Learning and Motivation **31,** 323–344 (2000)
doi:10.1006/lmot.2000.1056, available online at http://www.idealibrary.com on IDEAL®

Interference between Elementally Trained Stimuli Can Take Place in One Trial

Nuria Ortega and Helena Matute

Universidad de Deusto, Bilbao, Spain

Recent research has shown that the acquisition of a second cue–outcome association can interfere with responding appropriate to a previously acquired association between another cue and the same outcome, even if the two cues had never received compound training (Matute & Pineño, 1998a). This is similar to other results in the paired-associate literature but it is problematic for associative theories of learning because all of them assume that compound training is necessary for cues to interfere with each other. However, given several assumptions, a recent revision of Wagner's (1981) SOP model proposed by Dickinson and Burke (1996) could account for most of the data available on interference between elementally trained cues. According to the modified SOP model, the target cue that is paired with the outcome during Phase 1 could acquire an inhibitory association with the outcome during the Phase 2 trials in which the interfering cue is trained and the target cue is absent. This inhibitory association could be responsible for the weak responding observed to the target cue during testing because it could interfere with the excitatory association acquired during Phase 1. If this is true, interference should be weaker as the number of Phase 2 interfering trials is reduced. However, the three experiments reported here show that interference can occur even when only one interfering trial is given during Phase 2. The results of these experiments, along with other results in the literature, add support to the idea that interference between elementally trained cues occurs during retrieval and that it is not due to the formation of inhibitory associations between an absent cue and the outcome.

A well-known effect in predictive learning is that when two or more cues are trained together as predictors of the same outcome, they normally interfere with each other, so that if one of the cues is a good predictor of the

Support for this research was provided by Grants PI96/006 from Departamento de Educación, Universidades, e Investigación (Basque Government) and PB95-0440 from Dirección General de Enseñanza Superior (Spanish Ministry of Education) to H.M. and by an F.P.I. fellowship from the Basque Government (Ref. BFI97.102) to N.O. We thank Maria Teresa Bajo, Leyre Castro, Andrés Catena, Ottmar Lipp, Ralph Miller, Oskar Pineño, and Sonia Vegas for illuminating discussions, and two anonymous reviewers for their very helpful comments and suggestions on an earlier version of this manuscript.

Address correspondence and reprint requests to Helena Matute, Departamento de Psicología, Universidad de Deusto, Apartado 1, 48080 Bilbao, Spain. E-mail: matute@orion.deusto.es.

0023-9690/00 $35.00

outcome, humans and other animals normally disregard the capacity of the other cue or cues to predict the outcome. For example, in a typical blocking experiment (also known as forward blocking; Kamin, 1968), a cue, A, is trained as a predictor of an outcome (O) during the first phase of the study (i.e., A $\rightarrow$ O). If another cue, X, is trained in compound with A as predictor of the same outcome during Phase 2 (i.e., AX $\rightarrow$ O), responding to X at test will be weak as compared to responding in a control group which has not received the A $\rightarrow$ O trials of Phase 1.

According to traditional learning theories, the previously acquired A-O association interferes with the acquisition of the X-O association during Phase 2 (e.g., Rescorla & Wagner, 1972). According to response theories, however, the A-O association does not interfere with the acquisition of the X-O association, but with responding to X at test (e.g., Miller & Matzel, 1988). Thus, for example, if we invert the order of phases in a forward blocking experiment (i.e., backward blocking, which consists of AX $\rightarrow$ O training during Phase 1, followed by A $\rightarrow$ O training during Phase 2 and test on X, see Shanks, 1985), traditional acquisition theories will predict that the A $\rightarrow$ O training that occurs during Phase 2 cannot retrospectively affect the status of the X-O association that was already acquired during Phase 1, whereas response theories predict that the A $\rightarrow$ O training will still affect responding to X, regardless of the order in which A and X are trained.

Therefore, backward blocking, as well as several other retrospective revaluation effects in which the response potential of a given cue is altered on trials in which the cue itself is no longer present, has been regarded as one of the greatest challenges for traditional learning theories (see, e.g., Shanks & Dickinson, 1987; Wasserman & Berglan, 1998). Although backward blocking is a rather elusive effect (see Larkin, Aitken, & Dickinson, 1998; Miller, Hallam, & Grahame, 1990, for discussion), there are reports that show that it can be obtained, at least under particular conditions, in both humans (Shanks, 1985; Wasserman & Berglan, 1998) and other animals (Denniston, Miller, & Matute, 1996; Miller & Matute, 1996). Moreover, there are other related effects that also imply the retrospective processing of absent cues that, although not ubiquitous, have been reported in a variety of situations and species (Chapman, 1991; Dickinson & Burke, 1996; Kaufman & Bolles, 1981; Larkin et al., 1998; Matzel, Schachtman, & Miller, 1985; Van Hamme & Wasserman, 1994; Wasserman & Berglan, 1998). These results would have been taken until quite recently as clear evidence in favor of response, rather than acquisition theories of cue interference because, as previously mentioned, traditional acquisition theories were not equipped to deal with learning of absent cues. However, traditional theories of acquisition such as Rescorla and Wagner's (1972; hereafter R-W model) and Wagner's (1981) SOP model, have recently been modified so as to be able to account for cue interference effects that imply retrospective revaluation (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). Thus, nowadays, either acquisition or response theories can, once again, explain most of the available data on cue interference and retrospective revaluation effects. Nevertheless, and regardless of whether cue interference can better be explained as an acquisition or as a retrieval effect, both types of theories do agree in that compound training of the two cues is necessary for interference to take place. That is, presumably, only cues that are trained in compound with each other as predictors of the same outcome can interfere with the acquisition or with the retrieval of the association between the other cue and the same outcome (e.g., Larkin et al., 1998; Miller & Matzel, 1988; Wasserman & Berglan, 1998).

This prediction was recently challenged by Matute and Pineño (1998a, 1998b), who observed that, under certain circumstances, cues could interfere with each other even if they had not received compound training. In their experiments, two cues, A and X, were separately trained as predictors of the same outcome in different phases of the study. The design was similar to that of a blocking experiment except that the two cues did not receive compound training. That is, in the forward condition, subjects received A → O trials during Phase 1 followed by X → O trials during Phase 2 (instead of the typical forward blocking training consisting of A → O followed by AX → O training) and were then tested on X. In the backward condition the order of phases was reversed so that subjects received X → O followed by A → O (instead of the typical backward blocking training consisting of AX → O followed by A → O) and were also tested on X. The results showed that in the backward condition, the capacity of X to predict the outcome at test was weak as compared to several control groups that did not receive the critical A → O training during Phase 2.

