Learning and Motivation 39 (2008) 313–322

Contents lists available at ScienceDirect

Learning and Motivation

journal homepage: www.elsevier.com/locate/l&m

Overshadowed cues have reduced ability to retroactively interfere with other cues ☆

Miguel A. Vadillo *, Cristina Orgaz, Helena Matute

Departamento de Psicología, Universidad de Deusto, Apartado 1, 48080 Bilbao, Spain

ARTICLE INFO

Article history:
Received 4 January 2008
Revised 19 May 2008
Available online 22 August 2008

Keywords:
Associative
Interference
Cue competition
Overshadowing

ABSTRACT

The present series of experiments explores the interaction between retroactive interference and cue competition in human contingency learning. The results of two experiments show that a cue that has been exposed to a cue competition treatment (overshadowing) loses part of its ability to retroactively interfere with responding to a different cue that was paired with the same outcome. These results pose problems for associative models of contingency learning and are also difficult to explain in terms of current theories of causal reasoning. Additionally, it is proposed that in light of the interaction between interference and cue competition, interference could be used as an indirect measure for the study of cue competition effects.

© 2008 Elsevier Inc. All rights reserved.

Cue competition is one of the main topics explored by current research on associative learning. Since the early studies performed by Pavlov (1927), it is well known that pairing a cue, Y, with an unconditioned stimulus results in lower conditioned responding to Y, if another cue, A, has also been presented in compound with Y during training. This cue competition effect is known as overshadowing. Blocking (Kamin, 1968) and relative validity (Wagner, Logan, Haberlandt, & Price, 1968) are similar instances of cue competition: in all of them a target cue-outcome association is poorly learned or expressed because of its having been trained in compound with an alternative cue-outcome association.

Interestingly, evidence coming from verbal learning studies performed with paired-associates tasks suggests that cue-interaction effects similar to overshadowing and blocking can occur even when the

☆ Support for this research was provided by Dirección General de Investigación of the Spanish Government (ref. SEJ2007-63691/PSIC), Departamento de Educación, Universidades e Investigación of the Basque Government (ref. PI2008-9 and apoyo a grupos de investigación), and Dirección General de Investigación, Tecnología y Empresa of the Junta de Andalucía (ref. SEJ-406). We would like to thank Fernando Blanco, Pedro Cobos, Paco López, David Luque, Ralph Miller, Serban Musca, and Oskar Pineño for illuminating discussions concerning this experiment, and Gumersida Alonso for her assistance in running Experiment 2.

* Corresponding author. Fax: +34 94 4139089.
E-mail address: mvadillo@fice.deusto.es (M.A. Vadillo).

0023-9690/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.lmot.2008.05.005

competing cues are never trained in compound (see Slamecka & Ceraso, 1960, for a review). For instance, retroactive interference studies show that after X-outcome pairings, training of a novel cue, Y, with the same outcome can reduce the retrievability of the X-outcome association. Recent studies have shown that this retroactive interference effect can also take place in human contingency learning (Cobos, López, & Luque, 2007; Matute & Pineño, 1998a, 1998b) and animal conditioning (Escobar, Matute, & Miller, 2001).

Available evidence suggests that this retroactive interference effect shares some properties with cue competition effects, such as overshadowing and blocking (for a review, see Miller & Escobar, 2002). Experiments conducted with nonhuman animals indicate that manipulations known to affect cue competition effects also have a similar impact on interference. For example, manipulations of the spatial or temporal cue-outcome relationship that are known to impair cue competition (Amundson & Miller, 2007; Barnet, Grahame, & Miller, 1993) have similar effects on interference (Amundson & Miller, 2007; Escobar & Miller, 2003). Similarly, Wheeler and Miller (2007) recently found that other treatments known to counteract cue competition, such as degraded contingency (Urcelay & Miller, 2006) and massed training (Stout, Chang, & Miller, 2003), have the ability to counteract interference as well.

Two recent studies performed with human contingency learning tasks (Escobar, Pineño, & Matute, 2002; Vadillo, Castro, Matute, & Wasserman, 2008) suggest that cue competition and interference might share some basic mechanisms. Both of these studies show that at least part of the response deficit observed in backward blocking can be attributable to retroactive interference. According to these results, the similarities between cue competition and interference might be due to their common mechanisms.

In the present series of experiments we aimed at further investigating the mechanisms of cue competition and interference by exploring the potential interaction between both effects. As we will show later, this interaction between overshadowing and interference would pose important constraints for models designed to account for one or the other effect. Moreover, on the basis of this interaction novel procedures for the investigation of cue competition can be developed.

