Research Article

Psychological Science
23(1) 59–68

DOI: 10.1177/0956797611429577
http://pss.sagepub.com

Two Mechanisms of Human Contingency Learning

Daniel A. Sternberg and James L. McClelland
Stanford University

Abstract

How do humans learn contingencies between events? Both pathway-strengthening and inference-based process models have been proposed to explain contingency learning. We propose that each of these processes is used in different conditions. Participants viewed displays that contained single or paired objects and learned which displays were usually followed by the appearance of a dot. Some participants predicted whether the dot would appear before seeing the outcome, whereas other participants were required to respond quickly if the dot appeared shortly after the display. In the prediction task, instructions guiding participants to infer which objects caused the dot to appear were necessary in order for contingencies associated with one object to influence participants' predictions about the object with which it had been paired. In the response task, contingencies associated with one object affected responses to its pair mate irrespective of whether or not participants were given causal instructions. Our results challenge single-mechanism accounts of contingency learning and suggest that the mechanisms underlying performance in the two tasks are distinct.

Keywords

learning, reasoning, cognitive processes, associative processes

Received 10/27/10; Revision accepted 6/20/11

Understanding how people learn contingencies between events has been a focus of research for many years (Krechevsky, 1932; Pavlov, 1927; Tolman, 1948). In standard contingency-learning tasks, participants are presented with cues that are followed by outcomes and later have to predict outcomes for cues presented during a test block. Two prominent types of accounts have been offered to describe the process underlying performance in such tasks: pathway-strengthening accounts, which attribute performance to the strengthening of pathways that link representations of cues to representations of outcomes (Rescorla & Wagner, 1972), and inference-based accounts, which assume that an explicit reasoning process leads to inferences about the causal relations between cues and outcomes (De Houwer, 2009; Mitchell, De Houwer, & Lovibond, 2009).

Pathway strengthening has been proposed as the mechanism underlying the gradual learning of tendencies to respond to particular inputs. Stronger pathways promote faster, more automatic responses (Cohen, Dunbar, & McClelland, 1990), and pathway-strengthening models can account for the gradual speeding of responses in fast-paced sequence-learning tasks (Cleeremans & McClelland, 1991). However, considerable evidence from more recent research suggests that contingency learning may rely on a resource-intensive process of making explicit inferences. This evidence has led some researchers to propose that an explicit inference-based process can provide a complete account of contingency learning (De Houwer, 2009; Mitchell et al., 2009). We suggest that both types of processes exist and that the conditions of the contingency-learning task determine the degree to which either process underlies performance. In our investigation of this issue, we focused on a difference between the two accounts' predictions regarding how the instructions in contingency-learning tasks should influence what are known as *indirect effects*.

Indirect effects in contingency-learning tasks can arise when potential predictive cues can occur either alone or in combination. In some training events, two cues (X_1 and X_2) are presented together, and an outcome (O) occurs shortly afterward. In other training events, X_1 occurs without X_2. If X_1 occurs alone and is then followed by O, the learner's tendency to anticipate O when X_2 is presented alone will be reduced; this effect is called *blocking* (Kamin, 1969). However, if X_1 occurs alone and is not followed by O, the learner's tendency to anticipate O when X_2 occurs on its own will be increased; this effect is called *screening* or *reduced overshadowing* (Carr,

Corresponding Authors:
Daniel A. Sternberg, Jordan Hall, Building 420, Stanford University, 450 Serra Mall, Stanford, CA 94305
E-mail: sternberg@stanford.edu

James L. McClelland, Jordan Hall, Building 420, Stanford University, 450 Serra Mall, Stanford, CA 94305
E-mail: mcclelland@stanford.edu

1974). In both cases, the learner's response to X_2 is indirectly affected by what happened when X_1 was presented alone.

Pathway strengthening can explain these indirect effects if the changes in the strengths of the pathways that link cues to outcomes are proportional to the learner's prediction error (i.e., the difference between the strength of the prediction of the outcome, which is based on the current connection strengths between the present cues and that outcome, and the actual outcome). If there are trials on which X_1 occurs alone and is followed by O, X_1 will develop a strong connection to O. In that case, there will be little error on other trials in which X_1 is presented together with X_2 and followed by O, so there will be little strengthening of the association between X_2 and O (blocking). Conversely, if there are trials on which X_1 occurs alone and is not followed by O, the connection between X_1 and O will remain weak; this weak connection will lead to a large prediction error when O occurs after a presentation of X_1 and X_2 together, and therefore a strong association will develop between X_2 and O (screening). According to this account, the same process underlies both the learning of contingencies during training and the emergence of indirect effects at test, and the occurrence of indirect effects requires no extra processing steps. The earliest contingency-learning models to incorporate indirect effects in general and blocking in particular relied on this kind of mechanism (e.g., Rescorla & Wagner, 1972).