This effect of interference between individually trained cues is hard to explain in the framework of current theories because all of them assume the necessity of compound training. In order to explain interference between cues in the absence of compound training we could assume, for example, that the context in which A and X are trained plays the role of a cue that is trained in compound with X. In that case, the design would be similar to a traditional (i.e., compounded) blocking design. From this perspective, traditional theories of acquisition, such as R-W and SOP models, would expect blocking by context in the forward, but not in the backward, condition (recall that these theories do not allow for retrospective processing of information), which is contrary to the results obtained. On the other hand, response theories would predict a similar blocking by context effect in the forward and backward conditions, which is also contrary to the results, which showed interference in the backward but not in the forward condition.

The revised R-W (Van Hamme & Wasserman, 1994) and SOP models (Dickinson & Burke, 1996) also predict blocking in both the forward and the backward conditions, but the mechanisms suggested for each case are different. Thus, there is a possibility that the mechanism responsible for backward blocking could produce an effect of greater magnitude than the one responsible for forward blocking. In the forward condition, the mechanism should be identical to that predicted by the original R-W and SOP models; that is, forward blocking by context. In the backward condition, however, the predictions of these two theories differ from each other. According to the revised R-W model, in the backward condition (i.e., X → O followed by A → O and test on X), the common context in which X and A are trained can activate the representation of the absent cue, X, during the A → O trials of Stage 2. According to this model, the repeated activation of an absent cue, X, which is followed by the occurrence of the outcome, will result in the unlearning of the X-O association. This could in principle explain the weak responding to X observed at test by Matute and Pineño (1998a) in the backward condition. Thus, in order to test this potential explanation, Matute and Pineño (Experiment 3) manipulated the contexts in which the different phases of the study took place. Presumably, unlearning of the X-O association should occur in the backward condition only if A were trained in the same context in which X had been trained. However, weak responding to X was only observed when X was tested in the context in which A had been trained, regardless of whether A and X had been trained in the same or in different contexts. Moreover, even when A and X were trained in the same context, good responding to X was still observed as long as X was tested in a different context, which suggested that the X-O association had not been unlearned.

Thus, neither blocking by context, nor unlearning of X seemed to provide an explanation for the effect of interference between elementally trained cues. Instead, Matute and Pineño (1998a) suggested that the effect of interference in the absence of compound training is a retrieval effect that occurs when the interfering association, A-O, is more strongly activated than the target one, X-O, in the context in which the test of X takes place. Although the mechanism through which this occurs is not yet clear, it seems that, during testing, the target cue, X, has difficulty in activating the outcome representation if the outcome representation is already being activated by a different cue (i.e., A), and this interfering link (i.e., A-O) is more strongly activated in memory (e.g., because of recency or because of contextual similarity between the test context and the context in which the interfering association was trained).

There might be, however, at least one alternative way to explain those results (see Matute & Pineño, 1998b for discussion on some alternative possibilities). Wagner's revised SOP model (Dickinson & Burke, 1996; Larkin et al., 1998) could in principle account for the results of the backward condition by assuming that the representation of the absent cue, X, is activated by the common context during Stage 2 and that this causes X to acquire an inhibitory association with the outcome that occurs during the A → O trials of Stage 2. That is, according to this model, the X-O association that is learned during Phase 1 is not unlearned during Phase 2; however, X could acquire a new, inhibitory association during Phase 2, provided that its representation was activated by the common context and became associated with the outcome. This acquisition of inhibitory strength by X could explain the weak responding observed to X at test in the experiments in which A and X were trained in the same context (Experiments 1 and 2 in Matute and Pineño, 1998a), but it cannot explain the occurrence of the interference effect when A and X were trained in different contexts or the lack of interference that was obtained by simply changing the context in which X was tested (Experiment 3). However, according to Matute and Pineño (1998b), it should be possible to extend the revised SOP model to account for those results if we assume that inhibitory associations generalize less readily across different contexts than excitatory associations, and that, if a given cue acquires two opposite associations with the same outcome, the excitatory association will tend to interfere with the expression of the inhibitory association unless testing occurs in the context in which the inhibitory association was acquired (Bouton, 1993). This could explain why good responding to X occurred when X was tested in a context different from that in which A had been trained (which is the context in which the inhibitory association between the absent cue, X, and the outcome had presumably been acquired) as well as why weak responding to X was observed when X was tested in the context in which A had been trained. Nevertheless, this would still be insufficient to explain how an inhibitory association to X was acquired during Stage 2 in those cases in which A was trained in a context which did not have a within-compound association with X. But in any case, this idea deserves attention because it is, to our knowledge, the only theory that can reasonably explain almost all of the data available on this effect in terms of the acquisition of new knowledge about X when it was no longer present (and subsequent interference during testing) rather than as mere interference at the retrieval stage (see Matute & Pineño, 1998b, for further discussion).

One important problem of accepting this theory as an explanation for this effect, however, is that, to our knowledge, there are no published reports that have tested for such inhibitory learning of absent cues that were not trained in compound with the cue that is physically present. Moreover, our own attempts to provide a direct test for this view have been unsuccessful. For this reason, in the present research, we decided to use an alternative strategy to test it. According to the revised-SOP model, the strength of the inhibitory association that is presumably acquired by the absent cue during Phase 2 should be weaker as the number of A → O trials is reduced. Thus, if we give only one A → O trial during Phase 2, no inhibitory X-O association should be acquired and the effect should not occur.

