

Outcome predictability biases cued search

Oren Griffiths¹, May Erlinger¹, Tom Beesley² and Mike E. Le Pelley¹

1. School of Psychology, UNSW Sydney, Sydney, Australia

2. Department of Psychology, Lancaster University, Lancaster, UK

Page count: 21 pages (exc. references, abstract, notes)

Word count: 6891 words (exc. references, abstract, notes)

Corresponding author:

Dr Oren Griffiths,

Mail: School of Psychology, UNSW Sydney, Sydney, Australia, 2052

Phone: +61-2-93853020

Email: oren.griffiths@unsw.edu.au

Abstract

Within the domain of associative learning, there is substantial evidence that people (and other animals) select amongst environmental cues on the basis of their reinforcement history. Specifically, people preferentially attend to, and learn about, cueing stimuli that have previously predicted events of consequence (a predictiveness bias). By contrast, relatively little is known about whether people prioritize some (to-be-predicted) outcome events over others on the basis of their past experience with those outcomes (a predictability bias). The present experiments assessed whether the prior predictability of a stimulus results in a learning bias in a contingency learning task, as such effects are not anticipated by formal models of associative learning. Previously unpredictable stimuli were less readily learned about than previously predictable stimuli. This pattern is unlikely to reflect the use of strategic search processes or blocking of learning by the context. Instead we argue that our findings are most consistent with the operation of a biased learning mechanism.

Key words: Associative learning, uncertainty, contingency learning, visual search, prediction error

Encoding predictive (cue-outcome) relationships between mental representations of events in the environment is adaptive. It allows a person to predict the future occurrence of a stimulus or event (the “outcome”) on the basis of an earlier stimulus or event (the “cue”). For example, it has been argued that increases in gold price (the cue) typically precede increases in the value of the Australian dollar (the outcome; Aspergis, 2014). If true, learning such a relationship would be financially advantageous to the learner. The manner in which such cue–outcome beliefs are formed and updated is generally well predicted by formal associative learning theories (Pearce & Mackintosh, 2010; Mackintosh, 1975; Rescorla & Wagner, 1972; Le Pelley, 2004; Esber & Haselgrove, 2011), and their theoretical descendants in other domains (Kruschke, 2006; Dayan, Kakade & Montague, 2000). At their core, these theories are accounts of why some experiences of cue-outcome relationships produce stronger learning than others; that is, of why some events are prioritized over others for learning. Traditional theories posit two primary mechanisms in this regard. The first pertains to prior knowledge of cue-outcome relationships, and the second involves prior knowledge about the cueing stimulus. These are considered briefly in turn, before we consider a third possible source, which provides the focus of the present experiments.

The first mechanism is quite straightforward: If a cue (A) is already known to precede an outcome (X), then when this outcome follows the cue subsequently, learning about the cue-outcome relationship is de-prioritized. Formally, this is implemented by having learning modulated by *prediction error*, such that unpredicted (surprising) outcomes generate stronger learning than predicted outcomes (e.g., Rescorla & Wagner, 1972). For example, if a person already possesses a (perfect) mental representation of the A-X association, then that allows them to predict the consequences of cue A occurring (X). Should those predictions be confirmed (X actually occurs), then there will be no prediction error and consequently no new learning for the A-X association. Although the formal representations of this principle

differ between models, the guiding intuition is shared: A person must be surprised by some aspect of their experience for learning to take place. There is an abundance of empirical evidence to show that a person's prior knowledge (their A-X mental association, or lack thereof) shapes how much they will learn about the relationship between cue A and outcome X when these events are later observed (see Glimcher, 2011; Pearce, 2013 for reviews).

The second class of mechanisms uses prior knowledge specific to the cueing stimuli to modulate associative learning. Perhaps the simplest instantiation of this principle is the Mackintosh (1975) model, although this principle also underlies more complex, contemporary models of learning (Le Pelley, 2004; Esber & Haselgrove, 2011; Pearce & Mackintosh, 2011). Within the Mackintosh (1975) model, the cueing stimuli that have been shown to consistently predict important outcome events are subsequently prioritized during learning. Evidence consistent with this idea in humans comes from the *learned predictiveness* effect (e.g., Le Pelley & McLaren, 2003; Lochman & Wills, 2003; see Le Pelley, Mitchell, Beesley, George & Wills, 2016, for a review). In a typical demonstration, people must first learn to predict an outcome (X) using two cues (A and B). The task is arranged such that cue A reliably predicts outcome X, but cue B is uninformative as to whether outcome X will occur. Subsequently (in a second phase), people are found to be faster to learn an association between cue A and a novel outcome (Y) than between cue B and this same novel outcome, even if cues A and B are (objectively) equally good predictors of outcome Y. This suggests that people prioritize learning about cues that have been experienced as predictive (A) compared to cues previously experienced as uninformative (B).

