The Quarterly Journal of Experimental Psychology
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t716100704

Interference between cues of the same outcome depends on the causal interpretation of the events

Pedro L. Cobos [a]; Francisco J. López [a]; David Luque [a]
[a] University of Málaga, Spain

Online Publication Date: 01 March 2007

To cite this Article: Cobos, Pedro L., López, Francisco J. and Luque, David (2007) 'Interference between cues of the same outcome depends on the causal interpretation of the events ', The Quarterly Journal of Experimental Psychology, 60:3, 369 — 386

To link to this article: DOI: 10.1080/17470210601000961
URL: http://dx.doi.org/10.1080/17470210601000961

THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2007, 60 (3), 369–386

Psychology Press
Taylor & Francis Group

Interference between cues of the same outcome depends on the causal interpretation of the events

Pedro L. Cobos, Francisco J. López, and David Luque
University of Málaga, Spain

In an interference-between-cues design, the expression of a learned Cue A → Outcome 1 association has been shown to be impaired if another cue, B, is separately paired with the same outcome in a second learning phase. In the present study, we assessed whether this interference effect is mediated by participants' previous causal knowledge. This was achieved by having participants learn in a diagnostic situation in Experiment 1a, and then by manipulating the causal order of the learning task in Experiments 1b and 2. If participants use their previous causal knowledge during the learning process, interference should only be observed in the diagnostic situation because only there we have a common cause (Outcome 1) of two disjoint effects, namely cues A and B. Consistent with this prediction, interference between cues was only found in Experiment 1a and in the diagnostic conditions of Experiments 1b and 2.

In daily life causal relationships have to be inferred from data collected from a huge variety of different contexts (whether physical or temporal). In such circumstances, a crucial task for our cognitive system is to determine whether or not what we have learned in one context is valid in another one. In other words, our cognitive system has to evaluate whether the data sets coming from two (or more) different contexts have to be integrated to calculate the causal relationship between events or whether each data set has to enter into different computations to calculate causal relationships that are valid in one context but not in others. Thus, an important question in the causal learning field is to understand the circumstances in which people assume that what has been learned in one context is applicable in another one, as well as the circumstances in which people do not hold this assumption.

In our view, the study of retrospective interference in learning experiments can make important contributions to address the above questions. In an interference experiment, the expression of a previously learned relationship between a cue (or antecedent event), A, and an outcome (or consequent event), O1 (A-O1 relationship), is hindered if another relationship (e.g., A-O2, or B-O1) is learned in a later training phase[1].

[1] We should term this effect as retroactive interference to distinguish it from the proactive interference effect. Since we will only focus on the former, the term "interference" should be interpreted, hereafter, in the narrow sense of retroactive interference.

Correspondence should be addressed to Dr. Pedro L. Cobos, Departamento de Psicología Básica, Campus de los Teatinos, Universidad de Málaga, Málaga 29071, Spain. E-mail: p_cobos@uma.es

The research described here was supported by research grant from Junta de Andalucía (HUM-0105) from Spain. David Luque has been supported by F.P.I. fellowships also from Junta de Andalucía (Ref.: 4381).

We would like to thank Julián Almaraz, Fernando Blanco, Helena Matute, and Miguel A. Vadillo for their helpful comments on the experiments presented here.

http://www.psypress.com/qjep
DOI:10.1080/17470210601000961

There is extensive evidence, both from human and from animal learning experiments, that interference is not due to unlearning of the A-O1 relationship. It is well known that the expression of this relationship can recover almost completely if certain contextual cues (whether physical or temporal) are manipulated. For instance, if the A-O1 relationship is learned in Context X, and the second relationship is learned in Context Y, the expression of the former is recovered if participants are tested with A in Context X as well as in a new context (Bouton, 1993; García-Gutiérrez & Rosas, 2003a, 2003b; Pineño & Matute, 2000; Rosas & Bouton, 1998; Rosas, Vila, Lugo, & López, 2001; Vadillo, Vegas, & Matute, 2004). This recovery is also observed if a time delay is interpolated between the second training phase and the test phase (Bouton, 1993; Pavlov, 1927; Rosas & Bouton, 1998; Rosas et al., 2001). Thus, rather than unlearning the relationship learned in the first phase, people seem to behave as if what is learned in each phase were not valid in the other one. Moreover, contextual cues (whether physical or temporal) seem to play a crucial role in determining what relationship is pertinent at the present moment. But why is it that people tend to restrict the validity of each data set to different contexts in interference experiments?

At first glance, it could be thought that interference occurs because participants learn two conflicting predictive relationships. This could be the case in interference-between-outcomes treatments where participants learn that Cue A predicts Outcome 1 during the first training phase and Outcome 2 (or no outcome) during the subsequent training phase. These two conflicting relationships render Cue A ambiguous as it predicts either Outcome 1 or Outcome 2 but never both at the same time. But the idea of two conflicting predictive relationships does not apply so well in the case of interference between different cues of the same outcome (interference between cues hereafter). Table 1 shows a typical experimental design of interference between cues. This type of interference consists in the decrease of the expression of the A-O1 relationship learned in the first phase caused by the learning of a relationship between a different cue (i.e., Cue B) and the same outcome (i.e., Outcome 1) in a later phase. Because in this case there is no obvious conflict between the relationships learned –each cue only predicts one outcome– interference between cues has challenged current learning theories of interference and, thus, has recently attracted the attention of learning researchers.

Table 1. *Experimental design for interference between different cues of the same outcome*

	Training phases		
Experimental conditions	*Phase 1 Context X*	*Phase 2 Context Y*	*Test phase Context Y*
Same outcome	A → O1 C → O3	B → O1 C → O3	A?
Different outcome	A → O1 C → O3	B → O2 C → O3	A?

Note: Letters A-C and numbers stand for cues and outcomes, respectively; O refers to outcome; letters X and Y stand for different contexts.

Interference between cues, recently observed in predictive learning studies, had been found much earlier in the verbal learning tradition (Abra, 1967; Cheung & Goulet, 1968; Keppel, Bonge, Strand, & Parker, 1971; Schwartz, 1968). In predictive learning, the evidence for interference between cues comes mainly from human studies (see, for example, Escobar, Pineño, & Matute, 2002; Matute & Pineño, 1998; Ortega & Matute, 2000; Pineño, Ortega, & Matute, 2000), though there is also evidence for interference between cues from animal experiments (Escobar, Arcediano, & Miller, 2001; Escobar, Matute, & Miller, 2001; Escobar & Miller, 2003). As mentioned above for other forms of interference, interference between cues extinguishes or decreases whenever participants are tested with A either in a novel context or in the context in which the A-O1 relationship was learned (Matute & Pineño, 1998; Pineño et al., 2000; Pineño & Matute,

2000).[2] Thus, in an interference-between-cues treatment, rather than unlearning the A-O1 relationship, participants again seem to behave as if what is learned in one phase were not valid in the other.