Interestingly, some prior experiments in the literature suggest that the effect could still be obtained with just one A → O trial. For example, Siddle and his colleagues (Lipp, Siddle, & Dall, 1993; Packer & Siddle, 1989; Siddle, Broekhuizen, & Packer, 1990) have reported evidence of a *miscuing effect* which is similar to the effect observed by Matute and Pineño (1998a, 1998b) and that occurs with just one Phase 2 trial. In the miscuing experiments subjects typically receive X → O/A → no O discrimination training during Phase 1, and then, in Phase 2, subjects receive one trial in which the outcome is miscued by the cue that predicts its absence (i.e., one A → O trial in Phase 2). The outcome used by Siddle and his colleagues in these studies is normally a loud noise, and the miscuing effect has been evidenced as weak outcome expectation when X is tested after the A → O trial as well as through several other dependent variables. These results appear consistent with those obtained by Matute and Pineño with more Phase 2 trials, although they used a different preparation and dependent variable. Thus, it is possible that the one-trial effect reported by Siddle and his colleagues was specific to their preparation.

For this reason, in Experiment 1 we tested whether the miscuing effect with just one trial reported by Siddle and his colleagues (e.g., Lipp et al., 1993) could also be obtained using the preparation that had been used in Matute and Pineño's studies. In Experiment 2 we showed that this one-trial effect is not due to backward blocking by context. In Experiment 3 we tested both the miscuing effect and the effect reported by Matute and Pineño (i.e., using a neutral cue rather than a cue that predicted the absence of the outcome during Phase 2), using the same preparation as Matute and Pineño but with only one Phase 2 interfering trial. Using that preparation, Arcediano, Ortega, and Matute (1996) showed that it takes between four and eight trials of an A → O/B → no O discrimination to establish an excitatory association between A and the outcome, and Ortega and Matute (1999) have shown that a retardation test shows inhibition to cue B after 25 trials but not after 12 trials. Moreover, according to the SOP model (Wagner, 1981), inhibitory associations are normally slower to acquire than excitatory associations. Thus, we can safely assume that with this preparation, one trial is not sufficient to establish an inhibitory association between an absent cue and the outcome. Hence, if the formation of an inhibitory association between X and the outcome were responsible for the effect of interference between elementally trained cues reported by Matute and Pineño (1998a, 1998b), this effect should not be observed in the present experiments.

EXPERIMENT 1

The purpose of this experiment was to study whether the miscuing effect obtained by Siddle and his colleagues with just one interfering trial (e.g., Lipp et al., 1993) can generalize across different tasks and dependent variables. To these ends, we used the same procedure that had been used in the experiments described by Arcediano et al. (1996) and Matute and Pineño (1998a, 1998b).

Method

Subjects

Fifty-nine undergraduate students volunteered for the study. Random assignment of subjects resulted in 27 subjects in group Miscuing and 32 in group Control.

Apparatus

The study was run using personal computers, with subjects responding via keyboard. All subjects were run simultaneously. Subjects were separated from each other by about 1.5 m and each subject was exposed to a different experimental condition (and counterbalancing of stimuli) than the two subjects sitting next to him or her.

The experiments were superimposed on a video game in which the task of the subjects is to shoot a laser gun (the space bar) at Martians that are trying to invade the Earth (see Arcediano et al., 1996). In this preparation, Martians appear, one at a time, every 0.2 s, which gives a behavioral baseline of about five responses per second. Subjects are also told that, from time to time, the Martians may connect an antilaser shield (evidenced through a white flashing of the screen that lasts 0.4 s) and that, if they shoot while the shield is connected, a Martian's invasion will immediately occur (i.e., if a response is made during the 0.4-s white flashing screen, this white flashing remains on while Martians invade the screen during 5 s during which the subject is unable to do anything to stop their landing). Therefore, the subjects normally suppress their bar-pressing behavior when the white shield (outcome) is connected, even though Martians continue landing at their regular 0.2-s intervals while the shield is connected. Thus, the outcome is given motivational value through instructions and punishment.

Changes in the background color of the screen are used as cues that sometimes signal that the shield (O) is about to be connected and sometimes signal nothing. These changes occur while the Martians continue landing on the screen at 0.2-s intervals. For example, a change from black (the default color of the screen during the intertrial intervals) to blue may signal that the shield is about to be connected, whereas a change to yellow may signal nothing. During the intertrial intervals and during the presentation of the cues, bar-presses are reinforced with Martians exploding. However, bar-presses that occur while the white shield (O) is connected are punished with an invasion. Acquisition of a cue–outcome expectation is evident when subjects suppress their bar-pressing behavior in the presence of the cue that predicts the outcome.

Therefore, this preparation might be described as an anticipated discriminated punishment procedure (only responses that occur during the white flash are punished; but some cues predict when the white flash will appear and subjects end up suppressing their bar-pressing behavior not only when the white flash appears but also when the predictive cues are present). The suppression of bar-pressing that occurs while a predictive cue is present can be regarded, therefore, as an anticipatory reaction to the passive avoidance to the white flash that is required by the game instructions.

Procedure

After subjects read the instructions, the experimenter performed the pretraining phase on a computer that was connected to a large screen that could be seen by all subjects. Any questions were answered aloud. Next, each subject completed the pretraining phase. This phase consisted of 100 presentations of Martians. The purpose of this phase was to teach each subject to bar-press consistently using the space bar of the computer (see Arcediano et al., 1996). No cues or outcomes were presented at this stage. The instructions for this phase were as follows:

> Your task is to prevent Martians from landing. Every half second one new Martian will try to land. In order to destroy them, you must use your laser-gun (the space bar) before they can see you, that is, just before you can see them. But do not shoot too early because you only have one shot per Martian. At the end of this phase, we will tell you the percentage of Martians that you have killed. THE PLANET DEPENDS ON YOU!! DO NOT ALLOW THEM TO LAND!!

When all subjects had finished the pretraining phase, the following instructions were presented.

> Now the Martians have developed a powerful anti-laser shield. You must continue using your laser to prevent their landing. BUT BE CAREFUL because if you shoot your laser-gun when the shield is connected, your shot will reflect back to you, thousands of Martians will land safely immediately, and you no longer will be able to avoid that invasion. You will know that the shield is connected when you see a WHITE INTERMITTENT FLASHING on the screen.

