Transfer of control between causal predictive judgments and instrumental responding

CONCEPCIÓN PAREDES-OLAY, MARÍA J. F. ABAD, MATÍAS GÁMEZ, and JUAN M. ROSAS
University of Jaén, Jaén, Spain

Four experiments were conducted to explore outcome-specific transfer from causal predictive judgments to instrumental responding. A video game was designed in which participants had to defend Andalusia from navy and air force attacks. First, they learned the relationship between two instrumental responses (two keys on a standard keyboard) and two different outcomes (destruction of the ships or destruction of the planes). Then they learned to predict which of two different stimuli predicted which outcome. Finally, they had the opportunity of making either of the two instrumental responses in the presence of either stimulus. Transfer was shown as a preference for the response that shared an outcome with the current stimulus. The presentation of the stimulus during the test produced a decrease in the overall rate of response. Responding to a neutral stimulus in Experiments 2 and 3 suggested that this overall decrease in responding was due to a combination of the time needed to process the meaning of the stimulus and the activation of the representation of the outcome in the presence of the stimulus during the test. Transfer between predictive judgments and instrumental responding mirrors the outcome-specific Pavlovian instrumental transfer observed in conditioning studies with rats.

The study of causal and contingency learning has become common in the literature over the past 2 decades (e.g., Shanks, Holyoak, & Medin, 1996). Some of the authors interested in the study of human causal learning have proposed that causal learning is regulated by a statistical algorithm, such as delta P or one of its variations (e.g., Allan, 1980; Cheng, 1997; Cheng & Novic, 1990, 1992). However, the parallelism found between data from causal learning and data from classical and instrumental conditioning with animals has led other authors to suggest that these forms of learning may share the same underlying mechanism (e.g., Dickinson & Burke, 1996; Dickinson & Shanks, 1987; Shanks, 1993). Most of the authors who have taken this approach have suggested that the shared mechanism is associative (e.g., Dickinson, 2001; Shanks, 1995). However, the assumption of a common algorithm of learning does not necessarily imply that it is an associative one. Statistical models have also been used to explain animal conditioning (e.g., Gallistel, 1990). Similarly, the results of recent studies on memory in causal learning have revealed a strong parallelism with those of studies on memory in nonhuman animal conditioning, suggesting that some memory mechanism may be shared

across species (e.g., Paredes-Olay & Rosas, 1999; Rosas, Vila, Lugo, & López, 2001; Vila & Rosas, 2001a, 2001b).

The general aim of the experiments reported here was to explore this parallelism between animal conditioning and causal learning and extend it to the contents of learning. One technique that has been particularly useful for assessing the contents of learning in Pavlovian and instrumental conditioning is the transfer test. With this test, the potential of a stimulus to transfer its control to an instrumental response that is followed by the same outcome is evaluated (e.g., Colwill & Rescorla, 1988; Rescorla, 1994). For instance, Delamater (1996) trained animals on two instrumental responses with two different outcomes. Then he paired two different conditioned stimuli, one with each outcome. After different interference treatments, he ran an extinction test that involved a choice between the two instrumental responses in the presence of each conditioned stimulus. The animals responded more on the alternative that shared an outcome with the current stimulus, thereby revealing a transfer of control between the conditioned stimulus and the instrumental response. This result suggests that the representation of both the stimulus–outcome and the response–outcome associations includes some of the specific sensory features that differentiate the two outcomes.

The transfer test has been extensively used to explore the associative structure of instrumental learning (e.g., Colwill, 1993, 1994; Rescorla, 1991, 1992a, 1992b). One interesting aspect revealed by this analysis is that associations with an outcome are extremely resistant to manipulations, such as extinction, that are intended to undermine

This research was funded by the research group HUM642 from the Junta de Andalucía, Spain. We thank Juan Espinosa and Ana García for their help in running the experiments. Correspondence concerning this paper may be addressed to C. Paredes-Olay or J. M. Rosas, Departamento de Psicología, Universidad de Jaén, 23071 Jaén, Spain (e-mail: cparedes@ujaen.es or jmrosas@ujaen.es).

the strength of those associations. This has been found in both instrumental (Rescorla, 1991, 1993a, 1993b, 1996b) and classical (e.g., Delamater, 1996; Rescorla, 1996a) conditioning.

Taking in account how useful the transfer technique has been in the study of the contents of learning in animal conditioning and extinction, the aim of the experiments reported here was to examine a transfer procedure that could be used to study the contents of human causal learning. Participants were trained in a situation in which two instrumental responses were followed by two different outcomes (R1 +, R2*). Then they had to learn a predictive stimulus–outcome relationship between two stimuli and the outcomes previously used during the instrumental training (A +, B*). Finally, the participants received a transfer test in which they had the opportunity to perform both instrumental responses in the presence of either stimulus. In Experiment 1, this procedure was tested with the aim of finding the basic transfer effect. The purpose of Experiments 2, 3, and 4 was to replicate the transfer effect, in order to evaluate the specific features of transfer in this situation.

EXPERIMENT 1

Experiment 1 was conducted with the aim of testing a transfer technique for studying the contents of causal learning similar to techniques previously used to study the contents of animal learning. Participants performed a video game in which they had to defend Andalusia from air force and navy attacks by using two keys on a standard computer keyboard (R1 and R2). During this instrumental training, pressing one of the keys destroyed ships (R1 +), whereas pressing the other destroyed planes (R2*). Then the participants received training on a response-independent stimulus–outcome situation in which they had to predict which of two stimuli announced the destruction of either ships or planes (A +, B*). The presentation of the figures was followed by the appropriate outcome, independent of the judgment emitted by the participants. During the final test, the participants had the opportunity to choose between R1 and R2 in the presence of each of the stimuli. A transfer effect would be shown if the participants preferentially chose the response that had been reinforced with the same outcome as the one announced by

the current stimulus—that is, R1 rather than R2 in the presence of A and R2 rather than R1 in the presence of B.

