Memory & Cognition
2005, 33 (8), 1388-1398

Associative and causal reasoning accounts of causal induction: Symmetries and asymmetries in predictive and diagnostic inferences

FRANCISCO J. LÓPEZ, PEDRO L. COBOS, and ANTONIO CAÑO
University of Málaga, Málaga, Spain

Associative and causal reasoning accounts are probably the two most influential types of accounts of causal reasoning processes. Only causal reasoning accounts predict certain asymmetries between predictive (i.e., reasoning from causes to effects) and diagnostic (i.e., reasoning from effects to causes) inferences regarding cue-interaction phenomena (e.g., the overshadowing effect). In the experiments reported here, we attempted to delimit the conditions under which these asymmetries occur. The results show that unless participants perceived the relevance of causal information to solving the task, predictive and diagnostic inferences were symmetrical. Specifically, Experiments 1A and 1B showed that implicitly stressing the relevance of causal information by having participants review the instructions favored the presence of asymmetries between predictive and diagnostic situations. In addition, Experiment 2 showed that explicitly stressing the relevance of causal information by stating the importance of the causal role of events after the instructions were given also favored the asymmetry.

Inference making serves the important function of reducing uncertainty about the occurrence of events. In causal contexts, predictive and diagnostic inferences may be distinguished. When people make predictive inferences, they reason from causes to effects. For example, from knowledge that it is raining, we may anticipate that the pavement will be wet. Inversely, when people make diagnostic inferences, they reason from effects to causes. For example, from knowledge that the pavement is wet, we may infer that it has been raining. In both cases, people can reduce their uncertainty concerning one event by knowing of the presence of another event with which a causal relationship is known to exist.

The most influential accounts provided for causal learning and inference making in the last two decades are probably those based on either associative learning mechanisms or causal reasoning processes. Both types of accounts make very different assumptions about the nature of the processes through which inferences are made. According to associative theories, inferences are the output of informatively encapsulated, bottom-up processes whereby knowledge is evoked or triggered automatically by cues from the environment (see, e.g., Rescorla & Wagner, 1972). According to causal reasoning models, inferences are the output of elaborate, cognitively demanding reasoning processes. These are top-down processes sensitive to various forms of knowledge (as proposed, for example, by causal model theory; see Waldmann & Holyoak, 1992).

Numerous studies have been carried out to evaluate the predictions of both types of theories. One common strategy has focused on their predictions concerning differences between predictive and diagnostic inferences in cue-interaction situations. In these situations, individuals are presented with a scenario in which several cues occur simultaneously so that the predictive value of any one of them is relative to the predictive value of all the cues with which it has co-occurred. Typically, in studies in which this strategy is followed, predictive inferences were requested when, according to the instructions provided, cues played the role of causes and their outcomes played the role of effects. On the other hand, diagnostic inferences have been requested in situations in which cues played the role of effects and outcomes played the role of causes.

According to associative accounts, cue-interaction phenomena occur whenever various cues co-occur prior to a target outcome no matter what their causal role is—that is, cue interaction should be found in both predictive and diagnostic inferences. Conversely, according to causal reasoning models, cue-interaction phenomena should be found only when various causes co-occur and not when various effects co-occur—that is, they should be found in predictive but not in diagnostic inferences (for extensive discussions of the reasoning behind such predictions, see Cobos, López, Caño, Almaraz, & Shanks, 2002; De Houwer, Beckers, & Glautier, 2002; López, Cobos, Caño, &

The research described here was supported by Research Grants BSO2000-1216 from the Dirección General de Investigación and HUM-0105 from the Junta de Andalucía, Spain. We thank Julián Almaraz, David Luque, David R. Shanks, Michael Waldmann, and two anonymous reviewers for their helpful comments on previous versions of the manuscript. Correspondence concerning this article should be addressed to F. J. López, Departamento de Psicología Básica, Campus de los Teatinos, Universidad de Málaga, Málaga 29071, Spain (e-mail: frjlopez@uma.es).

Shanks, 1998; Waldmann, 2001; Waldmann & Holyoak, 1992).

Despite the numerous studies that have been carried out following this strategy, the results obtained are far from conclusive. In fact, some of these studies show cue-interaction effects in both predictive and diagnostic inferences, whereas others show cue interaction in predictive inferences only (for illustrations of these diverse results, see Cobos et al., 2002; De Houwer et al., 2002; Matute, Arcediano, & Miller, 1996; Price & Yates, 1995; Shanks & López, 1996; Tangen & Allan, 2004; Van Hamme, Kao, & Wasserman, 1993; Waldmann, 2000, 2001; Waldmann & Holyoak, 1992).

The objective of the present study is to evaluate a hypothesis that may help to explain the reasons for these conflicting results. We turn to Cobos et al. (2002, Experiment 4) and Waldmann (2001, Experiment 2) for examples of these conflicting results despite their use of similar experimental procedures. In both cases, overshadowing was the cue-interaction paradigm used. In this paradigm, participants are typically presented with two main trial types during a learning phase: AB $\rightarrow O_1$ (i.e., Cues A and B are presented together and followed by Outcome 1) and C $\rightarrow O_2$ (i.e., Cue C is followed by Outcome 2). Once this learning phase is completed, during an inference phase participants are asked to evaluate the relationships between cues and outcomes. The overshadowing effect, initially observed in animal conditioning by Pavlov, consists in the significant trend to perceive a weaker relationship between A (or B) and O_1 than between C and O_2 even though A (or B) has been paired with O_1 as many times as C has been paired with O_2. In other words, the fact that Cue A always co-occurs with another cue (B) associated with the same outcome (O_1) seems to undermine perception of the relationship between Cue A (or Cue B) and Outcome O_1.

Also in both cases, the cover story presented in the instructions served to manipulate the causal role that cues and outcomes played (although the cover story was different in the two experiments). Specifically, as was mentioned above, in the predictive group cues and outcomes played the roles of causes and effects, respectively, and the inferences requested were predictive inferences. On the other hand, in the diagnostic group cues and outcomes played the roles of effects and causes, respectively, and the inferences requested were diagnostic inferences.

