

Instrumental judgment and performance under variations in action–outcome contingency and contiguity

DAVID R. SHANKS

MRC Applied Psychology Unit, Cambridge, England

and

ANTHONY DICKINSON

University of Cambridge, Cambridge, England

The intentional theory of instrumental performance proposes that performance of an action is determined in part by a belief about its causal effectiveness in producing a desired outcome. At variance with this notion, previous implicit learning experiments appear to have yielded dissociations between subjects' performance and beliefs. In two experiments, subjects were given an opportunity to perform an action—pressing a key on a computer keyboard—which was associated with an outcome on the computer screen according to a free-operant contingency. The subjects in one group were asked to judge the effectiveness of the action in causing the outcome, while those in a second group were asked to maximize their points score under a payoff schedule. In the first study, the effect of varying the contingency between the action and outcome was examined by keeping the probability of an outcome contiguous with an action constant and varying the probability of an outcome in the absence of an action. Performance and judgments showed a comparable sensitivity to variations of the instrumental contingency. In the second study, the delay between the action and the resultant outcome was varied. Increasing the action–outcome delay from 0 sec up to 4 sec produced a systematic decline in both causal judgments and performance relative to noncontingent, control conditions. These results are in accord with the intentional theory of performance, but they present difficulties for the notion of implicit learning.

Cognitive or *intentional* theories of instrumental action have been commonplace in psychology and philosophy for many years. These theories assume that performance is mediated by a belief about the effectiveness of the action in causing a desired outcome. For instance, Tolman (1959) argued that experience of a positive instrumental relationship results in the acquisition of an expectancy or belief “that an instance of this *sort* of stimulus situation, if reacted to by an instance of this *sort* of response, will lead to an instance of that *sort* of further stimulus situation” (p. 113; the italics are Tolman's). Furthermore, he assumed that such an expectancy interacts with the current “valence” of the outcome (“that *sort* of further stimulus situation”) to determine performance. This is a view that has been endorsed, in one form or another, by subsequent statements of the intentional account (e.g., Bolles, 1972; Dickinson, 1989; Heyes & Dickinson, 1990; Irwin, 1971).

Although rarely stated explicitly in presentations of intentional theory, it is reasonable on this account to assume that the level of performance, such as the rate of

responding on a free-operant schedule, is determined, at least in part, by the agent's belief about the strength of the causal relationship between the action and the outcome.¹ An obvious implication of this account is that there should be a concordance, across variations in the critical parameters of an instrumental association, between (1) the subject's belief about the causal effectiveness of an action, and (2) their performance of that action.

Several authors, however, have recently defended a quite different theory of action, proposing that human instrumental performance on complex tasks is typically under the control of knowledge that is acquired *implicitly*—that is to say, in the absence of awareness and not available for verbal report. *Explicit* knowledge, which is available to awareness and can be verbalized, is functionally separate from implicit knowledge and is presumed to play no role in performance on complex tasks.

The reason for making this radical claim is that in a variety of learning situations, performance and awareness appear to be dissociable. Thus Berry and Broadbent (1984), Hartman, Knopman, and Nissen (1989), Hayes and Broadbent (1988), Lewicki (1986), Lewicki, Czyzewska, and Hoffman (1987), Nissen and Bullemer (1987), Reber (1989), Sanderson (1989), Stadler (1989), and Willingham, Nissen, and Bullemer (1989), among others, have all reported situations in which subjects' per-

Requests for reprints should be sent to David R. Shanks, MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge CB2 2EF, England.

formance on a task has improved in the absence of any corresponding change in reportable knowledge.

Of course, performance and verbal reports do not always dissociate; in Sanderson's (1989) experiments, they typically were associated rather than dissociated. Although associations do not necessarily undermine the implicit-explicit distinction, some researchers have questioned whether any of the evidence satisfactorily demonstrates dissociations between performance and verbal report (e.g., Brody, 1989; Green & Shanks, 1990; Perruchet, Gallego, & Savy, 1990; Perruchet & Pacteau, 1990, 1991; Sanderson, 1989, 1990). Thus, at present, the case for the implicit-explicit dissociation appears to be undecided.

