

Selectional processes in causality judgment

DAVID R. SHANKS

MRC Applied Psychology Unit, Cambridge, England

Two experiments illustrate the way in which competition between potential causes occurs when subjects are asked to judge the extent to which an action is the cause of an outcome. In the first experiment, it was found that introducing occurrences of the outcome in the absence of the action reduced causality judgments, but this effect was attenuated if these outcomes were signaled by another stimulus. In the second experiment, a delay between the action and the outcome reduced judgments, but this could be abolished by inserting a stimulus between the action and the outcome. The results are discussed in terms of a view of causality judgment that assumes that such judgments are based on associations between the mental representations of the action and the outcome.

Several correspondences between animal instrumental conditioning and human action-outcome causality judgment have recently been established. In a typical causality judgment experiment, a person is required to judge the extent to which performing some action (e.g., pressing a key) is the cause of a particular outcome (e.g., a light flashing). If the action is compared with the operant and the outcome with the reinforcer, then the causality judgment procedure resembles an animal instrumental conditioning preparation. In animal conditioning, the measure of interest is the response rate, whereas in its human analogue, it is the subject's judgment of causality, usually made on a rating scale.

The correspondences that have been established include the shape of the learning curve, or acquisition function (Shanks, 1987), the role of the temporal contiguity of the two events (Shanks & Pearson, 1987), and the sensitivity to *contingency* (e.g., Wasserman, Chatlosh, & Neunaber, 1983). The degree of contingency, or temporal correlation, between the action and the outcome is typically measured (Allan, 1980) by the difference (dP) between the conditional probability of the outcome given the action [$P(O/A)$] and the conditional probability of the outcome in the absence of the action [$P(O/-A)$]. Wasserman et al. (1983) found that causality judgments decreased as $P(O/-A)$ was raised, when $P(O/A)$ was held constant. Figure 1 illustrates two causal sequences in which the degree of contingency, dP , is reduced by increasing $P(O/-A)$ from Panel a to Panel b. One obvious conclusion from this sensitivity to contingency is that causality judgments are not based simply on the degree of contiguity between the action and the outcome, since this is unaffected by raising $P(O/-A)$.

As a result of the similarities between conditioning and causality judgment, Shanks and Dickinson (1987) have argued that the sorts of theoretical accounts that are applied to animal conditioning should also be applied to hu-

man causality judgment. In particular, they argued that an *associative* view of causality judgment should be adopted. With this view, causality judgments would be based on associations between the mental representations of causes and effects. Shanks and Dickinson presented a formal theory of causality judgment, on the basis of this associative view, that bears many similarities to one of the best-known animal learning models, the Rescorla and Wagner (1972) theory of Pavlovian conditioning.

As Figure 1 illustrates, a causal sequence occurs in the context of what we might call a "causal background," that is, a set of background stimuli that are constantly present when a causal sequence occurs and that represent a set of potential alternative causes of the outcome. The target cause, therefore, is not occurring in isolation, but is in competition with this background. It is assumed that there is a strict limit to the amount of associative strength that the outcome can support: An outcome that is perfectly predicted by a target cause is unable to enter into any more associations.

When there is a strong contingency between the target cause and the outcome, the background is unlikely to offer much competition since the target cause is so much more informative about when the outcome will occur. When the contingency is degraded, however, the causal background takes on greater significance. Referring again to Figure 1, occurrences of the outcome in the absence of the target cause in Panel b must be attributed to the background. In terms of the associative theory of Shanks and Dickinson (1987), the background in this case will become associated with the outcome.

The significance of this association between the background and the outcome is that it suggests an explanation for sensitivity to contingency. Consider the subsequent pairings of the action and the outcome after the background has already become associated with the outcome: These outcomes, being well predicted by the background, can only support a small amount of additional associative strength. Thus, the increase in the associative strength of the action will be minimal, relative to the situation in

The author's mailing address is MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge CB2 2EF, England.