However, the results obtained were theoretically unclear. Cobos et al. (2002) showed an equivalent overshadowing effect in both predictive and diagnostic inferences, as is predicted by associative theories; Waldmann (2001), in contrast, found the overshadowing effect in predictive but not in diagnostic inferences, in consistency with causal reasoning accounts.

At present, the origin of this difference in results is not known. In fact, the procedures adopted in the two studies differ in various ways (e.g., different cover stories were used, and the information was presented in different ways throughout the training phases), so it would be premature to attribute the difference in the results obtained to any one of these procedural differences in particular. Thus, our objective is not to evaluate exhaustively the procedural factor or factors that may explain the different results obtained, but rather to test a more general hypothesis that is related to a difference between the procedures adopted in the two experiments. According to this hypothesis, the use of causal knowledge predicted by causal reasoning models (i.e., considering the causal role that cues and outcomes play in inference making) is restricted to cases in which individuals notice that such knowledge is relevant to correctly performing the experimental task. The reason for this is that, according to causal reasoning models, the use of causal knowledge must involve the activation of costly cognitive processes (Waldmann & Holyoak, 1992). For example, it involves building up a causal mental model of the situation from the learning input provided or accessing this mental model flexibly to make inferences (see De Houwer & Beckers, 2003, or Waldmann & Hagmayer, 2001, for empirical evidence showing the large working memory demands involved in these mental operations). Thus, it follows from our hypothesis that individuals will tend to use less demanding cognitive strategies, such as those envisaged by associative models, unless they clearly perceive that the causal information provided is relevant to solving the task. Henceforth, we will refer to this hypothesis as the *perceived relevance hypothesis*.

In this sense, Waldmann (2001, Experiment 2) and Cobos et al. (2002, Experiment 4) differ in that only the former used a procedure including some features that participants could have understood as encouragement to consider the relevance of the causal information provided by the instructions. Specifically, Waldmann's (2001) participants were asked to explain the causal situation (i.e., the causal role that cues and outcomes played) by using a picture model that was handed to them after they had read the instructions, and corrective feedback was provided when needed. Thus, it is reasonable to assume that this procedure simultaneously focuses participants' attention on just that causal information that causal reasoning models regard as relevant to the extent that it is the only aspect of the instructions that is reviewed. Consequently, the participants could have regarded this procedural feature as an implicit demand to use the causal information provided, which would explain the asymmetries found between predictive and diagnostic inferences. On the other hand, no aspect of Cobos et al.'s (2002) procedure can be viewed as equivalent to this aspect of Waldmann's (2001) procedure.

The objective of the present study is to test this hypothesis empirically. The instructions of Experiment 1A cannot be viewed as focusing participants' attention on the causal information provided and, hence, as an invitation to use it. Accordingly, those cognitive processes that use this causal information will not be activated, and much simpler processes, such as associative processes, will come into play instead. Regarding the overshadowing effect, cue interaction should be obtained in both predictive and diagnostic inferences. In Experiment 1B, part of the instructions focused participants' attention on the causal information in a way similar (though not identical) to that of Waldmann (2001, Experiment 2). Thus, participants should perceive the relevance of this causal information, and, consequently, processes that use such information should be activated as predicted by causal reasoning models. Regarding the overshadowing effect, it should be obtained only in the predictive situation, and not in the diagnostic situation.

In Experiment 2, we tested the perceived relevance hypothesis in a different way. In our first two experiments, we tested the hypothesis by having participants review those aspects of the instructions that were related to the causal information provided. This has been interpreted as an implicit indication of the relevance of this information to solving the task. In Experiment 2, the relevance of the causal information to solving the task was made explicit without requiring a review of the instructions provided. Thus, our hypothesis was tested in a much more direct way.

EXPERIMENT 1A

The purpose of the experiment was to evaluate whether the causal order of the task (predictive vs. diagnostic) interacts with the overshadowing effect when no aspect of the instructions can be regarded as focusing participants' attention on the relevant causal information. That is, our goal is to determine whether or not, under these procedural circumstances, the overshadowing effect is equivalent in predictive and diagnostic inferences.

In order to make minimal changes to the original experiment of Waldmann (2001, Experiment 2) and, hence, to make the results we obtained comparable to Waldmann's (2001), in this experiment (and throughout the experimental series) we used the same causal scenario. In this scenario, participants learned about the functioning of an electrical box. On one side of the box (the "cause" side), there were switches, and on the other side (the "effect" side) there were bulbs. Participants learned how the switches and the bulbs were causally connected and were required to make either predictive (from causes to effects, or CE) or diagnostic (from effects to causes, or EC) inferences concerning these causal relationships (see the Procedure section of this experiment for further details concerning the causal scenario). Given that, unlike in Waldmann's (2001) procedure, the relevant causal information was not reviewed during the instructions, to know whether participants had correctly understood this causal scenario, at the end of the experimental task (i.e., once inferential judgments had been made) the participants were administered a comprehension questionnaire that focused on the causal role of the events involved.

Method

Participants and Apparatus. A total of 40 undergraduate psychology students from the University of Málaga took part in the present experiment for course credit. The task was performed on IBM-PC compatible computers in semi-isolated individual cubicles. A paper-and-pencil questionnaire was administered to evaluate the participants' understanding of the instructions provided.

Procedure. After the participants read the instructions on the computer screen and became familiar with the different parts of the experimental task through a pretraining task, the learning phase started. The participants had to learn causal relationships involving switches and bulbs in an electrical box. In the predictive (CE) condition, the front (or visible) side of the box included three lights of different colors and three different switches, each of which turned on its corresponding light. The back (or invisible) side of the box included two numbered bulbs. According to the causal scenario, the colored lights on the front side (cues) caused the illumination of the bulbs on the back side (outcomes).

In the diagnostic (EC) condition, the front (or visible) side of the box included three bulbs of different colors. The back (or invisible) side of the box included two different numbered lights and two different switches, each of which turned on its corresponding light. According to the causal scenario described, the illumination of the bulbs on the front side (cues) was caused by the lights on the back side (outcomes; see the Appendix for the specific instructions used in both CE and EC conditions).