In implicit learning experiments, researchers generally adopt the strategy of trying to demonstrate improvement on some measure of performance in a subset of subjects for whom there is no concomitant improvement (or even a reduction) in reportable knowledge. In the present experiments, we adopted a somewhat different strategy to examine the concept of implicit learning. Instead of simply classifying subjects as aware or unaware, we used a procedure in which subjects reported their beliefs about an action-outcome relationship on an ordinate scale. We then attempted to see whether certain factors known to affect instrumental performance have a similar effect on reported beliefs. Although it is generally accepted that successful conditioning is usually accompanied by an awareness of the instrumental relationship (see Brewer, 1974), there has been little work assessing the concordance between instrumental action and beliefs across the major parameters of the relationship. This is the issue addressed here.

In experiments on instrumental conditioning in animals, two major parameters of the instrumental relationship have been identified. The first is the temporal contiguity between the outcome and the action that caused it. If the occurrence of the outcome is delayed after the action, performance decreases systematically with the length of the delay (see Mackintosh, 1974, and Tarpay & Sawabini, 1974, for reviews). A second important determinant of performance is the action-outcome contingency. The effect of contingency can be dissociated from that of contiguity in a free-operant procedure by keeping the probability of a contiguous outcome [$p(O/A)$] constant and varying the probability of an outcome in the absence of an immediately preceding action [$p(O/-A)$]. It is now well established that as the causal effectiveness of the action is degraded by increasing $p(O/-A)$, performance of the instrumental action declines (see, e.g., Hammond, 1980).

The intentional account of instrumental action predicts that judgments of the causal effectiveness of an action and performance of that action should show a comparable sensitivity to variations in action-outcome contiguity and contingency. Chatlosh, Neunaber, and Wasserman (1985) presented extensive evidence for such a concordance

across variations of the instrumental contingency. Their human subjects performed a simple free-operant task, in which pressing a key was the instrumental action and a flash of a light the outcome. When the subjects were encouraged to maximize outcome occurrences by associating each outcome with the gain of a point, Chatlosh et al. (1985) found that both the probability of keypressing and judgments of the causal effectiveness of the action decreased systematically as $p(O/-A)$ was raised. Our first study was an attempt to replicate this observation, whereas in the second, we investigated the concordance between performance and judgment across variations in action-outcome contiguity.

EXPERIMENT 1

The design of our study replicated a subset of the conditions investigated by Chatlosh et al. (1985). In three of the four conditions, $p(O/A)$ per second was set at 0.875, and the contingency was varied across these conditions by setting $p(O/-A)$ at 0.125, 0.500, and 0.875 per second. Under the latter condition, keypressing had no causal effect, and the occurrence of the outcome was noncontingent. In addition, we presented the subjects with another noncontingent condition in which the outcome probability per second was 0.125 rather than 0.875. This low-frequency, noncontingent condition was included to investigate a potentially crucial dissociation between performance and judgments observed by Chatlosh et al. They reported that judgments under noncontingent schedules tended to increase with outcome probability, whereas performance decreased. Clearly such a dissociation would be problematic for an intentional account. Therefore, the reliability of this observation should be assessed.

We made two modifications to the procedure used by Chatlosh et al. (1985). They employed a within-subjects design in which performance and judgments were measured in the same subjects. The advantage of such a within-subjects design for assessing the concordance between performance and causal judgment is that the subjects' judgments are based on the actual experience of the instrumental contingency generated by their performance under this relationship. With this design, however, there is the possibility that requiring the subjects to make an explicit causality judgment biases them toward adopting the cognitive strategy of basing their performance on their current causal belief, a strategy that they may not spontaneously adopt when faced solely with the task of maximizing their points score. Perruchet and Baveux (1989) have criticized some earlier implicit learning research on much the same grounds. Consequently, we used a between-subjects design in which performance and judgments were assessed in separate groups. Second, our payoff schedule specified not only a points gain for outcomes but also a cost for keypressing. In the absence of such a cost, there is no penalty for overresponding on non-

contingent schedules and thus no explicit reason for regulating performance according to the action–outcome contingency.

Method

Subjects and Apparatus. The subjects were 28 members of the Applied Psychology Unit's subject panel, who were paid for their participation. The experiment was conducted on a BBC B/Torch Z80 microcomputer, connected to a video screen (white on black). The subjects were tested individually in a sound-proofed cubicle.

Procedure. At the beginning of the experiment, the subjects were given the following instructions on the computer screen:

Please read and follow the instructions very carefully. Take as much time as you like. There will be a triangle on the screen like this:

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

Now press the SPACE BAR and see what happens...