These two mechanisms exploit existing knowledge about (i) cue–outcome associations, and (ii) the predictiveness of the cues, to prioritize the allocation of learning resources to events that are mostly likely to be related to each other. Yet, as recently argued

by Griffiths & Thorwart (2017), a third source of potentially useful information is missed in these formal conceptualizations: information about the *outcome* stimulus or event.

Outcome-driven selectivity

It is also possible that people learn which aspects of the environment tend to be predictable (and are thus likely to be the “signal”) and which tend to be unpredictable (and are thus likely to be the “noise”), and prioritize their learning resources accordingly. Only one study has directly investigated this predictability hypothesis in humans. Griffiths, Mitchell, Bethmont & Lovibond (2015) used a causal learning task in which people were asked to learn which foods elicited allergic reactions in a fictional patient, Mr Smith. Participants were shown the different foods Mr Smith ate (the cues) and the type of allergic reaction that he experienced (the outcomes) after eating those foods. There were two stages to the task. Crucially, the cues (foods) differed between the two stages, but the outcomes (the allergic reactions) were common to both. In both stages, people needed to predict the type of skin reaction Mr Smith would have (itchiness, swelling or no reaction) and also the type of stomach reaction he would have (bloating, cramping or no reaction). They were given corrective feedback on every trial, and thereby learned the food-reaction associations. The task was arranged so that one type of reaction (e.g., skin reactions) was *predictable* during Stage 1, such that each time Mr Smith ate a particular food (e.g., lettuce) he would experience a particular symptom on the skin dimension (e.g., swelling). The other reaction dimension (stomach reactions) was *unpredictable* on the basis of the food cues shown in Stage 1. For example, after eating a second food (e.g., leek), it was equally likely that Mr X would experience any of the three values on the stomach reaction dimension (cramping, bloating or no reaction).

In Stage 2, a new set of foods was introduced (all were fruits, versus vegetables in Stage 1). However, the to-be-predicted allergic reactions were retained. People thus needed to

learn a new set of food-reaction associations between the novel Stage 2 foods and the existing reactions (skin and stomach reactions). Critically, in Stage 2 the foods and reactions were arranged so that both types of reaction were perfectly and equally predictable on the basis of the novel food cues; objectively, skin reactions could be predicted equally as accurately as stomach reactions. However, people more readily learned to predict the previously predictable outcome (skin reactions) than the previously unpredictable outcome (stomach reactions). This observed prioritization of learning for outcomes that were previously predictable, over those that were previously not (termed the *Outcome Predictability bias*) suggests that people do indeed make use of prior knowledge about outcome stimuli when selecting which stimuli to prioritize for learning.

However, some aspects of Griffiths et al.'s (2015) procedure limit the generalizability of this conclusion. First, people were required to make predictions about both outcomes (skin and stomach reactions) on every trial. This means that the observed bias for the previously predictable versus unpredictable outcome may have been a product of a within-trial interaction between the outcome stimuli, whereby one outcome event actively inhibited the capacity of the concurrent outcome to become associated with the target cue (analogous to "lateral inhibition" amongst cues; Kruschke, 2006, Griffiths & Le Pelley, 2009). Thus, it is not clear that this bias extends to instances in which outcomes must compete for learning across separate trials. If predictability is an encoded feature of the outcome's mental representation, then the Outcome Predictability bias ought not to depend upon both outcomes being present at the same time.

