A Production System Model of Causality Judgment

David R.Shanks
MRC Applied Psychology Unit, Cambridge

and

Susan M.Pearson
University of Cambridge

1. Introduction

Building an internal representation of the causal structure of the world is a critically important cognitive ability. Both in order to understand the relationships between events in the external world, and in order to control events in the world, it is necessary to be able to detect causality and to be able to perform causal reasoning. In psychology this ability has traditionally been seen as of considerable importance in understanding such diverse areas as perception (e.g. Michotte, 1963), decision-making (e.g. Nisbett and Ross, 1980) and psychopathology (e.g. Seligman, 1975). Recently too, researchers studying human-computer interaction have come to recognise the importance of causal knowledge (e.g. Lewis, 1986).

In fact such has been the concentration of effort in understanding causality judgment, that it is now possible to construct a fairly complete model of causality judgment: not only can the processes involved be described, but it is also possible to specify in reasonable detail what the form of representation is likely to be. In order to do this, we have chosen to use a production system architecture to capture the principal features of the human causality judgment mechanism.

Because the empirical studies are so important in determining what features a theory of causality judgment must have, we begin by describing three experiments which illustrate the importance of certain features of causal situations in determining the formation of causal knowledge. Then a recent model of causality judgment (Shanks and Dickinson, 1987) is described, followed by some extensions of this theory which cover the way causal knowledge might be represented. Finally, an implementation of this theory in terms of a production system is presented.

2. Contiguity and Contingency

It has been recognised at least since the time of Hume that judgments of causality depend on close temporal and spatial contiguity between the target cause and the effect. Although there have been some tests of this in causality judgment (e.g. Wasserman and Neunaber, 1986), no parametric data have ever been presented. The first experiment attempts to see whether subjects' judgments of the extent to which an action causes an outcome are reduced when delays of 4, 8 or 16 sec are inserted between the action and the outcome, relative to a condition in which there is no delay.

2.1. Experiment 1

In this experiment subjects were each given eight causality judgment problems in which they were required to judge the extent to which an action (pressing the space bar on a computer keyboard) caused an outcome (the flashing of a triangle for 0.1 sec) to occur on the video screen. Each condition lasted for 2 min and was divided into 1 sec time intervals. In the four experimental conditions, if the action occurred during a particular 1 sec interval, then the outcome followed the action with probability 0.75 and never occurred independently of the action. These four experimental conditions differed in the degree of temporal contiguity between the action and the outcome. Each time an action was performed and an outcome was programmed to follow it, a delay of either 0, 4, 8 or 16 sec was inserted before the outcome. During the delay the schedule proceded normally, so that further actions could set up further outcomes. Thus for these four conditions, the probability of the outcome given the action was constant; what differed was the temporal interval before this outcome occurred.

The remaining four conditions were all control conditions. Each experimental condition was immediately followed by a control condition in which the pattern of outcomes that occurred in the experimental condition was played back to the subjects independently of their actions. Thus the temporal distribution and frequency of outcomes is matched in the experimental and control conditions. This is important because if the number of actions (and hence the number of outcomes) differed across the experimental conditions, then this would be confounded with any differences in the subjects' judgments of the extent to which the action caused the outcome. The conditions were presented in pairs consisting of an experimental condition followed by its control, but the order of pairs of conditions with respect to the action-outcome delay was random.

If people are sensitive to contiguity in their causality judgments, we would expect judgments in the experimental conditions to be reduced as the delay is increased; the control conditions provide the baseline against which the judgment in each experimental condition can be assessed.

The subjects were 16 students who were tested individually. The instructions given to the subjects at the beginning of the experiment were similar to those used in the experiments decrsibed by Wasserman, Chatlosh and Neunaber (1983). At the end of each condition the subjects were asked to make a rating of the extent to which pressing the space bar caused the triangle to flash, using a scale from 0 to 100. 100 indicated that pressing the space bar always caused the triangle to light up, and zero indicated that pressing the space bar had no effect on whether or not the triangle lit up. After typing in a number, the next problem was presented.

2.2. Results

The principal results, the judgments of causality, are shown in Table 1. As the table shows, judgments were substantially reduced by increasing the delay between the action and the outcome from 0 to 4, 8 and 16 sec. An analysis of variance found a reliable difference between the experimental and control conditions, $F(1,15)=19.63$, a reliable overall effect of the delay, $F(3,45)=7.75$, and a significant interaction, $F(3,45)=7.80$. Individual tests found that there were significant differences between the experimental and control conditions at the 0 and 4 sec delays, t's$(15)=6.37$ and 1.90, but no differences at the 8 and 16 sec delays, t's$(15)<1.06$.

