patients to detect a causal relation between, say, a therapy and a beneficial effect.

Second, many laboratory experiments have shown that exposure to uncontrollable outcomes may result in superstitious behavior and illusion of control instead of resulting in learned helplessness and depression, as had been generally assumed within clinical psychology. Indeed, much of the human data that apparently support learned helplessness theory can be explained by alternative theories such as exposure to failure feedback or an absence of reinforcers. On the other hand, the conditions that led to the perception of uncontrollability (and that could lead, therefore, to unambiguous helplessness deficits) are not yet well established. But it seems that conditions leading to passivity and reduced responding (e.g., fatigue, punishment, extinction) probably favor subjects' detection of response-outcome independence, and thus the development of learned helplessness over superstitious behavior. Additionally, the existence of some alternative potential cause for the outcome (other than the subject's behavior) may also reduce the illusion of control. Nevertheless, it is possible that superstitious behavior and illusion of control may prevent the development of helplessness and depression. Thus, in many situations they should be encouraged when possible.

Third, the process of selective attribution and discounting of potential causes allows subjects to select the most likely cause for an outcome. This process readily explains why subjects may attribute their improvement to some (but not all) of the components of a given therapy (e.g., "What makes me sleep is the pill, not the relaxation exercise"). Additionally, it has been shown that subjects can reconsider these selective attributions should they receive new information suggesting that the initial attribution was in error. However, these selective attribution processes seemingly occur only in situations with relatively neutral events; biologically significant events (those that are important to the well-being of the organism) are apparently protected from discounting processes. For example, would it be possible to discount the role of a potential cause if this cause is a very significant event in the patient's life (e.g., the loss of a loved one)? Probably not. Discounting procedures in applied situations would *probably* require the addition of extinction (or other) procedures if the to-be-discounted cause is of inherent or acquired biological significance.

Finally, causes and effects do not appear to be substantially different in their susceptibility to competition. That is, under analogous conditions, multiple effects seem to compete or collaborate for diagnostic value just as multiple causes compete or collaborate for predictive value. As we mentioned, a patient may pay attention to just one symptom as *the* diagnostic cue indicating that the illness is coming back, but may also predict that, once the illness is established, all other possible symptoms will occur. These attributional biases sometimes may prove adaptive, and sometimes may not (e.g., the patient may be using a wrong diagnostic cue). The patient's recognizing that selective attributions can occur, and that under some conditions they can be rather biased, might prove helpful in obtaining more accurate and complete judgments of the actual causal factors.

ENDNOTES

Support for the preparation of this chapter was provided by Dirección General de Investigación Científica y Técnica (Spain) Grant PB95-0440, Grant PI96-6 from the Basque Government, and National Institute of Mental Health (USA) Grant 33881. We thank Pilar Antón, José Cáceres, Pilar Castro, and Ane Uribe for illuminating conversations suggesting potential clinical applications of phenomena that we described in an earlier draft, and for their critical comments that helped us turn that draft into a chapter that we hope will be more relevant to behavior therapy.

REFERENCES

Abramson, L. Y., & Alloy, L. B. (1981). Depression, nondepression, and cognitive illusions: Reply to Schwartz. *Journal of Experimental Psychology: General, 110,* 436–447.

Abramson, L. Y., Seligman, M. E. P., & Teasdale, J. D. (1978). Learned helplessness in humans: Critique and reformulation. *Journal of Abnormal Psychology, 87,* 49–74.

Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin, 114,* 435–448.

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General, 108,* 441–485.

Alloy, L. B., & Abramson, L. Y. (1988). Depressive realism: Four theoretical perspectives. In L. B. Alloy (ed.), *Cognitive processes in depression,* 223–265. New York: Guilford.

Alloy, L. B., & Clements, C. M. (1992). Illusion of control: Invulnerability to negative affect and depressive symptoms after laboratory and natural stressors. *Journal of Abnormal Psychology, 101,* 234–245.

Baker, A. G., & Mackintosh, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and US. *Animal Learning and Behavior, 5,* 315–319.

Buchwald, A. M., Coyne, J. C., & Cole, C. S. (1978). A critical evaluation of the learned helplessness model of depression. *Journal of Abnormal Psychology, 87,* 180–193.

Catania, A. C., & Cutts, D. (1963). Experimental control of superstitious responding in humans. *Journal of the Experimental Analysis of Behavior, 6,* 203–208.

Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17,* 837–854.

Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Effect of relative stimulus validity: Learning or performance deficit? *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 293–303.

Davison, G. C., Tsujimoto, R. N., & Glaros, A. G. (1973). Attribution and the maintenance of behavior change in falling sleep. *Journal of Abnormal Psychology, 82,* 124–133.

Denniston, J. C., Miller, R. R., & Matute, H. (1996). Biological significence as a determinant of cue competition. *Psychological Science, 7,* 325–331.

Dickinson, A. (1980). *Contemporary animal learning theory.* Cambridge: Cambridge University Press.

Dickinson, A., & Charnock, D. J. (1985). Contingency effects with maintained instrumental reinforcement. *Quarterly Journal of Experimental Psychology, 37B,* 397–416.

Dickinson, A., & Shanks, D. R. (1995). Instrumental action and causal representation. In D. Sperber, D. Premack, & A. J. Premack (eds.), *Causal cognition: A multidisciplinary debate,* 5–25. Oxford: Clarendon Press.

Esmoris-Arranz, F. J., Miller, R. R., & Matute, H. (1997). Blocking of subsequent and antecedent events. *Journal of Experimental Psychology: Animal Behavior Processes, 23,* 145–156.

Frankel, A., & Snyder, M. L. (1978). Poor performance following unsolvable problems: Learned helplessness or egotism? *Journal of Personality and Social Psychology, 36,* 1415–1423.

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 and Motivation, 23,* 225–249.

Hammerl, M. (1993). Blocking observed in human instrumental conditioning. *Learning and Motivation, 24,* 73–87.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

Hiroto, D. S., & Seligman,. M. E. P. (1975). Generality of learned helplessness in man. *Journal of Personality and Social Psychology, 31,* 311–327.

Hume, D. (1964). *Treatise on human nature* (ed., L. A. Selby-Bigge). London: Oxford University Press. (Originally published 1739.)

Kahneman, D., Slovic, P., & Tversky, A. (eds.) (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, UK: Cambridge University Press.

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M.R. Jones (ed.), *Miami symposium on the prediction of behavior: Aversive stimulation,* 9–31. Miami, Fl.: University of Miami Press.

Kaplan, P. S., & Hearst, E. (1982). Bridging temporal gaps between CSs and USs in autoshaping: Insertion of other stimuli before, during, and after CS. *Journal of Experimental Psychology: Animal Behavior Processes, 8,* 187–203.

Kaufman, M. A., & Bolles, R. C. (1981). A nonassociative aspect of overshadowing. *Bulletin of the Psychonomic Society, 18,* 318–320.

Kofta, M., & Sedek, G. (1989). Repeated failure: A source of helplessness or a factor irrelevant to its emergence? *Journal of Experimental Psychology: General, 118,* 3–12.

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311–328.

Leddo, J., Abelson, R. P., & Gross, P. H. (1984). Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology, 47,* 933–943.

Maier, S. F. (1989). Learned helplessness: Event covariation and cognitive changes. In S. B. Klein & R. R. Mowrer (eds.), *Contemporary learning theories: Instrumental conditioning theory and the impact of biological constraints on learning.* Hillsdale, N.J.: Erlbaum.

Maier, S. F., & Seligman, M. E. P. (1976). Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General, 105,* 3–46.

Matute, H. (1994). Learned helplessness and superstitious behavior as opposite effects of uncontrollable reinforcement in humans. *Learning and Motivation, 25,* 216–232.

Matute, H. (1995a). Human reactions to uncontrollable outcomes: Further evidence for superstitions rather than helplessness. *Quarterly Journal of Experimental Psychology, 48B,* 142–157.

Matute, H. (1996). Illusion of control: Detecting response-outcome independence in analytic but not in naturalistic conditions. *Psychological Sciences, 7,* 289–293.

Matute, H., Arcediano, F., & Miller, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 182–196.

Matzel, L. D., Schachtman, T. R., & Miller, R. R. (1985). Recovery of an overshadowed association achieved by extinction of the overshadowing stimulus. *Learning and Motivation, 16,* 398–412.

Michotte, A. (1963). *The perception of causality.* New York: Basic Books.

Miller, R. R., Hallam, S. C., & Grahame, N. J. (1990). Inflation of comparator stimuli following CS training. *Animal Learning & Behavior, 18,* 434–443.