After reading the instructions, the participants were invited to ask questions concerning any aspect of the instructions. Once all their questions were resolved, the training phase began.

A total of three trial types were included in the experiment: one per experimental condition (overshadowing, $AB \rightarrow O_1$; control, $C \rightarrow O_2$) plus a filler trial type ($\varnothing \rightarrow \varnothing$) in which the participants could see that no outcome occurred in the absence of a cue. On each trial, the participants could see the relevant cue or cues and had to predict the correct outcome. Once they had made a choice, feedback was provided: The participants could read at the bottom of the screen what the right choice was for that particular trial. Incorrect choices were accompanied by a beep. Specifically, on each trial the participants read, in the upper part of the screen, a question of the form: "In example '#' the following colored [CE, lights; EC, bulbs] have illuminated . . ."; then they read the cues programmed for that trial. After that, on the same screen but in a lower position, the participants were prompted to choose, from a list, among the two possible outcomes for that trial (numbered bulbs and numbered lights for the predictive and diagnostic conditions, respectively) and the absence of any outcome. The different cues were assigned different color names, identified by initial[1] (i.e., B for *blanco* [white], F for *fucsia* [fuchsia], and E for *esmeralda* [emerald]), and the outcomes were assigned different numbers. Specifically, B, F, and E were assigned to Cues A, B, and C, respectively. All possible assignments of color names were made for the overshadowing and control condition cues. Orthogonally, for half of the participants each outcome was numbered 1 or 2 (Outcome O_1 or Outcome O_2, respectively), and for the other half the outcomes were reversed.

In addition, from Trial 12 onward, each trial included explicit visual and numerical information about the percentage of correct responses across the last 12 trials. The percentage of correct responses was indicated by a yellow rectangle in the upper right-hand corner of the screen. The sequence of trials was randomized for each participant with the single restriction of being presented in blocks of 3 trials that included every trial type, although such blocks were not marked to the participants. A learning criterion was included so that the training phase ended whenever the participant had correctly responded on three complete blocks of trials.

Once the training phase had ended, the participants had to respond to a questionnaire that required predictiveness ratings of the target cues regarding the different outcomes. The questions were of the form, "To what extent does the colored [CE, light; EC, bulb '*X*'] predict the [CE, illumination of the different bulbs; EC, activation of the different lights]?" Although the participants were required to give predictiveness ratings for all cues regarding the different outcomes, only those judgments concerning the target $A \rightarrow O_1$, $B \rightarrow O_1$, and $C \rightarrow O_2$ relationships were analyzed. Due to the counterbalancing scheme followed, $A \rightarrow O_1$ and $B \rightarrow O_1$ ratings from the overshadowing condition were collapsed into a single rating, which will be referred to as the *$A \rightarrow O_1$ rating*. The participants gave these ratings on a scale of 0–100, divided in units of 10.

At the end of the experimental task, the participants were administered a questionnaire that evaluated their comprehension of the causal scenario (see the Appendix for details). Both questionnaires included a pictorial representation of the events involved in the box of lights and two questions, each with two alternative responses. The questions were designed specifically to determine whether or not the participants had understood the causal role (i.e., cause or effect) played by cues and outcomes within the scenario.

Results and Discussion

All the participants reached the learning criterion, and the mean numbers of trials used in the CE and EC groups did not differ (23.1 and 19.2, respectively). A 2 (overshadowing: $A \rightarrow O_1$ vs. $C \rightarrow O_2$) × 2 (causal order: CE vs. EC) ANOVA was performed on the participants' ratings. The main effect of overshadowing was significant [$F(1,38) = 27.11$, $MS_e = 461.15$, $p = .000$]. Neither the main effect of causal order nor the overshadowing × causal order interaction was significant (all Fs < 1). Thus, the overshadowing effect found was equivalent in both causal order conditions, contrary to the predictions of causal reasoning models and in keeping with those of the associative account.

However, the equivalence of the overshadowing effect found in both predictive and diagnostic inferences may simply reflect a lack of understanding of the causal scenario provided. Remember that, at variance with Waldmann (2001, Experiment 2), our instructions did not allow the participants to review the relevant causal information. Thus, the apparent insensitivity of inferences to causal knowledge may alternatively be explained by the participants' confusion about the causal scenario (e.g., by their confounding the causal roles played by cues and by outcomes in the box of lights). The inclusion of the final questionnaire allowed us to test this alternative hypothesis. In what follows, then, we reanalyze only the data of those participants who had correctly answered both of the questions included in the final questionnaire. Table 1 shows the mean ratings for this analysis. Again, the results showed a main effect of overshadowing [$F(1,24) = 21.32$, $MS_e = 548.72$, $p = .000$], and neither the main effect of causal order nor the overshadowing × causal order interaction was significant (both Fs < 1). This analysis shows that the participants' ratings were not sensitive to causal order even though they knew the roles played by cues and by outcomes during the task.

The analysis just described is important because it shows that the participants were able to remember the instructions provided concerning the causal scenario (as was shown by their performance on the comprehension questionnaire), although they did not use this causal knowledge when interpreting the learning input. Thus, within our paradigm, at least, participants do not seem to use their causal knowledge spontaneously during the task. We will turn to this again in the General Discussion section.

In addition, these data show that an identical number of participants was eliminated from the analysis in each causal order condition (i.e., 7 per condition), so the causal scenarios were of comparable levels of difficulty and the participants in both causal order groups remembered their respective causal scenario equally well.

In other words, in consistency with our hypothesis, there seems to be more to the ability to use the relevant causal information during the task than having this information. From the point of view of our hypothesis, it might be argued that if the relevance of its use had been suggested to the participants the results would have been different. This, in turn, might explain why Waldmann (2001, Experiment 2) revealed different judgments for predictive and diagnostic inferences. Waldmann (2001) used instructions that focused participants' attention on the relevant causal information, whereas our instructions did not.