These results confirm that the occurrence of a temporal delay between an action and an outcome can decrease judgments of the extent to which the action caused the outcome. Clearly, accounting for this sensitivity to contiguity is an essential requirement for any theory of causality judgment.

3. Sensitivity to Contingency

The preceding experiment shows that the degree of contiguity between the cause and the effect does affect causality judgments. Consider the difference between the experimental and control conditions with no delay between the action and the outcome. Exactly the same pattern of outcomes occurred in these two conditions, but in the control condition this pattern was non-contingent on the pattern of actions. The notion of the contingency between the action and the outcome is a second important determinant of the formation of causal knowledge.

It is not hard to see that contiguity alone is insufficient to account for causality judgments. Imagine the two conditions illustrated in Figure 1 in which the probability of the outcome given an action, $P(O/A)$, is the same, say 0.75. If the probability of the outcome given no action, $P(O/{-A})$, is zero [panel (a)], as was the case in the experimental conditions of the preceding experiment, there is little difficulty in detecting the causal relationship provided that the action-outcome delay is

Table 1
Results of Experiment 1

Mean judgements of causality in the Control and Experimental Conditions at each of the four action-outcome delays.

	Delay (secs)			
	0	**4**	**8**	**16**
Experimental	**69.8**	**33.8**	**34.7**	**22.7**
Control	**22.4**	**19.7**	**26.8**	**19.3**

Table 3
Results of Experiment 3

Mean judgements and mean actual contingency (dP x 100) for each of the conditions.

Condition	**50/0**	**50/50**	**50/50(S)**
Mean Judgement	**49.0**	**20.9**	**32.4**
Mean dP x 100	**48.3**	**0.6**	**-0.1**