Interference between cues has mainly been addressed by means of associative theories of learning. Although there are several associative explanations, most of them are inspired to a greater or lesser extent by Bouton's (1993) theory for interference between outcomes in animal learning (for more details, see Escobar et al., 2002; Matute & Pineño, 1998; Miller & Escobar, 2002). Without going into specific details concerning these explanations, overall, they build on the following two main ideas. First, interference occurs if two learned associations have a common element in a common temporal location (whether the cue or the outcome location), in which case such associations are learned separately. Second, the context in which a given association is acquired has the capacity of priming such association in detriment of others that share a common element with the previous one. Consequently, if two associations have a common cue or outcome, our cognitive system tends to diminish the validity of each association in the context in which the other was learned. According to these assertions, interference between cues occurs because: 1) the associations learned in the first and the second training phases have the outcome as a common element; 2) the context for the test is the same (or almost the same) as for the second training phase.[3]

Whereas almost all of the existing explanations for interference between cues originate from an associative framework, we propose that interference may be better understood from the viewpoint of what could be termed as causal reasoning theories, such as Causal Bayes Nets (Glymour, 2001; Gopnik et al., 2004; Sloman & Lagnado, 2004) and Causal Model Theory (Waldmann & Holyoak, 1992). This constitutes a novelty in the research on interference because, as far as we know, no attempt has been made to apply causal reasoning theories (CR theories hereafter) to explain interference, even though most of the experiments with humans have been framed within scenarios that are clearly susceptible of a causal interpretation.

Our CR account of interference between cues is based on two main assumptions: 1) people tend to give a causal interpretation to cues and outcomes; 2) they use their previous causal knowledge to infer the causal structure underlying the observed data. Regarding the causal interpretation, our hypothesis is that interference between cues is observed because cues and outcomes are interpreted as effects and causes, respectively, i.e., people perceive their task as a diagnostic learning task. According to this assumption, in an experiment of interference between cues, participants in the same-outcome condition learn that Cause 1 produces Effect A during the first phase and Effect B during the second phase (see the design in Table 1). Regarding the second assumption, we propose that people use structural knowledge about cause-effect relationships to infer the causal structure underlying the observed data. Such knowledge may be conceived of as a sort of intuitions about the probability distribution of events that should be expected given a specific causal model. If the expected distribution does not fit the observed probability distribution, the entertained causal model is rejected and substituted by

[2] Here, we use the term "context" in a wide sense to refer to a constant background or cue, a discrete cue immediately preceding a training trial, or to a temporal locus.

[3] Interference between cues has also been found without the use of different physical contexts or contextual cues to help participants differentiate between the two subsequent training phases (see, for example, Escobar et al., 2002). Based on this evidence, some authors have claimed that interference between cues could be explained by recency of the competing B-O1 association at the time of test (Ortega & Matute, 2000; Pineño et al., 2000). However, these results are also consistent with stating that the temporal context at the time of test, which is more similar to the temporal context of Phase 2 than to the temporal context of Phase 1, is priming the competing B-O1 over the A-O1 association.

another causal model. The application of this process may eventually lead to the postulation of hidden causes to improve the fit between the observed and the expected probability distribution.

What are the consequences of these assumptions in a design of interference between cues? If the learning task is given a diagnostic interpretation, the first thing to notice is that the probability distribution of events in the same-outcome condition is not compatible with a single-cause model (see Table 1). If Cause 1 were the only relevant cause explaining effects A and B, then, keeping constant the presence of Cause 1, there should be independence between the occurrence of A and of B. But this expectation is contradicted by the data since, given the presence of Cause 1, there is a highly negative correlation between its effects, i.e., A and B are disjoint effects. To explain this negative correlation, participants have to reject the single-cause model, and have to postulate the existence of a hidden causal factor that interacts with Cause 1, thereby producing either Effect A or Effect B (see also Hagmayer & Waldmann, 2007, in this issue for further information about inferences about unobserved causes from a causal Bayes net perspective). Such a causal factor works as an enabling or disabling condition that limits the capacity of Cause 1 to produce each effect. On this basis, the fact that Cause 1 systematically generates Effect A during Phase 1 and Effect B during Phase 2 can be accounted for by assuming that the hidden causal factor has one value in one phase and another value in the other phase. Given that there is no temporal or physical separation between the test trial and the trials of Phase 2, the hidden causal factor is assumed to be in the same state as in Phase 2. Such a state enables Cause 1 to produce Effect B, but not Effect A. Thus, when participants are tested with A, they are uncertain about the occurrence of Cause 1, even though they know of no other alternative cause that could have produced the effect.

When we say that participants in the same-outcome condition might be postulating a hidden causal factor we do not mean that they might necessarily be entertaining a specific hypothesis about what factor could explain the disjoint effects of Cause 1. People might simply be assuming that something unknown has changed from one phase to another that makes Cause 1 produce a different effect. This is likely to be the case in those circumstances in which people lack domain specific causal knowledge to imagine a plausible account of the data. Interference experiments are, in general, an example of such circumstances.

As a consequence of postulating a hidden causal factor in the same-outcome condition, contextual cues (whether temporal or physical) become causally relevant through the different status of the hidden causal factor that they signal. According to our CR account of interference between cues, contextual cues acquire causal significance because they are the only cues that inform participants about the state of the hidden causal factor. This could be the reason why contextual manipulations exert such a great influence on the extent to which interference is observed (for more on this, see the General Discussion).

Things are quite different in the control (different-outcome) condition: Cause 1 produces Effect A and Cause 2 produces Effect B during Phases 1 and 2, respectively. The data provided are consistent with single-cause models (i.e., one per causal relationship) so that no hidden causal factor needs to be assumed. Therefore, there is no reason to confine the validity of the calculations to a restricted training phase. Thus, when participants in the control condition are tested with A, they should be more confident that Cause 1 has occurred than participants in the same-outcome condition.

A critical prediction from our CR account is that interference between cues should not be observed in a predictive situation, i.e., if cues and outcomes are interpreted as causes and effects, respectively. In a predictive situation, participants in the same-outcome condition would learn that Cause A produces Effect 1 during Phase 1, and that Cause B produces the same effect during Phase 2 (see Table 1). Both relationships could be explained by single-cause models (one per relationship) because each cause only produces one effect. The fact that there are two events causing the same effect does

not require the postulation of a hidden causal factor which affects the capacity of each cause to produce the effect. Thus, there is nothing in the causal structure of the data that should prevent participants from generalising the knowledge acquired in Phase 1 to Phase 2. Consequently, participants' responses to Cue A at test in the same-outcome condition should not differ from those in the control condition.

A relevant fact consistent with our prediction is that there is no evidence (that we are aware of) of interference between cues in predictive (cause-effect) learning tasks. The available evidence of interference between cues comes, generally, from studies in which the learning task can reasonably be interpreted in the diagnostic direction (we will return to this in the General Discussion).

From an associative framework, no asymmetry between predictive and diagnostic situations is expected. Associative learning mechanisms are thought to be blind to the causal role played by cues and outcomes. Therefore, according to associative learning theories (see above), interference between cues should be observed both in predictive and in diagnostic situations.