Table 1
Design, Sample Size, Mean Ratings, and Standard Errors for Target Relationships in the CE and EC Groups of Experiments 1A and 1B

	Cue Interaction			
	Overshadowing ($AB \rightarrow O_1$)		Control ($C \rightarrow O_2$)	
Causal Order	*M*	*SE*	*M*	*SE*
Experiment 1A				
CE	63.85	7.38	92.31	7.69
EC	68.46	8.54	100.00	0.00
Experiment 1B				
CE	69.58	4.99	89.80	3.67
EC	88.56	3.24	97.33	2.09

Note—A, B, and C stand for the different cues, and O_1 and O_2 for the different outcomes used throughout the experiments. CE, cause–effect; EC, effect–cause. In Experiment 1A, $n = 13$ in each causal order condition. In Experiment 1B, $n = 45$ in each causal order condition.

EXPERIMENT 1B

Unlike the instructions of Experiment 1A, those used in Experiment 1B focused participants' attention on the relevant causal information, but in a way different from that of Waldmann (2001, Experiment 2). Participants did not have to explain to the experimenter the causal mechanism of the box of lights from the graphical representation provided. Instead, they had to respond to the same comprehension questionnaire as in our previous experiment. In addition, individual corrective feedback was provided if participants did not answer the questionnaire correctly. Thus, the procedure ensured that all participants had sufficiently understood the causal scenario before carrying out the experimental task.

To the extent that the questionnaire focused on one particular aspect of the instructions (i.e., the causal role of events), it may be argued that it implicitly suggested the relevance of this aspect. Thus, according to our hypothesis, participants should use the causal information throughout the task and, as in Waldmann (2001, Experiment 2), the overshadowing factor should interact with the causal order of the task.

Method

Participants and Apparatus. Ninety[2] undergraduate psychology students from the University of Málaga took part in the present experiment for course credit. The apparatus was the same as in our previous experiment.

Procedure. The procedural details were as described for Experiment 1A except for the fact that the questionnaire concerning the causal scenario was administered at the end of the instructions and, consequently, before the experimental task started. If the questionnaire was not correctly answered, the experimenter verbally interviewed the participant and any doubts or misunderstandings concerning the instructions, and particularly those concerning the causal order of the task, were resolved. The experimental task did not start until the experimenter considered that the participant correctly understood all the relevant procedural details.

Results and Discussion

All the participants reached the learning criterion, and the mean number of trials used in the CE and EC groups did not differ (17.8 in both groups). See Table 1 for mean ratings in the target overshadowing and control conditions.

A 2 (overshadowing: $A \rightarrow O_1$ vs. $C \rightarrow O_2$) $\times$ 2 (causal order: CE vs. EC) ANOVA was performed on the participants' ratings. The main effects of causal order [EC $>$ CE ratings, $F(1,88) = 8.50$, $MS_e = 930.18$, $p = .005$] and overshadowing [$A \rightarrow O_1 < C \rightarrow O_2$ ratings, $F(1,88) = 35.39$, $MS_e = 267.36$, $p = .000$] and, crucially, the overshadowing $\times$ causal order interaction [$F(1,88) = 5.51$, $MS_e = 267.36$, $p = .021$] were significant. The analysis of simple effects revealed that the overshadowing effect was significant in both causal order conditions [$F(1,88) = 34.41$, $MS_e = 267.36$, $p = .000$ and $F(1,88) = 6.48$, $MS_e = 267.36$, $p = .013$ for the CE and EC groups, respectively], although, as is indicated by the significant interaction, overshadowing was of greater magnitude in the CE than in the EC group. This interaction may also help account for the main effect of causal order obtained: Due to the greater overshadowing effect, ratings in the CE group were lower than those in the EC group.

Overall, the results were consistent with the predictions of causal reasoning models and contradicted those made by associative models, contrary to those obtained in Experiment 1A. Although the overshadowing effect found in the EC group is not anticipated by causal reasoning accounts, only this type of account offers a theoretical framework for understanding the influence of causal knowledge in inference making, as is shown by the significant overshadowing $\times$ causal order interaction.

Thus, the results showed that including in the instructions a comprehension questionnaire that focuses on the relevant causal information and then providing corrective feedback were enough to produce asymmetrical predictive–diagnostic inferences. In consistency with our hypothesis, implicitly stressing the relevance of causal information via the instructions somehow encouraged the participants to use their abstract causal knowledge during the learning and/or inferential judgment tasks.

Although our hypothesis is supported by the results from Experiments 1A and 1B, we have evaluated it in a rather indirect or implicit way, and this evaluation has involved procedural measures that may also be understood as favoring the participants' comprehension of the cover story before the experimental task started (i.e., the instructions included a comprehension questionnaire focusing attention on the causal information, and corrective feedback was provided only in Experiment 1B). Thus, the results obtained do not speak uniquely in favor of the perceived relevance hypothesis, but could also be explained by assuming different levels of comprehension of the cover story by participants from Experiments 1A and those from Experiment 1B while they performed the task. Note that this is an alternative hypothesis, even though only those participants that had correctly understood the cover story (as shown by their good performance on the comprehension questionnaire) were selected for data analysis in Experiment 1A. That is, even if the participants from both experiments had an effective understanding of the cover story, it could still be argued that the participants from Experiment 1B had a better comprehension of the cover story than did those from Experiment 1A and that such enhanced comprehension could facilitate their use of causal knowledge. Regardless of the credit one gives to this alternative explanation, this confound was avoided in our next experiment.

Therefore, the objective of Experiment 2 was twofold. First, we evaluated our hypothesis in a more direct or explicit way; second, we evaluated it in a way that avoided any conflation between stressing the relevance of the causal information and improving the participants' comprehension of such information.

EXPERIMENT 2

Specifically, we compared the participants' predictive and diagnostic inferences regarding the overshadowing effect in two different conditions. In the causal-cue condition, to encourage use of the causal information provided in the instructions, we included a sentence explicitly stating the relevance of such information to solving both the learning and the judgment tasks. Importantly, and at variance with Experiment 1B, the sentence did not involve a review of any specific aspect of the instructions that might be interpreted as enhancing participants' comprehension of the cover story (see below for details). In the no-causal-cue condition, no such sentence was included. At the end of the experiment, all the participants had to answer the same comprehension test as in the previous experiments. Thus, we could now know whether reading the target sentence as part of the instructions enhanced participants' comprehension of the cover story as evidenced by their performance on this test.