As in the pretraining phase, the experimenter explained these instructions using a large screen on a demonstration variant of the program. In this case, two trials consisting of one presentation of the outcome alone (i.e., with no cues) were performed by the experimenter. The purpose of this demonstration was to show the difference between a ‘‘short’’ outcome (i.e., that in which there is no punishment and just the white shield is presented for 0.4 s) and a ‘‘long’’ outcome (i.e., that in which a response is made during the 0.4-s white shield presentation and thus, it is followed by a 5-s Martian invasion which is no longer avoidable). When this demonstration was complete, the following instructions were presented:

> Some INDICATORS will help you predict when the shield is about to be connected, but there will also be some irrelevant cues. If you learn to distinguish between the valid and the irrelevant indicators, you will always be able to avoid the shield. Otherwise, every time that the Martians connect the shield (white flashing light) they will catch you shooting and thousands of them will invade. Remember, just one shot while the shield is connected (WHITE FLASH) and the Martians will invade.

TABLE 1
Design Summary of Experiment 1

Group	Treatment		Test
	Phase 1	Phase 2	
Miscuing	X → O, A → no O	A → O	X
Control	X → O, A → no O	A → no O	X

Note. A and X were yellow and blue, counterbalanced. Presentations of X were always followed by the outcome, whereas presentations of A were never followed by the outcome, except in the Phase 2 trial received by the Miscuing group.

Immediately after these instructions, the critical phases of the experiment occurred. Cues X and A were blue and yellow backgrounds, counterbalanced. The duration of the cues was 1 s. The background color remained black during the intertrial intervals (ITI). The duration of the ITI was pseudorandom with a range between 5 and 10 s and a mean of 7.5 s.

The design summary is presented in Table 1. Two groups of subjects were exposed during Phase 1 to X → O and A → no O trials. There were 12 trials for each cue. Their distribution was pseudorandom and identical for all subjects, starting with an A trial, and with no more than three trials of the same type in succession. After this phase was complete, Phase 2 consisted of one trial of A followed by O (i.e., A → O; miscuing trial) in the Miscuing group and one trial of A without being followed by O (i.e., A → no O) in the Control group. After that, the test phase was presented to all subjects and consisted of one presentation of X. In the test trial the cue was presented during 2 s. The different phases were presented without interruption and were separated from each other by a regular ITI.

Preanalysis treatment of the data. A suppression ratio of the form a/(a + b), where *a* is the number of responses during the presentation of the cue, and *b* is the number of responses in an equal period of time immediately preceding the occurrence of the cue, was computed for the test trial with X and was used as the dependent variable. A suppression ratio of zero would indicate that the subject has totally suppressed responding during the presentation of the cue, and hence, this indicates a high outcome expectation. By contrast, a suppression ratio of .5 indicates that the subject does not expect the outcome to occur when the cue is presented (the number of responses during the presentation of the cue and before the presentation of the cue is identical).

A selection criterion was used in order to ensure that subjects were paying attention to the experiment and had acquired the discrimination during Phase 1. According to this criterion, the suppression ratio to the nonreinforced cue (i.e., A) during the 12th trial in which it was presented in Phase 1 had to be higher (i.e., weaker conditioning) than the suppression ratio on the 12th trial for the reinforced cue in Phase 1 (i.e., X). Even though this criterion is quite light, we sometimes find that a high percentage of subjects has to be eliminated (particularly when we test too many subjects simultaneously, which probably increases distractions). In the present experiment, 59 subjects were run simultaneously, and 14 of them did not meet the selection criterion and were eliminated from the analyses (5 of them from the Miscuing group and 9 from the Control group).

Results and Discussion

Weak suppression of bar-pressing to X at test was observed in the Miscuing group as compared with its control. Thus, a clear miscuing effect was found in this experiment.

Mean suppression ratio for the target stimulus, X, was .32 (SE = .03) for the Miscuing group and .22 (SE = .02) for the Control group. Thus, control subjects suppressed more to X than miscuing subjects during the test phase, $t(43) = 2.55$, $p < .05$. This indicates that a miscuing effect was obtained during the presentation of X. This result replicates and extends the results of Lipp et al. (1993), who did observe a weak expectation of the outcome when X was presented. The present result shows that miscuing can also be evidenced through weak behavioral suppression to a cue that had been paired with the outcome during Phase 1. This also extends the results reported by Matute and Pineño (1998a, 1998b) by showing that a very similar effect can be obtained with their preparation using just one Phase 2 interfering trial.

EXPERIMENT 2

Experiment 1 replicated and extended the miscuing effect reported by Lipp et al. (1993), Packer and Siddle (1989), and Siddle et al. (1990), who had used a single trial during Phase 2. This suggests that the acquisition of an inhibitory association between the absent cue and the outcome is not an explanation for that effect.

Nevertheless, and even though Matute and Pineño (1998a) had discarded an explanation of their experiments in terms of backward blocking by context (see the introduction), there is a possibility that backward blocking by context could explain the results of the present experiments, in which only one Phase 2 trial is used. Thus, in order to provide additional evidence about the potential role of the context in this situation, we conducted an experiment that assessed directly the possible backward blocking by context explanation that could be derived from theories such as those proposed by Dickinson and Burke (1996), Miller and Matzel (1988), or Van Hamme and Wasserman (1994). According to these theories, the strength of the context-O association could increase during the A $\rightarrow$ O trial in Phase 2, and therefore, backward blocking of X by context is possible. Thus, in Experiment 2 we tested this hypothesis by adding a new group in which during Phase 2, the outcome was presented once without being preceded by a punctual cue.

Method

Subjects

Eighty-one undergraduate students volunteered for the study. Random assignment of subjects resulted in 28 subjects in the Miscuing group, 28 in the Context group, and 25 in the Control group.

Apparatus

For 36 subjects the apparatus was the same as in Experiment 1. The experiment was run for the remaining 45 subjects in a smaller laboratory that allowed for simultaneous testing of 5 subjects at a time in individual cubicles. In both testing conditions, the subjects were randomly distributed across the three experimental groups.

Procedure

The procedure was the same as in Experiment 1, except as noted otherwise. During the assessment trials used to ensure that subjects had acquired the discrimination of Phase 1 (i.e., the 1st and the 12th trials of Phase 1) the cues were presented over 2 s (instead of 1 s) in order to provide for a more reliable selection criterion.

Phase 1 was identical to that of Experiment 1 for all three groups. Then, during Phase 2, the Miscuing group received one A → O trial, the Context group received one trial in which the outcome was presented alone (i.e., without any cue preceding it), and the Control group received one A → no O trial. At test, all three groups were exposed to one presentation of X that lasted for 3 s (instead of 2 s). Table 2 summarizes this design.