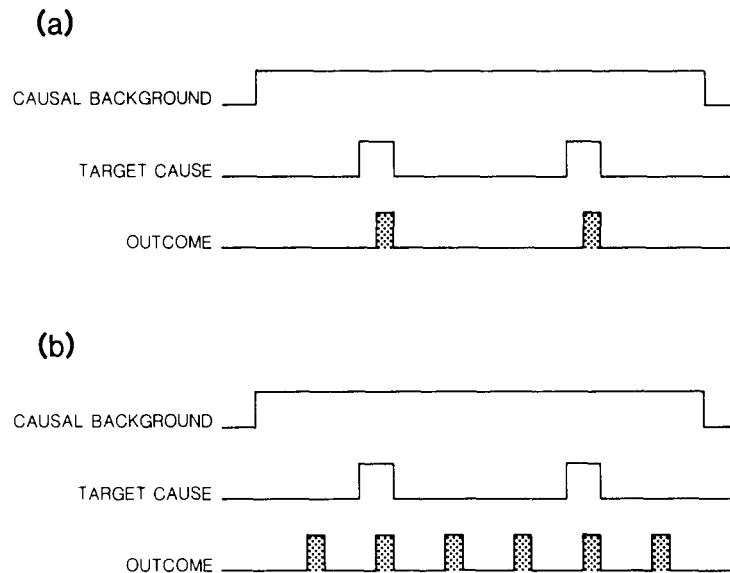


Figure 1. The sequence of events across time for a contingent (a) and a noncontingent (b) relationship between a target cause and an outcome.

which the background has not become associated with the outcome. If these outcomes are less likely to be associated with the action, then judgments of the extent to which the action caused the outcome will be reduced; this is exactly the finding when $P(O/-A)$ is increased.

Such an account of sensitivity to contingency assumes a crucial role for the causal background. It also implies that causal attribution is selective: Selections will be made among potential causes, in terms of how well established they already are as causes—in other words, in terms of how strongly associated with the outcome they are. The two experiments described here use variations on the same procedure to provide evidence for the role of associations in causality judgment. Experiment 1 attempts to show how the action and the background compete as potential causes.

EXPERIMENT 1

The analysis described above proposes that the impact of outcomes occurring in the absence of the target cause comes about because such outcomes are attributed to the background. But suppose this could be prevented: What would happen if such outcomes were attributed to some event other than the background? According to the above analysis, if this happened then the background would not become associated with the outcome and therefore would not be in strong competition with the action when subsequent outcomes occurred in the presence of the action. Thus, any procedure that prevents outcomes that occur in the absence of the action from being attributed to the background should elevate judgments about the action.

In Experiment 1, a straightforward way of doing this was employed. Every outcome occurring in the absence

of the action was preceded by another stimulus, called the "signal," that only occurred on those trials. When outcomes occurring in the absence of the action are attributed to the signal instead of to the causal background, we predict an elevation of judgments of the causal relatedness of the action and outcome. In fact, this signaling procedure is one that has been widely used in animal conditioning experiments to investigate the role of the causal background or context (e.g., Durlach, 1983).

This experiment used the same design as was used in two previous experiments (Shanks, 1986, Experiments 3 and 4), where it was found that a signal did increase judgments, as the associative analysis predicts. However, those experiments used a procedure very different from the free-operant procedure that has been used in most recent causality judgment experiments, including those of Wasserman and his colleagues. In particular, the experiments used a discrete trial procedure in which the subject's responding was highly constrained. The subject's task in those experiments was to judge the extent to which firing a shell at a tank caused the tank to blow up. The present experiment was conducted not simply to replicate a theoretically important result, but to replicate it with a procedure similar to the one used by Wasserman et al. (1983) that has motivated much of the recent theorizing about causality judgment.

Method

Subjects and Apparatus. The subjects, 24 members of the Applied Psychology Unit subject panel, were tested individually in a testing room. The entire experiment was controlled by a Torch Z80/BBC model B microcomputer system. The computer was connected to a Zenith video screen (green on a black background), measuring approximately 24×19 cm. The subjects were required to

respond to the instructions presented on the screen by typing on the computer keyboard placed immediately in front of them.