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 had occurred, the outline of the triangle remained on the screen as before. Then the instructions continued:

and press it a few more times... Then press the RETURN key to continue...

The subject could press the space bar and observe the triangle flashing a few more times. Then the instructions continued:

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. For the 12 subjects in the judgment group, the task was to make judgments about the action–outcome relationship. For this group, the instructions were as follows:

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.

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.

You will be given four different problems, each lasting for two 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.

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.

For the 16 subjects in the performance group, the measure of interest was their rate of pressing the space bar. The instructions given to subjects in this group were as follows:

Your task in this experiment is to score as many points as you can. Every time the triangle flashes, 3 points will be added to your score, but each press of the SPACE BAR will cost you 1 point, which will be deducted from your score.

Thus if pressing the SPACE BAR tends to cause the triangle to flash, then you will gain more points by pressing it. But if the triangle tends to flash of its own accord, you will simply lose points by pressing the SPACE BAR.

You will be given four different problems, each lasting for two 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.

At the end of each problem you will be told how many points you scored in that problem. You must try to get as many points as possible in each problem.

The first problem commenced after these instructions. Each problem lasted for 2 min. In each problem, the 2-min period was divided into 120 1-sec intervals. If a press occurred during a given interval, the outcome occurred at the end of that interval with probability $p(O/A)$. If no press occurred, the outcome occurred with probability $p(O/-A)$. Only the first press in an interval was recorded and had any programmed consequences.

In three conditions (875/125, 875/500, and 875/875), $p(O/A)$ was 0.875. The conditions differed in the value of $p(O/-A)$, which was increased from 0.125 (875/125) to 0.500 (875/500) to 0.875 (875/875). Finally, there was another noncontingent condition (125/125) in which $p(O/A)$ and $p(O/-A)$ were both 0.125. These probabilities acted as parameters for a software probability generator, so that the actual $p(O/A)$ and $p(O/-A)$ experienced by each subject could differ slightly from these nominal values. The order in which the conditions were presented was counterbalanced across subjects in such a way that each condition occurred equally as often as the first, second, third, and fourth conditions.

At the end of each problem, subjects in the judgment group gave a rating of the causal effectiveness of pressing the space bar. They were given the following instructions:

Type in a number below to indicate the extent to which pressing the SPACE BAR caused the triangle to light up. Use a scale from zero 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 judgement. Your judgement?

The subjects in the performance group, on the other hand, were simply told how many points they had scored in the preceding problem. This was calculated by adding 3 points for every outcome that occurred and subtracting 1 point for each press of the space bar.

Results and Discussion

The top panel of Figure 1 displays the mean presses per minute emitted by the performance group in the various conditions, whereas the mean ratings of the causal effectiveness of keypressing made by the judgment group are illustrated in the bottom panel. The reliability of the effects of varying action–outcome contingency was assessed against a criterion of 0.05 for rejecting the null hypothesis.

As can be seen from Figure 1, our results replicated the general concordance between performance and causal judgment observed by Chatlosh et al. (1985); both measures showed a systematic decline as the causal relationship between the action and outcome was degraded by increasing $p(O/-A)$. An overall analysis revealed a significant effect of contingency for both the judgments [$F(3,33) = 25.56$, $MS_e = 308.7$], and keypresses per minute [$F(3,45) = 8.09$, $MS_e = 134.8$]. Preplanned, nonorthogonal contrasts, assessed against the critical values for Bonferroni F , revealed significant differences between the 875/125 and 875/500 conditions for both ratings [$F(1,33) = 17.69$, $MS_e = 51.5$] and presses per minute [$F(1,45) = 6.47$, $MS_e = 16.8$], but not between the 875/500 and 875/875 conditions [$F(1,33) = 6.01$, $MS_e = 51.5$, and $F(1,45) = 1.14$, $MS_e = 16.8$, respectively].




Figure 1. Experiment 1: Mean number of presses per minute for the performance group (top panel), and mean causality ratings for the judgment group (bottom panel), across the different action-outcome contingencies.