Method

Participants

Sixteen undergraduate students at the University of Jaén participated in the experiment. They were between 18 and 25 years old and had no previous experience with this task. Approximately 75% were women, and 25% were men.

Apparatus

The participants were run individually in three adjacent isolated cubicles. Each cubicle had an IBM-compatible personal computer on which the task was presented. The procedure was implemented using the program SuperLab Pro (Cedrus Corporation). Keys F and J, covered with either a green or an orange sticker, were used as R1 and R2. Outcomes + and * were represented by a cross superimposed on the drawing of either the ships or the planes with the label “DESTROYED!” below. A yellow star and a circle (formed by green, blue, and yellow arrows) were used as Stimuli A and B. Stimuli, responses, and the position and type of the outcomes were fully counterbalanced. All actual words were, of course, in Spanish.

Procedure

The participants entered their cubicles and sat in front of the computer. The experimenter asked him or her to pay attention to the instructions on the computer screen. When the participant said that he or she was ready to begin the experiment, the experimenter left the cubicle.

The following instructions were presented in Spanish as successive screens. Instructions were presented using a black Times New Roman 28 font (bold and italic) against a white background. At the bottom of each screen, the sentence *press the space bar* indicated to the participants how to go to the following instructions screen. The sentence *press the space bar to start* was at the bottom of the last instruction screen. This started training. During instrumental training, the sentence *press B key to go on* appeared in the middle of the screen after each trial, and served as intertrial interval. The experiment was run in four phases (see Table 1).

Instrumental training 1. This phase began with the following instructions:

(Screen 1) Welcome! Andalusia is being attacked by sea and air. Your work will be to defend Andalusia by using the ORANGE and GREEN keys. One of the keys fires anti-aircraft missiles, whereas the other one fires anti-boat torpedoes. Your mission consists in destroying the ships and the planes before they reach the Andalusian coast. (Screen 2) Ships and planes have a protective shield that enables them to resist your attack, so sometimes you will need to shoot several times to destroy them. Remember, one key ONLY destroys SHIPS, and the other one ONLY destroys PLANES. The sooner you discover the functions of the keys, the more efficient your defense will be. Good luck! We're depending on you.

Table 1
Design of Experiments 1, 2, 3, and 4

Experiments	Instrumental 1	Stimulus–Outcome	Instrumental 2	Test
1 and 4	R1 +	A +	R1 +	A: R1 vs. R2
	R2*	B*	R2*	B: R1 vs. R2
2 and 3	R1 +	A +	R1 +	A: R1 vs. R2
	R2*	B*	R2*	B: R1 vs. R2
		C–		C: R1 vs. R2

Note—R1 and R2, orange and green keys; A, B, and C, yellow star, circle, and green arrow; + and *, destruction of the ship and destruction of the plane; –, no destruction. Responses, stimuli, and outcomes were counterbalanced across subjects. During the stimulus–outcome training, C was presented 12 times in Experiment 2 and 2 times in Experiment 3.

Twelve trials with each type of attack were then presented randomly intermixed. On each trial, a ship or a plane appeared in either the lower right or the lower left corner of the screen, with its position counterbalanced across trials. Reinforcement (destruction of the attacker) was delivered under a variable interval reinforcement schedule with a mean of 5 sec and a range of between 2 and 8 sec.

Stimulus–outcome training. The phase began with the following instructions:

(Screen 3) Your turn is finished, and you have done good work. Now it is the turn of another defender, and you can only observe what happens. A figure will appear in the screen that will allow you to discover whether the other defender is using the anti-aircraft missiles or the anti-boat torpedoes. Your mission in this part of the game will consist in discovering which figure represents the anti-aircraft missiles and which one represents the anti-boat torpedoes. (Screen 4) Each time a figure appears, you will be asked whether it announces the destruction of the ship, the destruction of the plane, or no destruction. The effect of the figure will appear immediately after your answer. At the beginning, your answer will be random, but little by little, you will discover what each figure means.

The participants received 12 trials, with each figure (A+ or B*) randomly intermixed. In each trial, the figure was presented in the upper center of the screen. The following instructions appeared below the figure:

- Choose a number according to your belief:
1. Destroys the plane.
 2. Destroys the ship.
 3. Does not destroy anything.

The participants had to make predictive judgments about the relationship between the figures and their outcomes. Immediately after a response and independently of which response was made, either a destroyed ship or a destroyed plane was presented, depending on the stimulus. The participants who did not reach a criterion of the last four consecutive trials with each stimulus without error were eliminated from the study.

Instrumental training 2. This phase was identical to the first instrumental training phase, with the exception that the participants received only four trials with each response and that this phase was preceded by the following instructions.

(Screen 5) It is again your turn as a defender. Now you have to use again the ORANGE and GREEN keys to defend the Andalusian coast. Remember, one key ONLY destroys SHIPS, and the other one ONLY destroys PLANES.

Test. The test started with the following screen:

(Screen 6) FINAL ATTACK!!! The enemy has improved its defense, and now it is MORE DIFFICULT to destroy their protective shields. Consequently, they have decided to mount their final attack. Since more torpedoes are now needed to destroy the ships and more missiles to destroy the planes, all the defenders on your side will come into service. (Screen 7) Your mission will be to use the ORANGE and GREEN keys (do not use numbers) to COOPERATE in the destruction of planes and ships. Remember that the MORE missiles a plane receives, and the MORE torpedoes a ship receives, the faster it will be destroyed. Good luck ... Andalusia is in your hands!!