Experiment 1

The purpose of Experiment 1 is to test whether a cue that has been exposed to an overshadowing treatment (probably, the simplest cue competition effect) is impaired to act as an interfering cue. As a means to explore this possibility, one cue, Y, received overshadowing treatment and, afterwards, a different cue, X, was paired with the same outcome as Y. If overshadowed cues lose part of their ability to interfere with other predictors of the same outcome, then additional presentations of cue Y should have little impact on responding to X, relative to a control group which did not receive overshadowing treatment.

Method

Participants and apparatus

Sixty-two undergraduate students from Deusto University volunteered to take part in the experiment. None of them had taken part in any related associative learning experiment. Random assignment of these participants resulted in 31 participants in each group. The experiment was performed in a large computer room, with each participant at least 1.5 m apart from the adjacent participant. The experimental task was implemented in a HTML document, using JavaScript functions to manage the presentation of the stimuli on the screen and to collect participants' responses.

Procedure and design

The experiment was conducted with a preparation and cover story that had already been used in several experiments performed in our laboratory (for a detailed description of the task and instructions, see Escobar et al., 2002). Participants were asked to imagine that they were soldiers whose task was to rescue some refugees that were hidden in a ramshackle building. On each trial, participants were given the opportunity to place a number of people in a truck and to take them to a safe place. They could enter people into the truck by pressing the space bar repeatedly. Participants could also enter larger numbers of people into the truck by keeping the space bar pressed instead of pressing it repeatedly. However, the refugees placed into the truck did not always arrive safely at their destination. In some trials, the road the truck had to move through contained mines that could explode. Participants could predict whether or not the road would be mined on a given trial by paying attention to a spy-radio installed in the truck, which appeared in the middle of the computer's screen. Certain colors in the spy-radio predicted that the road would be safe (and, therefore, that participants should place as many refugees as possible into the truck), whereas other colors predicted that the road would be mined (and, therefore, that participants should avoid placing refugees into the truck during those trials). Participants were not told which lights predicted which outcome but they could learn this by paying attention to what happened after the presentation of each color. Thus, the number of refugees placed in the truck on each trial when the light was on was taken as an index of the extent to which participants had learned that the cue (i.e., the color light) presented in that trial predicted that the road would be safe. Participants earned one point (a symbolic reward) for each refugee placed in the truck in the trials in which the road was safe and lost one point for each refugee placed in the truck in the trials in which the road was mined. Participants could only decide whether or not to put people in the truck and could enter their responses during the 3-s interval in which the color lights were present on each trial. The inter-trial interval was randomized, ranging from 3 to 7 s.

The design of the experiment is summarized in Table 1. During Phase 1, participants in the Experimental Group were exposed to two types of trials. In one of them, participants were presented with two simultaneous color lights in the spy-radio (playing the role of cues A and Y) which signaled the opportunity to earn points (outcome O1) by responding. The training conditions for participants in the Control Group were identical, except that only one color (cue Y) was presented in the trials in which participants could earn points by performing the response. During Phase 2, all participants were exposed to a novel trial type in which a new cue (X) was also followed by the opportunity to earn points. Finally, during Phase 3, all participants were exposed to presentations of cue Y alone signaling the opportunity to earn points. In order to avoid any tendency to respond indiscriminatively to all the cues, participants were also exposed to filler trials during phases 1–3. In these filler trials a different color (B) was presented and participants lost points (outcome O2) if they performed the response. At test, participants were again presented with the cue that had been trained in Phase 2, X. The number of trials presented on each training phase is shown in Table 1. Colors yellow, blue, red, and green played the role of cues A, B, X, and Y, counterbalanced in a latin square. Training proceeded without any disruption between phases, and all the trials within phases were presented in pseudorandom order.