In summary, according to our CR account of interference between cues, interference should be found in diagnostic but not in predictive situations. According to associative theories, equivalent interference should be obtained in predictive and diagnostic situations.

Interestingly, the manipulation of the causal role of events has been shown to have an influence on cue interaction effects, but in the opposite way (López, Cobos, & Caño, 2005; Tangen & Allan, 2004; Van Hamme, Kao, & Wasserman, 1993; Waldmann, 2000, 2001; Waldmann & Holyoak, 1992). Specifically, there is evidence showing cue interaction effects in predictive but not in diagnostic conditions. Such asymmetries are consistent with CR theories and are again not predicted by associative theories. Note that there is no conflict between the predictions derived from our CR account concerning interference between cues and the asymmetry predicted from CR theories concerning cue interaction effects because each prediction is based on different principles.

The aim of our study was to find out whether or not interference between cues is affected by the causal role of cues and outcomes. For this, we created a predictive condition in which cues and outcomes played the role of causes and effects, respectively, and a diagnostic condition in which cues and outcomes played the role of effects and causes, respectively. Observing a higher interference between cues in the diagnostic than in the predictive condition would lend support to our CR account. Alternatively, if the interference effects in the predictive and in the diagnostic conditions are not significantly different, the results would be consistent with associative theories.

Our first purpose was to replicate the interference between cues by using a learning task framed in a causal scenario susceptible to be defined in either a diagnostic or a predictive direction with minimal changes. To test whether our task was suitable to obtain interference between cues in the first place, in Experiment 1a all participants learned the task in the diagnostic direction, as both theoretical accounts predict interference between cues in this causal order condition. Since a reliable interference effect was obtained in Experiment 1a, Experiment 1b was conducted to find out whether such effect was modulated by the manipulation of the causal order of the learning task. Thus, in Experiment 1b participants were divided into two groups, one of which learned in a predictive condition whereas the other learned in a diagnostic condition. Finally, Experiment 2 was carried out to improve some methodological aspects of Experiments 1a and 1b, and to rule out alternative explanations to our CR account of the results.

EXPERIMENT 1A

The purpose of Experiment 1a was to find out whether our learning task was suitable to induce and replicate the interference between cues previously obtained with other learning tasks. According to Escobar et al. (2002), some tasks have proven to be rather insensitive to interference between cues. Thus, we tried to develop a sensitive

procedure to induce interference effects. One procedural measure that could help induce interference is providing a clear and different context for each training phase. Evidence coming from interference-between-outcomes experiments shows that the interference effect is stronger when contextual cues other than the temporal locus are added to help discriminate between the two training phases (Rosas & Bouton, 1998; Rosas et al., 2001; Vadillo et al., 2004). Thus, participants were trained in Context X during Phase 1 and in Context Y during Phase 2. Each context consisted of a series of cues that remained unchanged along each training phase. To favour the interference effect, the test trial took place in Context Y, i.e., in the same context as the interfering training phase. In addition, because both the associative and our CR accounts predict the interference effect in diagnostic situations, all participants in the present experiment learned in the diagnostic condition.

Method

Participants, design, and apparatus

Thirty eight psychology undergraduate students from University of Málaga took part in the present experiment for course credits. They were randomly assigned to one of the two experimental conditions shown in Table 1. The task was performed on IBM-PC compatible computers in semi-isolated individual cubicles.

Procedure

Participants started by reading the instructions on the computer screen placed in one of the ten cubicles in the laboratory. The instructions contained detailed explanations of the learning task described here. The learning task was a variation of the spy-radio task developed by Pineño et al. (2000) to obtain interference between cues. The main distinctive feature of our learning task concerns the causal scenario. We designed a causal scenario that could be defined in both causal order conditions with minimal changes. In our task, participants were asked to imagine that they were physicists working for the Red Cross organisation in a very poor area of Central Africa. Due to hunger, members of certain tribes had begun to eat a series of plants, some of which were poisonous. On each trial, participants had to guess whether or not a hypothetical patient had eaten a poisonous plant and, if so, they had to administer a certain amount of antidote to prevent the poisoning symptoms. If the patient had eaten a poisonous plant, then participants could gain as many points as antidote units were administered. Such units could be administered by repeatedly pressing the space bar or by keeping it pressed for a period of time. Rather than the poisonous plants, patients could have eaten either an anomalous or an innocuous plant. Administering the antidote to a patient who had eaten the anomalous plant caused her/him to be intoxicated, and, thus, participants lost as many points as antidote units were given. Eating this plant had no effect when the antidote was not administered. Finally, administering the antidote when the patient had eaten an innocuous plant had no effect, and, thus, participants neither gained nor lost points. Participants were encouraged to gain as many points as possible.

Instructions stated that the ingestion of each type of plant caused different changes in the patient's body pH that could be detected thanks to a special sort of litmus paper that could take on one of three different colours: brown, yellow, or blue. Thus, each litmus-paper colour was the effect of having eaten a different type of plant. Participants had to decide, on each trial, the amount of antidote that had to be administered from knowing what colour the litmus paper took on. After responding, participants received information about the outcome which consisted of three different messages: (a) a photo of the plant including its name and the category label (poisonous, anomalous, and innocuous); (b) the amount of points gained or lost on the current trial; (c) the points accumulated throughout the training session.

In the same-outcome condition, each litmus-paper colour played the role of one of the cues in the design shown in Table 1. Cues A and B indicated that the patient had eaten the same poisonous plant, i.e., Plant 1, whereas Cue C indicated that the patient had eaten an anomalous plant, i.e.,

Plant 3. The different-outcome condition was identical except that Cue B indicated that the patient had eaten an innocuous plant, Plant 2. Because Cue A indicated the ingestion of the poisonous plant, good acquisition would involve a high response rate (a great amount of antidote units) in the presence of Cue A. An interference effect would be observed if the number of responses at test were lower in the same-outcome condition than in the different-outcome condition.

The instructions stressed the importance of: (a) maintaining attention throughout the entire learning task; (b) gaining as much points as possible; (c) placing the hands on the keyboard in a specific fixed position to prevent participants from seeking the response key during the task. At the end of the instructions, participants were invited to ask the experimenter all possible questions about the task. Once all questions were resolved, participants went into the learning task which consisted of two learning phases plus an additional test trial at the end. In each training phase, participants were presented with hypothetical patients from one of two tribes: "ULUS" and "GANTUAS". These tribes played the role of the contexts X and Y displayed in Table 1. The assignment of these roles to each tribe was counterbalanced across participants. During Phase 1, all participants were exposed to ten A-O1 trials intermixed with ten C-O3 trials. Then participants received ten B-O1 trials intermixed with ten C-O3 trials during Phase 2 in the same-outcome condition, whereas participants in the different-outcome condition received ten B-O2 trials intermixed with ten C-O3 trials during Phase 2. The order of trial presentation was randomised with the constraint of disallowing more than two consecutive presentations of the same trial type. The test trial was not marked to participants and consisted of an additional A-O1 trial after the last trial of Phase 2. The context for the test trial was the same as for Phase 2, i.e., Context Y. The abstract cues A through C were assigned to the different litmus-paper colours according to a counterbalancing procedure. Outcomes 1, 2 and 3 were the poisonous, the innocuous, and the anomalous plant, respectively. Also according to a counterbalancing procedure, the three different plants were assigned corresponding plant photos labelled "Dobe", "Yamma", and "Kollin".