There were two trials with each stimulus. Each trial began with the presentation of a ship and a plane at the same time for 6 sec. This allowed us to assess baseline responding in the absence of either stimulus. Following baseline assessment, the stimuli were presented singly for 6 sec each while both the plane and the ship remained on the screen. No ships or planes were destroyed during this test. Trial order was ABBA for half of the participants and BAAB for the other half.

Dependent Variable and Statistical Analysis

Responses per minute in R1 and R2 during the two last trials of instrumental training and during the test were registered and eval-

uated with an analysis of variance (ANOVA). Planned comparisons were conducted using the methods discussed by Howell (1987, pp. 431-443). The rejection criterion was set at $p < .05$.

Results and Discussion

All the participants successfully reached the stimulus–outcome criterion and learned the instrumental response. The participants' choice of response during the test was determined by the current stimulus. They preferentially chose the instrumental response that had been reinforced with the same outcome as that previously announced by the current stimulus.

Instrumental training proceeded smoothly. The participants learned to use R1 and R2 to destroy ships and planes (+ and *). The mean correct responses per minute during the last two trials of instrumental training were 134.62 ($SE = 19.48$) and 126.71 ($SE = 17.17$) for R1 and R2, respectively. There were no significant differences between response alternatives ($F < 1$, $MS_e = 520.32$).

Transfer test results are shown in Figure 1, which presents the mean responding per minute for the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*) and for the alternative that had been reinforced with a different outcome (*different*), both in the absence (*pre*) and the presence (*stimulus*) of the stimulus during the test. The participants responded more on the *same* key than on the *different* key during presentation of the stimulus. This difference was much lower in the absence of either stimulus, suggesting a selective outcome-based transfer effect between the stimulus and the instrumental response. A 2 (stimulus) $\times$ 2 (outcome) ANOVA showed main effects of stimulus [$F(1,15) = 4.68$, $MS_e = 3,040.16$] and outcome [$F(1,15) = 5.50$, $MS_e = 7,687.55$].

Figure 1. Mean responses per minute on the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*) and on the alternative that had been reinforced with a different outcome (*different*), both in the absence ("Pre") and in the presence ("Stimulus") of the stimulus during the test in Experiment 1.

The stimulus $\times$ outcome interaction fell just short of significance [$F(1,15) = 3.94$, $MS_e = 2,633.26$, $p = .066$].

Although the interaction was not reliable, planned comparisons were run to explore the simple effect of outcome at each level of the stimulus. These analyses found that the simple effect of outcome was significant during the *stimulus* condition [$F(1,15) = 6.63$, $MS_e = 7,132.59$; $SE = 22.74$ and 12.84 , for *same* and *different*, respectively], but it was not significant during the *pre* condition [$F(1,15) = 1.69$, $MS_e = 3,188.28$; $SE = 23.50$ and 20.44 , for *same* and *different*, respectively]. An analysis run to explore the simple effect of the stimulus at each level of the outcome revealed a significant difference between the *pre* and the *stimulus* conditions when the outcome was *different* [$F(1,15) = 7.22$, $MS_e = 3,385.70$; $SE = 20.44$ and 12.84 , for *pre* and *stimulus* conditions, respectively], but not when the outcome was *same* ($F < 1$, $MS_e = 2,287.71$; $SE = 23.50$ and 20.44 , for *pre* and *stimulus* conditions, respectively).

The results of this experiment show that this procedure is effective in producing a selective transfer effect based on the outcome shared by the stimulus and the response. This result is consistently found in the animal conditioning literature (Colwill & Rescorla, 1988; Delamater, 1996; Kruse, Overmier, Konz, & Rokke, 1983; Rescorla, 1994). For example, Colwill and Rescorla reported outcome-specific transfer between a Pavlovian CS+ and an instrumental response. They found that the presentation of the Pavlovian stimulus led to a greater response on the alternative with which it shared a reinforcer than on the other alternative.

Similar to what has been found in animals, transfer was reflected in greater responding to the *same* alternative than to the *different* alternative. Because of the cover story used in this experiment, transfer could have appeared as a decrease in responding on the *same* alternative. Since the stimulus predicts the reliable occurrence of one of the outcomes (i.e., A destroys ships), a rational agent could have devoted his or her instrumental responding to trying to get the alternate outcome (i.e., responding on the instrumental alternative that destroys planes). This possibility was prevented by including the instruction prior to test that an increase in the response was required for plane and ship destruction (see Procedure, Screen 7). Pilot experiments had shown that, without this instruction, transfer could occur as either an increase or a decrease in the different response alternatives, depending on the idiosyncratic interpretation of the situation by the participant. Without greater specification, roughly half of the participants reported that they had interpreted the instructions to mean that they should collaborate with the other defender by trying to destroy the same attacker, and the other half had interpreted the instructions to mean that they should collaborate by destroying a different attacker. Under these conditions, mean performance at testing gave the false impression that transfer did not occur. The instructed increase in the requirement for destruction biased the participants who showed transfer to show transfer in the same direction.

The present transfer effect was seen as a decrease in responding on the alternative key that had been followed by an outcome *different* from the one signaled by the stimulus, rather than as an increase in responding on the *same* alternative. This result is somewhat surprising. One would have expected that transfer would be shown as a redistribution of responding in the presence of the stimulus, keeping constant the overall rate of response. That is, transfer was expected to be seen as an increase in responding to the *same* alternative, with respect to responding in the absence of the stimulus, equal to the decrease in responding on the *different* alternative. Instead, the overall rate of responding was lower during the stimulus than in the absence of the stimulus, as was reflected by the significant main effect of stimulus.