An inference-based process can also produce indirect effects if a learner treats X_1 and X_2 as potential independent causes of the outcome. For instance, imagine that the cues represent foods and the outcome is an allergic reaction. If no reaction occurs when food X_1 is eaten alone, but a reaction does occur when both food X_1 and food X_2 are eaten (a screening situation), a person could infer that food X_2 was the cause of the allergy. However, if eating food X_1 by itself results in an allergic reaction, then experiencing a reaction after food X_1 and food X_2 are eaten together (a blocking situation) will provide little evidence about food X_2. Thus, individuals who rely on an inference-based process and adopt the independent-cause assumption should predict an allergic reaction more strongly in the screening situation than in the blocking situation.

These alternative accounts of the indirect effects observed in contingency learning make contrasting predictions. According to the inference-based account, the learner's assumption that causes act independently is necessary to license the inferences that give rise to indirect effects (Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003). If the cues are not assumed to be independent causes of the outcome, the implications of events involving X_1 for events involving X_2 are ambiguous in both the blocking and the screening cases. For example, in the screening situation (O occurs after X_1X_2 but not after X_1 alone), there is no clear implication that O will follow X_2 alone; a person could suppose instead that X_1 and X_2 must occur together for O to occur. Evidence from contingency-learning tasks supports the importance of the independent-cause assumption: Giving participants a cover story that is consistent with the independent-cause assumption has been shown to increase indirect effects (Lovibond et al., 2003; Waldmann & Holyoak, 1992, 1997; Williams, Sagness, & McPhee, 1994).

According to inference-based accounts of contingency learning, indirect effects also require that the learner has enough time and free mental resources to make an inference on the basis of the evidence provided by the events observed during training. This claim is supported by the finding that introducing a distracting secondary task reduces indirect effects (De Houwer & Beckers, 2003). Thus, inference-based accounts predict that instructions encouraging learners to identify independent causes should promote indirect effects in contingency-learning tasks, whereas limitations on available resources (including the amount of time learners have to retrieve relevant information about training events from memory) should reduce them. In contrast, because pathway strengthening is thought to result from error-correcting learning and to operate in neural circuits that are not penetrated by verbal instructions, pathway-strengthening accounts predict that indirect effects may arise whether or not participants are given instructions that promote the independent-cause assumption.

We suggest that both inference-based and pathway-strengthening processes underlie contingency learning, but in different circumstances. When participants are given unlimited time in a contingency-learning task and receive instructions that encourage them to make inferences, they will do so; in that case, performance should conform to the predictions of inference-based accounts. However, when participants must respond quickly and thus have little time to reflect and make inferences, they will learn the contingencies through an error-correcting, pathway-strengthening process, which should lead to indirect effects whether or not participants receive instructions that promote the independent-cause assumption.

In order to test this hypothesis, we designed two tasks that exposed participants to the same contingencies between cues and outcomes. Participants who completed our untimed prediction task were given an unlimited amount of time to predict whether cues would be followed by a particular outcome; they were then given information about the outcomes that either confirmed or disconfirmed their predictions. Half of these participants were given *causal-framing* instructions, which encouraged them to learn which objects caused the outcomes. The other participants were given *object-framing* instructions, which did not guide them to make such inferences.

Our fast-paced reaction time (RT) task was performed by another two groups of participants. In the RT task, however, the outcome (if any) automatically occurred shortly after the objects had appeared, and participants were instructed to respond quickly when the outcome occurred and to refrain from responding when it did not. Half of the participants who completed the RT task received causal-framing instructions, and the other half received object-framing instructions.

Table 1 shows the competing predictions of the inference-based and pathway-strengthening accounts for performance in

Table 1. Competing Accounts' Predictions About Whether Indirect Effects Should Occur in the Four Combinations of Task and Framing Used in the Experiment

Account and task	Causal-framing condition	Object-framing condition
Inference-based account		
Untimed prediction task	Yes	No
Fast-paced RT task	?	No
Pathway-strengthening account		
Untimed prediction task	Yes	Yes
Fast-paced RT task	Yes	Yes
Two-process account		
Untimed prediction task	Yes	No
Fast-paced RT task	Yes	Yes

Note: An inference-based account predicts that indirect effects should depend on causal framing and may not occur at all in our fast-paced reaction time (RT) task. An account of pathway strengthening via error-correcting learning predicts indirect effects in all conditions. Our two-process account predicts that participants should make explicit inferences when they have enough time to do so but should rely on pathway strengthening under time constraints; indirect effects should thus depend on causal framing in the untimed prediction task but should occur under both framing conditions in the fast-paced RT task.

the four conditions of our experiment, as well as the predictions made by our two-mechanism account. According to our account, performance in the prediction task should be consistent with the predictions of the inference-based account, whereas performance in the RT task should be consistent with the predictions of the pathway-strengthening account.

Method

Participants

Forty-eight members of the Stanford psychology department's paid participant pool performed the prediction task, and 96 members of the pool performed the fast-paced RT task. Payment of participants was based in part on their performance: Participants earned \$0.02 for each correct response and lost \$0.01 for each incorrect response (for more information about payment, materials, and methods, see the Methods section in Additional Methods and Analyses in the Supplemental Material available online). We removed 1 participant in the RT task for failing to respond to one of the critical displays; we included the remaining 95 participants (47 in the causal-framing condition and 48 in the object-framing condition) in our analyses.