On each trial, the message "In the case of Patient #, the colour the litmus paper took on was" appeared at the top centre of the screen. Just below, a small rectangle filled with the corresponding colour was displayed. The colour was visible for 3.5 s only. After that time, the rectangle was filled with the same colour as the background, i.e., grey colour. Participants could administer the antidote by pressing the space bar while the litmus-paper colour was visible. After the colour disappeared, pressing the space bar had no effect on the amount of antidote given anymore. A scrollbar at the bottom centre of the screen indicated the amount of antidote that was being administered. If the spacebar was held pressed, the scrollbar face moved smoothly from left to right. The position of the face was translated into a number of antidote units from 1 through 100 displayed in a small textbox on the right of the scrollbar. The initial position of the face on each trial was the left extreme of the scrollbar, which corresponded with zero antidote units. Once the 3.5 s for the cue presentation had elapsed, participants received information about the outcome. The outcome included a photo of the plant with its corresponding name together with an indication of the type of plant (poisonous, anomalous, and innocuous) displayed in a small rectangle on the left of the scrollbar. Besides this, a message conveying the number of points won or lost was displayed in a small textbox just above the scrollbar. If the hypothetical patient was poisoned, the message displayed was "you win # points"; if the patient had eaten the anomalous plant, the message was "you lose # points"; if the patient had eaten the innocuous plant, the message was "you neither win nor lose points". The points referred to in the message were the amount of antidote units given by the participant. These points were added or subtracted to the accumulated points displayed in a small textbox located at the top right of the screen, at the same height as the litmus paper. The outcome information remained on the screen until participants pressed

the 'X' on the keyboard. Then all the outcome information was removed except the accumulated points, and the next trial followed 2 s later. Throughout, information about the context was displayed in a big rectangle as wide as the screen width at the centre of the screen. The message "ULUS (or GANTUAS) TRIBE" appeared in big size letters within the rectangle. The rectangle was filled either with red or with green colour depending on the training phase. The assignment of each colour to each phase was counterbalanced across participants. The context was displayed as a constant background that remained unchanged throughout the entire training phase.

Results and discussion

To ensure that only the data from those participants who paid attention to the task and showed a reasonable discriminative acquisition were include in the analysis, we decided to select those participants who met a certain learning criterion. Accordingly, participants in all the experiments were selected if the total responses to Cue C were less than 100 and the total score was higher than the mean minus two standard deviations at the end of Phase 1. As a result, 25 participants from the initial sample were selected for data analysis: 10 were from the same-outcome condition, and 15 were from the different-outcome condition.

Table 2 shows the mean response rate to Cue A at test. A one-tailed *t*-test was performed on participants' responses revealing a significantly lower response rate for the same-outcome group than for the different-outcome group, $t(23) = 3.8$, $p < .001$, indicating interference between cues.

The purpose of finding out whether our preparation was sensitive to interference between cues has been successfully achieved. It seems that the causal scenario together with the use of different contexts for each training phase have been particularly effective in favouring the interference effect. This is an interesting result because interference between cues has shown to be somewhat elusive. For example, Lipp and Dal Santo (2002) have reported several failures to observe interference between cues even when using a task that had produced interference before (Ortega & Matute, 2000). It remains to be shown, however, what the role of causal order is in inducing the interference effect.

EXPERIMENT 1B

After showing that our learning task can be used to induce interference between cues, Experiment 1b was conducted to test whether the effect can be modulated by the causal order of the task. Thus, we had two groups in this experiment: (a) a diagnostic group, similar to Experiment 1a; (b) a predictive group who received the information in a cause-effect order on each trial. According to our CR account, interference should only be observed in the diagnostic condition. According to associative theories, interference should be equivalent in both conditions.

Method

Participants, design, and apparatus

Seventy two psychology undergraduate students from University of Málaga took part in the present experiment for course credits. They were randomly assigned to four equally sized experimental groups ($N = 18$ for each group): a same-outcome predictive (SO-PR), a different-outcome

Table 2. *Means and standard errors from Experiments 1a, 1b, and 2*

	Exp. 1a	*Exp. 1b*		*Exp. 2*	
Causal order	Diagnostic	Diagnostic	Predictive	Diagnostic	Predictive
Group SO	23.7 (8.62)	45.71 (7.09)	57 (5.84)	31.17 (7.39)	54.95 (4.36)
Group DO	61.5 (5.76)	64.27 (3.18)	53.27 (3.76)	53.17 (4.42)	55.05 (4.08)

Note: SO and DO stand for same outcome and different outcome, respectively; standard errors are shown in parentheses.

predictive (DO-PR), a same-outcome diagnostic (SO-DG), and a different-outcome diagnostic (DO-DG) group. The task was performed using the same apparatus as for Experiment 1a.

Procedure

The procedure was identical to Experiment 1a except for some minimal changes regarding the predictive groups. Participants in these groups received the same causal scenario and the same cover story as participants in the diagnostic groups. However, they had to learn to decide the amount of antidote to be given from knowing what plant the patient ate (the cause of the patients' symptoms) instead of knowing what colour the litmus paper took on (the effect of the type of plant eaten). As in the diagnostic conditions, the cues, the contexts, and the outcomes were displayed at the top, at the centre, and at the bottom of the screen, respectively. The same plant photos and names that played the role of outcomes in Experiment 1a played the role of cues for the predictive groups. Specifically, each photo included a fictitious name for the plant, and was displayed, on each trial, for the same amount of time as cues in the diagnostic condition. The outcomes presented to participants were identical to those presented to the diagnostic groups except that a rectangle filled in with the corresponding litmus-paper colour was displayed on the left of the scrollbar instead of the plant photo. This rectangle was accompanied by the category label (poisonous, anomalous, or innocuous) for the plant eaten. Colours blue, yellow, and brown were used as litmus-paper colours. An equivalent counterbalancing scheme as in Experiment 1a was adopted for cues, outcomes, and contexts.

Results and discussion

Sixty three participants were selected for data analysis as a result of applying the learning criterion described above: 16 from the SO-PR group, 15 from the DO-PR group, 17 from the SO-DG group, and 15 from the DO-DG group. Meeting the learning criterion did not correlate with any of the manipulated factors, as the number of participants eliminated in each group was almost identical (2 from the SO-PR, 3 from the DO-PR, 1 from the SO-DG, and 3 from the DO-DG group).