Procedure. At the beginning of the experiment, the subjects were given the following instructions on the video screen (responses that the subjects were required to make are given in parentheses):

Please read and follow the instructions very carefully. Take as much time as you like. Your task in this experiment is to judge the extent to which you can cause something to happen on the computer screen. There will be a triangle on the screen like this:

The outline of a triangle with 5-cm-long sides was presented on the screen. The next instruction was as follows:

Now press the SPACE BAR and see what happens . . . (Space bar)

When the subject pressed the space bar, the whole triangle lit up for 0.1 sec, and the computer simultaneously generated a tone (middle C) for the same duration. These two events constituted the outcome. After the outcome occurred, the outline of the triangle remained on the screen as before. The instructions continued:

Then press the RETURN key to continue . . . (Return key)

The instructions continued as follows:

...the triangle lights up.

Sometimes the triangle may light up of its own accord, like this:

The triangle then flashed four times, at 2-sec intervals, independently of any response. Finally, the main part of the instructions was presented:

Your task in this experiment is to find out whether pressing the SPACE BAR has any effect on whether or not the triangle lights up. At any time you may choose whether or not to press the SPACE BAR. You can press it as often or as little as you like. However, because of the nature of the task it is to your advantage to press it some of the time and not to press it some of the time. (Return) Sometimes the triangle will flash when you press the SPACE BAR and sometimes it will flash of its own accord. You must judge the extent to which pressing the SPACE BAR is the cause of the triangle lighting up. (Return)

You will be given three different problems, each lasting for two or three minutes. The relationship between pressing the SPACE BAR and whether or not the triangle lights up will be constant within each problem but may well differ from one problem to the next. (Return)

At the end of each problem you will be asked to give an estimate on a rating scale of the extent to which you think that pressing the SPACE BAR caused the triangle to light up during that problem.

Each subject was given three conditions; each condition consisted of 180 1-sec intervals. The schedule operating in each problem was similar to that described by Wasserman et al. (1983): If the action occurred during a particular 1-sec interval, then the outcome occurred at the end of that interval with one probability, $P(O/A)$; if no action occurred, the outcome occurred with a different probability, $P(O/-A)$. Only the first response in each 1-sec interval was recorded. Subsequent responses had no programmed consequences; however, informal observation suggested that the subjects very rarely responded twice in any particular 1-sec interval. In one condition (50/0), there was a positive contingency between the action and the outcome: $P(O/A)$ was 0.5, while $P(O/-A)$ was zero. The first figure in the designation of each condition refers to $P(O/A) \times 100$ and the second to $P(O/-A) \times 100$. In the second condition (50/50), the contingency was reduced to zero by increasing $P(O/-A)$.

The critical condition was the one in which the signal was presented [50/50 (S)]. The contingency was identical to that in the noncontingent 50/50 condition except that all of the outcomes occurring in the absence of the action were preceded by the signal, which was a short tone (A below middle C) readily distinguishable from the tone that occurred as part of the outcome. The signal was not mentioned in the instructions to the subject.

For half of the subjects, the duration of the signal tone was 0.5 sec, and for the other half, it was 0.75 sec. If no response occurred during a particular 1-sec interval, and the outcome was scheduled to occur at the end of that interval, then the outcome was delayed at the end of the interval for an amount of time equal to the duration of the tone; this occurred for all of the conditions. Responses occurring during the tone were not recorded. The difference in the tone's duration had no effect on the results, which are therefore collapsed across this factor.

At the end of each condition, the subjects were asked to rate the extent to which pressing the space bar caused the triangle to light up:

Type in a number below to indicate your judgment of the extent to which pressing the SPACE BAR caused the triangle to light up. Use a scale from 0 to 100. 100 indicates that pressing the SPACE BAR always caused the triangle to light up, and zero indicates that pressing the SPACE BAR had no effect on whether or not the triangle lit up. Press the RETURN key after typing in your judgment. Your judgment?