Materials

The same stimuli and displays were used in both tasks. Cues consisted of clip-art objects from the OpenClipArt library (http://www.openclipart.org) and were randomly assigned to specific displays within the design; a different random assignment of clip-art objects to displays was used for each participant. Eleven displays were used during training; each display consisted of either one object (singleton display) or a pair of objects (compound display). For the outcomes, a dot either appeared after the event or did not appear. When the dot appeared, it was located either to the left or to the right of the side of the screen on which the objects were displayed.

We used two critical sets of displays (a blocking set and a screening set) to assess both the degree to which participants learned the contingencies and whether participants exhibited an indirect effect. Each set consisted of a compound training display in which two objects were shown as a pair, a singleton training display in which one object from the compound display appeared on its own, and a singleton test display in which the other object from the compound display appeared on its own (Table 2); only the training displays were used in the training phase, and both the training and the test displays were used in the test phase. The compound training displays in both the blocking set and the screening set (B_1B_2 and S_1S_2, respectively) were followed by the dot on approximately 90% of the trials on which they appeared. The blocking training singleton (B_1) was followed by the dot on approximately 90% of trials, whereas the screening training singleton (S_1) was followed by the dot on approximately 10% of trials. If the participants learned the training contingencies, then during the test block they should have expected the dot to appear following B_1 more strongly than they should have expected it to appear following S_1. If participants expected the dot to appear following S_2 more strongly than they expected the dot to appear following B_2, they would be demonstrating an indirect effect.

Two additional sets of displays were included to allow us to consider how participants generalized across the training displays in more detail (see Table 2): For the negative set, both the compound training display (N_1N_2) and the singleton training display (N_1) were followed by the dot on approximately 10% of the trials in which they appeared, whereas the control set consisted of only a compound training display (C_1C_2) that was followed by the dot on approximately 90% of the trials in which it appeared. Finally, we included four additional filler

Table 2. Stimuli in the Training and Test Sessions

Display set	Training compound display	Training singleton display	Test singleton display
Blocking set	B_1B_2+	B_1+	B_2
Screening set	S_1S_2+	S_1-	S_2
Negative set	N_1N_2-	N_1-	N_2
Control set	C_1C_2+	—	C_1, C_2

Note: The table shows the four sets of displays for both conditions of both tasks. Plus signs indicate displays that were usually followed by a dot; minus signs indicate displays that were usually not followed by a dot. In addition to the displays presented in the table, we included four filler displays (one compound display not followed by a dot, two singleton displays not followed by a dot, and one singleton display followed by a dot) in the training session to reduce the dot's overall rate of occurrence.

displays (one compound display not followed by a dot, two singleton displays not followed by a dot, and one singleton display followed by a dot). These fillers, which were shown only during the training phase, were used in order to reduce the overall occurrence rate of the dot, because pilot testing had shown that participants tended to overestimate the probability of the dot's occurrence.

Each participant was given one of two sets of framing instructions (causal framing or object framing; see Table 3) in addition to a set of task-specific instructions (prediction or RT). The causal-framing instructions emphasized that individual objects had the power to make the dot appear and encouraged participants to infer which objects had that power. The object-framing instructions emphasized that the computer's presentation of outcomes followed rules based on the objects that appeared in the displays but did not describe the nature of the rules. These instructions also encouraged participants to pay attention to the objects in each display without explicitly instructing participants to make inferences.

Table 3. Instructions for the Causal-Framing and Object-Framing Conditions of the Two Tasks

Causal-framing instructions: Some of the objects that you will see during the experiment have the power to make the dot appear, while others do not. The computer has randomly decided which objects have this power at the beginning of the experiment. You cannot tell just by looking at the object whether it has the power to make the dot appear. However, you can learn which objects have this power based on whether or not the dot appears when the object is inside the box. If two objects appear inside the box and at least one of them has this power, the dot will usually appear. Sometimes the box may malfunction, and the dot may occasionally fail to appear when it should, and may occasionally appear when it shouldn't. If you can determine which objects have the power to make the dot appear, this will help you to make predictions during the experiment.

Object-framing instructions: The computer follows rules to determine whether the dot will appear, based on the object(s) currently displayed on the screen. Sometimes the dot may fail to appear when it should, and may sometimes appear when it shouldn't. If you pay attention to the object(s) displayed on the screen, this will help to improve your performance during the experiment.

Procedure

Figure 1 shows an example sequence of displays and the time course for trials in both tasks. Objects appeared inside a rectangle. On each trial, either two objects appeared (compound display) or a single object appeared (singleton display). On compound-display trials, one object was randomly chosen to appear in the center of the upper half of the box, and the other object appeared in the center of the lower half of the box. On singleton-display trials, a single object appeared randomly in one of these two locations.

In the prediction task, after an object display appeared, participants were given an unlimited amount of time to predict whether the dot would appear; participants responded by pressing a "yes" or a "no" button. After a participant responded, the outcome was presented and was followed by visual and auditory feedback that indicated whether the response was correct and how many points the participant had gained or lost. During the training phase, each training display appeared 24 times, and the most likely outcome (dot or no dot) for each training display occurred on 22 of these 24 trials. At test, each display appeared 5 times, and no feedback about the outcome or the accuracy of the prediction was provided. Instead, after the participant responded, the word "Recorded" appeared in the center of the box.