Table 2 shows the mean response rate to Cue A at test for each group. A 2 (group: same-outcome group vs. different-outcome group) $\times$ 2 (causal order: predictive vs. diagnostic) analysis of variance yielded a significant interaction between group and causal order, $F(1, 59) = 4.26$, $p = .043$, $MSE = 457.68$. The remaining effects were non significant (largest $F = 1.9$, $p > .17$). A look at the means reveals that the difference between the same-outcome and the different-outcome group was considerably greater in the diagnostic than in the predictive condition. Moreover, the difference is in line with an interference effect in the diagnostic condition, whereas, if anything, there is a small trend in the opposite direction in the predictive condition. Thus, the interaction seems to mean that the interference between cues was observed in the diagnostic but not in the predictive condition. This interpretation was confirmed by an analysis of the simple effects of group in each level of causal order. The effect of the factor group was significant in the diagnostic condition, $F(1, 59) = 6.00$, $p = .017$, $MSE = 457.68$, but not in the predictive condition, $F(1, 59) < 1$, $MSE = 457.68$.

In summary, in Experiment 1b we have replicated the interference effect of Experiment 1a in the diagnostic group, and we have found evidence that the interference effect is influenced by the causal interpretation of the task. As predicted by our CR account, and at odds with associative theories' predictions, interference was found when cues and outcomes were interpreted as effects and causes, respectively, but not when they were interpreted as causes and effects.

Although the results of Experiment 1b suggest that interference between cues can be influenced by people's previous causal knowledge, some aspects of the procedure could raise doubts about our CR explanation. First, it could be argued that participants in the predictive group were not informed about the absence of Cue A in B-O1 trials, whereas this information was implicit for

participants in the diagnostic group because presenting the litmus paper in one colour necessarily implies that it did not take on the other colours. This difference could inspire different explanations of why interference was only found in the diagnostic group. For example, participants might have assumed the presence of Cue A in B-O1 trials unless otherwise indicated. Since the presence of Cue A was implicitly negated in the diagnostic condition, only participants in the predictive condition would have actually made this assumption. This would have had a consequence similar to training the A-O1 relationship in Phase 2 (or in Context Y), which explains why interference was not observed. From an associative point of view, we could say that the representation of Cue A could have been elicited in B-O1 trials by way of a backward association with Outcome 1, resulting in a sort of mediated learning in Context Y analogous to mediated conditioning in animals.

A second criticism could be that, in both causal order conditions, the outcome included information about the category label of the plant. Such category label could be easily understood as an essential property of the plant. Thus, for the predictive group, this category label informed participants about an essential property of the cue. This could be viewed as an aid to generalise responding to A from Phase 1 to the test phase because the essential properties of plants or whatever object are not supposed to change from one context to another. This would explain the absence of interference in the predictive group. Contrarily, in the diagnostic condition, the category label included as part of the outcome did not refer to an essential property of the cue (the litmus paper) but to an essential property of another part of the outcome (the plant). Thus, participants in the diagnostic condition were not induced to generalise from Phase 1 to the test phase to the same extent as participants in the predictive group.

EXPERIMENT 2

Experiment 2 was conducted to achieve two main purposes. First, we changed some aspects of the procedure to overcome the shortcomings referred to above. Second, we intended to replicate the asymmetry observed in Experiment 1b to obtain more compelling evidence of the influence of the causal interpretation of the task on the interference effect.

To overcome the problem of implicitly informing participants about the absence of Cue A in B-O1 trials in the diagnostic but not in the predictive condition, we made two changes. First, participants in both causal order conditions were explained through instructions that there were three chemical reagents of different colours, instead of a unique litmus paper, that could react as a consequence of the type of plant eaten. Thus, each trial informed about the reagent altered in a hypothetical case, and participants had to learn which reagent was altered as a consequence of eating each type of plant. Second, the instructions for the diagnostic group stated that, on each trial, participants would be presented with the only reagent altered. The reagents not presented were not altered. The same strategy was used regarding the information about the plant eaten in the predictive condition. Instructions stated that, on each trial, participants would be presented with the only plant eaten. The plants not presented were not eaten by the hypothetical patients. All this was explicitly stated on each trial. This way, both causal order groups were equivalent regarding the information received about absent cues on each trial.

To solve the problem related to presenting the category label, we decided to eliminate this information from the outcome in both causal order conditions. Instead, information about the state of the patient as a consequence of having eaten the plant was added as part of the outcome.

Method

Participants, design, and apparatus

Ninety nine psychology undergraduate students from University of Málaga took part in the present experiment for course credits. They were randomly assigned to four groups resulting in the following distribution: $N = 26$ for group

SO-PR, $N = 24$ for group DO-PR, $N = 24$ for group SO-DG, and $N = 25$ for group DO-DG. The task was performed on the same apparatus as in the previous experiments.

Procedure

The procedure was identical to Experiment 1b except for the following aspects. As said before, each trial informed about the only chemical reagent that had been altered, and about the only plant eaten by the hypothetical patient. Thus, on each trial, participants in the diagnostic condition could read the message "In the case of Patient # the only reagent altered was" at the top of the screen. As in the previous experiments, a rectangle filled with the corresponding reagent colour was displayed below the message. The information about the plant eaten was displayed at the bottom left of the screen and consisted of a plant photo and a fictitious name at the top of the photo. In the predictive condition, participants could read the message "In the case of Patient #, the only plant eaten was" at the top of the screen on each trial. Below the message, the corresponding plant photo was displayed, accompanied by a fictitious name. The rectangle representing the reagent was displayed at the bottom left of the screen below the message "Reagent altered". In both causal order conditions, a message indicating the state of the patient as a consequence of having eaten the plant was added as part of the outcome. This message was displayed in the same small textbox as the points won or lost by the participant. If the patient had eaten a poisonous plant, the message was "poisoned patient"; if the plant eaten was anomalous, the message was "sensitised patient"; and if the plant eaten was innocuous, the message was "healthy patient". The meaning of these terms was clearly stated to participants through instructions.

Results and discussion

Seventy eight participants were selected for data analysis according to the same learning criterion as in the previous experiments: 22 from the SO-PR group, 21 from the DO-PR group, 17 from the SO-DG group, and 18 from the DO-DG group. The number of participants whose data did not enter the analysis was distributed across groups as follows: 4 from the SO-PR, 3 from the DO-PR, 7 from the SO-DG, and 7 from the DO-DG group. Because the percentage of participants selected from the predictive group was higher than from the diagnostic group, we performed a chi-square test for independence between causal order and being selected or not. The analysis revealed that the distribution of participants per cell did not significantly depart from an independence distribution [$\chi = 3.1439 < \chi_{.05}(1) = 3.841$]. In other words, the probability of being selected did not significantly differ between participants in the predictive and in the diagnostic groups.

Table 2 shows the mean response rate to Cue A at test for each group. A 2 (group: same-outcome group vs. different-outcome group) × 2 (causal order: predictive vs. diagnostic) analysis of variance yielded a significant effect of group, $F(1, 74) = 4.75$, $p = .033$, $MSE = 495.31$, a significant effect of causal order, $F(1, 74) = 6.41$, $p = .013$, $MSE = 495.31$, and a significant interaction between group and causal order, $F(1, 74) = 4.67$, $p = .034$, $MSE = 495.31$. A look at the means reveals a clear trend toward an interference effect in the diagnostic group, whereas the means in the predictive group are almost identical. This was confirmed by an analysis of the simple effects of group within each causal order condition. The effect of group was significant in the diagnostic condition, $F(1, 74) = 8.54$, $p = .005$, $MSE = 495.31$, but not in the predictive condition, $F(1, 74) < 1$, $MSE = 495.31$. The significant and robust interference effect found in the diagnostic condition clearly explains the main effects of group and causal order.