Second, in Griffiths et al.'s (2015) procedure the clearest evidence for learning biases was seen in a self-reported test phase that occurred after training, which explicitly asked people to judge the likelihood that the outcome would occur given the presence of a cue. This raises the possibility that the observed bias in this task was not reflective of a bias in learning

per se. Instead, people may have correctly learned all the relevant cue-outcome contingencies during training, but causally discounted some of these learned contingencies when forming explicit likelihood judgments at test (Baker, Mercier, Vallee-Tourangeau, Frank & Pan, 1993; Fugelsang & Thompson, 2001). This method makes it difficult to determine the relative contributions of any bias that occurred during learning, as distinct from any bias in causal reasoning that may have occurred during the test stage. Thus, the present procedure was designed to minimize the *requirement* to engage in this form of causal reasoning by having participants perform a task (visual search) in which perfect response accuracy could be achieved with no explicit causal knowledge regarding the cues and outcomes (see also Beesley & Le Pelley, 2010). In addition, in contrast to Griffiths et al. (2015), the primary dependent variable used in these experiments (trial-by-trial response latency) did not require people to explicitly reflect upon their earlier learning.

Experimental series

The present experiments sought to test whether previously predictable outcomes were more readily learned about than previously unpredictable outcomes. For this purpose, a visual search task was used in which the location of the target stimulus was predictable when it occurred in some locations, but was unpredictable when it occurred in others. Then, in the second stage, the target stimulus was rendered predictable both when it was located in a previously predictable location, and also when it appeared in a previously unpredictable location. The key question, and hypothesis tested in Experiment 1, was whether participants would be biased in their capacity to learn to predict the location of the target stimulus when it appeared in a previously predictable location, versus a previously unpredictable location. This is distinct from the possibility that people learn to generally orient towards previously predictable locations (a biased search), which is addressed in Experiment 2.

Experiment 1

On each trial in Experiment 1, participants were required to respond as rapidly as possible according to the orientation of an arrow that could appear in one of eight possible locations in a search array. Each array was preceded by a cue stimulus: a line drawing of a common object (e.g., a carrot). Some cue stimuli (cues A and B) perfectly predicted the location of the target arrow in the upcoming search array (but not its orientation). The locations where the target occurred on these trials (e.g., upper and lower locations) are referred to as the *predictable locations*. Some cue stimuli (cues C and D) preceded the target arrow appearing in two different locations (e.g., left and right positions), each with 50% probability. These were the *unpredictable locations*, as the exact location of the target arrow could not be known in advance. Other filler cues were included to balance various design features of the experiment, such as making sure that the target could appear in all possible locations in the search array.

Then in a second stage, a new set of cue images was shown. Critically, these new cueing stimuli allowed the appearance of the target arrow to be reliably predicted during Stage 2 in both the previously predictable locations (upper and lower positions), and also in the previously unpredictable locations (left and right positions). If people were biased to learn about previously predictable locations (over previously unpredictable locations), then they ought to learn to locate the target arrow more rapidly when it was cued to appear in a previously predictable location, than in a previously unpredictable location.

Method

Participants, apparatus, and stimuli. Forty-one undergraduate students (mean age = 19.7 years, 17 were female) participated in exchange for course credit. The study was approved by UNSW Human Research Ethics Advisory Panel, and was conducted in

accordance with the ethical standards laid down in the Declaration of Helsinki. Stimuli were presented on a widescreen monitor (1920 × 1080 resolution, 60 Hz) at a distance of approximately 50cm. Stimulus presentation was controlled by MATLAB (Mathworks, 2012) using Psychophysics Toolbox extensions (Kleiner, Brainard, Pelli, Ingling, Murray & Broussard, 2007).

All stimuli were presented on a black background. Cue stimuli were ten images taken from Snodgrass & Vanderwart's (1980) database (airplane, banana, bell, book, butterfly, cake, carrot, cat, chair, and chicken), presented centrally at a size of $7.6^\circ \times 8.9^\circ$ visual angle. The coloured shape icons that appeared in the search array each measured $5.4^\circ \times 6.3^\circ$. The set of 60 coloured shapes consisted of every possible combination of 6 colours (red, yellow, blue, orange, green, purple) and 10 shapes (squares, circles, vertical ovals, horizontal ovals, vertical rectangles, horizontal rectangles, pentagons, hearts, diamonds, and triangles). There were four types of arrow-like stimuli that could appear on these shapes: left-facing target arrow, right-facing target arrow, horizontal line distractor image, bidirectional arrow distractor image (all appear in Figure 1).

[Figure 1 about here.]