We failed to replicate the apparent dissociation between judgment and performance observed by Chatlosh et al. (1985). In our study, neither the causal ratings [$F(1,33) = 2.34$, $MS_e = 51.5$] nor the rate of keypressing [$F(1,45) = 1.19$, $MS_e = 16.8$] differed significantly under the two noncontingent schedules. The direction of the difference was in agreement with Chatlosh et al.'s observations for the judgments, in that the noncontingent schedule with a higher outcome frequency (875/875) tended to produce higher ratings than did the low-frequency schedule (125/125), a finding that has been reported in several other studies (see Shanks, 1991, for a review). In contrast to Chatlosh et al.'s (1985) findings, however, the difference was in the same direction for the judgments and for the rate of keypressing in the performance group. The reason for this discrepancy is not clear, but it should be noted that the subjects in Chatlosh et al.'s study were exposed to negative contingencies as well as the positive and noncontingent schedules employed in the present experiment.

The task demands affected performance differentially in the two groups; the rate of keypressing was significantly affected by variations in the action-outcome contingency only in the performance group. An overall analysis of the rate of keypressing revealed a significant group \times contingency interaction [$F(3,78) = 3.18$, $MS_e = 124.3$], and an analysis of the simple main effect failed to reveal a significant effect of contingency in the judgment group [$F(3,33) = 2.79$, $MS_e = 110.0$]. The mean rates of keypressing in this group under the various contingencies were 875/125: 25.7; 875/500: 16.0; 875/875: 15.5; and 125/125: 22.8 presses per minute. An interaction between the effect of the task instructions and the action-outcome contingency was also reported by Chatlosh et al. (1985); as in our study, variations in contingency failed to affect the performance of their subjects unless they were requested to maximize the number of outcomes. Presumably, this difference reflected the fact that the performance strategy for maximizing the points score differs from that for determining the contingency. For instance, the judgment strategy of sampling both $p(O/A)$ and $p(O/-A)$ is appropriate under each contingency, but it is clearly at variance with maximizing the points score.

Finally, it should be noted that the actual contingencies—as measured by the difference between $p(O/A)$ and $p(O/-A)$ experienced by the two groups—were comparable. The differences between these two conditional probabilities for the judgment group were 875/125: 0.725; 875/500: 0.386; 875/875: -0.019 ; and 125/125: 0.003; and for the performance group, they were 875/125: 0.734; 875/500: 0.373; 875/875: -0.053 ; and 125/125: -0.023 . An overall analysis revealed neither a significant effect of group nor a significant group \times contingency interaction ($F < 1$ in both cases).

In conclusion, this study replicated the concordance observed by Chatlosh et al. (1985) between causal judgments and performance across variations in action-outcome contingency. In fact, the concordance observed in our study was somewhat better, in that we failed to detect any dissociation between the two measures under noncontingent schedules differing in outcome frequency. Moreover, this concordance was observed in the absence of any explicit requirement that the subjects in the performance group make a cognitive judgment. In the next study, we assessed the effect of varying the action-outcome contiguity rather than contingency.

EXPERIMENT 2

The general design of this study was based on that of an experiment reported by Williams (1976). Williams compared the performance of pigeons under different delays of reinforcement with that on a yoked, noncontingent schedule delivering the same rate and temporal pattern of reinforcement. With this design, the ability of the subjects to detect the causal relationship between an action and the outcome, a reinforcer, under different de-

lays can then be determined by comparing performance on the contingent schedule with that under the yoked, non-contingent condition. Williams reported that a delay of 3 sec or more produced rates that were only marginally higher than those of the noncontingent controls.

We have already reported that human causality judgments showed a similar sensitivity to outcome delay in the free-operant task employed in Experiment 1. Shanks, Pearson, and Dickinson (1989) found that in judging the causal effectiveness of an action, subjects discriminated significantly between the contingent and noncontingent conditions with delays of 0 and 2 sec, but not with delays of 4 sec or more. What remains unknown, however, is whether human performance would show a comparable sensitivity to temporal contiguity.

In Experiment 2, two groups of subjects were given the opportunity to press a key that was associated with the outcome on either contingent schedules with outcome delays of 0, 2, and 4 sec or on yoked, noncontingent schedules. After experiencing each condition, one group of subjects was asked to rate the causal effectiveness of the action on the immediately preceding schedule. By contrast, the second performance group was not required to make a causal judgment but simply to maximize their points score under a payoff schedule similar to that used in Experiment 1. The intentional theory anticipates that causal judgments and performance should show a comparable sensitivity to variations in action–outcome contiguity.

Method

Subjects and Apparatus. The subjects were a further 24 members of the APU subject panel, and the apparatus was the same as that employed in Experiment 1.