According to these results, we have achieved the main purposes pursued in this experiment. First, we have replicated the interaction between causal order and group found in Experiment 1b. Like in that experiment, the interference effect was observed in the diagnostic but not in the predictive condition. Second, we have shown that neither of

the two procedural aspects of Experiment 1b that were changed in Experiment 2 can explain the interaction between group and causal order. Specifically, the interaction did not vanish due to (1) making explicit the absence of cues in the two causal order conditions, nor by (2) eliminating the category label for the plant from the outcome. Thus, the results found in this experiment add more compelling evidence for the influence of the causal interpretation of the task and of previous causal knowledge on interference between cues. Such results cannot be accounted for by associative learning theories because, as previously mentioned, these theories do not differentiate between predictive and diagnostic processing of the cue-outcome relationships.

GENERAL DISCUSSION

Summary and comments on alternative explanations

According to our CR account, interference between different cues of the same outcome takes place as a consequence of the causal interpretation of the learning task together with the use of previous causal knowledge. If the task is interpreted as a diagnostic learning task, participants need to postulate the existence of some hidden causal factor with different states for each training phase to explain why Cause 1 produces Effect A in Phase 1 and Effect B in Phase 2. The state of the hidden cause would enable or disable the capacity of Cause 1 to produce either A or B. Interference would occur because participants assume, when tested on Cue A, that the hidden causal factor is in a state that does not allow Cause 1 to produce Effect A. From our CR account it is predicted that interference between cues should not be observed if the task is given a predictive interpretation, i.e., if cues and outcomes are interpreted as causes and effects, respectively. Experiments 1a through 2 were intended to provide evidence consistent with this prediction. The purpose of Experiment 1a was to find out whether interference between cues could be obtained with a new learning task framed within a causal scenario that allowed for both predictive and diagnostic learning with minimal changes. For this, all participants learned in a diagnostic condition, and different contextual cues were used for each training phase. Since we were successful at inducing the interference effect, the causal order was manipulated in Experiment 1b to test whether the interference effect is influenced by the causal interpretation of the task. The data analysis yielded a significant interaction between causal order and interference that resulted from the interference effect being significant in the diagnostic but not in the predictive condition. Finally, the results of Experiment 1b were replicated in Experiment 2 after changing two aspects of the procedure to rule out some alternative explanations. Experiment 2 provided additional evidence of the influence of causal knowledge on interference between cues. Overall, the results of Experiments 1b and 2 are consistent with our CR account, and cannot be accounted for by associative learning theories.

A possible objection to Experiments 1b and 2 could be that the stimuli used as cues and outcomes in the diagnostic condition were different to those used in the predictive condition. For example, in the diagnostic condition, the cues consisted of coloured rectangles, whereas in the predictive condition, they consisted of plant photos labelled with fictitious names. It could be argued that learning to discriminate between familiar colours is not the same as learning to discriminate between unknown plant-photos labelled with fictitious names. For one thing, in the latter case, participants have to discriminate between more complex stimuli consisting of a greater amount of features than in the former case. But if this difference were important enough to explain participants' responses at test, it should also have a detectable influence on, for example, the acquisition of discrimination over the course of training. After all, the test trial was one more training trial at the end of the learning task. The data concerning response acquisition in Experiment 2 contradict this prediction. Figure 1 shows the response rate per trial for each cue-outcome pairing in

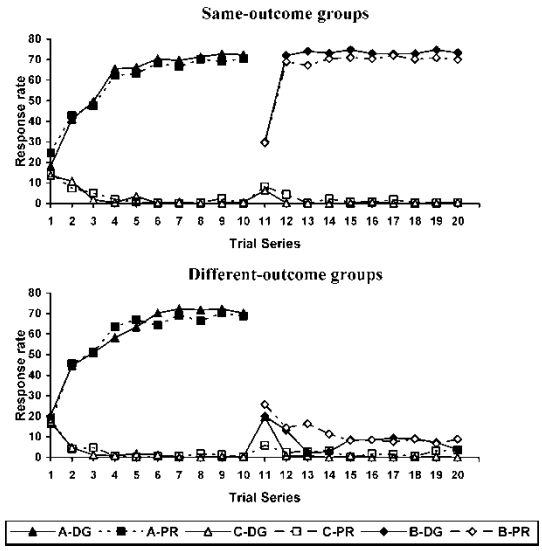

Figure 1. *Response-rate curves for each group and for each cue along the training phases of Experiment 2.*

each training phase. It shows that the acquisition curve in the diagnostic group was virtually identical to that in the predictive group. The results of the statistical analyses confirmed this impression. We first averaged responses to each cue along each training phase in three trial blocks. The first, second, and third block comprised trials 1 through 4, 5 through 7, and 8 through 10, respectively.[4] We conducted separate analyses for each phase and each cue. In Phase 1, for Cue A, a 2 (group: same-outcome group vs. different-outcome group) × 2 (causal order: predictive vs. diagnostic) × 3 (trial block 1 vs. trial block 2 vs. trial block 3) analysis of variance with trial block as a within-subjects factor yielded a significant effect of trial block, $F(2, 148) = 179.31$, $p < .001$, $MSE = 91.31$. None of the remaining effects were significant (all Fs < 1). The same analysis of variance on participants' responses to Cue C yielded a significant effect of trial block, $F(2, 148) = 50.00$, $p < .001$, $MSE = 18.79$. None of the remaining effects were significant (all Fs < 1). Consequently, the predictive and the diagnostic groups were not significantly different regarding the acquisition of the discrimination between Cues A and C along Phase 1. Because responses to Cue C did not differ between the predictive and the diagnostic groups in Phase 1, the same statistical analysis was only performed on responses to B in Phase 2. This analysis yielded a significant effect of group, $F(1, 74) = 545.60$, $p < .001$, $MSE = 363.35$, a significant effect of trial block, $F(2, 148) = 3.41$, $p < .036$, $MSE = 72.15$, and a significant interaction between group and trial block, $F(2, 148) = 25.47$, $p < .001$, $MSE = 72.15$. All of the remaining effects were non-significant (largest $F = 1.50$, $p > .24$). The interaction between trial block and group is hardly surprising if we take into account the different meaning of Cue B for each group. What really matters, however, is that causal order was not involved in any of the significant effects found. Thus, the predictive group did not differ from the diagnostic group regarding the acquisition of responses to B along Phase 2. In summary, the use of different stimuli as cues and outcomes in the predictive and the diagnostic groups did not produce any detectable influence on participants' response rates to any cue along the training phases of Experiment 2.[5] Moreover, the difference in stimuli did not influence participants' responses to A at test in the different-outcome group either. Consequently, given that the use of different stimuli did not significantly affect participants' responses in any of the multiple cases in which it could potentially have had an influence, it is very unlikely that such use of different stimuli could account for the interaction found in Experiments 1b and 2 between causal order and interference.