According to our hypothesis, predictive and diagnostic asymmetries regarding the overshadowing effect should arise only in the causal-cue groups, since these are the only groups that received an explicit indication of the relevance of the causal information to solving the task.

Method

Participants and Apparatus. Ninety-three psychology undergraduates from the University of Málaga volunteered to take part in the present experiment for course credit. The apparatus was the same as in our previous experiments.

Procedure. The procedural details were as in Experiment 1A. After reading the instructions, the participants were invited to resolve any doubts they may have had before the learning phase started. In the causal-cue groups, once all questions had been answered the participants could read, "To solve the task correctly, it is VERY IMPORTANT to take into account what you have just read in the instructions. Most importantly, in order to solve the different examples of the task and to answer the questionnaire, bear in mind that the lights on one side of the box cause the illumination of the lights on the other side of the box." Then, the learning phase started for these participants. In the no-causal-cue groups, the training phase began immediately after all doubts had been resolved. All the participants had to answer the comprehension questionnaire once they had made their inferential ratings.

Results and Discussion

All the participants reached the learning criterion, and the mean numbers of trials used in the CE and EC groups did not differ (16.2 and 17.9, respectively).

To facilitate the interpretation of the results, their statistical analysis was performed on the magnitude of the overshadowing effect (i.e., the difference between the control and the overshadowing target cues; see Table 2 for the ratings on these cues in each experimental group). A 2 (causal cue vs. no causal cue) $\times$ 2 (causal order: CE vs. EC) ANOVA was performed on the magnitude of the overshadowing effect. The mean magnitude of the overshadowing effect was significantly greater than zero [$F(1,89) = 105.97$, $MS_e = 686.63$, $p = .000$], which should be interpreted as a robust overshadowing main effect. The main effects of the factors of causal order [$F(1,89) = 4.77$, $MS_e = 686.63$, $p = .032$], causal cue [$F(1,89) = 5.27$, $MS_e = 686.63$, $p = .024$], and, importantly, the causal order $\times$ causal cue interaction [$F(1,89) = 8.61$, $MS_e = 686.63$, $p = .004$] were all significant. Since the interaction was significant, we performed a simple analysis to find the specific effect of the causal order factor in the causal-cue and no-causal-cue groups. According to this analysis, the causal order factor was significant in the causal-cue but not in the no-causal-cue groups [$F(1,89) = 13.58$, $MS_e = 686.63$, $p = .000$ and $F(1,89) = 0.20$, $MS_e = 686.63$, $p = .654$, respectively]. This pattern of results is consistent with our predictions, according to which the target predictive–diagnostic asymmetry should be found only in the causal-cue condition. Overall, the results obtained were equivalent to those obtained in Experiment 1B, showing that inclusion of this target sentence had an effect similar to that of inclusion of the comprehension test in the instructions.

Table 2
Design, Sample Size, Mean Ratings, and Standard Errors for Target Relationships in the Different Groups of Experiment 2

	Cue Interaction			
	Overshadowing ($AB \rightarrow O_1$)		Control ($C \rightarrow O_2$)	
Causal Order	*M*	*SE*	*M*	*SE*
Causal Cue				
CE	56.67	6.66	92.92	3.97
EC	84.78	5.83	92.61	5.19
No Causal Cue				
CE	62.61	6.74	94.78	4.40
EC	64.35	4.75	100.00	0.00

Note—A, B, and C stand for the different cues, and O_1 and O_2 for the different outcomes used throughout the experiment. CE, cause–effect; EC, effect–cause. For each group, $n = 23$, except for the causal-cue CE group, where $n = 24$.

In addition, one of the objectives of Experiment 2 was to test the perceived relevance hypothesis in a way that would allow us to detect any possible improvement in the participants' comprehension of the causal information provided. For this, a chi-square analysis was performed to test for the independence of the causal-cue factor and responding correctly on the comprehension test (i.e., giving a correct response to both test questions). The analysis yielded no significant dependence between the two variables [$\chi^2(1) = 0.53$, $p > .4$]. Thus, on the basis of the participants' responses on the comprehension test, the inclusion of the causal cue did not significantly improve their comprehension of the causal mechanism involved in the box-of-lights scenario.

GENERAL DISCUSSION

Using the same overshadowing paradigm, Waldmann (2001, Experiment 2) and Cobos et al. (2002, Experiment 4) obtained different results concerning the asymmetry between predictive and diagnostic inferences. Specifically, only Waldmann's (2001) results showed this asymmetry and, consequently, were consistent with participants' use of abstract causal knowledge in inference making. According to the hypothesis presented here, because learning processes that use abstract causal knowledge are cognitively demanding (for empirical evidence, see De Houwer & Beckers, 2003, or Waldmann & Hagmayer, 2001), individuals tend to use less demanding, associative learning processes in performing the learning and inference tasks. This, in turn, would lead them to ignore the relevance of the causal information provided during the instructions (i.e., information concerning the causal role of cues and outcomes). However, should the experimental procedure include any indication (either explicit or implicit) that focuses participants' attention on the relevance of this information, they would use this abstract causal knowledge, resulting in the predictive–diagnostic asymmetries predicted by causal reasoning accounts. In our view, Waldmann's (2001) procedure included such an indication. Specifically, participants had to explain to the experimenter the causal mechanism of the box of lights before the learning phase started, and corrective feedback was provided if necessary. The experiments reported here, guided by our hypothesis, helped us evaluate whether stressing the relevance of the causal information favored the observation of asymmetries between predictive and diagnostic situations.