Design. The design of Experiment 1 is outlined in Figure 2. Every trial began with the presentation of a cue stimulus, which was followed after 1.5s by a search array of 8 coloured shapes arranged in a ring around the cue stimulus (at 6.5° eccentricity). An arrow-like icon was displayed on each of these coloured shapes: one shape contained a target arrow (a single-headed arrow pointing either left or right); all other shapes contained a distractor icon (a horizontal line or a bidirectional arrow, determined randomly for each shape). The critical locations throughout the experiment were the cardinal locations (upper, lower, left, right) in the search array. These locations were assigned to roles in a counterbalanced fashion. For half of participants, appearance of the target arrow in upper and lower locations was *predictable*

(on the basis of the cue stimulus) during Stage 1, but appearance of the target in left and right locations was *unpredictable*. For the remaining participants these assignments were reversed. Consider a participant for whom upper/lower locations were predictable and left/right locations were unpredictable. For this participant, during Stage 1, one image (cue A) predicted the location of the target image in the upper location with 100% probability, and a second cueing image (B) predicted the location of the target image in the lower position with 100% probability. Two further cues, C and D, predicted that the target arrow would appear in either the left or right positions, each with 50% probability. On the remaining trials, a fifth cue (E) was shown. On these trials the target arrow appeared equally often in each of the four diagonal locations (for all participants). Cue E was included so that the target appeared in all locations, and thus all locations were task-relevant. This cue was otherwise irrelevant to our hypotheses and is not discussed further.

[Figure 2 about here.]

In Stage 2, five new cues (F-J) were shown. Cues F-I each reliably indicated that the target would appear in a particular cardinal location (upper, lower, left, right). Two of these cardinal locations had been well predicted (by cues A and B) in Stage 1 and two of these locations had been somewhat unpredictable (following cues C and D) in Stage 1. The final cue, cue J, preceded the target appearing in one of the four diagonal locations.

The assignment of each image (e.g., carrot, plane) to each cue role (e.g., A, B) was randomly determined for every individual. The particular coloured shape (e.g., blue square for cue A) upon which the target arrow appeared was randomly determined for each participant, with the constraint that both the shape and colour of each stimulus differed between trial-types (e.g., if the blue square appeared on cue A trials, the shape that appeared on cue B trials could neither be blue nor a square). Within each trial, the colours and shapes of the irrelevant distractor stimuli were randomly determined, with the caveat that no

distractor image could be identical to the coloured shape in which the target arrow appeared (e.g. there were no blue square distractors on cue A trials).

Procedure. People were instructed that their task was to find a unidirectional arrow in the search array, and to press either the left (“1” key) or right (“9” key) button, to indicate the orientation of that arrow, as rapidly as possible. They were also told that the arrow would appear in an array of similar-looking icons, and that the pictures shown in the middle of the screen (i.e. the cueing images) would “help you find the target arrow faster.”

The procedure of each trial is summarized in Figure 1. Each trial commenced with the presentation of a white fixation cross for 1s. The cross was then replaced by the cue image. After 1.5s, the search array appeared, and participants could make their response. If the response was correct, a chime sounded and the word “CORRECT” appeared for 0.5s. If the response was incorrect, a buzzer sounded and “INCORRECT” appeared for 2.5s. The screen then blanked, and the next trial began after 1s.

The five trial-types (indicated by letters A-E in Figure 2) were arranged into blocks of twenty trials (four per trial-type), and the trial order was randomized within each of these blocks. People completed 5 such training blocks (100 trials in total) in Stage 1, and then proceeded to Stage 2. The transition between stages was not signalled. The five trial-types shown in Stage 2 (indicated by letters F-J) were arranged into blocks of twenty trials (four per trial-type), and the trial order was randomized within each of these blocks. People completed 5 such training blocks (100 trials in total) in Stage 2.

Results

Trials with erroneous responses were excluded (2.07% of all trials), as were trials with exceptionally long response latencies that were indicative of distraction from the task; any latencies more than 5 standard deviations from the overall mean were removed (0.33%). All

data from one participant were excluded as this participant achieved only 53% accuracy on the task. The average error rate for the remaining participants was 2.07% (SD = 1.71).

Mean response latencies were separated according to the location of the target on that trial: locations that were 100% predictable in Stage 1 (termed predictable), locations that were 50% predictable in Stage 1 (termed unpredictable). The mean response latencies for targets appearing in each of these locations across Stages 1 and 2 are shown in Figure 3. The response latency data were analysed separately for each Stage using a 2×2 ANOVA, with factors of predictability (predictable versus unpredictable in Stage 1) and training (first half, second half). Familywise error rates were controlled at .05.