However, these results are not completely new in the animal literature (Colwill & Motzkin, 1994; Colwill & Rescorla, 1988). For example, Colwill and Rescorla (Experiment 3) found that the Pavlovian stimulus left unchanged the response that had earned the *same* reinforcer, whereas it depressed performance of the response that earned the *different* reinforcer. The authors suggested the possibility that a Pavlovian stimulus may elicit competing responses, such as the conditioned response of approaching the magazine, decreasing the opportunity for instrumental responding. Potentially, our procedure allowed for competing responses that played a role similar to Pavlovian conditioned responses. Since the participants were trained to use numbers 1, 2, and 3 during the stimulus–outcome training, the use of the number keys could have competed with the instrumental response during the transfer test. We attempted to minimize this possibility by instructing the participants not to use the number keys during the test. In fact, the participants only used the instrumental keys during the test.

Colwill and Rescorla (1988) suggested that a Pavlovian stimulus can activate the representation of the reinforcer during the test, interfering with the potential of the response to activate other reinforcer representations. In our procedure, the activation of the outcome associated with the *same* response could have interfered with the activation of the alternative outcome, thereby causing the observed decrease in responding to the *different* alternative.

Alternatively, the observed decrease in responding on the *different* alternative may have been an artifact caused by an overall decrease in responding during the test. A plausible explanation of this overall decrease in responding comes from the specific features of our procedure. Since we used a short test trial duration (6 sec), it is possible that the necessary time to process the stimulus and make a decision about which response alternative to choose led the participants to decrease their overall rate of responding. Experiments 2 and 3 were designed to test these two alternative explanations.

EXPERIMENT 2

Experiment 1 yielded a technique that was successful in producing an outcome-specific transfer effect between

causal predictive judgments and instrumental responding. As in previous studies with rats (Colwill & Motzkin, 1994; Colwill & Rescorla, 1988), transfer was revealed as a decrease in responding to the *different* response alternative, rather than as an increase in responding to the alternative with which the stimulus shared an outcome. That is, overall responding decreased between the baseline and the stimulus. We proposed two alternative explanations for this decrease in the overall response rate. The stimulus may have activated a representation of the outcome during the test that interfered with the potential of the response to activate the other outcome representation, or the participants may have needed some time to process the meaning of the stimulus. Either possibility would have led to a decrease in the overall rate of responding.

Experiment 2 was designed to replicate the transfer result and to make it possible to decide between the two alternative explanations for the decrease in the overall instrumental responding between the absence and the presence of the stimulus. We introduced a new stimulus (C) that was presented without the outcome during stimulus–outcome training. Otherwise, the design was identical to the one used in Experiment 1. If the decrease in overall responding during the test was due to the time needed to process the stimulus, responding to C should fall approximately at the middle point between responding to *same* and responding to *different*. That is, overall responding in the presence of C should decrease as much as the mean response to *same* and *different*.

On the other hand, if the overall response decrease during the stimulus occurred because the stimulus activated the outcome representation, preventing the activation of the alternative outcome, mean responding to C should be similar to mean baseline responding, since C was never paired with any of the outcomes.

Method

Participants and Apparatus

Eighteen students with characteristics similar to those of the participants in Experiment 1 took part in the experiment. The apparatus was the same as that used in Experiment 1. A new stimulus was added—specifically, a figure of a green arrow oriented like a V. Responses, outcomes, and stimuli were fully counterbalanced.

Procedure

The design of this experiment is presented in Table 1. The procedure was identical to the one used in Experiment 1, except for the following. A new stimulus, C, was presented 12 times during stimulus–outcome training, followed by the outcome *does not destroy anything*, randomly intermixed with 12 A+ and 12 B* trials.

The test was conducted with the three stimuli, A, B, and C. Each one was presented three times. One third of the participants received the test in the order ABCBCACAB, one third received trials BCACABABC, and the last third received CABABCBCA. Thus, trial order during the test was counterbalanced within and between participants. Responding to C was registered as the mean response to the orange and the green keys.

Results and Discussion

The participants responded selectively to the *same* alternative, rather than to the *different* alternative, in the

presence of A and B stimuli, thereby replicating the results of Experiment 1. Moreover, mean responding to C fell between responding to *same* and *different*, decreasing with respect to baseline responding.

During instrumental training, the participants successfully learned to use R1 and R2 to destroy ships and planes. The mean number of correct responses per minute on the two final training trials was 152.04 ($SE = 21.19$) and 155.93 ($SE = 22.20$) to R1 and R2, respectively. There was no significant difference between them ($F < 1$, $MS_e = 273.61$).

Figure 2 presents mean responses per minute with the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*), with the alternative that had been reinforced with a different outcome (*different*), and to stimulus C, both in the absence (*pre*) and the presence (*stimulus*) of the stimulus during the transfer test. Responding to *same* was higher than responding to *different* in the presence of the stimulus, whereas responding to C fell between the two. In the absence of the stimuli, there were no differences in responding. A 2 (stimulus) $\times$ 3 (outcome) ANOVA showed main effects of stimulus [$F(1,17) = 19.15$, $MS_e = 4,092.61$] and outcome [$F(2,34) = 6.53$, $MS_e = 2,351.02$]. Most important, the stimulus $\times$ outcome interaction was also significant [$F(2,34) = 11.66$, $MS_e = 1,819.43$].