TABLE 2
Design Summary of Experiment 2

	Treatment		
Group	Phase 1	Phase 2	Test
Miscuing	X → O, A → no O	A → O	X
Context	X → O, A → no O	O	X
Control	X → O, A → no O	A → no O	X

Note. A and X were yellow and blue, counterbalanced. Presentations of X were always followed by the outcome, whereas presentations of A were never followed by the outcome, except in the Phase 2 trial received by the Miscuing group. The Context group received one presentation of the outcome alone during Phase 2.

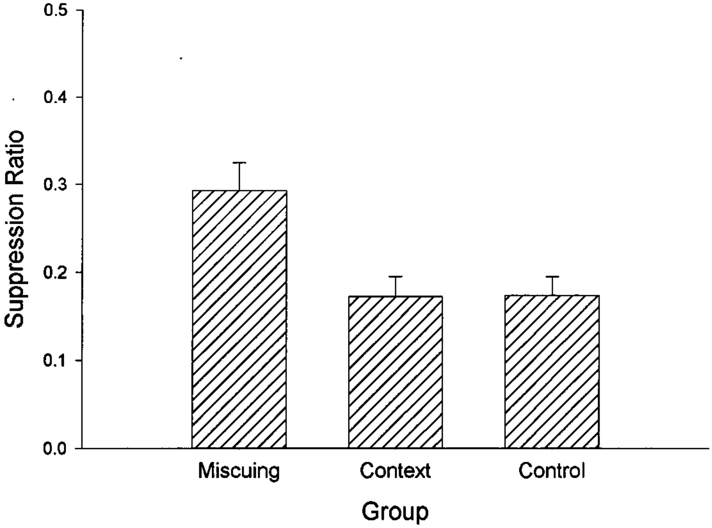

FIG. 1. Outcome expectation during the test trial of Experiment 2. The dependent variable was mean suppression ratio. Thus, a lower value represents stronger outcome expectation. Error bars represent standard errors of means.

Preanalysis treatment of the data. The data from five subjects (three from the Miscuing group, one from the Context group, and one from the Control group) were eliminated from the analysis according to the data-selection criterion described in Experiment 1.

Results and Discussion

This experiment replicated the difference obtained in Experiment 1 between the Miscuing group and the Control group. Moreover, a similar difference was found between the Miscuing group and the Context group. Therefore, cue A, which had been previously trained as a predictor of the absence of the outcome interfered with X after being paired once with the outcome in the Miscuing group, and this was not due to backward blocking by context. Both the Context group and the Control group showed good suppression to X at test. These results are depicted in Fig. 1.

A one-way ANOVA on suppression ratios during the test phase showed an overall group effect, $F(2, 73) = 7.22$, $p < .01$, and pairwise comparisons confirmed that interference occurred not only in the Miscuing group as compared to the Control group, $F(1, 73) = 10.44$, $p < .01$, but also in the Miscuing group as compared to the Context group, $F(1, 73) = 11.29$, $p < .01$. No differences were observed between the Context and the Control groups $F(1, 73) = 0.001$, $p > .05$.

In Experiment 1 we had recorded suppression ratios directly, and hence, there was a possibility that differences in context conditioning and base rate behavior could be responsible for the observed results. Thus, in the present experiment we assessed the number of responses during the precue interval at test, in addition to suppression ratios, in order to allow assessment of this potential problem. An analysis of the number of responses during the 3-s precue interval at test showed similar response rates for the Miscuing group ($M = 12.92$, $SE = 0.47$), the Context group ($M = 12.22$, $SE = 0.57$), and the Control group ($M = 12.33$, $SE = 0.62$), $F(2, 73) = 0.43$, $p > .05$, thereby indicating that differences during the precue interval were not responsible for the differential suppression to X at test.

Therefore, the results of this experiment suggest that backward blocking by context cannot provide an explanation for the miscuing effect.

EXPERIMENT 3

Experiment 1 replicated and extended the miscuing effect reported by Siddle and his colleagues (Lipp et al., 1993; Packer & Siddle, 1989; Siddle et al., 1990) using just one Phase 2 trial. Experiment 2 added evidence against the explanation of this effect as backward blocking by context. In principle, there is a possibility that the miscuing effect could be explained as a rule-learning effect in which subjects learn that the contingencies change from one phase to another. However, this could not account for the results of Matute and Pineño (1998a) in which a neutral cue, rather than an inhibitory cue, was what predicted the outcome during Phase 2. Thus, the purpose of the present experiment was to test whether the one-trial effect observed in Experiments 1 and 2 could also be obtained when a neutral cue, rather than an inhibitory cue, was used during Phase 2 to predict the outcome. If rule learning were responsible for the miscuing effect, we should still be able to replicate the miscuing effect in this one-trial experiment, but not Matute and Pineño's results.

Therefore, in this experiment, we tested whether a single trial with a new, neutral cue, introduced in the second phase had the same influence as a single trial with a stimulus that had been previously trained as a predictor of the absence of the outcome. For this purpose we added a new group which, during Phase 2, received a new stimulus followed by O (e.g., B $\rightarrow$ O) instead of A $\rightarrow$ O. In the control group we changed the second phase and instead of giving one trial of A without the outcome, we gave no cues or outcomes (subjects simply kept shooting at Martians during the corresponding period of time). This control condition was used in order to make sure that the results of Experiments 1 and 2 were not due to the use of an A $\rightarrow$ no O control, which according to some theories (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994) could result in increased responding to X and consequently would not be a neutral control condition.

TABLE 3
Design Summary of Experiment 3

Group	Treatment		Test
	Phase 1	Phase 2	
Miscuing	X → O, A → no O	A → O	X
New Cue	X → O, A → no O	B → O	X
Control	X → O, A → no O	—	X

Note. A, X, and B were yellow, blue, and purple, counterbalanced. Presentations of B and X were always followed by the outcome, whereas presentations of A were never followed by the outcome, except in the Phase 2 trial received by the Miscuing group. The Control group did not receive any cue or outcome during Phase 2 but was exposed to the experimental context for the corresponding amount of time.