Procedure. The instructions given to the subjects were the same as those presented in Experiment 1, except that they indicated that there would be six 2-min problems. For the 12 subjects in the judgment group, the task was to rate the effectiveness of keypressing the space bar in causing the outcome, whereas the remaining 12 subjects in the performance group were asked to maximize their points score using the payoff schedule employed in Experiment 1.

In the three contingent conditions, every press was recorded and set up the outcome with probability 0.75. Outcomes were never set up in the absence of a response. What differed across the conditions was the delay between the action and the outcome, which was 0, 2, or 4 sec. Each contingent condition was followed by its noncontingent, control condition. In these conditions, the temporal pattern of outcomes that had been delivered in the immediately preceding contingent condition was played back to each subject independently of the subject's responding. Presses were recorded, but they had no programmed consequences. The three pairs of problems were presented in a counterbalanced order across subjects, so that each pair occurred equally as often as the first, second, and third pairs.

Results and Discussion

The top panel of Figure 2 displays the mean number of keypresses per minute in the various conditions emitted by the performance group, whereas the mean ratings of the causal effectiveness of keypressing by the judgment group are illustrated in the bottom panel. The performance and judgment data were analyzed separately.




Figure 2. Experiment 2: Mean number of presses per minute for the performance group (top panel), and mean causality ratings for the judgment group (bottom panel), across the different action–outcome delays.

The pattern of judgments essentially replicated that observed by Shanks et al. (1989). The subjects clearly discriminated between the contingent and noncontingent conditions when the outcome was immediate, since the contingent condition yielded high ratings and the noncontingent condition low ratings. The ratings decreased with increasing delay in the contingent conditions, whereas this factor had little effect on judgments when keypressing did not affect the outcome. An overall analysis revealed a significant effect of contingency [$F(1,11) = 74.10$, $MS_e = 478.0$], but not of delay [$F(2,22) = 3.05$, $MS_e = 402.4$], and a significant contingency \times delay interaction [$F(2,22) = 9.89$, $MS_e = 374.2$]. Separate preplanned orthogonal comparisons revealed that the contingent condition yielded significantly higher judgments than did the corresponding noncontingent condition at each delay [$F(1,11) > 6.41$, $MS_e = 79.7$], in all cases. However, the ability to discriminate the contingency at each delay does not mean that the subjects were insensitive to the delay in the contingent condition. Separate one-way analyses revealed a significant effect of delay in the contingent conditions [$F(2,22) = 7.14$, $MS_e = 641.3$], but not in the noncontingent conditions [$F(2,22) = 2.60$, $MS_e = 135.3$].

The rate of keypressing in the performance group showed essentially the same pattern as that seen for judg-

ments. Although neither the main effect of contingency [$F(1,11) = 4.11$, $MS_e = 442.1$], nor that of delay [$F(2,22) = 1.27$, $MS_e = 515.9$], was significant, these two factors did interact [$F(2,22) = 5.30$, $MS_e = 267.7$]. As in the case of the judgments, subjects discriminated between the contingent and noncontingent conditions under a 0-sec delay [$F(1,11) = 9.82$, $MS_e = 73.7$] by keypressing at a high rate with contingent outcomes, but not with noncontingent outcomes. In contrast to the judgments, however, the rate of keypressing failed to reveal discrimination between the contingent and noncontingent conditions with delays of 2 and 4 sec ($F < 1$ in both cases).

The effects of delay cannot be explained in terms of fortuitous differences in the probability of an outcome, given a keypress. These probabilities in the contingent conditions at the 0-, 2-, and 4-sec delays were 0.75, 0.75, and 0.72, respectively, for the judgment group, and 0.74, 0.77, and 0.71, respectively, for the performance group. An overall analysis failed to yield significant effects of group ($F < 1$), or delay [$F(2,44) = 1.66$, $MS_e = 79.0$] on these probabilities and no significant interaction between these factors ($F < 1$).

In a free-operant procedure, such as that employed in the present experiments, the delay between the action that causes the outcome and the actual occurrence of the outcome need not correspond to the experienced delay between the outcome and the immediately preceding action, since the subject is free to keypress during the delay period. Table 1 shows the mean delays between each outcome and the immediately preceding action, regardless of whether this was the action that caused the outcome. There was no evidence that the experienced delays differed significantly between the groups. The effect of group and all interactions involving this factor failed to reach significance in analyses of the experienced delays in both the contingent conditions under the 2- and 4-sec delays and the noncontingent conditions at all delays [$F(1,22) < 4.27$, $MS_e = 11.5$, and $F(2,44) = 2.18$, $MS_e = 42.0$, in all cases].