Results and discussion

Before analyzing the results a data selection criterion was used in order to remove from the sample those participants who had not paid enough attention to the experimental task. As in other experiments we have performed with this task, we removed participants who, by the end of each training phase, failed to show appropriate discrimination between the cues that predicted the opportunity

Table 1
Design summary of Experiment 1

Group	Phase 1	Phase 2	Phase 3	Test
Experimental	4 AY-O1 2 B-O2	8 X-O1 4 B-O2	4 Y-O1 2 B-O2	1 X
Control	4 Y-O1 2 B-O2	8 X-O1 4 B-O2	4 Y-O1 2 B-O2	1 X

Note: Cues A–Y are color lights in the spy-radio that can predict two different outcomes: O1 (the participants gain points if they perform the response) and O2 (participants lose points if they perform the response). Numbers before each trial type denote the number of trials of that type that were presented in a given phase.

to earn points and the filler cue that predicted a lost of points if the response was performed. Specifically, we removed participants that responded equal or more during the last B-O2 trial of each phase than during the last reinforced trial of the respective training phase. Following this selection criterion data from seven participants in the Experimental Group and six participants in the Control Group were removed from subsequent analyses. Therefore, the following analyses were performed with 24 participants from the Experimental Group and 25 participants from the Control Group.

Mean responding of the remaining participants during training and at test is shown in Fig. 1. As expected, these participants learned to respond to the appropriate cues in few trials in each learning phase. A 2 (group) × 4 (trial) mixed analysis of variance performed on participants responses in AY/Y trials during Phase 1 showed a main effect of trial, $F(3,141) = 64.41$, $p < .001$, but no effect of group, $F(1,47) = 1.05$, $p = .31$, and no group × trial interaction, $F(3,141) = 1.002$, $p = .39$. Similarly, a 2 (group) × 8 (trials) mixed ANOVA performed on responses in X trials during Phase 2 showed a main effect of trial, $F(7,329) = 77.82$, $p < .001$, but again the main effect of group and the group × trial interaction were nonsignificant, $F(1,47) = 524.81$, $p = .19$, and $F(7,329) = 0.558$, $p = .79$. These analyses show that learning proceeded uneventfully in both groups and without significant differences between them.

Most importantly, responding to cue Y in the first trial of Phase 3 is significantly lower in the Experimental Group, which suggests that cue Y was successfully overshadowed during the Phase 1 compound training. A 2 (group) × 4 (trial) mixed ANOVA performed on participants' responses during the four Y-O1 trials in Phase 3, yielded a main effect of trial, $F(3,141) = 16.88$, $p < .001$, and a significant trial × group interaction, $F(3,141) = 6.11$, $p < .005$; the main effect of group failed to reach statistical significance, $F(1,47) = 2.43$, $p = .13$. Planed comparisons showed that the number of responses during the first Y-O1 trial in Phase 3 was lower in the Experimental than in the Control Group, $t(47) = 2.65$, $p < .05$. However, there were no significant differences in the remaining Y-O1 trials of Phase 3, largest $t(47) = 0.72$. These analyses confirm that compound training of cues A and Y during Phase 1 in the Experimental Group proceeded as expected, resulting in overshadowed responding to Y during Phase 3 relative to the Control Group.

The next step was to assess whether the overshadowed cue was able to interfere at test with responding to cue X. Consistent with the hypothesis that the presentation of an overshadowed cue would cause little interference in responding to cue X at test, the level of responding was higher in the Experimental Group than in the Control Group. But this difference failed to reach statistical significance, $t(47) = 1.10$, $p = .28$. However, visual inspection of Fig. 1 suggests that the level of responding to cue X increased in the Experimental Group relative to the end of Phase 2, suggesting that no interference took place in this group, whereas it decreased in the Control Group, suggesting that interference might have taken place. Thus, in order to further assess interference at test, interference scores were

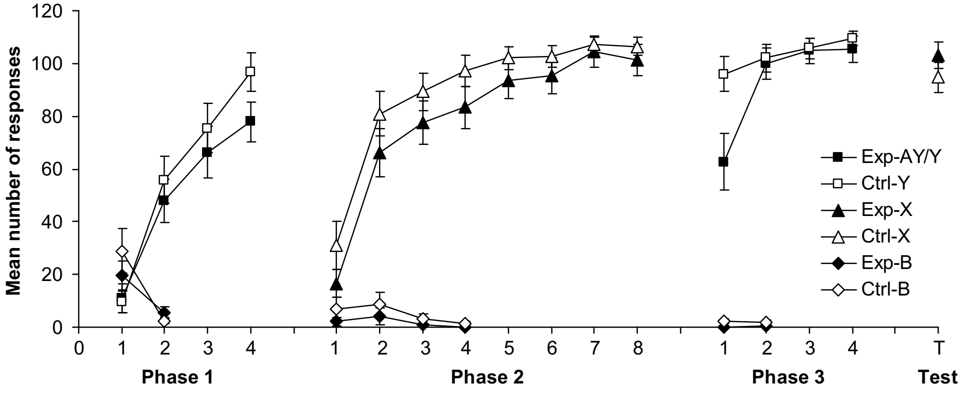

Fig. 1. Mean responses during training and test in Experiment 1. Data series depicted in black refer to the Experimental Group and those depicted in white refer to the Control Group. Error bars denote the standard error of the mean.