In the fast-paced RT task, the dot automatically did or did not appear 350 ms after the object display appeared. Participants were required to respond to the appearance of the dot within a response window that began at the dot's onset (see Fig. 1). Participants received visual and auditory feedback after each response. Feedback indicating a correct response was provided if participants responded within the response window on trials in which the dot occurred or avoided responding on trials in which the dot did not occur. The duration of the response window was initially 400 ms and decreased by 12.5 ms every 10 trials for the first 200 training trials and by 2.5 ms every 10 trials for the next 100 trials. For the remaining training trials and during the test, the response window was 250 ms long. Each training display occurred 72 times during the training block, and the most likely outcome for each display occurred on 65 of these 72 trials. The test consisted of the same set of displays used in the prediction task. All training

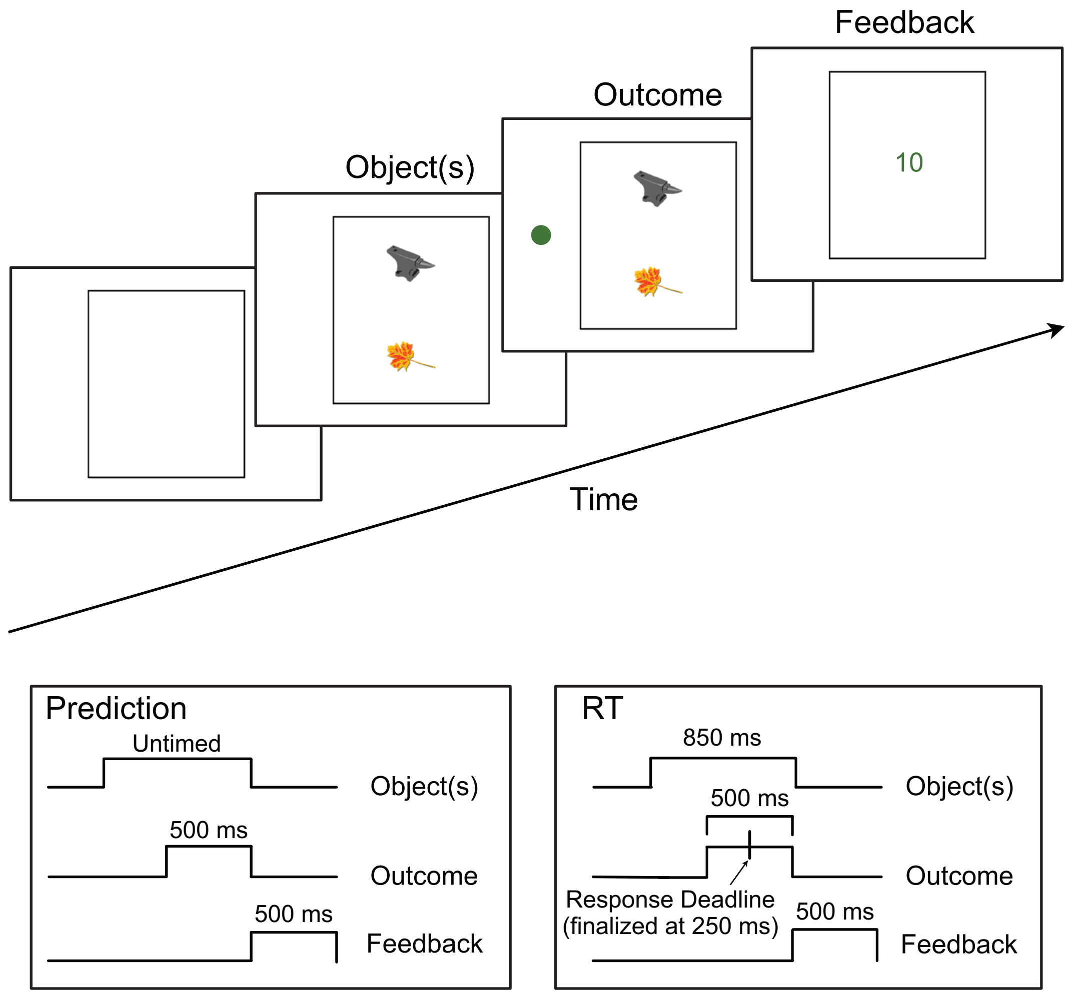

Fig. 1. Example of the sequence of displays (top panel) and timing of trials (bottom panel) in the prediction and reaction time (RT) tasks. In each task, the appearance of either one object (singleton display) or two objects (compound display) was followed by an outcome (dot or no dot) and feedback. In the prediction task, after the object or objects appeared, participants had unlimited time to predict the outcome. After making their prediction, they were shown the outcome: Either the dot appeared for 500 ms, or nothing happened. The outcome was followed by visual feedback (10 points for a correct prediction, –5 points for an incorrect prediction) and auditory feedback (a pleasant sound for a correct prediction, a buzzer for an incorrect prediction) for 500 ms. In the RT task, the dot automatically did or did not appear 350 ms after the object or objects had appeared; the outcome was displayed for 500 ms. Participants had to respond to the appearance of the dot within a response window; after the response window ended, participants received the same feedback for correct and incorrect responses as in the prediction task. The duration of the response window was initially 400 ms and decreased in small steps until it reached its final value of 250 ms.

and test displays occurred 16 times during the test block. So that we could obtain sufficient data on reaction times for the critical items, all singleton displays were followed by the dot on 8 of the 16 test trials. The *dot-likely* compound displays (the B_1B_2, S_1S_2, and C_1C_2 pairs) were followed by the dot on 14 of the 16 trials, whereas the *dot-unlikely* compound display (the N_1N_2 pair) was followed by the dot on 2 of the trials.