The three conditions were presented in a random order. Although the randomization procedure did not ensure that each condition would occur equally often in each position, this was, in fact, approximately what happened. Because the outcome on each trial was determined by a software random number procedure, the actual probabilities, $P(O/A)$ and $P(O/-A)$, could differ slightly from the nominal probabilities. The actual probabilities for each subject were recorded and used in the analysis in the form of dP , the difference between $P(O/A)$ and $P(O/-A)$.

Results and Discussion

The subjects responded a mean of 55.1 ($SD = 31.1$) times in Condition 50/0, 48.2 (33.5) times in Condition 50/50, and 62.4 (36.9) times in Condition 50/50 (S). A Friedman test found no overall significant difference between these scores ($\chi^2 = 3.06, p > .1$).

Figure 2 presents the main results of Experiment 1. The actual contingency, as expected, was close to 0.5 in Condition 50/0 ($SD = 0.075$), and close to zero in Conditions 50/50 ($SD = 0.104$) and 50/50 (S) ($SD = 0.118$). An overall Friedman test found a reliable difference between the conditions ($\chi^2 = 36.08, p < .001$). There was, of course, a significant difference between dP in Conditions 50/0 and 50/50 [Wilcoxon $T(24) = 0$], but no difference between Conditions 50/50 and 50/50 (S) [$T(24) = 147, p > .1$].

The principal results are the causality judgments, shown in Figure 2. Standard deviations for Conditions 50/0, 50/50, and 50/0 (S) were 28.5, 25.1, and 23.3, respectively. An overall Friedman test found a reliable difference between the judgments in the three conditions ($\chi^2 = 9.52, p < .01$). Judgments were reduced in Condition 50/50 relative to Condition 50/0 by increasing $P(O/-A)$ [$T(23) = 50, p < .05$]; this corroborates previous findings (e.g., Wasserman et al., 1983) that contingency is a strong determinant of causality judgments.

The critical result, however, is that the mean judgment in Condition 50/50 (S) was greater than that in Condition 50/50 [$T(21) = 48, p < .02$]. Thus, in accordance with the associative account of Shanks and Dickinson (1987), signaling the outcomes that occur in the absence of the action elevates judgments and, hence, at least par-




Figure 2. Mean actual contingency ($dP \times 100$) and mean causality judgment in each condition of Experiment 1. In Condition 50/0, there was a positive contingency, whereas in Condition 50/50, the contingency was reduced to zero by increasing $P(O/A)$. In Condition 50/50 (S), all outcomes occurring in the absence of the action were signaled by a brief tone.

tially reverses the effect of those outcomes on causality judgments.

Judgments were not elevated to the level of those in Condition 50/0. This implies that the signal was not capable of completely preventing the background from acquiring associative strength.

One possible criticism of the results of Experiment 1 is that the signal, a relatively salient stimulus, occurred only in Condition 50/50 (S). It might be, for example, that simply presenting an additional event in this causality judgment context brought about an elevation in judgments via some nonassociative process. This possibility can be ruled out, however, by considering an additional condition included by Shanks (1986, Experiment 4). In that condition [SIG (O/A)], the signal occurred immediately before each outcome that had been preceded by the action, rather than before each outcome that had occurred in the absence of the action. Thus, the signal occurred in this additional control condition as it did in the normal signal condition [SIG (O/-A)]. The difference between the two conditions was that it was only in the normal signal condition that the signal had a chance to reduce the associative strength of the background. Since judgments in the control condition [SIG (O/A)] were reduced relative to the noncontingent condition, it can be argued that the signaling effect is not simply a nonassociative effect of the presence of the signal.

The signaling finding exactly parallels results from animal instrumental conditioning experiments (e.g., Colwill & Rescorla, 1986; Dickinson & Charnock, 1985;

Hammond, 1980), and therefore provides further evidence for the claim that causality judgment in humans and conditioning in animals should be accounted for in terms of similar mechanisms (Shanks & Dickinson, 1987). Dickinson and Charnock (1985), for example, trained thirsty rats to leverpress for a saccharin reinforcer, and then reduced the response-reinforcer contingency by introducing free presentations of the saccharin when the animal was not emitting any responses. This had the effect of reducing the response rate, which indicates that the animals were sensitive to the response-reinforcer contingency; but this response-rate reduction was attenuated when the free saccharin presentations were signaled by a light stimulus.