[Figure 3 about here.]

Stage 1. A main effect of training was observed, $F(1,40) = 58.73, p < .001, \eta^2_P = 0.59$, with response times decreasing across training. Importantly, there was a main effect of predictability, $F(1,40) = 20.38, p < .001, \eta^2_P = 0.34$, with faster responses when the target location was predictable (cue A and B trials) than when it was unpredictable (cue C and D trials). While there appeared to be a trend towards the effect of predictability increasing across training, predictability did not significantly interact with the training factor, $F(1,40) = 3.68, p = .06, \eta^2_P = 0.08$.

Stage 2. A main effect of training, $F(1, 40) = 83.93, p < .001, \eta^2_P = 0.68$, was observed. Response times decreased across training. Most importantly the main effect of prior predictability was significant, $F(1,40) = 6.84, p = .01, \eta^2_P = 0.15$, such that people responded more rapidly to the target in the previously predictable locations than in the previously unpredictable locations. Although the effect of prior predictability appeared to increase across training, the training \times prior predictability interaction was not significant, $F(1, 40) = 3.76, p = .06, \eta^2_P = 0.09$.

Discussion

People readily learned to predict the location of the target outcome during Stage 1, and (unsurprisingly) were able to do so more readily when that location was completely predictable than when the location was somewhat unpredictable. Crucially, this difference in the probability of target location had an effect on responses in Stage 2, which used a new and independent set of cueing images. That is, even though the new cues F-I each perfectly (i.e. with 100% conditional probability) predicted the location of the target stimulus on each trial, people more rapidly located the target arrow when it appeared in a location that was previously predictable than when it appeared in a location that was previously unpredictable. This difference suggests that people were biased to learn more readily about the new cueing relationships that involved locations that were previously predictable than those that were previously unpredictable.

There is an alternative interpretation of these data. In the prior analysis, it was assumed that prior learning about the predictability of a target location allows the participant to more readily associate that location with a novel cue in Stage 2. An alternative possibility is that people learned which locations were predictable during Stage 1, and then preferentially searched those locations in the future irrespective of the cue shown on that trial (see also effects of “search history” on visual search, Awh, Beloposky & Theeuwes, 2012). Such an effect could result in faster responding to the previously predictable outcome locations (than the previously unpredictable outcome locations), in the absence of any cue-location learning during Stage 2, and thereby undermines any conclusions drawn about the influence of prior predictability on differential learning rates. This is because it is possible that the observed differential response latencies between locations in Experiment 1 did not reflect cue-location learning during Stage 2, but instead reflected a cue-independent general bias to preferentially

look for the target in some locations over others. This possibility is addressed in Experiment 2.

Experiment 2

The key point of difference between a *biased learning* versus *biased search* account is the dependence of the bias on the cue shown prior to the search array. This is because the *biased search* hypothesis predicts that people will search previously predictable locations first, irrespective of the cue shown on that trial: it does not depend on new cue-outcome learning. Thus, it predicts that the bias will be just as evident on trials in which a completely uninformative (or *invalid*) cue is shown, as when an informative (or *valid*) cue is shown. By contrast, under a *biased learning* account, people would more readily learn to associate novel cues with previously predictable (as opposed to previously unpredictable) locations. That is, this account relies on a difference in the rate at which people learn cue–location associations during Stage 2, with these associations allowing them to form expectations about the likely location of the target. Therefore this account states that the predictability bias (faster responding to the target when it appears in a previously predictable location) should manifest only when the target is preceded by a valid cue – since it is only under these circumstances that participants can generate the appropriate expectations. In contrast, no bias should be observed on trials in which the cueing stimulus is absent or uninformative (i.e. those trials with an *invalid* cue).

To discriminate between these accounts, Experiment 2 replicated Experiment 1, but now included an *invalid* cue in each stage (cue E in Stage 1, cue J in Stage 2). This cue was uninformative as to where the target would appear, because it preceded the target appearing equally often in all of the 8 possible locations (with each occurring 12.5% of the time). If people biased their search for the previously predictable locations in general in Stage 2, then

people should more readily find the target when it appeared in a previously predictable location (than in a previously unpredictable location) both on validly cued trials (i.e. after cues F-I) and also following the non-informative invalid cue J. By contrast, if people's cue-location learning was biased in Stage 2, then the facilitated detection of the target arrow should occur only on validly cued trials (F-I), and not following the non-informative cue J.