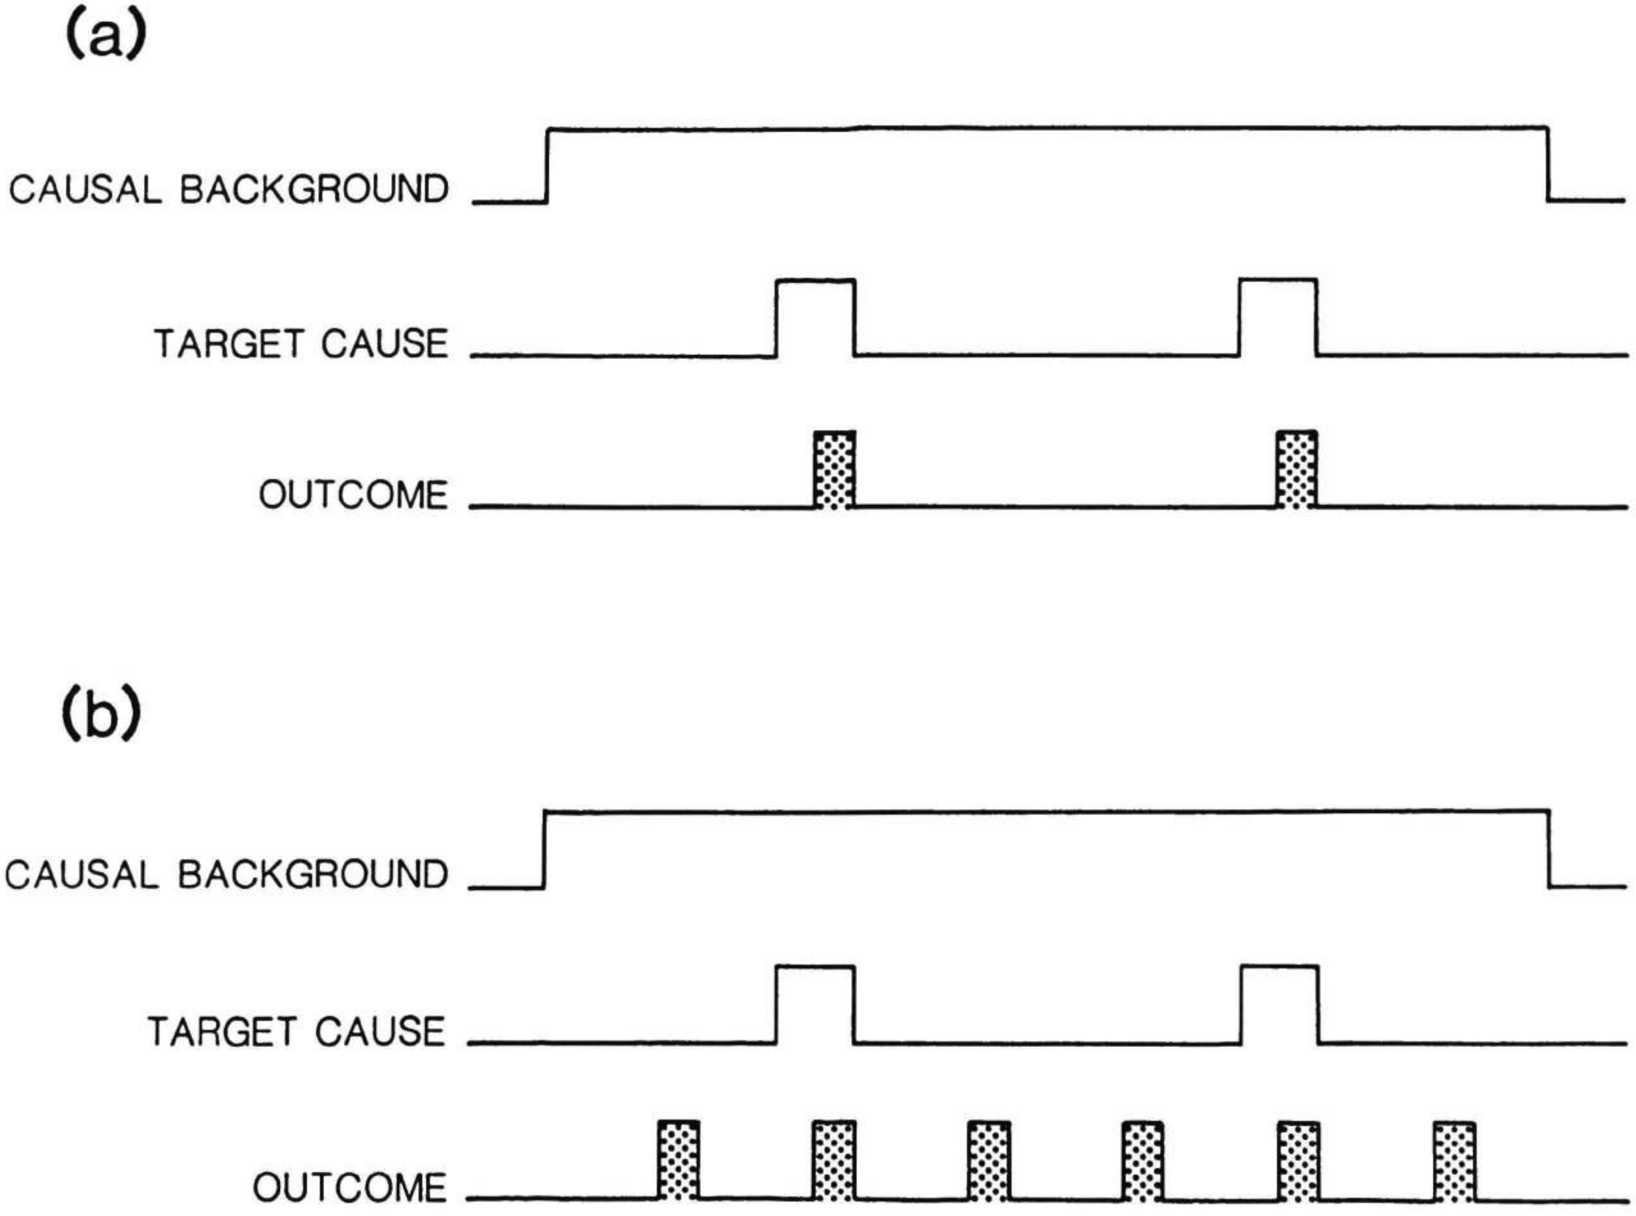
(a)
CAUSAL BACKGROUND
TARGET CAUSE
OUTCOME
(b)
CAUSAL BACKGROUND
TARGET CAUSE
OUTCOME

Figure 1

small. But suppose that in the second condition P(O/-A) is the same as P(O/A), as shown in panel (b). These two conditions have the same degree of temporal contiguity between the action and the outcome, but in the second condition the outcome is just as likely to occur independently of any action as in its presence. That is, in this second condition, the action and the outcome are noncontingent: the action does not cause the outcome, and hence subjects should judge that there is no causal relation. This is in fact exactly what happens (e.g. Wasserman et al., 1983). Clearly, sensitivity to the contingency between the action and the outcome is independent of the degree of contiguity between them. Traditionally, contingency has been defined statistically by dP, which is the difference between P(O/A) and P(O/-A).