As it has been argued above that the signaling effect in causality judgment demonstrates the interaction between the context and the action, so signaling is regarded as illustrating the competition between the context and the conditioned stimulus (CS) in Pavlovian conditioning (e.g., Durlach, 1983) and between the context and the response in instrumental conditioning (Dickinson & Charnock, 1985).

EXPERIMENT 2

One obvious feature of the results of Experiment 1 is the fact that the contingency between the action and the outcome was identical in Conditions 50/50 and 50/50 (S). From a normative, statistical point of view the degree to which the action caused the outcome was actually identi-

cal in these two conditions, and yet the subjects gave different judgments for the two conditions. The associative account offers the explanation that despite the identical statistical relationship between the cause and the effect, the strength of the association formed between them differed in the two conditions.

The purpose of Experiment 2 was to give an additional illustration of this. Suppose that a short delay intervenes between the action and the outcome; we would expect this delay to reduce causality judgments relative to a situation in which there was no delay—in fact, this has been shown to be the case with the procedure used here (Shanks & Pearson, 1987). Now suppose that we fill this delay with a salient stimulus, such as a tone. Once again, the insertion of this tone between the action and the outcome has no effect on the normative action–outcome contingency. However, many associative learning experiments have found that responding is elevated by filling the delay between a CS and an unconditioned stimulus (US) or between an instrumental response and a reinforcer, as a result of an increase in the strength of the association between the two stimuli. In causality judgment, we might therefore predict that such a gap-filling stimulus will bring about a substantial increase in judgments, despite leaving the action–outcome contingency unaffected.

Method

Subjects and Apparatus. The subjects were 16 students who participated as part of an undergraduate psychology practical class. The experiment was conducted on three BBC B+ microcomputers, placed adjacent to one another on a table in the classroom. Each computer was connected to a video screen (green on a black background), measuring approximately 24×19 cm.

Procedure. In Experiment 2, the subjects were given three conditions, again in a random order. Procedural details were similar to those of Experiment 1. The conditions were reduced in duration to consist of 120 1-sec intervals. In one condition, 75/25, there was a positive contingency between the action and the outcome. In a second condition, 75/25 (D), a 2-sec pause occurred after each response. Responses occurring during this pause had no programmed consequences and were not recorded. In the final condition, 75/25 (T), this pause was filled by a tone (A below middle C) generated by the computer. This tone was readily distinguishable from the tone that occurred as part of the outcome. The expectation was that the delay in Condition 75/25 (D) would reduce judgments relative to Condition 75/25, but that the tone would counteract this effect and would elevate judgments in Condition 75/25 (T) relative to Condition 75/25 (D).

Results and Discussion

The subjects responded a mean of 46.6 ($SD = 21.0$) times in Condition 75/25, 32.6 (15.9) times in Condition 75/25 (D), and 36.7 (20.8) times in Condition 75/25 (T). An overall Friedman test found a significant difference between the numbers of responses in the three conditions ($\chi^2 = 9.09, p < .05$). There were more responses in Condition 75/25 than in Condition 75/25 (D) [$T(16) = 16, p < .01$], but Conditions 75/25 (D) and 75/25 (T) did not differ [$T(16) = 48.5, p > .1$].

The actual contingency as measured by dP was identical in the three conditions, as Figure 3 illustrates. The standard deviations were 0.065, 0.085, and 0.085 for Conditions 75/25, 75/25 (D), and 75/25 (T), respectively. An overall Friedman test found no significant difference between the values of dP in the three conditions ($\chi^2 = 0.38, p > .1$).