An alternative explanation of our results is that the calculation of the conditional Δp (Cheng & Holyoak, 1995; Cheng & Novick, 1990) in the same-outcome condition gives a value of 1 for

[4] The first block was allowed to include one more trial than the other two blocks to compensate for the greater response variability that is usually observed during the first training trials.

[5] The same analysis on participants' responses in Experiment 1b yielded virtually the same results.

the Cause A → Effect 1 causal relationship in the predictive condition, and a value of .5 for the corresponding Cause 1 → Effect A causal relationship in the diagnostic condition. In contrast, in the different-outcome condition, the value for Δp is 1 in both causal order conditions. This could explain the interaction found between causal order and the interference effect. Notice that, as our CR account, this explanation assumes that interference between cues is influenced by the causal interpretation of events because the conditional Δp is always computed in the cause-effect direction regardless of the order in which information is presented. The main difference between both accounts is that the latter assumes that, in the same-outcome diagnostic condition, B-O1 trials of Phase 2 are taken into account to compute the Cause 1 → Effect A causal relationship. In contrast, our CR account states that participants should estimate that Phases 1 and 2 differ with respect to a hidden causal factor. Thus, neither the B-O1 trials should enter into the calculation of the Cause 1 → Effect A causal relationship, nor vice versa. In other words, participants should estimate that what is learned in Phase 2 does not generalise to Phase 1, and that what is learned in Phase 1 does not generalise to Phase 2. The main problem with the alternative Δp explanation is that it turns out to be implausible in the light of the evidence provided in other studies. Specifically, such explanation is incompatible with studies showing that interference between cues vanishes when certain contextual manipulations are made (Matute & Pineño, 1998; Pineño & Matute, 2000; Pineño et al., 2000), or when a delay is introduced between the second training phase and the test (Pineño et al., 2000). There is also evidence that interference between cues vanishes when A-O1 and B-O1 trials are intermixed in a single training phase (Pineño et al., 2000). Even though this result is not necessarily incompatible with the alternative Δp explanation, it poses an additional difficulty. In contrast, as we will see in the next paragraphs, our CR account is highly consistent with these findings.

Explaining the main findings of previous studies on interference between cues from the CR account

An interesting point of the CR explanation developed in this article is that it provides a parsimonious account for the main findings in research on interference between cues. In this respect, there are three points worth mentioning. First, the data obtained to date are consistent with the prediction that interference between cues should be found in the diagnostic rather than in the predictive direction. Apart from our learning task, there are only three tasks that have been used to obtain interference between cues, even though with different frequency and success. The most frequently and successfully used is the spy-radio task (Castro, Ortega, & Matute, 2002; Escobar et al., 2002; Pineño & Matute, 2000; Pineño et al., 2000). In this task, participants have to rescue as many refugees as possible in a war zone plagued with hidden mines by repeatedly pressing the space bar to place the refugees in a series of trucks. The coloured lights of a spy radio, which play the role of cues, tell the state of the road (free of mines, mined, or closed) on each trial. In this case, the causal scenario is very likely to induce participants to interpret the coloured lights of the spy radio as effects of the state of the road on some kind of detector device. After all, people are very familiar with the existence of mine detectors and devices that can detect metals or other sort of materials. In any case, it is very unlikely for participants to have attributed any causal power to the illumination of the lights to produce one road state or another. A similar analysis could be made regarding the Martians task (Lipp & Dal Santo, 2002; Matute & Pineño, 1998; Ortega & Matute, 2000), and, to a lesser extent, the Air-Force task (Escobar et al., 2002: Experiment 1). In any case, there is no evidence for interference in learning tasks framed in a clear predictive causal scenario.

A second feature of interference between cues is its tendency to vanish when some contextual manipulations are made. As said in the introduction, in whatever sense the term "context" may

be used (see Footnote 2), results from different studies converge in that interference between cues decreases (or disappears) if the test takes place in a different context than that of the second training phase (Matute & Pineño, 1998; Pineño & Matute, 2000; Pineño et al., 2000). From our CR account, the reason for this is that the context comes to be causally relevant in the same-outcome condition when the learning task is given a diagnostic interpretation. As said in the introduction, the causal relevance of the context lies in its informative value for the different states of the hidden causal factor that enables Cause 1 to produce either Effect A or B. Thus, if the test phase takes place in the same context as Phase 2, participants infer that the hidden causal factor is in the state which enables Cause 1 to produce Effect B rather than A. If the context changes from Phase 2 to the test phase, there are two possibilities: (a) some participants might infer that the hidden causal factor has also changed to the state which enables Cause 1 to produce effect A; (b) others might be uncertain about the state of the hidden causal factor. In both cases, however, an increase in the response rate to A at test should be expected with respect to when participants are tested in the same context as for the second training phase.

Finally, a third point regarding previous studies on interference between cues that deserves our attention is, as mentioned above, the absence of interference when A-O1 and B-O1 trials are intermixed in a single training phase (see Pineño et al., 2000, Experiment 2). According to our CR account, if participants learn in a diagnostic situation, they should also infer the existence of a hidden causal factor to explain why Cause 1 produces Effects A and B in a disjoint manner. However, contrary to what is the case when participants are trained in two different phases, the context would be the same for both trial types, and, thus, it would not convey any information about the state of the hidden causal factor. At the same time, because of the training regime, participants would get used to expect rapid successive changes in the state of the hidden causal factor. Thus, when participants are tested with A, they can be sure that Cause 1 has occurred, and no interference effect should be observed.

In light of these considerations, it seems that our CR account is highly consistent with the main findings on interference between cues. However, this does not mean that the effects of interference between cues found in previous studies are necessarily caused by the processes envisaged by CR theories. It could well be that associative and CR processes are both responsible for interference between cues. In fact, interference between cues has also been found in studies with animals, who are not supposed to enjoy the competencies postulated by CR theories (but see Beckers, Miller, De Houwer, & Urushihara, 2006). At the same time, the idea that both associative and CR processes could be responsible for a common learning phenomenon is gaining increasing acceptance as a consequence of evidence obtained in other experimental paradigms (Cobos, López, Caño, Almaraz, & Shanks, 2002; Le Pelley, Oakeshott, & McLaren, 2005; López et al., 2005; Price & Yates, 1995; Tangen & Allan, 2004). Consequently, rather than invalidating the associative explanations, the experiments reported here show the viability of alternative accounts of interference between cues based on CR processes. In future studies, both theoretical approaches should be borne in mind when offering an account of the empirical phenomena obtained.