Throughout the experimental series, we manipulated the causal order of the task and used the same overshadowing design and causal scenario as did Waldmann (2001, Experiment 2). The instructions of Experiment 1B included a comprehension questionnaire about the specific causal mechanism involved in the scenario, which has been interpreted as an indication to participants of the relevance of this causal information. In Experiment 1A, the same comprehension questionnaire was administered at the end of the task (i.e., once the inference ratings had been requested) and, thus, no indication was made to the participants regarding the relevance of any kind of causal information. In consistency with the perceived relevance hypothesis, the results showed an asymmetry between predictive and diagnostic inferences in Experiment 1B but not in Experiment 1A, even though the participants in Experiment 1A had also correctly understood the relevant causal information (as was shown by their performance on the comprehension questionnaire). In Experiment 2, we evaluated our hypothesis more directly in a way that avoided any conflation between stressing the relevance of the causal information and improving the participants' comprehension of such information. Specifically, Experiment 2 showed that including (or not including) an explicit sentence stating the relevance of the target causal information had an effect on the results obtained equivalent to the effect of including (or not including) the comprehension test as part of the instructions. That is, the group that could read the explicit sentence evidenced asymmetrical predictive–diagnostic inferences, as predicted by causal reasoning accounts; in contrast, the group for which the sentence was not included showed equivalent predictive and diagnostic inferences, as predicted by an associative account. Furthermore, the results showed that reading the target sentence during the instructions did not produce any advantage in the comprehension of the causal mechanism involved in the scenario, as was evidenced by the participants' performance during the comprehension test. Thus, the results from Experiment 2 served to show that it was the perceived relevance of causal knowledge to solving the task, and not an enhanced comprehension of the instructions provided, that led to the asymmetrical inferences found. It is important to note that these results do not rule out the relevance of enhancing the comprehension of the cover story as a factor favoring the use of causal knowledge, and, in fact, this was not our objective. On the other hand, what these results showed is that enhancing the perceived relevance of causal knowledge has an independent effect of favoring the use of causal knowledge.

As was mentioned in the introduction, the causal learning literature aimed at evaluating the relative merits of causal reasoning and associative accounts does not seem to agree on the processes involved in causal learning and inference making. Part of the reason, we argue, is that most, if not all, of this research has been conducted under the more or less implicit assumption that only one type of process determines causal learning and inference making in the experiments reported. The results obtained in the literature and those obtained here seem to converge on a rather different idea: Even under specific circumstances, causal learning may be mediated by more than one process and, depending on subtle procedural features, it is one or the other type of process that mediates individuals' performance on these tasks. In this sense, it seems more relevant to ask which factors facilitate the activation of associative process and which facilitate the activation of the more complex process envisaged by causal reasoning accounts. The answer to these questions is relevant not only on its own but also for an adequate comprehension of those associative and learning processes guided by causal knowledge, both seemingly involved in causal learning situations.

Converging evidence on the issue of the different processes involved in causal learning may be obtained by taking a closer look at the individual performance of the participants in our experimental series. Specifically, the present analysis will focus on the diagnostic conditions of the different experiments, since these are the conditions from which the predictions of causal reasoning and those of associative accounts diverge, in that only the associative accounts predict the overshadowing effect. On a descriptive level, Figure 1 shows that, across all the experiments reported, the participants fell into two main groups in terms of the magnitude of the overshadowing effect (i.e., the difference between the control and the overshadowing target cues): those who rated all target cues similarly (an overshadowing effect of 0) and those who showed an overshadowing effect with a magnitude of 50. These two groups included 77%, 93%, and 87% of all participants in Experiments 1A, 1B, and 2, respectively. Henceforth, then, we will concentrate our analysis on these two groups. We could refer roughly to the first group as the *causal reasoners group* and to the second as the *associative group*, meaning only that their performance was consistent with one or the other causal learning account. On a more theoretical level, Figure 1 also allows us to examine how these two groups of participants were redistributed in terms of the different experimental manipulations carried out. For example, if we compare Experiments 1A and 1B, we can see the effect of including instructions that focused on the relevance of causal knowledge on the number of participants in these two groups: Whereas in Experiment 1A the causal reasoners group and the associative group were of exactly the same size (i.e., 38.46% of the whole sample in both cases), in Experiment 1B the sizes of the two groups substantially varied (i.e., the causal reasoners group rose to 77.77% and the associative group diminished to 15.56% of the whole sample). Thus, focusing the participants' attention on the relevance of causal knowledge through the comprehension questionnaire increased the size of the causal reasoners group and diminished that of the associative group. Similarly, we can examine more directly the effect of suggesting the relevance of causal knowledge by comparing how these groups changed in the causal-cue and no-causal-cue conditions of Experiment 2. Specifically, in the no-causal-cue conditions, the associative and causal reasoners groups included 56.52% and 26.10% of the whole sample, respectively; in the causal-cue conditions, the associative and causal reasoners groups included 13.04% and 78.26% of the whole sample, respectively. Thus, again, focusing participants' attention on the relevance of causal knowledge altered the number of participants in the two main groups of Experiment 2 in the way