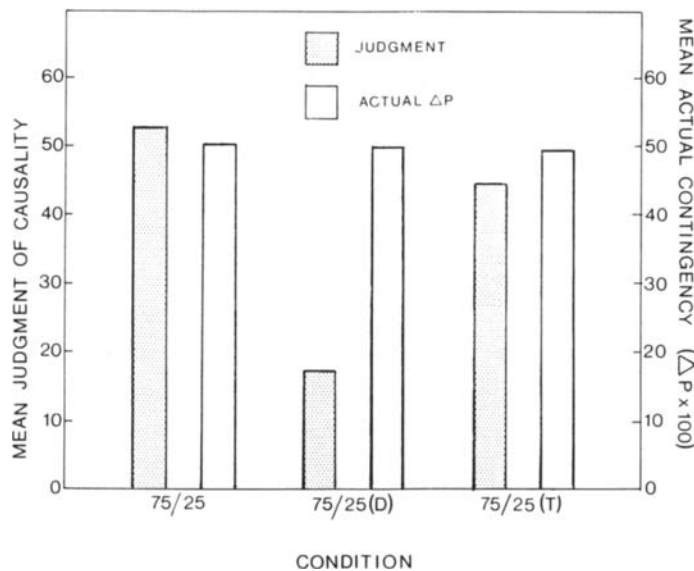


Figure 3. Mean actual contingency ($dP \times 100$) and mean causality judgment in each condition of Experiment 2. In each condition, there was a positive contingency. In Condition 75/25 (D), a 2-sec delay was inserted between the action and the outcome. In Condition 75/25 (T), a tone filled this delay.

Figure 3 shows the mean judgment of causality in each condition. The standard deviations for Conditions 75/25, 75/25 (D), and 75/25 (T) were 29.1, 24.5, and 29.6, respectively. Judgments were substantially reduced in Condition 75/25 (D) relative to Condition 75/25 by the interpolation of the 2-sec delay between the response and the outcome. More important, however, is the fact that judgments were elevated when this delay was filled by the tone in Condition 75/25 (T). A Friedman test found an overall reliable difference between the judgments ($\chi^2 = 10.97, p < .01$). Judgments were reliably lower in Condition 75/25 (D) than in Condition 75/25 [T(16) = 18, $p < .01$], and were greater in Condition 75/25 (T) than in Condition 75/25 (D) [T(15) = 15, $p < .01$].

These results replicate the findings of earlier experiments (Shanks & Pearson, 1987; Wasserman & Neunaber, 1986), revealing that a temporal delay between an action and an outcome reduces causality judgments. In Experiment 2, responses occurring during the delay were not recorded and had no programmed consequences; and so it could be argued that if there were many such responses, then the actual contingency between the action and the outcome differed in Conditions 75/25 and 75/25 (D). Thus, the reduction in judgments in Condition 75/25 (D) might not have been specifically due to the delay, but rather to the fact that the schedule might have been such as to alter the contingency between the action and the outcome. Against this argument, however, is the fact that Shanks, Pearson, and Dickinson (1988) obtained very similar results when they used a schedule in which responses occurring during the delay were recorded. In their experiment, delays of 4, 8, and 16 sec led to progressively lower judgments of causality.

Shanks et al. (1988) also found, as in the present experiment, that a delay between the action and the outcome tended to reduce the subject's response rate, and a comparable result occurs in animal instrumental conditioning experiments (e.g., Williams & Heyneman, 1982). In Experiment 2, there was no difference between the number of responses in Conditions 75/25 (D) and 75/25 (T), and so the effect of the tone on judgments cannot be attributed to any response rate differences. However, the fact that the subjects responded more in Condition 75/25 than in Condition 75/25 (D) means that, in the latter condition, the subjects actually experienced fewer pairings of the action and the outcome, which may well have biased their judgments. But Shanks et al. (1988) argue against this interpretation. They found that a delay still reduced judgments when the delay and control conditions, instead of lasting for a fixed 2-min period, were terminated when 25 responses had been performed, thus equating the number of action-outcome pairings in the two conditions. In a further experiment, delays of 0, 4, 8, and 16 sec produced progressively lower causality judgments in the absence of response rate differences.