A subsequent analysis exploring this interaction showed that the simple effect of outcome was significant in the presence of the stimulus [$F(2,34) = 10.46$, $MS_e = 3,471.87$; $SE = 16.71$, 15.40, and 19.38 for *same*, *different*, and C, respectively], whereas it was not significant in the absence of the stimulus ($F < 1$, $MS_e = 698.59$; $SE = 22.03$, 19.07, and 20.83 for *same*, *different*, and C, respectively). The simple effect of outcome during the stimulus came from higher responding to *same* than to *different* [$F(1,17) = 14.08$, $MS_e = 4,814.79$], thereby replicating the transfer effect found in Experiment 1. Responding to *same* was also higher than to C [$F(1,17) = 9.63$, $MS_e = 3,768.52$]. However, there were no differences on responding between C and *different* [$F(1,17) = 2.67$, $MS_e = 1,832.30$]. Moreover, the simple effect of stimulus was not significant for the *same* outcome ($F < 1$, $MS_e = 1,115.88$), but it was for *different* [$F(1,17) = 22.8$, $MS_e = 3,508.69$] and C [$F(1,17) = 13.13$, $MS_e = 3,106.90$]. That is, there was a decrease in responding between the *pre* and the *stimulus* conditions both in the *different* alternative and in overall responding to Stimulus C.

To summarize, Experiment 2 replicated the transfer effect found in Experiment 1. First, the participants preferentially chose the key that had been followed by the same outcome as the one that had been followed by the different stimulus. Causal predictive judgments that had been established independently of responding transferred their control to instrumental responses that were followed by the same outcomes. And second, transfer was seen as a decrease in responding to the *different* alternative, rather than as an increase to the *same* alternative. Moreover, there was a decrease in responding to a stimulus that had not been followed by an outcome (C), suggesting that the

Figure 2. Mean responses per minute on the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*), on the alternative that had been reinforced with a different outcome (*different*), and to Stimulus C, both in the absence ("Pre") and in the presence ("Stimulus") of the stimulus during the transfer test in Experiment 2.

decrease in the overall responding to the combination of *same* and *different* in the presence of the stimulus may have been due to the time needed to process the stimulus, rather than to an activation of a competing representation of the outcome by the stimulus.

However, it should be noted that the design used in this experiment allows for a different explanation for the decrease in the overall responding to Stimulus C. Although we had assumed that C remained a neutral stimulus after stimulus–outcome training, C had, in fact, received differential inhibition training (12 trials of C–, intermixed with 12 trials of A+ and 12 of B*) that could have made C a conditioned inhibitor for both outcomes. Consequently, it could be claimed that the lower responding to C during the test actually reflected transfer of inhibition to the instrumental response. There is nothing in our data that allows us to rule out this possibility.

To assess this alternative explanation of our results, we performed a new experiment in which Stimulus C was not presented before the test. This experiment essentially replicated the results reported here for Experiment 2 (with rate of responding to C being between the rates of responding to *same* and *different*), suggesting that the specific pattern of transfer in these experiments was due to the time necessary to process the stimulus during the test. However, although the design of this experiment prevented C from becoming a conditioned inhibitor, it left open an alternative explanation for the decrease in responding in the presence of C. The participants did not experience C before the test. Thus, the decrease in responding to C during the stimuli could have been due to the fact that C was a novel, unexpected stimulus, rather than to a hypothetical need for time to process stimuli in this situation.

Experiment 3 was designed to test the effect of a familiar stimulus in a situation that minimized the possibility that the stimulus was an inhibitor. This constituted a clearer

test between the two alternative explanations for the decrease in the overall responding in the presence of the stimulus.

EXPERIMENT 3

We noted that the design used in Experiment 2 could have made C a conditioned inhibitor, so that the decrease in responding in the presence of Stimulus C could have been due to transfer of inhibition, rather than to processing time. To differentiate between these two possible explanations for the decrease in the overall response to C, Experiment 3 repeated the design used in Experiment 2, with the exception that C was presented without either outcome only two times during stimulus–outcome training. We assumed that this would minimize the possibility of C's becoming an inhibitor, while keeping it familiar to the student. Not presenting C before the test could have caused a decrease in responding, because of the surprise that could have been evoked by the presentation of a completely new stimulus.

Method

Participants and Apparatus

Eighteen students with characteristics similar to those of the participants in the previous experiments took part in this experiment. The apparatus was the same as that described above.

Procedure

The design of the experiment is presented in Table 1. The procedure was identical to the one used in Experiment 2, with the exception that C was presented only two times without being followed by an outcome during the stimulus–outcome training.

Results and Discussion

Transfer between predictive judgments and instrumental responses was replicated. The participants responded

selectively in the presence of each stimulus, reducing pressing of the key that did not share an outcome with the stimulus. However, the decrease in responding to Stimulus C was lower than that observed in Experiment 2.

The instrumental training proceeded smoothly. The participants learned to use R1 and R2 to destroy ships and planes successfully. The means for correct responses per minute during the last two final trials of training were 185.31 ($SE = 16.61$) and 203.28 ($SE = 14.62$) to R1 and R2, respectively, with no significant differences between them [$F(1,17) = 2.39$, $MS_e = 1,213.25$].

Figure 3 presents mean responses per minute with the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*), with the alternative that had been reinforced with a different outcome (*different*), and to Stimulus C, both in the absence (*pre*) and in the presence (*stimulus*) of the stimulus during the test. Responding to *same* was higher than responding to *different* in the presence of the stimulus, whereas responding to C was very close to responding to *same*. There were no differences in the absence of the stimuli. These results were confirmed by statistical analysis. A 2 (stimulus) $\times$ 3 (outcome) ANOVA showed a main effect of stimulus [$F(1,17) = 17.84$, $MS_e = 3,031.86$] and outcome [$F(2,34) = 10.75$, $MS_e = 1,451.31$]. The stimulus $\times$ outcome interaction fell just short of significance [$F(2,34) = 3.23$, $MS_e = 2,478.59$, $p = .052$].