Results

Results from the test phase of the experiment indicated that participants in all four conditions learned the contingencies they had been exposed to during the training phase (see Table 4 for a summary of results from the test phase). Participants in both the causal-framing and the object-framing conditions of the prediction task were more likely to predict the dot after the B_1 item than after the S_1 item at test—causal-framing condition: Wilcoxon $z = 4.61$, $p < .0001$; object-framing condition: Wilcoxon $z = 4.34$, $p < .0001$ (see Fig. 2). Participants in both conditions of the RT task responded faster on dot trials for the B_1 item than on dot trials for the S_1 item at test—causal-framing condition: $t(46) = 10.9$, $p < .0001$; object-framing condition: $t(47) = 10.14$, $p < .0001$ (see Fig. 2). See Figure S1 in the Supplemental Material for learning curves for both tasks.

Table 4. Test Results: Prediction of the Dot and Reaction Times in the Untimed Prediction Task and the Fast-Paced Reaction Time Task

Task, condition, and display type	Training compound display	Training singleton display	Test singleton display
Untimed prediction task			
Causal-framing condition			
Blocking	1.0	.983	.583
Screening	.917	.066	.867
Negative	.042	.042	.050
Control	.950	—	.750
Object-framing condition			
Blocking	.942	.942	.692
Screening	.900	.108	.666
Negative	.092	.042	.192
Control	.950	—	.721
Fast-paced RT task			
Causal-framing condition			
Blocking	171.2 (64.0)	198.7 (62.9)	268.9 (48.5)
Screening	203.7 (54.8)	325.7 (49.8)	245.8 (62.6)
Negative	332.6 (59.4)	352.9 (52.6)	341.9 (49.1)
Control	187.1 (59.9)	—	270.1 (44.9)
Object-framing condition			
Blocking	171.1 (49.9)	215.8 (58.4)	269.4 (51.1)
Screening	197.7 (52.1)	312.8 (45.2)	239.1 (59.9)
Negative	309.9 (71.3)	339.8 (41.4)	321.2 (41.7)
Control	177.7 (61.5)	—	255.4 (49.5)

Note: For the prediction task, we report the proportion of test trials on which participants predicted that the dot would appear. For the reaction time (RT) task, we report the mean of participants' median reaction times (in milliseconds; standard deviations are shown in parentheses).

In the untimed prediction task, an indirect effect appeared in the causal-framing condition but not in the object-framing condition. In the prediction task, participants demonstrated an indirect effect if they predicted the dot more often for the S_2 item than for the B_2 item. Participants who received the causal-framing instructions showed this indirect effect, $z = 2.47$, $p < .05$, whereas participants who received the object-framing instructions did not, $z = -0.32$, $p = .78$ (Fig. 2). Results of Mann-Whitney U tests showed that this difference between the two groups was significant ($z = 2.10$, $p < .05$).

In contrast, in the RT task, both groups showed the expected indirect effect. In this task, participants demonstrated an indirect effect if they responded faster to the S_2 item than to the B_2 item. This effect was observed both in the causal-framing condition, $t(46) = 2.60$, $p < .05$, and in the object-framing condition, $t(47) = 3.13$, $p < .01$. The size of the effect did not differ by group, $t(93) = 0.54$, $p = .59$ (Fig. 2). See the Mixed Regression Analyses section in Additional Methods and Analyses in the Supplemental Material for results from mixed-effects regression analyses that provide further confirmation of these findings.

The fact that an indirect effect was not observed in the prediction task in the object-framing condition led us to examine the pattern of responses to all test singleton displays in this condition. In order to determine whether participants' predictions for these objects were influenced by the contingencies for the compound displays in which the objects appeared during training, we compared the average of each participant's predictions for the C_1, C_2, B_2, and S_2 objects, which had occurred in dot-likely compound displays during training, with the average of that participant's predictions for the N_2 object, which had occurred in a dot-unlikely pair. A Friedman test revealed that participants in the object-framing condition indeed predicted the dot more often for the dot-likely objects than for the N_2 item, $\chi^2(1, N = 24) = 9.78$, $p < .01$. However, there were no significant differences in the probabilities of positive dot predictions for the different dot-likely objects, $\chi^2(2, N = 24) = 0.222$, $p = .89$; this finding is consistent with the view that participants' predictions about objects in dot-likely pairs were not affected by the contingencies of the objects' pair mates. These findings indicate that participants in the prediction task who were given object-framing instructions learned the co-occurrence relations between compound training displays and outcomes and relied on these relations to make predictions about the test objects that had appeared in these compound displays; however, these participants did not take the further step of making inferences based on outcomes associated with training displays in which the pair mates of the test objects occurred alone.