Method

Participants. Sixty-five undergraduate students (mean age = 19.1 years, 49 were female) participated in exchange for course credit.

Design. The design was similar to Experiment 1, but with one change. The target arrow could appear in any of the eight locations (each with equal probability: 12.5%) following presentation of cue E (in Stage 1) and cue J (in Stage 2).

Procedure. The procedure was similar to that of Experiment 1. In order to improve learning of the Stage 1 contingencies, the trial-types were initially clustered (see Mitchell, Griffiths, Seetoo & Lovibond, 2012) such that 8 trials involving cues A and B were shown in one block, then 8 trials involving cues C and D were shown in a second block, and then a third block consisted of 8 trials involving cue E. These 3 blocks (one with cues A and B, one with cues C and D, and one with cue E) were each repeated four times, for a total 96 trials. Whether participants received blocks of predictable (cues A and B) or unpredictable (cues C and D) trials first in this sequence was counterbalanced across individuals. Once these clustered blocks were completed, participants were given a further 144 trials of Stage 1 training in which all of the trial-types were randomly intermixed. These Stage 1 trials were organized into blocks of 12 trials: Two repetitions each of trial-types A-D (which each used a valid cue) and four repetitions of the invalidly cued trial-type (cue E). In Stage 2, 20 blocks of six trials each were presented, for a total of 120 trials. Each block consisted of one

repetition of each trial-type involving a valid cue (F-I), and two repetitions of the invalidly cued trial-type (cue J). The trials were randomized within each block, for both Stage 1 and 2.

Results

As in Experiment 1, response latencies for incorrect responses were omitted, as were any response latencies that were more than 5 standard deviations from the individuals' overall mean. This resulted in the removal of 1.98% of the data.

Mean response latencies were separated into four trial-types that were demarcated by (i) whether the target appeared in one of the predictable locations versus one of the unpredictable locations, and (ii) whether that trial featured a 'valid' cue stimulus (cues A-D in Stage 1, cues F-I in Stage 2) or an 'invalid' cue stimulus (cue E in Stage 1, cue J in Stage 2). The mean response latencies for each of these trial-types across Stages 1 and 2 are shown in Figure 4. The response latency data were analysed separately for each Stage using a $2 \times 2 \times 2$ ANOVA, with factors of predictability (predictable versus unpredictable location in Stage 1), cue (validly or invalidly cued), and block (first half of training versus second half of training stage).

[Figure 4 about here.]

Stage 1. A main effect of training was observed, $F(1,64) = 18.74, p < .001, \eta^2_P = 0.23$, with faster responses in the second half of Stage 1 than in the first half. Similarly, responses were faster when the target arrow was preceded by a valid cue (A-D) than when it was preceded by the invalid cue (E), $F(1,64) = 85.24, p < .001, \eta^2_P = 0.57$. Averaged across validly and invalidly cued trials, there was a numeric benefit for the predictable locations, but this did not reach significance, $F(1,64) = 3.67, p = .06, \eta^2_P = 0.05$. However, there was a cue \times predictability interaction, $F(1,64) = 21.48, p < .001, \eta^2_P = 0.25$. Simple effect analyses revealed that responses were faster when the target appeared in a predictable location than an unpredictable location following a valid cue, $F(1,64) = 29.05, p < .001, \eta^2_P = .31$, but the

trend towards faster responses to the unpredictable locations than the predictable locations following the uninformative cue (E) did not reach significance, $F(1,64) = 3.14, p = .08, \eta^2_P = .05$. There was no significant three-way interaction, $F(1,64) = 3.90, p = .05, \eta^2_P = 0.06$.