Method

Subjects and Apparatus

Fifty-three students volunteered for the study. None of them had participated in Experiment 1 or 2. Random distribution of subjects resulted in 18 subjects in the Miscuing group, 14 subjects in the New Cue group, and 21 subjects in the Control group. The apparatus was the same as in Experiment 1.

Procedure

The procedure was identical to that in Experiment 2, except otherwise noted. The duration of the ITI was pseudorandom with a range between 6 and 12 s and a mean of 9 s. The three cues used in this experiment were blue, yellow, and purple backgrounds, counterbalanced. Phase 1 was identical to that of Experiment 1 for all three groups. Then, during Phase 2, group Miscuing received one A → O trial, group New Cue received one trial in which a new cue was paired with the outcome (i.e., B → O), and the Control group received no cues or outcomes (i.e., this group kept shooting at Martians during that time). At test, all three groups were exposed to one presentation of X. Table 3 summarizes this design.

Preanalysis treatment of the data. Even though we used assessment trials that lasted 2 s rather than 1 s to ensure that subjects had acquired the discrimination during Phase 1, the data from 18 subjects (7 from the Miscuing group, 1 from the New Cue group, and 10 from the Control group) had to be eliminated from the present analyses according to the data selection criterion described in Experiment 1. As in Experiment 1, the high number of subjects eliminated probably reflects that many subjects do not pay enough attention to the experiments when tested in large groups.

Results and Discussion

A clear stimulus interference effect was found in the Miscuing and in the New Cue groups as compared to the Control group. Thus, as in Experiments 1 and 2, cue A, which had previously been trained as a predictor of the absence of the outcome, interfered with X after being paired once with the outcome in the Miscuing group. Moreover, a neutral cue, B, associated with the outcome during Phase 2, yielded a similar impairment on responding to X in the New Cue group. The Control group showed good suppression to X in the test trial. These results are depicted in Fig. 2.

A one-way ANOVA on suppression ratios during the test phase showed an overall group effect, $F(2, 32) = 6.19$, $p < .05$, and pairwise comparisons confirmed that interference occurred not only in the Miscuing group as compared to the Control group, $F(1, 32) = 5.66$, $p < .05$, but also in the New Cue group as compared to the Control group $F(1, 32) = 11.93$, $p < .05$. No differences were observed between the Miscuing and the New Cue groups, $F(1, 32) = 0.95$, $p > 0.1$.

An analysis of the number of responses during the 3-s precue interval at test showed similar response rates for group Miscuing ($M = 14.18$, $SE = 0.37$), group New Cue ($M = 14.07$, $SE = 0.49$), and group Control ($M = 14.27$, $SE = 0.35$), $F(2, 32) = 0.05$, $p > 0.5$, thereby indicating that differences during the precue interval were not responsible for the differential suppression to X at test.

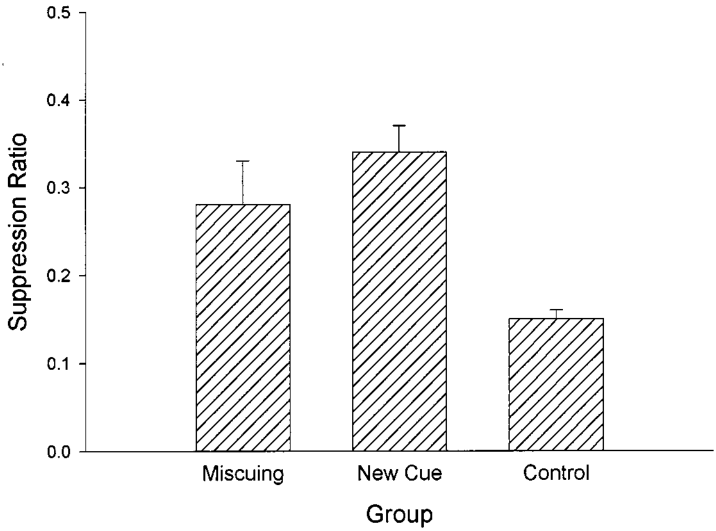

FIG. 2. Outcome expectation during the test trial of Experiment 3. The dependent variable was mean suppression ratio. Thus, a lower value represents stronger outcome expectation. Error bars represent standard errors of means.

Thus, the results of this experiment show that the miscuing effect is not an isolated effect, but rather, it is an instance of the more general effect of interference between individually trained stimuli. Most importantly, the results show that interference between elementally trained cues can occur in just one trial, regardless of whether a neutral cue or a cue which predicts the absence of the outcome is used in that trial to predict the outcome.

GENERAL DISCUSSION

Matute and Pineño (1998a, 1998b) showed that cues that are elementally trained to the same outcome in different phases can interfere with each other, and the miscuing studies reported by Siddle and his colleagues (Lipp et al., 1993; Packer & Siddle, 1989; Siddle et al., 1990) suggested that this interference effect could occur in just one trial, at least when the outcome was miscued, during the interfering trial, with a cue that predicted its absence. The results of Experiment 1 show that the one-trial miscuing effect can also be observed in a preparation different from that used by Siddle and his colleagues, and the results of Experiment 2 allow us to discard backward blocking by context as an explanation for this effect. In addition, the results of Experiment 3 suggest that the miscuing effect is a special case of a more general phenomenon of interference between elementally trained cues in that it is not necessary to miscue the outcome with a cue that predicts its absence in order to obtain the effect using just one interfering trial. Indeed, the effect can be observed even when a neutral cue is what is paired just once with the outcome during Phase 2 (Experiment 3).

As mentioned in the introduction, the results of interference between elementally trained cues cannot be explained by traditional theories of learning (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981), response theories (e.g., Miller & Matzel, 1988), or retrospective revaluation theories (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). However, Matute and Pineño (1998b) suggested that, given certain assumptions, the modified Wagner's (1981) SOP model (Dickinson & Burke, 1996; Larkin et al., 1998) could account for those results. More specifically, if we assume that an inhibitory association is formed between X and the outcome during the interfering A-O trials of Phase 2 in which X is absent (as long as X and A are trained in the same context; Dickinson & Burke, 1996), and if we also assume that inhibitory associations generalize less readily across contexts than excitatory associations (Bouton, 1993), then we could explain both the impaired responding to X during testing that occurs when X and A are trained in the same context as well as the absence of impairment observed when X is tested in a context that is different from that in which A was trained (see Matute & Pineño, 1998a, 1998b).