computed by subtracting each participant's mean responding to X during the last X-O1 trial in Phase 2 from the number of responses to cue X at test. Negative interference scores indicate that responding to X at test were reduced relative to baseline responding to X by the end of Phase 2. Therefore, negative scores were taken as indicative of interference at test. As expected, the interference scores showed that the Control Group exhibited stronger interference at test than the Experimental Group. A *t* test showed that the interference score was more negative (and, therefore, interference was stronger) for participants in the Control Group ($M = -11.76$, $SEM = 5.07$) than for participants in the Experimental Group ($M = 1.75$, $SEM = 4.31$), $t(47) = 2.02$, $p < .05$. Therefore, the overshadowing training received by cue Y in the Experimental Group did reduce the ability of cue Y to serve as an interfering cue during Phase 3, although this effect was evident only when interference scores were analyzed.

Experiment 2

The results of Experiment 1 provided partial support for the hypothesis that an overshadowed cue loses its ability to produce interference with other cues of the same outcome. On the one hand, the analysis performed on interference scores shows that no interference took place in the Experimental Group, whereas some interference could be observed in the Control Group. On the other hand, the direct analysis of the absolute level of responding at test failed to show significant differences between both groups.

Some features of the experimental design used in Experiment 1 might have been responsible for our failure to observe differences in interference at test when the absolute level of responding was taken as the dependent variable. Specifically, the design of the previous experiment does not control for potential differences in the amount of proactive interference that training during Phase 1 may exert on learning during Phase 2. During Phase 1 the Experimental Group learnt that two cues, A and Y, could predict the outcome. However, in the Control Group only one cue, Y, predicted the outcome. This means that during Phase 2, the learning of the X-outcome association was affected by two sources of proactive interference in the Experimental Group, but only by one source of interference in the Control Group. This difference in the amount of proactive interference might have been responsible for the slightly higher level of responding to cue X during Phase 2 in the Control Group. This differential level of responding to cue X might in turn have overshadowed the target differences at test between both groups.

In order to better control for the differential levels of proactive interference, Experiment 2 used a modified experimental design in which both groups were exposed to identical training conditions during training phases 1 and 2. The design of Experiment 2 is outlined in Table 2. As can be seen, all participants receive both the overshadowing treatment (AY-O1 trials) and the control treatment (Z-O1 trials) during Phase 1. This assures that the levels of proactive interference exerted by this training are equal in both groups. During Phase 2, all participants receive training with the X-O1 association. During Phase 3, participants in the Experimental Group are further trained with the overshadowed cue, Y, whereas participants in the Control Group are trained with the control cue, Z.

Table 2
Design summary of Experiment 2

Group	Phase 1	Phase 2	Phase 3	Test
Experimental	4 AY-O1	8 X-O1	4 Y-O1	1 X
	4 Z-O1			
	6 B-O2	6 B-O2	6 B-O2	
Control	4 AY-O1	8 X-O1	4 Z-O1	1 X
	4 Z-O1			
	6 B-O2	6 B-O2	6 B-O2	

Note: Cues A–Z are color lights in the spy-radio that can predict two different outcomes: O1 (the participants gain points if they perform the response) and O2 (participants lose points if they perform the response). Numbers before each trial type denote the number of trials of that type that were presented in a given phase.

Method

Participants and apparatus

The sample was composed by 36 psychology students from the University of the Basque Country who volunteered to take part in the experiment. Nineteen were assigned to the Experimental Group and 17 to the Control Group. These participants performed the experiment under conditions similar to those of Experiment 1. Apart from this students' sample, an additional sample was recruited over the Internet. This sample was composed by 45 Internet users who visited the website of our virtual laboratory, http://www.labpsico.com, and voluntarily decided to participate in the experiment. Twenty-three were assigned to the Experimental Group and 22 to the Control Group.

Procedure and design

All the procedural details were kept unchanged with respect to Experiment 1. Only minor changes were done to adapt the experimental task to the design of Experiment 2. Colors yellow, blue, red, green, and purple, all counterbalanced in a latin square, were used as color lights in the spy-radio, playing the role of cues A, B, X, Y, and Z.