Because this sensitivity to contingency has been a crucial factor in the formulation of theories of causality judgment, some further data will be described on this issue. It is not just important to look at the degree of sensitivity alone; for a variety of reasons which will become apparent it is also profitable to look at the learning curves for different contingencies. The following experiment therefore examines the acquisition functions for causality judgments under different contingencies, that is, the way in which judgments change as more and more information about the causal relationship is provided.

3.1. Experiment 2

In this experiment subjects were each given 4 min experience under each of six different contingencies, and in each condition they were required to make regular judgments of causality during the 4 min period. The contingencies used were as follows, where the first figure refers to P(O/A) x 1000 and the second to P(O/-A) x 1000: in three of the conditions P(O/A) was held constant as P(O/-A) was raised (conditions 875/125, 875/500 and 875/875); the last of these conditions was in fact a noncontingent one, and so another noncontingent condition was included for which the frequency of the outcome was lower, 125/125. Finally, judgments for the 875/875 condition could be compared with judgments for two other conditions for which P(O/-A) was held constant while P(O/A) was raised (conditions 125/875 and 500/875). In these latter conditions dP is in fact negative: the action actually prevents the outcome from occurring or reduces its likelihood.

There was no contiguity manipulation in this experiment. If the subject responded during a particular 1 sec interval, then the outcome occurred at the end of that interval with probability P(O/A). If there was no response during the interval, the outcome occurred at the end of the interval according to P(O/-A).

The subjects, who were 16 another members of the APU panel, were asked to make causality judgments after 10, 20, 30, 60, 90, 120, 150, 180, 210 and 240 sec in each condition. Essentially, the instructions were the same as in the preceding experiment except that here a rating scale going from -100 to + 100 was used. Negative judgments were indications that the action to some extent prevented the outcome.

3.2. Results

Since the events on each trial were determined by a software random number procedure, it is important to check that the actual contingency experienced by each subject was close to the nominal contingency. The deviations of dP from the nominal contingency were very minor, and more importantly, there was no change in dP across trials, F< 1, nor was there an interaction between trials and conditions, F(45,661)= 1.07. The actual contingency was, of course, different across the conditions, F(5,75)= 586.88. The actual contingencies (dP x 100) calculated across the whole 4 min period for each condition were 76.6, 37.5, 1.4, -1.0, -34.8 and -74.5 for conditions 875/125, 875/500, 875/875, 125/125, 500/875 and 125/875, respectively.

An indication of the extent to which P(O/A) and P(O/-A) were both sampled is given by the rate of responding in each condition. Out of a possible 240 responses, the mean number of times the subjects pressed the space-bar was 106.5, indicating that the two probabilities were almost

Table 2

Results of Experiment 2

Mean actual judgements of causality across trials under each contingency. The two numbers describing each condition refer to P(O/A) x 1000 and P(O/-A) x 1000 respectively. The predicted figures are from the simulation described in the final section.