Stage 2. There was a main effect of training, $F(1,64) = 5.23, p = .03, \eta^2_P = 0.08$, with response latencies generally decreasing over the course of Stage 2. There was also a main effect of cue, $F(1,64) = 45.47, p < .001, \eta^2_P = 0.42$, with faster responses when the target was preceded by a valid cue than an invalid cue. There was no significant main effect of prior predictability, $F < 1$, but as in Stage 1, there was a significant cue \times predictability interaction, $F(1,64) = 6.50, p = .01, \eta^2_P = 0.09$. There was also a three-way interaction, $F(1,64) = 4.11, p = .047, \eta^2_P = 0.06$. Simple effect analyses found no significant differences between the predictable and unpredictable locations (either when validly cued or when preceded by an uninformative cue) during the first half of Stage 2, all F s < 1 . However, response latencies to these locations did differ in the second half of training. Specifically, responses were faster for previously predictable locations than previously unpredictable locations following a valid cue (cues F-I), $F(1,64) = 6.04, p = .02, \eta^2_P = .09$, but were faster for previously unpredictable locations than previously predictable locations following the invalid cue J, $F(1,64) = 4.44, p = .04, \eta^2_P = .06$.

Discussion

The primary finding of Experiment 1 was replicated: participants responded more rapidly to the previously predictable locations, than to the previously unpredictable locations during Stage 2. However, this bias in responding was only obtained when a valid cueing stimulus preceded the presentation of the target. On trials that included an invalid cue (J) there was no benefit in detection of the target stimulus in the previously predictable location over the previously unpredictable location. In fact, there was a significant effect in the opposite

direction, whereby responses to the target were faster when it appeared in previously unpredictable locations after the uninformative cue. This effect was not anticipated, and possible explanations are discussed in the General Discussion. Overall, however, these data strongly suggest that biased cue-outcome learning is the cause of the Outcome Predictability bias in the present task, rather than adoption of a preferential search strategy that favoured previously predictive locations irrespective of the cue that was presented prior to the target.

General Discussion

Across two experiments, participants more readily learned to detect a target stimulus in previously predictable locations than in previously unpredictable locations. This Outcome Predictability bias was not a direct product of previously learned cue-location associations, as the crucial second stage of the procedure used novel cues, so any associations between the cues shown in Stage 1 and the target locations would not have aided performance in Stage 2. Instead, people had to learn new cue-location associations in Stage 2, and this process was reliably biased towards learning about the locations that were previously predictable.

Notably, our study extends previous work on the Outcome Predictability bias (Griffiths et al, 2015) by demonstrating this effect in a procedure in which participants were not *required* to learn cue–outcome associations at all. In Griffiths et al.’s (2015) study, participants were required to make an explicit outcome prediction on every trial; in contrast, in the current study participants could achieve perfect response accuracy without learning anything about the predictive status of the cue stimuli. The observation of an Outcome Predictability bias under these conditions of ‘incidental’ learning suggests a more general influence of this bias that operates even when there is no explicit requirement to learn. Perhaps more importantly, our studies demonstrate an influence of the Outcome Predictability bias during the process of learning (i.e., on participants’ responding during the focal task), rather than in a distinct test of explicit knowledge that occurred only once

learning was complete (as used by Griffiths et al., 2015). This shows that the Outcome Predictability bias has an ‘online’ effect on the formation of cue–outcome beliefs, rather than merely influencing the way in which these beliefs are translated into explicit judgments on test (e.g., through differences in causal discounting).

The current studies also allow us to rule out certain alternative accounts of the Outcome Predictability bias. Firstly, Experiment 2 tested whether people favoured previously predictable locations during their search in Stage 2 (a biased search account), against the possibility that people more readily learned cueing relationships involving previously predictable locations (a biased learning account). In this experiment, the Outcome Predictability bias occurred only on trials in which a valid cue was shown, suggesting that the facilitated responding for some locations over others was conditional on the cue shown on that trial. The observation that the bias was conditional on the cue shown on each trial strongly suggests that it was the specific learning of cue-outcome relationships during Stage 2 that facilitated responding for previously predictable locations, rather than a more generalized bias to search in these locations first.

Secondly, in Griffiths et al.’s (2015) study of the Outcome Predictability bias, two cues and two outcomes were presented within the same trial. One possible interpretation of these experiments was that the concurrent presentation of a previously predictable outcome directly inhibited any learning about the previously unpredictable outcome. This explanation cannot account for bias observed in the present data, however, as the “outcomes” in this task (the location of the target stimulus) were mutually exclusive: the target appeared in only one location on each trial. This means that stimulus-specific, direct inhibition from previously predictable to previously unpredictable outcomes is unlikely to be the source of the observed effects. Similarly, the restriction of one cue and one outcome per trial in the present task also excludes a class of explanations (suggested by Griffiths et al.) in which the observed bias in

responding is a product of people applying a “rule” whereby each cue can only elicit a single outcome. Such a rule is unlikely to be learned or used in the present experiments, because several cues (C-E and J) were each unambiguously associated with multiple (2 or more) locations across different trials.