Subsequent planned comparisons performed to explore the simple effects of stimulus and outcome showed that the simple effect of outcome was significant in the presence of the stimulus [$F(2,34) = 7.61$, $MS_e = 2,837.98$; $SE = 17.32, 21.47$, and 23.74 for *same*, *different*, and C, respectively], but not in its absence [$F(2,34) = 1.86$, $MS_e = 1,091.91$; $SE = 20.73, 21.63$, and 19.38 for *same*, *different*, and C, respectively]. Responding to *same* and

C was higher than responding to *different* [$F_s(1,17) = 8.98$, $MS_e = 1,708.45$], whereas there were no significant differences between responding to *same* and responding to C [$F(1,17) = 1.88$, $MS_e = 3,627.58$]. Similarly, the simple effect of stimulus was significant for the *different* outcome [$F(1,17) = 15.95$, $MS_e = 2,844.73$] and for stimulus C [$F(1,17) = 8.13$, $MS_e = 2,870.86$], but it was not significant for the *same* outcome ($F < 1$, $MS_e = 2,273.44$).

Transfer between causal judgments and instrumental responding was again replicated in this experiment. It was shown as a decrease in responding on the alternative that had not shared an outcome with the current stimulus, rather than as an increase in responding to the alternative that shared an outcome with the stimulus. We previously suggested that this specific pattern of results could have been due to the time needed by the students to process the stimulus during the test, rather than to the activation of a representation of the outcome previously related to the stimulus. This conclusion relied on the result found in Experiment 2, in which responding to a supposedly neutral stimulus was lower than responding in the absence of the stimulus. However, the design of Experiment 2 allowed the neutral stimulus to become a conditioned inhibitor.

Even though, in this experiment, the possibility of the neutral stimulus C's becoming an inhibitor was minimized, there was still a significant decrease in responding to C, with respect to responding in the absence of the stimulus. This suggests that the time needed to process the stimuli may be partially responsible for the decrease in the overall rate of responding.

However, it should be noted that, in this experiment, responding to C during the test, although it was lower than responding during the *pre* period, was not different from

Figure 3. Mean responses per minute on the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*), on the alternative that had been reinforced with a different outcome (*different*), and to Stimulus C, both in the absence ("Pre") and in the presence ("Stimulus") of the stimulus during the transfer test in Experiment 3.

responding to *same*. That is, the decrease in responding to C, with respect to the absence of the stimulus, seemed to be smaller when conditions were arranged so as to facilitate keeping the meaning of the stimulus neutral. This result suggests that the decrease in the overall response rate during transfer with this procedure and, consequently, the decrease in responding with the *different* alternative may be due to the activation of the outcome in the presence of the stimulus, probably combined with the time needed to process the stimulus when it appears during the test.

EXPERIMENT 4

Experiments 1–3 showed that a stimulus that was trained to predict a specific outcome reduced responding on a response alternative that had been followed by a different outcome. We have interpreted these results as a demonstration of transfer between predictive judgments and instrumental conditioning, a situation akin to the outcome-specific Pavlovian instrumental transfer observed in studies with rats (e.g., Delamater, 1996). However, as has been pointed out by an anonymous reviewer, this is not the only explanation for our results. The situation was designed so that during stimulus–outcome training, the stimulus was paired not only with the destruction of the respective attacker, but also with the invasion by a ship or a plane. Under these circumstances, selective transfer could have been mediated by an association either between a stimulus and the instrumental outcome (the destruction of the event) or between the stimulus and the discriminative stimulus in the instrumental phase (the picture of the ship or the plane).

Experiment 4 was run with the aim of assessing the possibility of participants' establishing a stimulus–discriminative-stimulus association during the stimulus–outcome training. The design was the same as that in Experiment 1, with one exception. Both discriminative stimuli, the plane and the ship, were presented together during stimulus–outcome training, so that a specific predictive stimulus was equally related to the invasion by both ships and planes; thus, it was equally related to both discriminative stimuli, but it was associated with the destruction of only one of the invaders, either a plane or a ship. Transfer under these circumstances would be mediated by the stimulus–outcome association established during predictive training.

Method

Participants and Apparatus

Twenty-four students with characteristics similar to those of the participants in the previous experiments took part in this experiment. The apparatus was the same as that described above.

Procedure

The experimental design is presented in Table 1. The procedure was identical to the one used in Experiment 1, with the exception that the plane and the ship were both presented below Stimulus A or B during stimulus–outcome training. Except for the instructions regarding responding, the screen used during stimulus–outcome training was identical to the one used later on the transfer test.

Results and Discussion

The transfer between predictive judgments and instrumental response found in Experiments 1, 2, and 3 was replicated. The participants chose the response alternative that had shared the outcome with the current stimulus, rather than the other response alternative.

Instrumental training proceeded uneventfully. The participants successfully learned to use R1 and R2 to destroy ships and planes. The mean of correct responses per minute during the last two final trials of training was 184.97 ($SE = 15.29$) and 193.67 ($SE = 14.84$) to R1 and R2, respectively, with no significant differences between them [$F(1,23) = 1.63$, $MS_e = 552.51$].