	Trial									
	10	**20**	**30**	**60**	**90**	**120**	**150**	**180**	**210**	**240**
Condition										
875/125 actual	**40.9**	**58.1**	**62.5**	**63.7**	**64.7**	**70.3**	**72.4**	**76.8**	**71.6**	**73.3**
875/125 predicted	34.3	52.5	60.2	68.1	68.2	68.8	69.5	67.4	66.2	66.9
875/500 actual	**26.4**	**36.9**	**33.4**	**41.9**	**39.7**	**41.3**	**34.1**	**43.8**	**48.1**	**44.6**
875/500 predicted	25.7	39.0	44.5	45.2	45.6	45.5	46.5	47.0	45.9	46.3
875/875 actual	**6.8**	**25.6**	**13.9**	**14.8**	**6.0**	**7.3**	**1.1**	**3.3**	**-0.6**	**0.6**
875/875 predicted	15.3	20.3	17.4	8.5	2.2	0.8	0.4	0.2	0.9	0.5
125/125 actual	**-6.1**	**-27.0**	**-28.9**	**-34.6**	**-39.4**	**-41.0**	**-38.4**	**-31.1**	**-37.9**	**-36.6**
125/125 predicted	1.7	0.8	0.5	0.5	-0.1	0.5	0.2	-0.3	0.0	0.1
500/875 actual	**-12.4**	**-19.0**	**-17.1**	**-13.4**	**-17.1**	**-23.0**	**-18.4**	**-24.4**	**-22.1**	**-26.4**
500/875 predicted	-2.7	-12.4	-22.2	-38.0	-42.7	-45.4	-45.9	-45.7	-44.4	-46.5
125/875 actual	**-30.6**	**-47.2**	**-55.3**	**-67.1**	**-64.1**	**-69.8**	**-77.2**	**-72.9**	**-60.0**	**-64.6**
125/875 predicted	-17.2	-33.4	-43.8	-61.3	-65.9	-67.3	-68.3	-67.9	-68.0	-68.6

equally sampled. There was no difference in the number of responses in each condition, $F < 1$, nor was there any trials by conditions interaction, $F < 1$.

Table 2 shows the mean judgments of contingency for each condition across trials. Just taking the terminal judgments in each condition, we can see that judgments were reduced from condition 875/125 to condition 875/500 to condition 875/875 as P(O/-A) was increased while P(O/A) was held constant. Similarly, judgments were increased (became less negative) from condition 125/875 to condition 500/875 to condition 875/875 as P(O/A) was increased while P(O/-A) was held constant. In addition, judgments were biased in the noncontingent conditions by the overall frequency of the outcome: judgments were substantially greater in condition 875/875 than in condition 125/125 where judgments were in fact strongly negative. This confirms previous claims (e.g. Alloy and Abramson, 1979) that in noncontingent situations judgments depend on the rate of occurrence of the outcome.

Statistically, there was a significant difference between the conditions, $F(5,75) = 39.04$, and a significant trials x conditions interaction, $F(45,675) = 2.54$. There was no overall effect of trials, $F < 1$. The main feature of the results, however, is the nature of the changes across trials under each contingency. When there is a positive contingency, judgments increase across trials [comparing the first and last judgments in condition 875/125, $t(15) = 3.52$], and when there is a negative contingency, judgments decrease [comparing the first and last judgments in condition 125/875, $t(15) = 3.71$]. These changes, furthermore, appear to be negatively accelerated as the judgment approaches asymptote. For the intermediate contingencies (875/500 and 500/875) these changes are less dramatic.

The experiment demonstrates two important features of causality judgments. The first is that people are highly sensitive to the actual degree of contingency as measured by the normative metric dP This applies both to positive and to negative contingencies: people make equally reliable judgments when the action prevents the outcome as when it causes it. The second feature of the results is that these accurate judgments are derived by an incremental (or decremental) learning process. The functions represented in Table 2 are in fact simply learning curves; the subjects' judgments increase across trials under a positive contingency, towards the asymptote of the actual contingency, and decrease across trials under a negative contingency.

4. Selection Amongst Potential Causes

The two experiments described so far have illustrated two factors that have strong effects on the formation of causal knowledge, namely contingency and contiguity. The next experiment looks at a further factor, which can again have a potent effect on judgments. This is the status of other potential causes which are present at the same time as the target cause (the action in this case).

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 which are constantly present in that situation and which 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. 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. But when the contingency is degraded, 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 can be attributed to either the background or to the action. But since the background is already associated with the outcome, surely subsequent occurrences of the outcome are now more likely to be attributed to the background than to the action, relative to the situation in which the background has not become associated with the outcome. If these outcomes are less likely to be attributed to 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 amongst potential causes in terms of how well established they already are as causes. The next experiment attempts to provide support for this claim.

4.1. Experiment 3

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 other event rather than to the background ? By the above analysis, if this happened then the background would not become associated with the outcome and hence would not be in strong competition with the action when subsequent outcomes occurred in the presence of the action. Thus any procedure which prevents outcomes occurring in the absence of the action from being attributed to the background should elevate judgments about the action. In this experiment a straightforward way of doing this is employed: every outcome occurring in the absence of the action is preceded by another stimulus, called the 'signal', which only occurs on those trials.