In order to test this view, we used only one interfering trial during Phase 2 in the present experiments. Larkin et al. (1998) suggested that at least more than three interfering trials should be needed to establish an inhibitory association between an absent cue and an outcome (p. 1350). Although they used a different preparation, inhibitory associations are generally assumed to be slower to acquire than excitatory ones (e.g., Wagner, 1981), and as previously mentioned, Arcediano et al. (1996) showed that with the preparation that we used in the present experiments, between four and eight trials were needed to establish an excitatory association between a cue and an outcome. Thus, for all these reasons, we can be confident that no inhibitory association should be responsible for the effect that we observed with just one trial in the present experiments. Interference was observed with just one trial, regardless of whether the interfering Phase 2 trial consisted of pairing a neutral cue with the outcome or a cue that predicted that the outcome would not occur. These results, along with the observation that interference could occur even when A and X were trained in different contexts (Matute & Pineño, 1998a, Experiment 3), suggest that the formation of inhibitory associations between an absent cue and the outcome is probably not the best explanation for this effect.

The hypothesis suggested by Matute and Pineño (1998a, 1998b) to explain their results could still account for the present ones. According to this hypothesis, this effect occurs during the retrieval stage. If the interfering association is more strongly activated in memory during testing than the target association (e.g., because of recency or because of contextual manipulations), it will interfere with the retrieval of the target association during testing, thus producing weak responding to X. This could explain why interference was observed in the present experiments, even though the associative strength of the interfering association was necessarily weaker than that of the target association. The contextual effects described above (i.e., interference occurs even when A and X are trained in different contexts but only if X is tested in the context in which A was trained), as well as a recent research by Pineño, Ortega, and Matute (2000), also add support to the view that the effect takes place during retrieval, and only when the interfering association is more strongly activated than the target association. To test this view, Pineño et al. manipulated the relative activation of the two associations at testing in different ways: By inserting a retention interval between the interfering A $\rightarrow$ O trials and the test trial; by randomizing the X $\rightarrow$ O and A $\rightarrow$ O trials throughout the experiment rather than separating them in two different phases; and by inserting, just before testing, a novel cue that, like the retention interval, had the purpose of separating the test trial from the interfering Phase 2 trials. Thus, all these three manipulations had the purpose of preventing the A-O association from being more strongly activated than the target association at testing. And all these three manipulations prevented the occurrence of the interference effect that was otherwise observed. Those results suggest not only that the effect occurs at the retrieval stage, but also that it is probably a short-term memory effect. Quite possibly, when the target cue, X, is presented at testing, the interfering association, which is more strongly activated (because of recency or other reasons), and the need for a very rapid response (or suppression of responding in this case), makes it difficult to access the relevant information (in this case, the X-O association).

The present series of experiments was not specifically aimed to test whether the effect was a short-term memory effect or whether it was dependent, as shown by Pineño et al., upon the relative activation of the associations, but it allows us to discard an explanation that was, to our knowledge, the only one that could account for much of the data available on this effect in terms of acquisition [e.g., the effect of a retention interval could be explained as the acquisition of an inhibitory association during Phase 2 (Dickinson & Burke, 1996) which does not generalize when testing occurs after a retention interval (Bouton, 1993)]. This is important in that once the effect does not seem to depend on the acquisition of inhibitory associations, then the research questions can be aimed at clarifying the type of retrieval process that is responsible for the effect (e.g., whether it depends on the relative activation of the associations at testing or on other mechanisms).

Another question that still needs to be addressed in future research is why the strong activation of the A-O association during the test phase does not produce a response. Matute and Pineño (1998b) had suggested that one potential explanation was that an inhibitory association between X and A is acquired during Phase 2 (because when A is present, X is absent, see Espinet, Iraola, Bennett, & Mackintosh, 1995), and this inhibitory association then prevents responding to the representation of the absent cue, A, which is activated at testing when X is presented (Bennett, Scahill, Griffiths, & Mackintosh, 1999; Mackintosh & Bennett, 1997). However, the one-trial effect observed in the present experiments suggest that the formation of inhibitory associations between the absent cue X and another event (be it the outcome or cue A) during Phase 2 is not a plausible explanation for the effect of interference between elementally trained cues.

An alternative possibility is to assume that if the test context is strongly activating the A-O association [i.e., context-A-O or context-(A-O)] when X is presented at test, the occurrence of X might be processed as ‘‘no A.’’ If this were the case, a conflicting association (context—no A) might be activated, thus interfering with the occurrence of the response. Moreover, we could also assume that the representation of cues that are present in a given trial differ from the representations of cues that are absent (e.g., Wagner, 1981) and that subjects do not respond to cues that are absent unless they expect them to immediately occur, such as, for example, in sensory preconditioning experiments. The data available so far suggest that these are plausible explanations that need to be explored.

Our interpretation of the effect as interference for retrieval of the target association when the interfering association is more strongly activated at testing can be contrasted to a rule-based account, according to which subjects might be learning that only one cue can be associated with the outcome.

Thus, if they first learn that X is correct and then they learn that X is not necessary to predict the outcome, they might infer that X is no longer correct and this would be enough to observe weak suppression to X. However, several of the experiments that we have conducted allow us to discard this view. For example, the results of group Context in Experiment 2 suggest that the mere occurrence of the outcome in the absence of X is not sufficient to produce interference. Moreover, the retention-interval experiment mentioned above (Pineño et al., 2000), in which the interference effect was alleviated by the mere insertion of a retention interval before testing cannot be explained as a rule-learning effect and suggests that manipulating the relative activation of the associations at testing was what determined the presence or absence of the effect in that experiment. Apparently, if two cues are associated to the same outcome, and one of these cue–outcome associations is more strongly activated than the other one, the weaker association will not be retrieved or, at least, its retrieval will be impaired. Nevertheless, our hypothesis needs to be further developed and tested in order to see, for example, whether the same type of account can be applied to explain interference between compounded cues. As noted in the introduction, it is not yet clear whether the more traditional forms of interference between cues (i.e., compound training) are better explained as retrieval (e.g., Miller & Matzel, 1988) or as acquisition effects (e.g., Dickinson & Burke, 1996; Rescorla & Wagner, 1972; Van Hamme & Wasserman, 1994; Wagner, 1981).