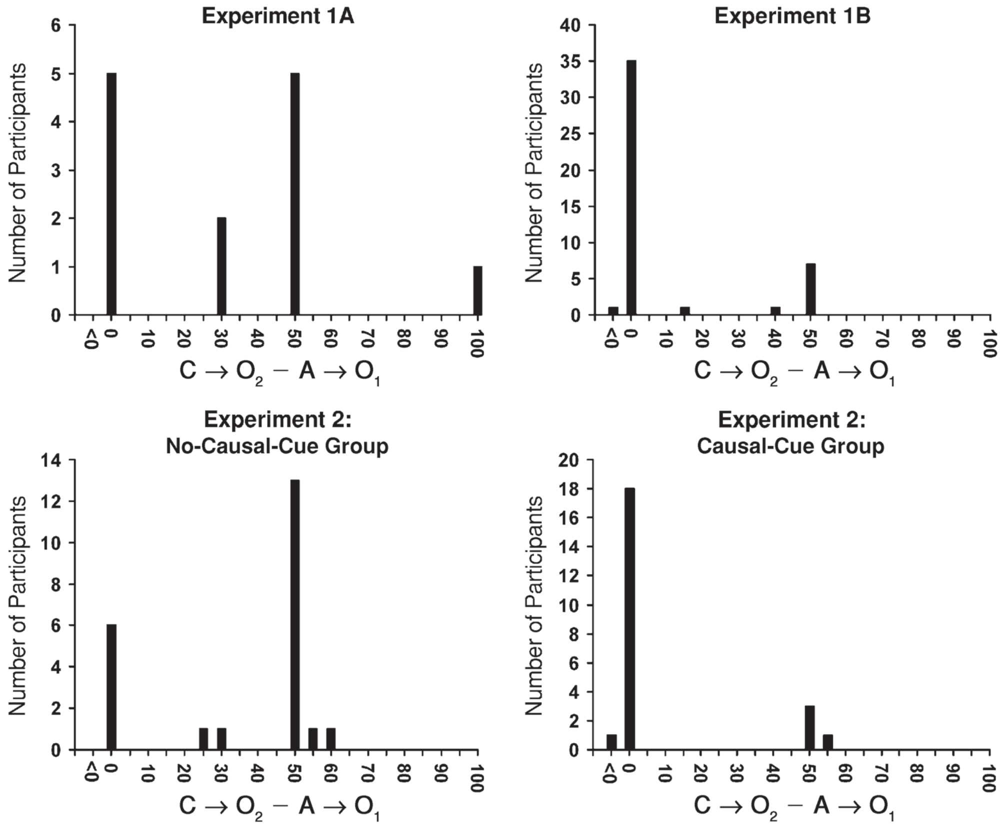

Figure 1. The distributions of participants across the different magnitudes of the overshadowing effect ($C \rightarrow O_2 - A \rightarrow O_1$) found in the EC conditions of the series of experiments.

put forth by our hypothesis. In general, then, the analysis carried out on individual performance has revealed two main groups of participants that may be readily defined in terms of the accounts proposed, and has converged on the same conclusions extracted from the analysis of the participants' inferential ratings.

Previous suggestions have been made about the relevance of factors such as the cognitive effort demanded by the experimental task, the number of trials, the familiarity with the causal scenario involved, and the credibility of the causal scenario as potential explanations for the different results obtained in the literature. Nevertheless, very few attempts have been made to systematically evaluate the relevance of such factors (see, e.g., Tangen & Allan, 2004; Waldmann & Hagmayer, 2001). This strategy would entail obtaining results consistent with both associative and causal reasoning accounts under identical circumstances except for the target factor whose relevance is being evaluated. This is, in our view, the most important objective that the present experimental series has satisfied. Under very similar circumstances, the results obtained were consistent with either the associative or the causal reasoning account. Specifically, the participants made use of their abstract causal knowledge only when the relevance of the causal role played by the events involved in the task was indicated to them either explicitly or implicitly. Otherwise, the participants' performance was consistent with an associative account of causal learning, even when they showed good comprehension of the relevant information.

Consequently, the results reported here may have important implications for the theoretical interpretation of previous results. For example, a number of studies have failed to report predictive–diagnostic asymmetries regarding cue interaction phenomena or have shown cue interaction effects in diagnostic situations (e.g., Cobos, Caño, López, Luque, & Almaraz, 2000; Cobos et al., 2002; Gluck & Bower, 1988; Le Pelley & McLaren, 2001; Matute et al., 1996; Price & Yates, 1995; Shanks, 1990; Shanks & López, 1996). Although these results have commonly been interpreted as showing the inability of individuals to use abstract causal knowledge, according to the results reported here such an inability may be interpreted as a consequence of the failure to perceive the relevance of the causal role played by the events involved in the tasks. This interpretation is consistent with the fact that none of these researchers used a procedure that stressed the relevance of causal information to solving these tasks. Of course, there may be additional reasons for which the participants in these studies did not use causal knowledge when making inferences. On the other hand, as has previously been argued, to say that people have the competence to use abstract causal knowledge does not suffice to explain the asymmetries already found in the literature (see, e.g., Waldmann, 2001, Experiment 2). They must also perceive the relevance of this causal information, a fact that seems to fall outside of the explanatory scope of causal reasoning models as currently formulated. Another example that shows the relevance of the present results relates to those reported by Tangen and Allan (2004). They identified the number of trials as a factor that demarcates the circumstances under which associative or causal reasoning accounts apply. Specifically, they showed that associative processes mediate inferential judgments only if the number of trials is large and, thus, they called into question the generalizability of Cobos et al.'s (2002) results to experimental procedures that lack such a large number of trials. The results reported in the present study show that a large number of trials is not necessary to demonstrate results consistent with associative models.

An important question that arises in this context is why participants apparently have difficulties in perceiving the relevance of the causal information provided during the instructions of the experiments. There are a number of possible answers to this question, and, unfortunately, the results obtained thus far do not help to discriminate between them. Throughout the present article, we have argued that the use of causal knowledge demands important cognitive resources that individuals initially try to avoid having to use by the activation of less costly processes, such as associative processes. This being the case, the tendency to ignore the causal information found in the experimental series reported here may generalize to extra-experimental contexts. However, it can also be argued that ordinary experimental situations have some features that lead participants to ignore the causal information provided in the instructions. For example, the experimental tasks used in the laboratory are usually of no practical interest and, consequently, the use of more demanding cognitive processes is not especially encouraged. In addition, it may be argued that the abstract and artificial formats in which the information is presented throughout the task hinder the use of abstract causal knowledge. In this sense, it should be noted that, in most of the experiments carried out, direct experience with the target causal events of the task is substituted by a series of verbal labels that appear on a computer screen. These different possibilities should be explored in future research.

REFERENCES

Cobos, P. L., Caño, A., López, F. J., Luque, J. L., & Almaraz, J. (2000). Does the type of judgement required modulate cue competition? *Quarterly Journal of Experimental Psychology*, **53B**, 193-207.

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.

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., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, **55A**, 965-985.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227-247.

Le Pelley, M. E., & McLaren, I. P. L. (2001). Retrospective revaluation in humans: Learning or memory? *Quarterly Journal of Experimental Psychology*, **54B**, 311-352.