The experiment involves the same basic procedure as in the previous experiments. The subjects (24 members of the APU subject panel) were each given three conditions, presented in a random order. 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. In the second condition (50/50) the contingency was reduced to zero by increasing P(O/-A) to 0.5. 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, but in this case all of the outcomes occurring in the absence of the action were preceded by the signal, which was a short tone. For half of the subjects the duration of the tone was 0.5 sec, while 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. The difference in the tone's duration had no effect on the results, which are therefore collapsed across this factor.

4.2. Results

The subjects responded a mean of 55.1 times in condition 50/0, 48.2 times in condition 50/50, and 62.4 times in condition 50/50 (S). The difference between the numbers of responses in the latter two conditions was in fact reliable, Wilcoxon $W = 62.5$, $p < 0.05$. Table 3 presents the main results of the experiment.

The actual contingency, as expected, was close to 0.5 in condition 50/0 and close to zero in conditions 50/50 and 50/50 (S). Judgments were reduced in condition 50/50 relative to condition 50/0 by increasing P(O/-A), $W = 50$, $p < 0.05$; this corroborates the findings of the previous two experiments that contingency is a strong determinant of causality judgments. The critical result, however, is the greater mean judgment in condition 50/50 (S) than in condition 50/50, $W = 48$, $p < 0.02$. Thus signaling those outcomes that occur in the absence of the action has the effect of elevating judgments, and hence of at least partially reversing the effect of those outcomes on causality judgments.

5. A Production System Model of Causality Judgment

If a temporal delay is interpolated between an action and an outcome, judgments of the extent to which the action is the cause of the outcome are reliably reduced, even though the delay might have no effect on the actual probability of the outcome. In Experiment 1 a short delay between the action and the outcome substantially reduced judgments. In the second experiment it was found that causality judgments follow growth functions whereby under a positive contingency judgments are incremented trial-by-trial towards an asymptote, whilst under a negative contingency they are decremented towards the asymptote. Finally, in the last experiment the importance of the causal background was illustrated in a situation where causal selection on the basis of the prior association between the background and the outcome could be prevented.

The account of causality judgment that will now be described proposes a crucial role for such associations. According to associative theories of learning, the occurrence of an outcome increments the associative strength of stimuli present at that time. A crucial aspect of such theories is the role they assign to the causal background, that set of stimuli which is always present in that context. In the model described by Shanks and Dickinson (1987), an outcome occurring in the presence of the target cause increments the associative strength of the compound of the target cause and the background according to the equation:

$$dV_{AB}=\alpha_{AB}\beta(\lambda-V_{AB}) \qquad (1)$$

where dV_{AB} is the increment in the associative strength of the compound of the action and the background, α_{AB} is a learning rate parameter for the compound, β is a learning rate parameter for the outcome, λ is the asymptote of associative strength and V_{AB} is the associative strength the compound already has. The asymptote λ is usually set to 1.0 on trials on which the outcome occurs and to zero when the outcome does not occur.

If the outcome occurs in the absence of the action, then the associative strength of the background, V_B is incremented according to the equation:

$$dV_B=\alpha_B\beta(\lambda-V_B) \qquad (2)$$

where dV_B is the increment and α_B is a rate parameter for the causal background. At the end of a series of occurrences of the action and the outcome, there will be two associative strengths, V_{AB} and V_B. The causality judgment, J_A, is then based on an inference step in which the difference between these two associative strengths is determined:

$$J_A=V_{AB}-V_B \qquad (3)$$

The model can account for sensitivity to contingency since increasing P(O/-A) will increase V_B and hence reduce judgments. In fact the model can readily reproduce the pattern of acquisition functions seen in the second experiment. Table 2 gives the results of a simulation of the model run under the same contingencies as the actual experiment. Each figure represents the mean from 1000 simulated subjects with the parameters as follows: $\alpha_B=0.1$, $\alpha_{AB}=0.2$, β for the outcome= 0.3 and β for the nonoccurrence of the outcome= 0.8, and assuming that V_B and V_{AB} start at zero. All of the main features of the actual results are reproduced, with the exception that the negative judgments in the 125/125 condition do not emerge (see Shanks and Dickinson, 1987, for a discussion of this discrepancy).