Extending the CR approach to interference between different outcomes of the same cue

A straightforward consequence of our account is that the same CR processes responsible for the interference between cues found in our experiments might also be responsible, at least in some circumstances, for interference between different outcomes of the same cue. In fact, the causal structure of the data for the same-outcome diagnostic condition used in our experiments is the same as for the interference groups of experiments on interference between outcomes. The most typical studies on interference between outcomes in humans have used causal learning tasks in which a given cause, Cue A, is paired with one effect

during Phase 1 (Outcome 1), and with another effect (or the absence of the previous one) during Phase 2 (García-Gutiérrez & Rosas, 2003a, 2003b; Matute, Vegas, & De Marez, 2002; Rosas et al., 2001; Vadillo et al., 2004). Thus, according to our CR account of interference between cues, participants should infer the existence of a hidden causal factor that interacts with the target cause, Cue A, to produce either O1 (during Phase 1) or O2 (during Phase 2). Again, the context present in each phase would serve as a signal of the different states of the hidden factor. Interference would occur if the context present at test is the same as (or similar to) the context of Phase 2 because, in such a case, participants would infer that the hidden causal factor is in the state that enables Cue A to produce O2 but not O1. Because associative theories also predict interference between outcomes, it turns out that, to extend our CR account to this kind of interference, we need to empirically discriminate between these different explanations. We are currently working on a design of interference between outcomes that will allow for such discrimination.

CONCLUDING COMMENTS

The results obtained in our study are consistent with studies showing the influence of causal knowledge on cue-interaction effects in causal learning experiments. Some of these studies, as in our case, have searched for asymmetries between predictive and diagnostic situations regarding cue-interaction effects as a means to provide evidence supporting CR theories (López et al., 2005; Tangen & Allan, 2004; Waldmann, 2000, 2001; Waldmann & Holyoak, 1992; Waldmann & Walker, 2005). In this sense, our research may be viewed as an extension of these strategies and findings to the area of interference between cues. But at the same time, our study differs in a relevant respect that goes beyond the experimental paradigm used or the phenomenon tackled here. Most of the studies on cue-interaction effects have focussed on how people distinguish between true and spurious causal relationships. At odds with this, our study has focussed on how people calculate causal structure from the probability distribution of events, and how the inferred causal structure determines whether or not the data collected in one context generalise to another one. Some advocates of causal Bayes nets have claimed that calculating causal structure is at the core of causal learning and may be viewed as the main use for causal knowledge (Griffiths & Tenenbaum, 2005; Sloman & Lagnado, 2004).

Finally, we hope to have shown the relevance of the interference between cues phenomenon to the study of causal learning. The main idea we have tried to pursue throughout the experimental work reported here was to show that understanding causal learning requires an understanding of how people assess whether or not the evidence gathered in one context generalises to other contexts. Our experiments have shown that interference paradigms may be a useful tool for such understanding. Thus, we think that future causal learning research will greatly benefit from more empirical work on interference effects. In line with this, a relevant research question would be to evaluate whether similar results to those reported here are obtained using a different interference paradigm like that involved in the effect of interference between different outcomes.

REFERENCES

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

Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: Forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference. *Journal of Experimental Psychology: General*, *135*, 92–102.

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

Castro, L., Ortega, N., & Matute, H. (2002). Proactive interference in human predictive learning. *International Journal of Comparative Psychology*, *15*, 55–68.

Cheng, P. W., & Holyoak, K. J. (1995). Complex adaptive systems as intuitive statisticians: causality, contingency, and prediction. In J. A. Meyer & H. Roitblat (Eds.), *Comparative approaches to cognition* (pp 271–302). Cambridge, MA: The MIT Press.

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545–567.

Cheung, C. G., & 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.

Cobos, P. L., López, F. J., Caño, A., Almaraz, J., & Shanks, D. R. (2002). Mechanisms of predictive and diagnostic causal induction. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 331–346.

Escobar, M., Arcediano, F., & Miller, R. R. (2001). Conditions favoring retroactive interference between antecedent events (cue competition) and between subsequent events (outcome competition). *Psychonomic Bulletin & Review*, *8*, 691–697.

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.

García-Gutiérrez, A., & Rosas, J. M. (2003a). Context change as the mechanism of reinstatement in causal learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *29*, 292–310.

García-Gutiérrez, A., & Rosas, J. M. (2003b). Empirical and theoretical implications of additivity between reinstatement and renewal after interference in causal learning. *Behavioural Processes*, *63*, 21–31.

Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychology Review*, *111*, 1–30.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology, 51,* 354–384.

Hagmayer, Y & Waldmann, M. R. (2007). Inferences about unobserved causes in human contingency learning. *Quarterly Journal of Experimental Psychology*, *60*, 330–355.

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.

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.

Lipp, O. V., & Dal Santo, L. A. (2002). Cue competition between elementary trained stimuli: US miscuing, interference, and US omission. *Learning & Motivation*, *33*, 327–346.

López, F. J., Cobos, P. L., & Caño, A. (2005). Associative and causal reasoning accounts of causal induction: Symmetries and asymmetries in predictive and diagnostic inferences. *Memory & Cognition*, *33*, 1388–1398.

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

Matute, H., Vegas, S., & De Marez, P. J. (2002). Flexible use of recent information in causal and predictive judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 714–725.

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.

Ortega, O., & Matute, H. (2000). Interference between elementally trained stimuli can take place in one trial. *Learning & Motivation 31*, 323–344.

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

Pineño, O., & Matute, H. (2000). Interference in human predictive learning when associations share a common element. *International Journal of Comparative Psychology*, *13*, 16–33.

Pineño, O., Ortega, N., & Matute, H. (2000). The relative activation of the associations modulates interference between elementally trained cues. *Learning and Motivation*, *31*, 128–152.

Price, P. C., & Yates, J. F. (1995). Associative and rule-based accounts of cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1639–1655.

Rosas, J. M., & Bouton, M. E. (1998). Context change and retention interval have additive, rather than interactive, effects after taste aversion extinction. *Psychonomic Bulletin & Review*, *5*, 79–83.

Rosas, J. M., Vila, N. J., Lugo, M., & López, L. (2001). Combined effect of context change and retention interval on interference in causality judgment. *Journal of Experimental Psychology: Animal Behavior Proceses*, *27*, 153–164.

Schwartz, M. (1968). Effect of stimulus class on transfer and RI in the A–B, A–C paradigm. *Journal of Verbal Learning & Verbal Behavior*, *7*, 189–195.

Sloman, S., & Lagnado, D. (2004). Causal invariance in reasoning and learning. In Brian H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 44, pp. 287–325). San Diego, CA: Academic Press.

Tangen, J. M., & Allan, L. G. (2004). Cue-interaction and judgments of causality: contributions of causal and associative processes. *Memory & Cognition*, *32*, 107–124.

Vadillo, M. A., Vegas, S., & Matute, H. (2004). Frequency of judgment as a context-like determinant of predictive judgments. *Memory & Cognition*, *32*, 1065–1075.

Van Hamme, L. J., Kao, S.-F., & Wasserman, E. A. (1993). Judging interevent relations: from cause to effect and from effect to cause. *Memory & Cognition*, *21*, 802–808.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 53–76.

Waldmann, M. R. (2001). Predictive versus diagnostic causal learning: evidence from an overshadowing paradigm. *Psychonomic Bulletin and Review*, *8*, 600–608.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222–236.

Waldmann, M. R., & Walker, J. M. (2005). Competence and performance in causal learning. *Learning & Behavior*, *33*, 211–229.