Figure 4 presents the mean responses per minute with the key that had been reinforced with the same outcome as that for the current stimulus (*same*) and with the alternative that had been reinforced with a different outcome (*different*), both in the absence (*pre*) and the presence (*stimulus*) of the stimulus during the transfer test. There were no differences in responding in the absence of the stimuli. During presentation of the stimulus, however, the participants chose the response alternative that shared an outcome with the present stimulus, rather than the other response alternative. A 2 (*stimulus*) $\times$ 2 (*outcome*) ANOVA showed a main effect of stimulus [$F(1,23) = 5.65$, $MS_e = 948.57$]. The main effect of outcome fell just short of significance [$F(1,23) = 3.57$, $MS_e = 3,609.06$, $p = .071$]. Most important, the stimulus $\times$ outcome interaction was statistically significant [$F(1,23) = 6.04$, $MS_e = 3,634.88$].

A subsequent analysis performed to explore the stimulus $\times$ outcome interaction showed that the simple effect of outcome was significant only during the stimulus [$F(1,23) = 5.21$, $MS_e = 6,567.13$; $SE = 12.01$ and 15.02 for *same* and *different*, respectively], in which responding

Figure 4. Mean responses per minute on the alternative that had been reinforced with the same outcome as that for the current stimulus (*same*) and on the alternative that had been reinforced with a different outcome (*different*), both in the absence (“Pre”) and in the presence (“Stimulus”) of the stimulus during the test in Experiment 4.

for the *same* alternative was higher than that for the *different* alternative. There were no differences in the absence of the stimulus ($F < 1$, $MS_e = 676.81$; $SE = 6.51$ and 5.71 for *same* and *different*, respectively). In contrast, the simple effect of stimulus was significant on the *different* alternative [$F(1,23) = 11.12$, $MS_e = 2,202.81$; $SE = 5.71$ and 15.02 , for the *pre* and the *stimulus* conditions, respectively], but it was not significant on the *same* alternative [$F(1,23) = 1.18$, $MS_e = 2,380.64$; $SE = 6.51$ and 12.01 for the *pre* and the *stimulus* conditions, respectively]. Thus, responding on the *different* alternative decreased between the absence and the presence of the stimulus, whereas responding on the *same* alternative did not change.

These results replicate those found in Experiments 1, 2, and 3, but in a situation in which an association between the stimulus used during the stimulus–outcome training and the discriminative stimulus used for the instrumental response (the plane or the ship) was prevented. In summary, these experiments suggest that outcome-specific predictive judgment instrumental transfer in humans parallels the Pavlovian instrumental transfer observed in conditioning studies in which rats have been used.

GENERAL DISCUSSION

The four experiments presented in this paper showed a consistent transfer effect based on the outcome shared by stimuli and responses, replicating the results previously found in the animal research literature (Colwill & Motzkin, 1994; Colwill & Rescorla, 1988; Kruse et al., 1983; Rescorla, 1994).

The transfer effect was shown consistently across the four experiments reported here as a decrease in responding to the alternative that did not share an outcome with the current stimulus, rather than as an increase in responding to the alternative that shared an outcome announced by the stimulus. This pattern of results replicates almost exactly what has been found in the animal literature with a procedure on which the present research was based (Colwill & Motzkin, 1994; Colwill & Rescorla, 1988; but see Delamater, 1996, Kruse et al., 1983, and Rescorla, 1994, for partial exceptions to the rule).

The animal research literature offers explanations for this pattern of results in terms of competing responses being elicited by the Pavlovian stimulus during the transfer test (e.g., Colwill & Motzkin, 1994; Colwill & Rescorla, 1988). However, this explanation did not seem to fit our results. It is unlikely that the stimuli used in these experiments would have led to a conditioned response that would compete with the instrumental responses used in this task. The possibility of competition between judgments and instrumental responses in this task would arise because the same keyboard was used to emit judgments and to emit instrumental responses. In fact, in a pilot experiment run before this series, we found that some of the participants responded during the test with the same numbers as those previously used to emit predictive judgments. This possibility was precluded in this exper-

imental series by specifically instructing the participants not to use numbers during the test.

Another explanation for this pattern of results, which also comes from the animal literature, suggests that the decrease in overall responding in the presence of the stimulus may arise from an activation of the representation of a outcome by the presence of the stimulus. This representation may then interfere with the potential of the instrumental response to activate the representation of the other outcome (Colwill & Rescorla, 1988). As an alternative explanation, we suggested that the specific features of our procedure could have led to a decrease in responding being caused by the time needed by the participants to process the stimuli during the test. These two alternatives were tested in Experiments 2 and 3 by testing responding to a stimulus that was not followed by outcomes during the stimulus–outcome training. In both experiments, we found that responding to the new stimulus was lower than responding in the absence of it, suggesting that the overall decrease in responding during presentation of the stimulus may have been caused by the time needed to process its meaning. However, it should be noted that when conditions were established that prevented the new stimulus from becoming a conditioned inhibitor and, thus, prevented the possibility of inhibition-based transfer that would have led to a decrease in overall responding, this decrease was quite low, and responding to the new stimulus was identical to responding on the *same* alternative. Overall, this suggests that the specific pattern of transfer found in these experiments is mainly the result of the activation of a representation of the outcome by the stimulus, perhaps combined with the time needed to process the meaning of the stimulus on this task.