Results and discussion

Six participants in the Experimental Group (five from the Internet sample and one from the students' sample) and six participants in the Control Group (three from the Internet sample and three from the students' sample) failed to meet the same data selection criteria that was used in Experiment 1 and were, consequently, removed from subsequent analyses. Therefore, the following analyses were conducted with 36 participants from the Experimental Group and 33 from the Control Group. Preliminary analyses showed that the origin of the sample (college students vs. Internet users) did not interact with the other factors studied in the present experiment. Therefore, this factor was removed from subsequent analyses.

Mean responding of the participants during training and at test is depicted in Fig. 2. As shown by the following analyses, training during Phases 1 and 2 proceeded as expected without remarkable differences between groups. A 2 (group) × 4 (trials) ANOVA performed on participants responses to AY during Phase 1 yielded a main effect of trial, $F(3,201) = 79.49$, $p < .001$. The main effect of group and the group × trial interaction were nonsignificant, all Fs < 1. An equivalent ANOVA performed on responses to Z during Phase 1 also showed a main effect of trial, $F(3,201) = 72.87$, $p < .001$, but no effect of group or group × trial interaction, all Fs < 1. Another 2 (group) × 8 (trials) ANOVA performed on responses to X during Phase 2 yielded a main effect of trial, $F(7,469) = 28.08$, $p < .001$, and an unexpected

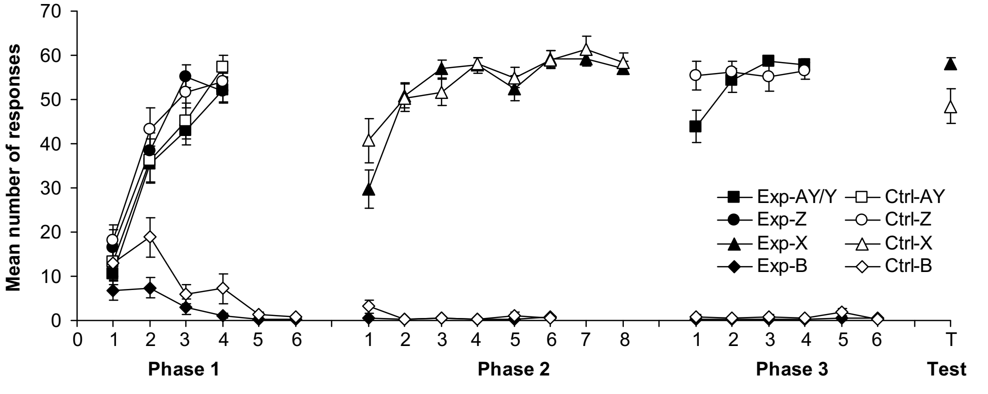

Fig. 2. Mean responses during training and test in Experiment 2. Data series depicted in black refer to the Experimental Group and those depicted in white refer to the Control Group. Error bars denote the standard error of the mean.

group × trial interaction, $F(7,469) = 2.26$, $p < .05$. The main effect of group was not significant, $F < 1$. Planned comparisons showed that the interaction was due to a marginally significant difference between both groups in the number of responses given in the first X-O1 trial of Phase 2, $t(67) = 1.69$, $p = .096$; there were no differences between both groups in the remaining X-O1 trials of Phase 2, largest $t(67) = 1.53$.

The pattern of responses during Phase 3 shows an overshadowing effect in the first trial. A 2 (group) × 4 (trial) ANOVA performed on responses to cues Y/Z during Phase yielded a main effect of trial, $F(3,201) = 6.01$, $p < .005$, and, most importantly, a group × trial interaction, $F(3,201) = 5.36$, $p < .005$. The main effect of group did not reach statistical significance, $F < 1$. Planned comparisons showed that responses to cue Y in the Experimental Group were lower than responses to Z in the Control Group in the first trial, $t(67) = 2.32$, $p < .05$, but not in the remaining trials, largest $t(67) = 1.08$. These analyses confirm that training of cues AY during Phase 1 resulted in overshadowed responding to Y during Phase 3 relative to the control cue Z.