The overall similarity between the pattern of judgments and performance across variations in outcome delay obviously accords with the general prediction of intentional theory. We did, however, detect one apparent dissociation. Although increasing the delay produced a decline in causal ratings, the subjects were still capable of detecting the presence of an action–outcome relationship with delays of 2 and 4 sec. By contrast, the instrumental

relationship had no significant effect on performance at these delays. Note that this is opposite to the usual finding in implicit learning experiments, since here the judgments, not the subjects' performance, is the more sensitive measure.

A reason for this dissociation might lie in the differing task demands faced by the two groups. It is possible that the performance subjects, in fulfilling the task demand of maximizing their points score, generated a performance profile that did not allow them to detect a causal relationship under delayed outcomes. We have already seen in the first study that the differing task demands of the performance and judgment conditions can lead to different performance profiles.

It is unlikely, however, that this account explains the dissociation observed in the present study because the two groups exhibited similar levels of keypressing under the various conditions. Delays of 0, 2, and 4 sec in the judgment group produced mean rates of 48.2, 21.1, and 22.8 presses per minute, respectively, in the contingent conditions and of 21.7, 22.0, and 17.6 presses per minute, respectively, in the noncontingent conditions. In an overall analysis in which the rates of keypressing in the two groups were compared, the main effect of groups ($F < 1$), and all interactions of the group factor with those of contingency and delay failed to reach significance [$F(2,44) < 1.05$, $MS_e = 175.2$], in all cases. An analysis of the rate of keypressing in the judgment group alone revealed a significant effect of delay [$F(2,22) = 24.03$, $MS_e = 65.9$], an effect of contingency [$F(1,11) = 14.21$, $MS_e = 133.8$], and a significant contingency \times delay interaction [$F(2,22) = 14.99$, $MS_e = 82.7$]. Individual comparisons attributed these effects to a significant difference between the contingent and noncontingent conditions only at the 0-sec delay [$F(1,11) = 31.50$].

We doubt, however, that the dissociation between causal ratings and performance under 2- and 4-sec delays constitutes a serious embarrassment for an intentional account, given the actual payoff schedule for the performance group under which each press cost 1 point and each outcome gained 3 points. Under this schedule, one can only achieve a net gain of points by keypressing if a press increases the probability of the outcome by more than 0.33. Thus, even if one believes that an action has a weak causal influence on an outcome, the intentional account is not required to assume that such beliefs are automatically expressed in action. Presumably, the subjects in the judgment group also attached some "value" to an outcome and were operating under the constraint of some implicit payoff rule. If so, we should not anticipate that a belief about a weak causal relationship would necessarily be manifest in performance. The dissociation that would cause problems for intentional theory would be a demonstration that performance yields discrimination between contingent and noncontingent schedules under conditions in which judgments do not.

Finally, the present results bear upon an interpretation considered by Shanks et al. (1989) of the sensitivity of causal judgments to action–outcome contiguity. They

Table 1
Mean Delay (in Seconds) Between an Outcome
and the Preceding Keypress

Group	Delay Condition		
	0 sec	2 sec	4 sec
Judgment			
Contingent		1.69	2.33
Noncontingent	9.65	5.18	8.95
Performance			
Contingent		1.61	2.65
Noncontingent	16.79	9.99	8.45

noted that subjects asked to make causality judgments may interpret the word "cause" in the context of the present task and instructions to mean "is followed immediately by" on the basis of their past experience of the immediacy of the cause-effect relation in electronic devices such as computers. Shanks et al. pointed out, however, that this interpretation could not explain why causal judgments under the same task demands are sensitive to contingency when the probability of an immediate outcome is kept constant, as in Experiment 1.

But even if we were to ignore this difficulty, the results from the performance group of the present experiment cannot be reconciled with the account. In applying the account to the performance group, one would have to presuppose that performance is mediated by a causal belief about the action, but in this case one based on the immediacy of the outcome rather than the actual causal efficacy of pressing the space bar. But such a belief is, of course, not appropriate under the performance task demands, because the actual causal effectiveness of keypressing is not affected by the outcome delay; subjects could have enhanced their points scores by keypressing to the same extent under all delays.

GENERAL DISCUSSION

In these experiments, we have demonstrated a general concordance between performance of an instrumental action and judgments about the causal effectiveness of the action in producing an outcome, across variations in two of the major parameters of an instrumental relationship. Such a concordance stands in contrast to the dissociations between performance and verbal reports that have been cited in the implicit-explicit learning literature.