A production system model of causality judgment which incorporates the linear operator equations of Shanks and Dickinson (1987) has recently been implemented in the computer language OPS-5 (Brownston, Farrell, Kant and Martin, 1985). It consists of two principal elements, namely the production rules and a working memory. The production rules consist of simple if...then... statements which operate when their conditions are satisfied by the contents of working memory. There will be two production rules relevant to causality judgment. The first will specify that if the causal background is represented in working memory, then the representation of the outcome should be placed in working memory, and equation (2) should be selected. This production rule has a measure of belief associated with it, which is the current value of V_B. The second rule states that if the action and the background are in working memory, then the representation of the outcome should be put in working memory and equation (1) should be selected. This rule also has a measure of belief associated with it, this time determined by V_{AB}.

Notice that the production rules do not incorporate the concept of causation. When the subject comes to make a causality judgment, the model proposes that this comes about by an inferential process which involves propostional knowledge from other sources as well as what is incorporated in the production rules. In causality judgment experiments, of course, it is assumed that any extra-experimental knowledge the subjects might bring to the situation will be minimal. In more realistic settings, though, it seems highly likely that prior knowledge about causal relationships might have a strong effect on causality judgments in particular circumstances. In terms of the model, such knowledge is brought in to the inference represented by equation (3).

How can such a model account for contiguity effects ? So far, the model does not specify what constitutes a co-occurrence of the action and the outcome [in which case equation (1) applies] and what constitutes an occurrence of the outcome in the absence of the action [in which case equation (2) applies]. But the crucial determinant is likely to be simply whether or not the action is still represented in working memory, above a certain threshold, at the time of the outcome. If we assume that the action is fully represented immediately after its occurrence but then its representation decays, then this can be captured in the model by a reduction in α_{AB} , the salience of the action-background compound, as time elapses from the occurrence of the action. Although there is no obvious reason to support one decay function over another, an exponential decay curve in the model gives predictions for contiguity effects that approximate those found experimentally. If α is being reduced, then the increment in the associative strength dV_{AB} will likewise be reduced as the action-outcome interval increases. Note that this does not affect the associative strength of the background, V_B: in accordance with the results of Shanks and Dickinson (1987); a reduction in the salience of the action brought about by an interval between the action and the outcome will reduce judgments of causality by reducing V_{AB} , but will leave V_B unaffected.

Obviously, at some point the action will no longer be represented in working memory and therefore an occurrence of the outcome must then increment V_B and not V_{AB}. A threshold below which the level of representation of the action in working memory causes equation (3) to be selected instead of equation (2) can readily be incorporated into the model.

The model can also explain the signaling effect of Experiment 3 because the introduction of the signal means that outcomes occurring in the absence of the action do not increment V_B; they increment the associative strength of the signal-background compound, which does not figure in equation (3).

In summary, the model in its present form is capable of accounting for the main features of causality judgment: first, it can explain sensitivity to contingency; secondly, it can cover contiguity effects; and thirdly, it can account for the selectional effects seen in Experiment 3. Because it is specified in the precise terms of a production system, it makes concrete predictions about a variety of causality judgment situations and can therefore readily be tested.

References

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

Brownston,L., Farrell,R., Kant,E. & Martin,N. Programming Expert Systems in OPS-5: An

Introduction to Rule-Based Programming. Reading, Mass.: Addison-Wesley, 1985.

Lewis,C. Understanding what's happening in system interactions. In D.A.Norman & S.W.Draper (Eds.), User Centered System Design: New Perspectives on Human-Computer Interaction. Hillsdale, N.J.: Erlbaum, 1986.

Michotte,A. The Perception of Causality. London: Methuen, 1963.

Nisbett,R.E. & Ross,L. Human Inference: Strategies and Shortcomings of Social Judgment. Englewood Cliffs: Prentice-Hall, 1980.

Seligman, M.E.P Helplessness. San Francisco: W.H.Freeman, 1975.

Shanks,D.R. & Dickinson,A. Associative accounts of causality judgment. In G.H.Bower (Ed.), The Psychology of Learning and Motivation (Vol.21). New York: Academic Press, 1987.

Wasserman,E.A., Chatlosh,D.L. & Neunaber,D.J. Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. Learning and Motivation, 1983, 14, 406-432.

Wasserman,E.A. & Neunaber,D.J. College students' responding to and rating of contingency relations: The role of temporal contiguity. Journal of the Experimental Analysis of Behavior, 1986, 46, 15-35.