The principal significance of Experiment 2 is that once again it presents support for the associative view of causality judgment. The fact that filling the temporal delay between the target cause and the outcome with a stimulus

increases judgments of causality shows again that factors other than simple temporal contiguity have a role in causality judgment. In conditions 75/25 (D) and 75/25 (T), the temporal relationship between the action and the outcome was identical, and yet judgments were greater in Condition 75/25 (T). The difference between judgments in these two conditions cannot have been due to any effects of contiguity. Instead, it is suggested that it stems from a difference in the extent to which associations between the action and the outcome were formed in the two conditions.

The effect of the tone in Experiment 2 parallels findings from animal conditioning experiments. For example, Rescorla (1982) found that when there was a 10-sec delay between the CS and the US in a pigeon autoshaping experiment, the rate of responding was very low. However, when this interval was filled with another stimulus, the response rate was substantially elevated. In an instrumental conditioning procedure, Williams and Heyneman (1982) found that animals' response rates for a food reinforcer were reduced when a 3-sec delay intervened between the criterion response and the reinforcer, but were elevated again when this delay was filled with a light stimulus. There are two main theoretical accounts of how the intervening stimulus affects the formation of associations involving the CS and the US, or instrumental response and reinforcer. Perhaps the most obvious account is that the gap-filling stimulus becomes a conditioned reinforcer, that is, becomes associated on its own with both the action and the outcome. It is possible, on the other hand, that the stimulus has its effect by virtue of "catalyzing," or "potentiating," the association between the target cause and the outcome (e.g., Rescorla, 1982). Whichever of these accounts is in fact correct (and, of course, they may both be true at the same time), it seems clear that the gap-filling stimulus has its effect by influencing the subject's associative representation of the causal situation.

The relationship between the signaling effect of Experiment 1 and the gap-filling effect of Experiment 2 is a theoretically interesting one, for they are clearly not unrelated. We might well anticipate, for instance, that instead of increasing causality judgments, the gap-filling stimulus may act as an alternative cause of the outcome and in fact reduce attribution of the outcome to the action, in much the same way as the signal prevented attribution of the outcome to the background in Experiment 1. It would seem that the precise temporal arrangement between the action, the tone, and the outcome will be of critical importance; in fact, this is borne out in animal conditioning experiments. Williams and Heyneman (1982), using a situation in which pigeons were keypecking for food, found that increasing the delay between the keypeck and the food reinforcer reduced the response rate, but whether an intervening light stimulus increased or decreased responding depended on exactly where in the response-reinforcer interval the stimulus occurred. If we take the pigeons' response rate as an index of their judgment of the causal relation between their responding and

the reinforcer, then we might expect comparable results with humans. This is obviously an important area for further investigation in the analogous human causality judgment situation.

GENERAL DISCUSSION

The two experiments reported here have provided additional confirmation for the claim that human causality judgment and animal conditioning share fundamental underlying processes. Experiment 1 showed that signaling outcomes occurring in the absence of the action reduced the impact of those outcomes on judgments of causality, which implies that competition between the action and the causal background is a basic element of causality judgment situations. In Experiment 2, filling the delay between the action and the outcome elevated judgments. Both of these results can readily be incorporated into an associative account of causality judgment that proposes a fundamental role for associations formed between causes and effects.

At the theoretical level, the results of these two experiments can be related to other recent attempts to describe the mechanism by which causality judgments are made. First, in neither experiment do the results imply that people are sensitive merely to the normative statistical relationship, dP , between the action and the outcome. In Condition 50/50 (S) of Experiment 1, the introduction of the signal had no effect on the degree of contingency between the action and the outcome, relative to Condition 50/50. Similarly, in Experiment 2, the introduction of the tone in Condition 75/25 (T) did not affect the contingency relative to Condition 75/25 (D). Thus, it appears that judgments cannot be accounted for on a simple covariational, or contingency, basis.