López, F. J., Cobos, P. L., Caño, A., & Shanks, D. R. (1998). The rational analysis of human causal and probability judgment. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 314-352). Oxford: Oxford University Press.

Matute, H., Arcediano, F., & Miller, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 182-196.

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

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. (1990). Connectionism and the learning of probabilistic concepts. *Quarterly Journal of Experimental Psychology*, **42A**, 209-237.

Shanks, D. R., & López, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, **24**, 511-522.

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

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, & Cognition*, **26**, 53-76.

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

Waldmann, M. R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition*, **82**, 27-58.

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.

NOTES

1. Initials refer to Spanish color names. All the experiments were conducted in Spanish.

2. Experiment 1B includes the samples of two identical experiments. The sample of the original experiment consisted of 33 participants (17 and 16 participants in the CE and EC groups, respectively). The results were as described in the main text. Specifically, the main effects of causal order [$F(1,31) = 7.21$, $MS_e = 815.41$, $p = .012$], overshadowing [$F(1,31) = 8.73$, $MS_e = 221.48$, $p = .006$], and, crucially, the overshadowing × causal order interaction [$F(1,31) = 4.42$, $MS_e = 221.48$, $p = .044$] were significant. The analysis of simple effects revealed that the overshadowing effect was significant only in the CE condition [$F(1,31) = 13.18$, $MS_e = 221.48$, $p = .001$] and not in the EC condition ($F < 1$). Due to the novelty of this result in our laboratory, we decided to replicate it in an independent experiment. Since the procedures were identical in all respects, we decided to collapse the results of both experiments and describe them here as the results of a single experiment. Thus, the different results obtained in Experiment 1B cannot be attributed to differences in experimental sensitivity, since the sample sizes of Experiments 1A and 1B are noticeably different.

APPENDIX
Translation of Instructions for Experiments 1A–2

CE Condition

You are now going to learn how a box of lights works. On the front side of the box are three colored lights (white, fuchsia, and emerald), and below each of the lights is a switch that activates its corresponding light. On the back side of the box, there are two bulbs (Bulb 1 and Bulb 2) that come on depending on the activation of the colored lights from the front side. The bulbs on the back side are never illuminated simultaneously.

Your objective is to learn the causal relationship between the activation of the colored lights (causes) and the illumination of the bulbs on the back side (effects). For this, you are going to see a series of examples. In each of them, you will be told which colored light(s) is activated and you will have to guess which bulb on the back side (Bulb 1, Bulb 2, or neither) will illuminate. In addition, after you respond, the screen will give you the right answer. Use this information to learn the relationship between the different colored lights on the front side and the bulbs on the back side.

Once you have examined all the examples, you will have to evaluate on a questionnaire to what extent each of the different colored lights predicts the illumination of the different bulbs.

EC Condition

You are now going to learn how a box of lights works. On the back side of the box are two lights (Light 1 and Light 2), and below each of the lights is a switch that activates its corresponding light. The lights on the back side are never activated simultaneously. On the front side of the box are three colored bulbs (white, fuchsia, and emerald) that come on depending on the activation of the lights from the back side.

Your objective is to learn the causal relationship between the activation of the lights on the back side (causes) and the illumination of the colored bulbs (effects). For this, you are going to see a series of examples. In each of them, you will be told which colored bulb(s) came on and you will have to guess which light on the back side (Light 1, Light 2, or neither) had been activated. In addition, after you respond, the screen will give you the right answer. Use this information to learn the relationship between the different colored bulbs on the front side and the lights on the back side.

Once you have examined all the examples, you will have to evaluate on a questionnaire to what extent each of the different colored bulbs predicts the activation of the different lights.

Test of Comprehension of the Instructions (Experiments 1A–2)

Questionnaire About the Functioning of the Box of Lights

Finally, to evaluate your comprehension of how the box of lights works during the task, you will have to answer the following questions. Take as much time as you need before giving an answer.

To facilitate your task, we provide a drawing representing the box of lights.

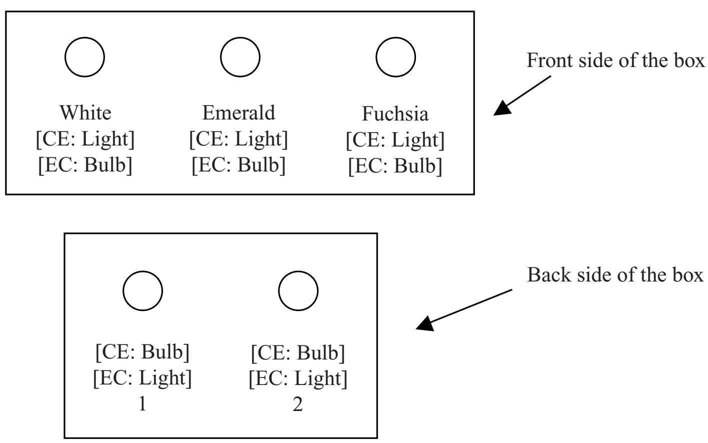

Question 1. What would you have to do to [CE, activate the colored lights; EC, illuminate the colored bulbs] on the front side? (Draw a circle around the correct choice.)

A. Since each colored [CE, light; EC, bulb] has its own switch, you only have to activate one of the three switches.

B. Since each numbered [CE, bulb; EC, light] has its own switch, you have to activate one of the [CE, bulbs; EC, lights] via one of the two switches in order to [CE, activate the colored lights; EC, illuminate the colored bulbs] indirectly.

APPENDIX (Continued)

Question 2. What would you have to do to [CE, illuminate the bulbs; EC, activate the lights] on the back side? (Draw a circle around the correct choice.)

A. Since each numbered [CE, bulb; EC, light] has its own switch, you only have to activate one of the two switches.

B. Since each colored [CE, light; EC, bulb] has its own switch, you have to activate one of the [CE, lights; EC, bulbs] via one of the three switches in order to [CE, illuminate the numbered bulbs; EC, activate the numbered lights] indirectly.

(Manuscript received June 21, 2004;
revision accepted for publication December 3, 2004.)