The first experiment replicated the concordance observed by Chatlosh et al. (1985) when the instrumental contingency was degraded by increasing the likelihood of the outcome in the absence of a contiguous action. However, it failed to reproduce the apparent dissociation that Chatlosh et al. observed under noncontingent schedules. Although varying the outcome probability had no significant effect on judgments and performance in our experiment, if anything, both measures tended to increase as outcome probability increased.

In the second experiment, it was found that delaying the outcome reduced both performance and judgments across a comparable temporal range. This general agreement between performance and judgment can be explained in terms of an intentional theory of instrumental action, in which it is assumed that performance is mediated by a belief about the causal efficacy of the action in producing the outcome.

Of course, the relationship between judgment and performance demonstrated in these experiments is only correlational in nature, and therefore it cannot strictly warrant the assumption that a common process is involved in the two measures. However, this conclusion receives increasing support with converging evidence for the corre-

lation. Especially compelling are concordances in cases in which performance and judgments are not veridical, for it is not clear why parallel processes should be subject to the same illusion. Thus the comparable effects of varying the action-outcome delay are particularly instructive, because, of course, the actual causal effectiveness of the action in augmenting the number of outcomes is unaffected by the delay. Nonveridical performance and judgments can also be induced by contingency manipulations. Both Dickinson and Charnock (1985) and Hammond and Weinberg (1984) found that signaling noncontiguous reinforcers under a noncontingent schedule augmented the instrumental performance of rats. A corresponding illusion of control was induced in human causality judgment by Shanks (1989), using the same signaling operation.

It will not have escaped notice that the latter concordance is between animal performance and human judgment. Indeed, both the experiments reported in this paper follow designs initially employed in studying instrumental conditioning in animals. To the extent that similarities between causal judgments and performance provide support for an intentional account, they also suggest that such a theory can be applied to animal performance during instrumental conditioning (see Dickinson & Shanks, 1985). And, of course, Tolman's (1932) original development of intentional theory was based on animal rather than human performance, a feature that is also true of a number of subsequent restatements of this position (e.g., those of Bolles, 1972; Dickinson, 1989).

Finally, what are we to conclude about the claims made in the implicit-explicit learning literature (e.g., Berry & Broadbent, 1984; Hartman et al., 1989; Hayes & Broadbent, 1988; Lewicki, 1986; Lewicki et al., 1987; Nissen & Bullemer, 1987; Reber, 1989; Stadler, 1989; Willingham et al., 1989) that behavior can adapt to complex contingencies without subjects' being "aware" of these contingencies? Leaving aside the problems involved in assessing "awareness" in such tasks (see, e.g., Brody, 1989), our results would appear to be at variance with claims about the status of implicit learning.

However, we do not wish to argue that the implicit-explicit distinction fails to capture any genuine feature of human learning. It is significant that cognitive theorists have always recognized a role for nonintentional processes in the control of behavior (Tolman, 1932, p. 300), a claim that Dickinson (1989) has endorsed on the basis of studies of the effect of goal or outcome devaluation studies. Dickinson reports a variety of studies in which animals were trained to press a lever for a food or liquid reinforcer. Subsequently, they were presented with some more of the same reinforcer (in the absence of the lever) but were poisoned when they consumed it, and hence an aversion to the reinforcer was established. According to the intentional account, when they were again given access to the lever, the animals should have suppressed their responding. Instead, Dickinson reports that in some circumstances they continued inappropriately to press the lever. Such

a finding provides convincing evidence for nonintentional factors in behavior, but this conclusion is not based on assessments of the concordance between cognition and action. If such a dual-process theory of action is accepted, the interesting question becomes that of assessing the conditions that determine the relative contributions of intentional and nonintentional processes to the control of action.