The analysis of the absolute number of responses at test showed that participants in the Experimental Group gave more responses than participants in the Control Group. Unlike Experiment 1, this difference was now statistically significant, $t(67) = 2.46$, $p < .05$. Additionally, interference scores were computed as in Experiment 1. Again, interference scores were more negative in the Control Group ($M = -9.97$, $SEM = 3.65$) than in the Experimental Group ($M = 1.08$, $SEM = 1.14$), $t(67) = 2.99$, $p < .005$. These analyses confirm that interference was significantly stronger in the Control Group than in the Experimental Group. In other words, exposure to an overshadowed cue failed to interfere with latter responding to another predictor of the same outcome, whereas such interference was readily observed in the Control Group. These results suggest that the failure to find significant differences at test in Experiment 1 can be attributable to the above mentioned problems in its design.

General discussion

Implications for associative models

It is quite difficult to account for both interference effects and cue competition effects within a common theoretical framework. Although many associative models can readily explain cue competition effects such as overshadowing (Miller & Matzel, 1988; Rescorla & Wagner, 1972; Wagner, 1981), most of these models fail to account for interference phenomena.

One family of associative models that can potentially account for both overshadowing and the interference effects is retrospective revaluation models. For instance, Van Hamme and Wasserman (1994) proposed an extension of the Rescorla and Wagner (1972) model which allowed them to account for changes in the associative status of absent cues. According to this extended model, absent (but expected) cues change their associative strengths in the direction *opposite* to that of the present cues. This feature allows the model to explain the basic interference effect described in the Introduction: if Y-outcome trials are provided after initial X-outcome training, the Y-outcome association is strengthened, but the X-outcome association is correspondingly weakened, due to the absence of cue X during Y-outcome trials.[1]

Interestingly, as a descent of the Rescorla–Wagner model, this model is also able to account for cue competition effects, such as overshadowing. However, the model fails to account for the interaction between overshadowing and interference that we have observed in this series of experiments. According to this model, the amount of associative change of absent cues is related to the amount of associative change of present cues: the more the associative status of the present cue changes, the more the associative status of the absent cues should change *in the opposite direction*. In our Experiment 1, Y-outcome training during Phase 3 should have resulted in different amounts of learning in the Experimental and Control groups. In the latter, Y-outcome training during Phase 1 should have resulted in asymptotic or nearly asymptotic learning of the Y-outcome association. Thus, Y-outcome training during Phase 3

[1] Of course, this explanation can only be applied to interference provided that cue X is, for any reason, expected during Y-Outcome trials (either because of the common context used in both phases or because of the common outcome shared by X and Y).

should have resulted in little additional learning of the Y-outcome association. In accordance, the absent target cue, X, should have suffered little, if any, associative change during Phase 3 in the Control Group. However, in the Experimental Group, the Y-outcome association was overshadowed and, therefore, should not have reached the learning asymptote in Phase 1. Thus, Y-outcome training during Phase 3 should have resulted in substantial learning of the Y-outcome association until the learning asymptote was reached. In this case, there should have been a proportional decrease in the associative strength of the absent cue, X, in Phase 3. As can be seen, these predictions are just the opposite to the actual results observed in the experiment. Other retrospective revaluation models based on similar assumptions (e.g., Dickinson & Burke, 1996) face similar problems to account for the present results.

In a very different vein, Miller and Escobar (2002) proposed an extension of the comparator hypothesis (Miller & Matzel, 1988) that can potentially account for the overshadowing and interference effects and for their interaction. The comparator hypothesis explains overshadowing as a retrieval deficit induced by the simultaneous activation of the representation of the outcome from two different sources during the test phase. On the one hand, the overshadowed cue is able to activate a direct representation of the outcome by virtue of its being paired with the outcome. On the other hand, the overshadowed cue also activates a representation of its associate overshadowing cue (i.e., the comparator cue); this representation of the overshadowing cue can in turn activate another (indirect) representation of the outcome. According to this framework, little responding is observed at test because the direct activation of the outcome elicited by the overshadowed cue is counteracted by the indirect activation of the outcome elicited by its comparator cue (i.e., the overshadowing cue).

Miller and Escobar (2002) proposed that comparator cues might also be primed at test even when they have not been paired with the target cue: either recent training or the presentation of the context in which the comparator cues were trained might result in the activation of the comparator representation at test. This mechanism would allow the model to account for the interference effect: after X-outcome training, Y-outcome training would result in strong priming of cue Y during testing, and this cue would therefore act as a strong comparator stimulus for X at test. In our experiment, this means that responding to cue X at test would be modulated by the associative status of cue Y, which would be primed because of its recent training.