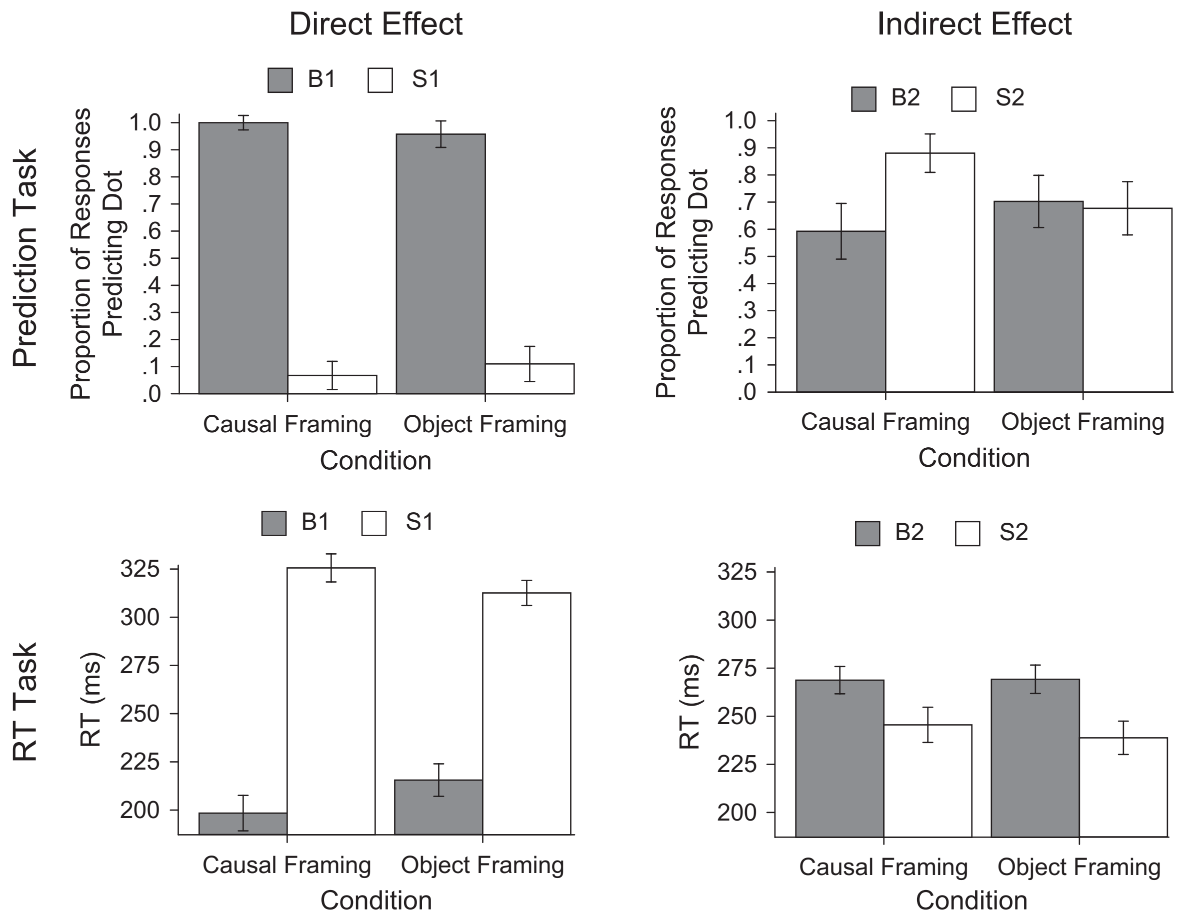

Fig. 2. Direct and indirect effects in the prediction task (top row) and the reaction time (RT) task (bottom row), as a function of condition. In the top row, the proportion of responses predicting the appearance of the dot is shown for the critical training displays (S_1 and B_1; direct effect), on the left, and for the test displays (S_2 and B_2; indirect effect), on the right. In the bottom row, RTs for the critical training displays are shown on the left, and RTs for the test displays are shown on the right. Error bars indicate standard errors of the mean.

Discussion

We observed a striking dissociation in our experiment. On the one hand, when participants made untimed predictions, causal-framing instructions led to an indirect effect, but object-framing instructions did not. This result is consistent with previous findings showing that indirect effects in explicit contingency-learning tasks are dependent on framing. Without causal framing, predictions about objects that had occurred only in pairs during training appeared to be based only on the outcomes of the relevant compound displays and not on outcomes associated with separate appearances of the object's pair mate. On the other hand, when participants had to respond under high time pressure, indirect effects occurred in both framing conditions, as one would expect if such effects resulted from pathway strengthening via error-correcting learning. This dissociation is consistent with our hypothesis that indirect effects are driven by an inference-based process in the untimed prediction task but rely on pathway strengthening in the fast-paced RT task.