The present and other related findings (e.g., Lipp, Siddle, & Dall, 1993; Matute & Pineño, 1998a, 1998b; Packer & Siddle, 1989; Siddle, Broekhuizen, & Packer, 1990) could seem to suggest that similar results should also be observed in animal research as well if one tested it carefully. However, it has been shown that other forms of retrospective revaluation using compound training (e.g., backward blocking) are hard to obtain in rats because, among other reasons, the stimuli that are used as outcomes in animal research are generally stimuli of high biological significance (i.e., unconditioned stimuli or USs) rather than the neutral stimuli used as outcomes in human research (Denniston et al., 1996; Miller & Matute, 1996). Following these ideas, Escobar, Matute, and Miller (in press) have recently tested whether the effect of interference between elementally trained cues reported here and elsewhere could also be obtained in rats using as outcomes either USs or neutral stimuli (i.e., through the use of sensory preconditioning). As expected, the interference effect was obtained in rats only when neutral stimuli, rather than USs, were used as outcomes.

Even though the present results might appear quite new and inconsistent with current theories, it is worth mentioning that older studies in the paired associate tradition had also shown interference between cues trained apart (Abra, 1967; Cheung & Goulet, 1968; Johnston, 1968; Keppel, Bonge, Strand, & Parker, 1971). Moreover, recent research in the interference literature is also showing that interference effects do not depend on the amount of Phase 2 training (e.g., Bäuml, 1996), which is consistent with the present results and also points to a retrieval rather than to an acquisition explanation. Those studies on paired associate learning have generally been overlooked by current theories of predictive learning, perhaps because they were generally concerned with language-related issues rather than with predictiveness. However, the observation of similar effects of interference between elementally trained cues in situations as different as paired-associate learning, miscuing effects using loud noises as outcomes, or the present research suggests that these are not isolated results and calls for a reformulation of theories that attempt to explain how we acquire and retrieve associative information and, more specifically, how cues interfere with each other to predict the events in our environment.

REFERENCES

Abra, J. C. (1967). Time changes in the strength of forward and backward associations. *Journal of Verbal Learning and Verbal Behavior,* **6,** 640–645.

Arcediano, F., Ortega, N., & Matute, H. (1996). A behavioural preparation for the study of human Pavlovian conditioning. *Quarterly Journal of Experimental Psychology B,* **49,** 270–283.

Bäuml, K. H. (1996). Revisiting an old issue: Retroactive interference as a function of the degree of original and interpolated learning. *Psychonomic Bulletin & Review,* **3,** 380–384.

Bennett, C. H., Scahill, V. L., Griffiths, D. P., & Mackintosh, N. J. (1999). The role of inhibitory associations in perceptual learning. *Animal Learning & Behavior,* **27,** 333–345.

Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin,* **114,** 80–99.

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

Cheung, C. C., & Goulet, L. R. (1968). Retroactive inhibition of R-S associations in the A-B, B-C, C-B paradigms. *Journal of Experimental Psychology,* **76,** 327–328.

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

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology B,* **49,** 60–80.

Escobar, M. C., Matute, H. & Miller, R. R. (in press). Cues trained apart compete for behavioral control in rats: Convergence with the associative interference literature. *Journal of Experimental Psychology: General*.

Espinet, A., Iraola, J. A., Bennett, C. H., & Mackintosh, N. J. (1995). Inhibitory associations between neutral stimuli in flavor-aversion conditioning. *Animal Learning & Behavior,* **23,** 361–368.

Johnston, W. A. (1968). Bidirectional interference in an A-B, C-B paradigm. *Journal of Verbal Learning and Verbal Behavior,* **7,** 305–311.

Kamin, L. J. (1968). ‘‘Attention-like’’ processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9–31). Miami, FL: Univ. of Miami Press.

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

Keppel, G., Bonge, D., Strand, B. Z., & Parker, J. (1971). Direct and indirect interference in the recall of paired associates. *Journal of Experimental Psychology,* **88,** 414–422.

Larkin, M. J. W., Aitken, M. R. F., & Dickinson, A. (1998). Retrospective revaluation of causal judgments under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **24,** 1331–1352.

Lipp, O. V., Siddle, D. A. T., & Dall, P. J. (1993). Effects of miscuing on Pavlovian conditioned responding and on probe reaction time. *Australian Journal of Psychology,* **45,** 161–167.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review,* **82,** 276–298.

Mackintosh, N. J., & Bennett, C. H. (1997). *Inhibitory associations in taste aversion conditioning.* IX Congreso de la Sociedad Española de Psicología Comparada, Salamanca (Spain).

Matute, H., & Pineño, O. (1998a). Stimulus competition in the absence of compound conditioning. *Animal Learning & Behavior,* **26,** 3–14.

Matute, H., & Pineño, O. (1998b). Cue competition in the absence of compound training: Its relation to paradigms of interference between outcomes. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 38, pp. 45–81). San Diego, CA: Academic Press.

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

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

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

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

Ortega, N., & Matute, H. (1999). *A parametric study of inhibition in a behavioral preparation with humans.* XI Meeting of the Spanish Society of Comparative Psychology, Baeza (Spain).

Packer, J. S., & Siddle, D. A. T. (1989). Stimulus miscuing, electrodermal activity and the allocation of processing resources. *Psychophysiology,* **26,** 192–200.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review,* **87,** 532–552.

Pineño, O., Ortega, N., & Matute, H. (2000). The relative activation of the associations modulates interference between elementally trained cues. *Learning and Motivation,* **31,** 127–151.

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

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology B,* **37,** 1–21.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H.

Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229–261). San Diego, CA: Academic Press.

Siddle, D. A. T., Broekhuizen, D., & Packer, J. S. (1990). Stimulus miscuing and dishabituation: Electrodermal activity and resources allocation. *Biological Psychology,* **31,** 229–243.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue interference in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation,* **25,** 127–151.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum.

Wasserman, E. A., & Berglan, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgement: The role of within-compound associations. *Quarterly Journal of Experimental Psychology B,* **51,** 121–138.

Received March 25, 1999
Revised February 15, 2000