Interestingly, in recent developments of the comparator hypothesis (Denniston, Savastano, & Miller, 2001; Stout & Miller, 2007) it is assumed that comparator cues themselves are subject to a comparator process that modulates their ability to interfere with responding to other cues. In our experiment, the ability of cue Y to interfere with responding to X would be modulated by comparator cues of cue Y. Given that in the Experimental Group, cue Y has a strong comparator cue, A, the ability of Y to interfere with X would be somewhat reduced, relative to the Control Group, in which cue Y has no comparator cue. Therefore, this integration of the extended comparator hypothesis (Denniston et al., 2001; Stout & Miller, 2007) and the assumption that comparator cues can be primed either by the context or by recent training (Miller & Escobar, 2002) can explain the present results.

Implications for theories of interference based on reasoning processes

Some researchers have explained previous demonstrations of interference between cues in our experimental task in terms of higher-order reasoning processes. According to Cobos et al. (2007; see also Luque, Cobos, & López, 2008), participants can perceive some contingency learning tasks, such as the spy-radio preparation used in this experiment, as a diagnostic task in which the goal is to diagnose the presence of an event (in this case, the safety of the road) on the basis of the visible effects of that event on the environment (in this case, the onset of the color lights in the spy-radio panel). According to this interpretation, in an interference design, participants observe that the safety of the road can produce two different effects on Phases 1 and 2: in Phase 1, color light X is on when the road is safe; however, in Phase 2, a different color light, Y, is on when the road is safe. According to Cobos et al. (2007), participants perceive a contradiction in this evidence, because a single cause (the safety of the road) can give rise to two different and mutually exclusive effects (the onset of color lights X and Y). This contradiction and participants' attempts to explain it would be responsible for the observed interference effect at test. Critically, this effect would not appear if participants perceived the predictive cues as potential causes (and not effects) of the state of the road.

Although this explanation of the interference effect is an appealing one and accounts for most of the evidence available in the literature, we think that it fails to account for the result of the present series of experiments. Both participants in the Experimental and Control groups were exposed to information that could be seen as contradictory in causal terms: all of them could be regarded as seeing that a single cause (the safety of the road) can give rise to different and mutually exclusive color cues. If anything, the information received by participants in the Experimental Group of Experiment 1 should result more counterintuitive for participants because participants see that a single event can produce three different configurations of cues (AY, X, and Y alone), whereas participants in the Control Group see that a single event can produce only two different configurations of cues (Y alone and X). Thus, if anything, this account would predict more interference at test in the (causally more complex) experimental condition than in the control condition.

Methodological implications

Regardless of the theoretical implications of the present study, we think that the interaction we observed between interference and cue competition has important methodological implications for the study of human associative learning processes. Although early studies on human contingency learning were generally accounted for in terms of lower-order psychological mechanisms (e.g., Dickinson, Shanks, & Evenden, 1984; Shanks, 1985; Wasserman, 1990), recent work suggests that most of the available evidence might be better explained in terms of higher-order cognitive processes related to causal and deductive reasoning (for a review, see De Houwer, Vandorpe, & Beckers, 2005).

According to this view, in order to explore the role of lower-level mechanisms potentially involved in contingency learning, it is necessary to use experimental procedures that preclude the intervention of higher-order cognitive processes. For example, as a means to impair these processes, researchers have tried to overload the participants cognitive system by using very complex experimental designs (e.g., Dickinson & Burke, 1996; Le Pelley, Oakeshott, & McLaren, 2005) or by asking participants to perform a complex secondary task while they are also performing the contingency learning task (De Houwer & Beckers, 2003). In other studies, the impact of higher-order processes on the experiments' results has been reduced by using implicit measures of participants' learning, instead of asking them to give subjective judgments or to perform a predictive behavior (Vandorpe & De Houwer, 2006).

In our opinion, the interaction between interference and cue competition suggests that interference can be used as an implicit measure for the study of human associative learning. If one wants to check whether a given cue has been overshadowed in a contingency learning experiment, one can either measure this overshadowing effect directly by paying attention to participants' responses to that cue, or, alternatively, one can indirectly measure the overshadowing effect by checking whether or not the potentially overshadowed cue has lost its ability to interfere with responding to another cue. This second, indirect measure of overshadowing is more likely to be unaffected by higher-order cognitive processes than the former one. This procedure assesses participants' learning of one cue-outcome association by asking them to respond to a different cue, which makes very difficult for them to become aware of what is really being assessed in the experiment. The same general procedure could be used to explore any contingency learning effect, such as blocking, inhibition, etc. An important advantage of this indirect measure of learning over alternative implicit measures, such as affective priming (Fazio, Sanbonmatsu, Powell, & Kardes, 1986) or the implicit association test (Greenwald, McGhee, & Schwartz, 1998), is that this procedure can be easily implemented on most preparations for the study of human contingency learning currently used by researchers (the only requirement is that trials for the measure of retroactive interference are included in the design), whereas standard implicit measures usually have to be implemented on specific preparations. Of course, this proposal of using interference as an implicit measure is only tentative and further research with other learning effects, apart from overshadowing, has to be done in order to assess the validity of this procedure.