The difference we observed in the effect of instructions in the two tasks supports our predictions and challenges unitary approaches to contingency learning. An inference-based approach nicely accounts for the difference between the two conditions in our prediction task and is consistent with previous work showing that causal framing is necessary to produce indirect effects in such tasks. However, although this explanation confirms the importance of causal framing, it is unclear how the inference-based approach could account for an indirect effect in the object-framing condition in the RT task. Inference-based models assume that indirect effects result from a time- and resource-dependent inference process, so participants should be less likely to make inferences in the RT task, in which time and cognitive resources are constrained, than in the prediction task. Thus, the presence of an indirect effect in the object-framing condition of the RT task but not in the object-framing condition of the explicit prediction task cannot be easily explained by a unitary inference-based account.

The observed pattern of results is also difficult to explain with a unitary pathway-strengthening account, because such an account predicts that indirect effects should occur in all four conditions. Some proponents of associative accounts of contingency learning have addressed framing effects by positing that causal framing promotes a shift from configural to elemental representations of compound cues (Melchers, Shanks, & Lachnit, 2008; Williams et al., 1994). If the representation of a compound cue is completely distinct from the representations of the elements that make it up, a pathway-strengthening process based on error-correcting learning should not result in indirect effects. However, incorporating the idea that object-framing instructions promote configural representations into the view that pathway strengthening underlies performance in both the prediction task and the RT task implies that indirect effects should disappear with object-framing instructions in both tasks. Again, the presence of indirect effects with object-framing instructions in the RT task but not in the prediction task cannot be easily explained even with this elaborated version of the pathway-strengthening account.

There is a third approach to understanding contingency learning that we have not yet explored in this article. According to this approach, learners attempt to determine the causal relations between events by performing Bayesian inference over a set of hypotheses represented as causal graph structures (Gopnik et al., 2004; Griffiths & Tenenbaum, 2005; Sobel, Tenenbaum, & Gopnik, 2004; Tenenbaum & Griffiths, 2003). Although this account is similar to an inference-based account because it invokes the concept of inference, it is framed at what Marr (1982) called the *computational level* rather than at the level of processes or mechanisms. That is, this account refers only to learners' hypothesis spaces (i.e., the range of alternative hypotheses about contingencies) and the extent to which evidence provided by experience provides a basis for inferring the probability of each hypothesis, without describing the process or mechanism through which the information is used to determine these probabilities.

This approach might predict that the framing instructions should make a difference, given that the object-framing instructions allow for a wide range of alternative possible hypotheses about how cues relate to outcomes, whereas the causal-framing instructions are far more restrictive. Like the other unitary accounts we have considered in this section, however, the computational-level inference-based account would have to be able to explain why causal-framing instructions were necessary for an indirect effect in the prediction task but not in the RT task. It is possible that the independent-cause assumption is a default assumption that participants relied on when resources were limited in the fast-paced RT task; however, if this assumption is indeed a default, it is unclear why it would not have been relied upon by participants in the object-framing condition of the explicit prediction task. It is also possible that participants' default assumptions differed between the two tasks because the two tasks engage different mechanisms, one that allows for a wide range of alternative hypotheses and another in which the hypothesis space is more restricted. Elaborating the computational-level account in this way acknowledges our key point, which is that different processes underlie indirect effects in untimed prediction tasks and tasks requiring fast-paced responding.

Our proposal can be seen as a specific version of a two-process account that makes specific predictions about the patterns of results that should be observed in the different task conditions we employed. Even if the evidence seems to favor a two-process account, the pattern of responses observed in the object-framing condition of the explicit prediction task remains to be explained. In this condition, did participants employ an inference-based strategy—one that failed to produce an indirect effect because participants' hypothesis spaces were not constrained in a way that would promote the necessary inferences? It is unclear what alternative hypotheses about the contingencies would be consistent with both the training data and participants' predictions at test. Training experience with the screening compound display (S_1S_2), which was usually followed by the dot, and with the screening singleton display (S_1), which was usually not followed by the dot, would rule out the hypothesis that the outcome occurs unless one or more objects in the display prevents it. In addition, our finding that participants' predictions for singleton test objects were related to the dot probabilities for the compound training displays in which those objects had occurred contradicts the idea that participants considered hypotheses that treated compound and singleton displays completely independently.

To explain performance in the object-framing condition of the explicit prediction task, we propose a simple, similarity-based generalization account based on the exemplar model (e.g., Medin & Schaffer, 1978) as an alternative to an inference-based account. Our account assumes that participants in the object-framing condition responded to test displays by first comparing each display to representations of displays encountered during training and then computing a probability for the occurrence of the dot on the basis of a similarity-weighted average of the dot probabilities associated with each of these displays (for more details, see the Exemplar-Based Context Model section in Additional Methods and Analyses in the Supplemental Material). Because it assumes this similarity-based weighting, this account predicts that exact matches to training displays should have greater weight than partial matches. For example, the blocking singleton object (B_1) that was presented during training exactly matches a representation of itself and partially matches the representation of the blocking compound display (B_1B_2), whereas the object from this compound event that was not presented by itself during training (B_2) does not exactly match any item presented during training but does partially match the representation of the compound display.