Wasserman and Neunaber (1986) proposed a theory of causality judgment that attempts to account for sensitivity to contingency and for the effects of temporal contiguity. Briefly, this theory claims that judgments are based on the *relative contiguity* of the action and the outcome. Relative contiguity is defined as the difference between the mean delay, d_o , between the action and the outcome and the mean delay, d_n , between competing causes and the outcome, such that relative contiguity is equal to $d_n - d_o$. In this theory, d_n represents the mean delay between alternative actions that the person is performing and the outcome, rather than referring to the causal background of events, which is constant in the experimental situation.

Wasserman and Neunaber's (1986) theory can account for contiguity effects such as those found in Experiment 2, because a delay between the action and the outcome increases d_o and thus reduces the relative contiguity. Similarly, contingency effects such as those shown in Experiment 1 can be accounted for, since increasing $P(O/-A)$ reduces d_n and hence reduces judgments.

What this theory cannot account for, however, are the subtler effects of the tone stimulus in each of the experi-

ments reported here. In neither of these experiments did the introduction of the tone affect the degree of relative contiguity, and so should have had no effect on judgments. Wasserman and Neunaber's theory, in other words, has no mechanism for selectional processes amongst potential causes.

The alternative view of causality judgment holds that such judgments are based on associations between the mental representations of the action and the outcome. This view clearly owes a great deal to conceptions of animal conditioning, but the correspondences between causality judgment and conditioning are striking (see Shanks, 1986). As such, the associative view borrows many ideas from the tradition of associationism in animal conditioning. While the formal details of an associative model of causality judgment will not be spelled out here, it is possible for such a model to account quantitatively, as well as qualitatively, for a wealth of causality judgment data (Shanks & Dickinson, 1987; Shanks & Pearson, 1987).

This view attributes a fundamental role to competition between the potential causes present. It accounts for contingency effects by virtue of the fact that increasing $P(O/-A)$ will increase the associative strength of the causal background and hence will allow it to overshadow the action. Contiguity effects are viewed as resulting from a basic property of associations, namely, that increasing the delay between the action and the outcome reduces the trace of the action in memory and hence leads to the formation of a weaker association. And in addition, associative theories can account for the sorts of selectional processes that have been demonstrated here (see Shanks & Dickinson, 1987). Clearly, such experiments as have been described here offer one informative way of differentiating between alternative views of causality judgment.

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.
- COLWILL, R. M., & RESCORLA, R. A. (1986). Associative structures in instrumental learning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 20). New York: Academic Press.
- DICKINSON, A., & CHARNOCK, D. J. (1985). Contingency effects with maintained instrumental reinforcement. *Quarterly Journal of Experimental Psychology*, 37B, 397-416.
- DURLACH, P. J. (1983). Effect of signaling intertrial unconditioned stimuli in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 374-389.
- HAMMOND, L. J. (1980). The effect of contingency upon the appetitive conditioning of free operant behavior. *Journal of the Experimental Analysis of Behavior*, 34, 297-304.
- 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., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research*. New York: Appleton-Crofts.
- SHANKS, D. R. (1986). Selective attribution and the judgment of causality. *Learning & Motivation*, 17, 311-334.

- SHANKS, D. R. (1987). Acquisition functions in contingency judgment. *Learning & Motivation*, **18**, 147-166.
- SHANKS, D. R., & DICKINSON, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21). New York: Academic Press.
- SHANKS, D. R., & PEARSON, S. M. (1987). A production system model of causality judgment. *Proceedings of the Ninth Annual Cognitive Science Society Conference*. Hillsdale, N.J: Erlbaum.
- SHANKS, D. R., PEARSON, S. M., & DICKINSON, A. (1988). *Temporal contiguity and the judgment of causality*. Manuscript submitted for publication.
- WASSERMAN, E. A., CHATLOSH, D. L., & NEUNABER, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning & Motivation*, **14**, 406-432.
- WASSERMAN, E. A., & 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., & HEYNEMAN, N. (1982). Multiple determinants of "blocking" effects on operant behavior. *Animal Learning & Behavior*, **10**, 72-76.

(Manuscript received May 1, 1987;
revision accepted for publication February 29, 1988.)