References

Amundson, J. C., & Miller, R. R. (2007). Similarity in spatial origin of information facilitates cue competition and interference. *Learning and Motivation, 38*, 155–171.

Barnet, R. C., Grahame, N. J., & Miller, R. R. (1993). Temporal encoding as a determinant of blocking. *Journal of Experimental Psychology: Animal Behavior Processes, 19*, 327–341.

Cobos, P. L., López, F. J., & Luque, D. (2007). Interference between cues of the same outcome depends on the causal interpretation of the events. *Quarterly Journal of Experimental Psychology, 60*, 369–386.

De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology, 56B*, 345–357.

De Houwer, J., Vandorpe, S., & Beckers, T. (2005). On the role of controlled cognitive processes in human associative learning. In A. Wills (Ed.), *New directions in human associative learning* (pp. 41–63). Mahwah, NJ: Lawrence Erlbaum.

Denniston, J. C., Savastano, H. I., & Miller, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65–117). Hillsdale, NJ: Erlbaum.

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

Dickinson, A., Shanks, D. R., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology, 36A*, 29–50.

Escobar, M., Matute, H., & Miller, R. R. (2001). Cues trained apart compete for behavioral control in rats: Convergence with the associative interference literature. *Journal of Experimental Psychology: General, 130*, 97–115.

Escobar, M., & Miller, R. R. (2003). Timing in retroactive interference. *Learning & Behavior, 31*, 257–272.

Escobar, M., Pineño, O., & Matute, H. (2002). A comparison between elemental and compound training of cues in retrospective revaluation. *Animal Learning & Behavior, 30*, 228–238.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229–238.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74*, 1464–1480.

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.: University of Miami Press.

Le Pelley, M. E., Oakeshott, S. M., & McLaren, I. P. L. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 56–70.

Luque, D., Cobos, P. L., & López, F. J. (2008). Interference between cues requires a causal scenario: Favourable evidence for causal reasoning models in learning processes. *Learning and Motivation, 39*, 196–208.

Matute, H., & Pineño, O. (1998a). 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.

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

Miller, R. R., & Escobar, M. (2002). Associative interference between cues and between outcomes presented together and presented apart: An integration. *Behavioural Processes, 57*, 163–185.

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.

Pavlov, I. P. (1927). *Conditioned reflexes*. London: Oxford University Press.

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-Century-Crofts.

Shanks, D. R. (1985). Continuous monitoring of human contingency judgment across trials. *Memory & Cognition, 13*, 158–167.

Slamecka, N. J., & Ceraso, J. (1960). Retroactive and proactive inhibition of verbal learning. *Psychological Bulletin, 57*, 449–475.

Stout, S. C., Chang, R., & Miller, R. R. (2003). Trial spacing is a determinant of cue interaction. *Journal of Experimental Psychology: Animal Behavior Processes, 29*, 23–38.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review, 114*, 759–783.

Urcelay, G. P., & Miller, R. R. (2006). Counteraction between overshadowing and degraded contingency treatments: Support for the extended comparator hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes, 32*, 21–32.

Vadillo, M. A., Castro, L., Matute, H., & Wasserman, E. A. (2008). Backward blocking: The role of within-compound associations and interference between cues trained apart. *Quarterly Journal of Experimental Psychology, 61*, 185–193.

Vandorpe, S., & De Houwer, J. (2006). *The implicit association test as a measure of acquired associative strength between cause and effect in human causal learning*. Paper presented at the XVIII annual meeting of the Spanish Society for Comparative Psychology. Málaga, Spain.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition 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.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76*, 171–180.

Wasserman, E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science, 1*, 298–302.

Wheeler, D. S., & Miller, R. R. (2007). Interactions between retroactive-interference and context-mediated treatments that impair Pavlovian conditioned responding. *Learning & Behavior, 35*, 27–35.