Miller, R. R., & Matute, H. (1996). Biological significance in forward and backward blocking: Resolution of a discrepancy between animal conditioning and human causal judgment. *Journal of Experimental Psychology: General, 125,* 370–386.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (ed.), *The psychology of learning and motivation,* Vol. 22: 51–92. San Diego, Calif.: Academic Press.

Overmier, J. B., & Seligman, M. E. P. (1967). Effects of inescapable shock upon subsequent escape and avoidance learning. *Journal of Comparative and Physiological Psychology, 63,* 28–33.

Pavlov, I. P. (1927). *Conditioned reflexes.* London: Clarendon Press.

Pearce, J. M., & Hall, G. (1978). Overshadowing the instrumental conditioning of a lever-press response by a more valid predictor of the reinforcer. *Journal of Experimental Psychology: Animal Behavior Processes, 4,* 356–367.

Peterson, C., Maier, S. F., & Seligman, M. E. P. (1993). *Learned helplessness: A theory for the age of personal control.* New York: Oxford University Press.

Piaget, J., & Inhelder, B. (1951). *La genese de l'ideé de hasard chez l'enfant* [The origin of the idea of chance in the child]. Paris: Presses Université France.

Price, P. C., & Yates, J. F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition, 21,* 561–572.

Reed, P. (1992). Effect of a signalled delay between an action and outcome on human judgment of causality. *Quarterly Journal of Experimental Psychology, 44B,* 81–100.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66,* 1–5.

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. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43,* 151–160.

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

Schwartz, B. (1981a). Does helplessness cause depression, or do only depressed people become helpless? Comment on Alloy and Abramson. *Journal of Experimental Psychology: General, 110,* 429–435.

Schwartz, B. (1981b). Helplessness, illusions and depression: Final comment. *Journal of Experimental Psychology: General, 110,* 448–449.

Schweitzer, L., & Green, L. (1982). Reevaluation of things past: A test of the "retrospective hypothesis" using a CER procedure in rats. *Pavlovian Journal of Biological Science, 17,* 62–68.

Sedek, G., & Kofta, M. (1990). When cognitive exertion does not yield cognitive gain: Toward an informational explanation of learned helplessness. *Journal of Personality and Social Psychology, 58,* 729–743.

Seligman, M. E. P. (1975). *Helplessness.* San Francisco: Freeman.

Seligman, M. E. P., & Maier, S. F. (1967). Failure to escape traumatic shock. *Journal of Experimental Psychology, 74,* 1–9.

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. (1993). 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: 229–261. San Diego, Calif.: Academic Press.

Shanks, D. R., Holyoak, K. J., & Medin, D. L. (eds.) (1996). *The psychology of learning and motivation, Vol. 34: Causal judgment.* San Diego, Calif.: Academic press.

Shanks, D. R., & Lopez, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition, 24,* 511–522.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgment of causality by human subjects. *Quarterly Journal of Experimental Psychology, 41B,* 139–159.

Skinner, B. F. (1948). Superstition in the pigeon. *Journal of Experimental Psychology, 38,* 168–172.

Staddon, J. E. R.,.& Simmelhag, V. L. (1971). The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review, 78,* 3–43.

Thornton, J. W., & Jacobs, P. D. (1971). Learned helplessness in human subjects. *Journal of Experimental Psychology, 87,* 367–372.

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. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation, 25,* 127–151.

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

Wasserman, E. A. (1990a). Detecting response-outcome relations: Toward an understanding of the causal texture of the environment. In G. H. Bower (ed.), *The psychology of learning and motivation,* Vol. 26: 27–82. San Diego, Calif.: Academic Press.

Wasserman, E. A. (1990b). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science, 1,* 298–302.

Wasserman, E. A. (1993). Comparative cognition: Toward an understanding of cognition in behavior. *Psychological Science, 4,* 156–161.

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.

Williams, B. A. (1975). The blocking of reinforcement control. *Journal of the Experimental Analysis of Behavior, 24,* 215–225.

Wright, J. C. (1962). Consistency and complexity of response sequences as a function of schedules of noncontingent reward. *Journal of Experimental Psychology, 63,* 601–609.

Young, M. E. (1995). On the origin of personal causal theories. *Psychonomic Bulletin & Review,* 2, 83–104.