This model fits the data well (see Fig. S2 in the Supplemental Material); in particular, it captures the fact that participants predicted the dot almost equally often for the blocking test singleton, the screening test singleton, and the two control test singletons but predicted the dot far less frequently for the negative test singleton. The probabilities of outcomes reflected by these predictions are less extreme (less close to 1 for the dot-likely displays and less close to 0 for the dot-unlikely displays) than the probabilities associated with displays participants actually saw during training; the exemplar model predicts this pattern of results because the test items only partially matched the representations of previously experienced displays stored in memory.

In line with this account, it is possible that participants in the object-framing condition did not engage in an explicit inference process at all; instead, these participants may have simply stored representations of the displays together with their associated outcomes in memory and relied on a similarity-based activation process to make predictions about the various test singletons. The activation of related items in memory might be an important component of an explicit inference-based process, as making an inference may also depend on the activation of related events; however, in our view, making an explicit inference about the causal power of an object requires additional processes beyond the activation of related items in memory.

Our findings are broadly consistent with other models that incorporate separate implicit and explicit learning mechanisms (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Sun, Slusarz, & Terry, 2005). Such a distinction already has considerable support in research on category learning (see Ashby & Maddox, 2005, for a review of some of the earlier findings). Previous attempts to separate implicit and explicit learning processes in contingency-learning tasks have often relied on participants' self-reports of awareness of contingencies. Such reports are notoriously sensitive to probing by experimenters (Maia & McClelland, 2004), and explicit awareness might arise in parallel with implicit, pathway-strengthening-based learning (Cleeremans, 1993; Cleeremans & McClelland, 1991); these issues make it difficult to use measures of awareness as definitive evidence for the existence of two separate mechanisms.

Another type of evidence that has been used to support a two-mechanism account is based on the finding that in a contingency-learning task in which a simple, abstract rule underlies a complex set of contingent relationships, learners who perform well tend to extend the rule to new events, whereas learners who perform poorly tend to rely on cue similarity even when it conflicts with the underlying rule (Shanks & Darby, 1998). However, these differences can be explained by an inference-based account that attributes them to participants' reliance on different hypotheses rather than different mechanisms (De Houwer, 2009). Our finding that causal-framing instructions affected performance in an untimed prediction task but not in a fast-paced RT task provides a new kind of evidence to support two-mechanism accounts, and this evidence may challenge inference-based approaches more strongly than previous evidence has.

It remains possible that some form of associative or pathway-strengthening process is at work in all forms of learning and that performance in our two tasks reflects associations or pathways in two distinct systems, one system that underlies the rapid formation of explicit memories for items, hypotheses, and explicitly formulated inferences and another system that underlies the gradual strengthening of connections in processing pathways underlying stimulus-driven responding (McClelland, McNaughton, & O'Reilly, 1995). Ultimately, learning in all contingency-learning tasks may thus in some sense depend on a pathway-strengthening process (Shanks, 2010). However, this underlying similarity should not obscure the differences in how these systems function and how they contribute to task performance. Further research is needed to fully delineate the characteristics of these quite different kinds of learning systems.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by Air Force Research Laboratory Grant FA9550-07-1-0537.

Supplemental Material

Additional supporting information may be found at http://pss.sagepub.com/content/by/supplemental-data

References

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178.

Carr, A. F. (1974). Latent inhibition and overshadowing in conditioned emotional response conditioning with rats. *Journal of Comparative and Physiological Psychology*, *86*, 718–723.

Cleeremans, A. (1993). Attention and awareness in sequence learning. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 330–335). Hillsdale, NJ: Erlbaum.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop task. *Psychological Review*, *97*, 332–361.

De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, *37*, 1–20.

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

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–384.

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment* *and aversive behavior* (pp. 279–296). New York, NY: Appleton-Century-Crofts.

Krechevsky, I. (1932). Hypotheses in rats. *Psychological Review*, *39*, 516–532.

Lovibond, P. F., Been, S., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition*, *31*, 133–142.

Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the Iowa gambling task. *Proceedings of the National Academy of Sciences, USA*, *101*, 16075–16080.

Marr, D. (1982). *Vision: A computation investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Melchers, K. G., Shanks, D. R., & Lachnit, H. (2008). Stimulus coding in human associative learning: Flexible representations of parts and wholes. *Behavioural Processes*, *77*, 413–427.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral & Brain Sciences*, *32*, 183–246.

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

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–69). New York, NY: Appleton-Century-Crofts.

Shanks, D. R. (2010). Learning: From association to cognition. *Annual Review of Psychology*, *61*, 273–301.

Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 405–415.

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian inference in preschoolers. *Cognitive Science*, *28*, 303–333.

Sun, R., Slusarz, P., & Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review*, *112*, 159–192.

Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, & K. Obermeyer (Eds.), *Advances in neural information processing systems 15* (pp. 35–42). Cambridge, MA: MIT Press.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189–208.

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., & Holyoak, K. J. (1997). Determining whether causal order affects cue selection in human contingency learning: Comments on Shanks and Lopez (1996). *Memory & Cognition*, *25*, 125–134.

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 694–709.