Aside from the explanation of the specific pattern of transfer found with this procedure, transfer between predictive causal judgments and instrumental responses was consistently replicated across experiments. As was stated above, the transfer technique has been shown to be an important tool with which to study the contents of learning. For instance, Delamater (1995) has observed different amounts of transfer displayed by stimuli that differ in associative strength. In one of these experiments, differential training (A^* , B^+) was followed by intertrial presentations of one of the reinforcers (A^* , $*$, B^+). As a result of this manipulation, transfer was stronger to Stimulus B than to Stimulus A. Along another line of work, Delamater (1996) found that transfer between Pavlovian stimuli and instrumental responses was not affected by different procedures that were designed to undermine the CS–US association, such as extinction, interference, or presentation of the outcome by itself after acquisition. Similar results have been found by Rescorla (e.g., 1991, 1993a, 1993b), using transfer between discriminative stimuli and instrumental responses. Overall, these results suggest that extinction and other interference treatments do not affect the stimulus–outcome association originally learned, thereby allowing for a better understanding of the learning that occurs during extinction (e.g., Res-

corla, 1993a). The present transfer technique, applicable to the study of the contents of causal learning in human beings, will allow us to explore the kind of learning that occurs after extinction within this paradigm in future research.

It should be noted that the results reported here with predictive judgments are practically identical to the results previously reported for animal conditioning (e.g., Colwill & Rescorla, 1988). As was noted in the introduction, this parallelism between the results from animal conditioning studies and those from causal learning research has been pointed out by several authors (e.g., Dickinson, 2001; Rosas et al., 2001; Shanks, 1993). The results reported here add to this evidence, thereby suggesting a shared mechanism of learning in human causal learning and animal conditioning.

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.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, **104**, 367-405.
- CHENG, P. W., & NOVIC, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality & Social Psychology*, **58**, 545-567.
- CHENG, P. W., & NOVIC, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, **99**, 365-382.
- COLWILL, R. M. (1993). An associative analysis of instrumental learning. *Current Directions in Psychological Science*, **2**, 111-116.
- COLWILL, R. M. (1994). Associative representations of instrumental contingencies. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 31, pp. 1-72). San Diego: Academic Press.
- COLWILL, R. M., & MOTZKIN, D. K. (1994). Encoding of the unconditioned stimulus in Pavlovian conditioning. *Animal Learning & Behavior*, **22**, 384-394.
- COLWILL, R. M., & RESCORLA, R. A. (1988). Associations between the discriminative stimulus and the reinforcer in instrumental learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **14**, 155-164.
- DELAMATER, A. R. (1995). Outcome-selective effects of intertrial reinforcement in a Pavlovian appetitive conditioning paradigm with rats. *Animal Learning & Behavior*, **23**, 31-39.
- DELAMATER, A. R. (1996). Effects of several extinction treatments upon the integrity of Pavlovian stimulus-outcome associations. *Animal Learning & Behavior*, **24**, 437-449.
- DICKINSON, A. (2001). Causal learning: An associative analysis (The 28th Bartlett Memorial Lecture). *Quarterly Journal of Experimental Psychology*, **54B**, 3-26.
- DICKINSON, A., & BURKE, J. (1996). Within-compound associations mediate the retrospective reevaluation of causality judgments. *Quarterly Journal of Experimental Psychology*, **49B**, 60-80.
- DICKINSON, A., & SHANKS, D. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229-261). San Diego: Academic Press.
- GALLISTEL, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- KRUSE, J. M., OVERMIER, J. B., KONZ, W. A., & ROKKE, E. (1983). Pavlovian conditioned stimulus effects upon instrumental choice behavior are reinforcer specific. *Learning & Motivation*, **14**, 165-181.
- PAREDES-OLAY, C., & ROSAS, J. M. (1999). Within-subjects extinction and renewal in predictive judgments. *Psicológica*, **20**, 195-210.
- RESCORLA, R. A. (1991). Associations of multiple outcomes with an instrumental response. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 465-474.
- RESCORLA, R. A. (1992a). Associations between an instrumental discriminative stimulus and multiple outcomes. *Journal of Experimental Psychology: Animal Behavior Processes*, **18**, 95-104.
- RESCORLA, R. A. (1992b). Response-independent outcome presentation can leave instrumental R-O association intact. *Animal Learning & Behavior*, **20**, 104-111.
- RESCORLA, R. A. (1993a). Inhibitory associations between S and R in extinction. *Animal Learning & Behavior*, **21**, 327-336.
- RESCORLA, R. A. (1993b). Preservation of response-outcome associations through extinction. *Animal Learning & Behavior*, **21**, 238-245.
- RESCORLA, R. A. (1994). Control of instrumental performance by Pavlovian and instrumental stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **20**, 44-50.
- RESCORLA, R. A. (1996a). Preservation of Pavlovian associations through extinction. *Quarterly Journal of Experimental Psychology*, **49B**, 245-258.
- RESCORLA, R. A. (1996b). Response-outcome association remains functional through interference treatments. *Animal Learning & Behavior*, **24**, 450-458.
- ROSAS, J. M., VILA, N. J., LUGO, M., & LÓPEZ, L. (2001). Combined effect of context change and retention interval upon interference in causality judgments. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 153-164.
- SHANKS, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, **84**, 319-354.
- SHANKS, D. R. (1995). *The psychology of associative learning*. Cambridge: Cambridge University Press.
- SHANKS, D. R., HOLYOAK, K. J., & MEDIN, D. L. (Eds.) (1996). *The psychology of learning and motivation: Vol. 34. Causal learning*. San Diego: Academic Press.
- VILA, N. J., & ROSAS, J. M. (2001a). Reinstatement of acquisition performance by presentation of the outcome after extinction in causality judgments. *Behavioural Processes*, **56**, 147-154.
- VILA, N. J., & ROSAS, J. M. (2001b). Renewal and spontaneous recovery after extinction in a causal learning task. *Mexican Journal of Behavior Analysis*, **27**, 79-96.

(Manuscript received February 6, 2002;
revision accepted for publication May 13, 2002.)