REFERENCES

- BERRY, D. C., & BROADBENT, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, **36A**, 209-231.
- BOLLES, R. C. (1972). Reinforcement, expectancy, and learning. *Psychological Review*, **79**, 394-409.
- BREWER, W. F. (1974). There is no convincing evidence for operant or classical conditioning in adult humans. In W. B. Weimer & D. S. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 1-42). Hillsdale, NJ: Erlbaum.
- BRODY, N. (1989). Unconscious learning of rules: Comment on Reber's analysis of implicit learning. *Journal of Experimental Psychology: General*, **118**, 236-238.
- CHATLOSH, D. L., NEUNABER, D. J., & WASSERMAN, E. A. (1985). Response-outcome contingency: Behavioral and judgmental effects of appetitive and aversive outcomes with college students. *Learning & Motivation*, **16**, 1-34.
- DICKINSON, A. (1989). Expectancy theory in animal conditioning. In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theory: Pavlovian conditioning and the status of traditional learning theory* (pp. 279-308). Hillsdale, NJ: Erlbaum.
- 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. (1985). Animal conditioning and human causality judgment. In L.-G. Nilsson & T. Archer (Eds.), *Perspectives on learning and memory* (pp. 167-191). Hillsdale, NJ: Erlbaum.
- GREEN, R. E. A., & SHANKS, D. R. (1990). *On the existence of independent learning systems: An examination of some evidence*. Manuscript submitted for publication.
- 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.
- HAMMOND, L. J., & WEINBERG, M. (1984). Signaling unearned reinforcers removes the suppression produced by a zero correlation in an operant paradigm. *Animal Learning & Behavior*, **12**, 371-377.
- HARTMAN, M., KNOPMAN, D. S., & NISSEN, M. J. (1989). Implicit learning of new verbal associations. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1070-1082.
- HAYES, N. A., & BROADBENT, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, **28**, 249-276.
- HEYES, C., & DICKINSON, A. (1990). The intentionality of animal action. *Mind & Language*, **5**, 87-104.
- IRWIN, F. W. (1971). *Intentional behavior and motivation: A cognitive theory*. Philadelphia: Lippincott.
- LEWICKI, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **12**, 135-146.
- LEWICKI, P., CZYZEWSKA, M., & HOFFMAN, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 523-530.
- MACKINTOSH, N. J. (1974). *The psychology of animal learning*. London: Academic Press.
- NISSEN, M. J., & BULLEMER, P. (1987). Attentional requirements of learning: Evidence from performance. *Cognitive Psychology*, **19**, 1-32.
- PERRUCHET, P., & BAVEUX, P. (1989). Correlational analyses of explicit and implicit memory performance. *Memory & Cognition*, **17**, 77-86.
- PERRUCHET, P., GALLEGO, J., & SAVY, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. *Cognitive Psychology*, **22**, 493-516.
- PERRUCHET, P., & PACTEAU, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, **119**, 264-275.
- PERRUCHET, P., & PACTEAU, C. (1991). The implicit acquisition of abstract knowledge about artificial grammar: Some methodological and conceptual issues. *Journal of Experimental Psychology: General*, **120**, 112-116.
- REBER, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, **118**, 219-235.
- SANDERSON, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 729-747.
- SANDERSON, P. M. (1990). Implicit and explicit control of a dynamic task: Empirical and conceptual issues (Tech. Rep. No. EPRL-90-02). University of Illinois at Urbana-Champaign, Engineering Psychology Research Laboratory.
- SHANKS, D. R. (1989). Selectional processes in causality judgment. *Memory & Cognition*, **17**, 27-34.
- SHANKS, D. R. (1991). *Actions and the judgment of causality*. Manuscript submitted for publication.
- SHANKS, D. R., PEARSON, S. M., & DICKINSON, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *Quarterly Journal of Experimental Psychology*, **41B**, 139-159.
- STADLER, M. A. (1989). On learning complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1061-1069.
- TARPY, R. M., & SAWABINI, F. L. (1974). Reinforcement delay: A selective review of the last decade. *Psychological Bulletin*, **81**, 984-997.
- TOLMAN, E. C. (1932). *Purposive behavior in animals and men*. New York: Irvington.
- TOLMAN, E. C. (1959). Principles of purposive behavior. In S. Koch (Ed.), *Psychology: A study of a science* (Vol. 2, pp. 92-157). New York: McGraw-Hill.
- WILLIAMS, B. A. (1976). The effects of unsignalled delayed reinforcement. *Journal of the Experimental Analysis of Behavior*, **26**, 441-449.
- WILLINGHAM, D. B., NISSEN, M. J., & BULLEMER, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1047-1060.

NOTE

1. This simple assumption is only reasonable for schedules in which the probability of an outcome does not vary with the time since the last response or outcome. This is true of the random ratio schedules employed in the present experiments.

(Manuscript received April 10, 1990;
revision accepted for publication December 3, 1990.)