Finally, the current data are also inconsistent with an account of the Outcome Predictability bias in terms of ‘blocking’ (Griffiths & Le Pelley, 2009; Kamin, 1969; Shanks, 1985). This account runs as follows. During Stage 1, the unpredictable outcomes were not well predicted by the cues that (inconsistently) preceded them (C and D). Hence these outcomes might have formed significant associations with the stimuli representing the experimental context: in effect participants learn that these outcomes cannot be accurately predicted by the preceding cues, but instead are more generally predicted by the context. In contrast, the predictable outcomes were well predicted by the cues that preceded them in Stage 1 (A and B), and hence would presumably have formed strong associations with these cues and *not* with the context. Importantly, the contextual stimuli are also present during Stage 2. This would mean that—at the start of Stage 2—the previously unpredictable outcomes were less surprising (since they were partially expected on the basis of the contextual stimuli) than the predictable outcomes (which were not). According to models that view surprise as the crucial determinant of learning (e.g. Rescorla & Wagner, 1972; Wagner, 1981), learning about the new cue–outcome associations for the previously unpredictable outcomes would therefore be blocked by this prior learning about the context. This could generate the Outcome Predictability bias observed in Experiment 1 (and by Griffiths et al., 2015). However, this approach has difficulty with the findings of Experiment 2. The blocking account anticipates that, since previously unpredictable locations are already expected on the basis of contextual stimuli during Stage 2, these locations should be (relatively) blocked from forming associations with *any* novel cues. However, data from the invalid cue trials of

Experiment 2 suggest that participants formed *stronger* associations between this invalid cue (J) and previously unpredictable locations, relative to previously predictable locations. This finding is therefore inconsistent with the blocking account of the Outcome Predictability bias.

This finding—that the Outcome Predictability bias was reversed on the trials in which the cueing stimulus (the invalid cue) provided no useful information—was unexpected and interesting. It is not immediately clear why the bias in responding reversed on these trials. One possibility is that people more readily associated a previously unpredictable outcome, than a previously predictable outcome, with non-predictive cues (see also Griffiths et al., 2015, for a similar pattern of results). If so, this suggests the intriguing hypothesis that people may seek to “match” the degree of predictiveness of the cueing stimuli with the predictability of the outcome stimuli, and thereby facilitate learning for those cue-outcome pairings that match in this manner. Such an approach would be consistent with a *causal model* view of cue-outcome learning, whereby people construct a model of the manner in which cues are related to outcomes (e.g. Waldmann, 1996; De Houwer, 2009), which includes information about the degree of predictiveness/predictability, and then preferentially search for new components (new cues or new outcomes) that match the properties of that inferred model. At a computational level, this might consist of learning and parameterizing a mental (causal) model (Griffiths & Tenenbaum, 2005, 2009; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008) which, by the end of Stage 1, would place highest credibility on high values (~ 1) of the “strength” parameter for the predictable outcomes, and most credibility on middling values (~ 0.5) for the unpredictable outcomes. Yet this implementation also requires an assignment of parameters that is atypical in the literature. Specifically, because only the outcomes are shared between Stages 1 and 2, the relevant parameter shaping belief updating in Stage 2 cannot refer to a particular cue (its “causal strength” parameter) or a hypotheses relating a

particular cue to a particular outcome (i.e. the likelihood of a given causal structure), but must instead be associated with the outcome stimulus (akin to a “predictability” parameter).

Theoretical implications of the Outcome Predictability bias

The key finding of the current study—that novel learning is systematically biased by prior experience of the predictability of an outcome—is important because this finding lies beyond seminal models of associative learning (Rescorla & Wagner, 1972; Mackintosh, 1975; Pearce & Hall, 1980) and more contemporary versions (Le Pelley, 2004; Le Pelley et al, 2016; Esber & Haselgrove, 2011; Pearce & Mackintosh, 2010). None of these existing models allow for a learned property of the outcome stimulus, such as its prior predictability, to affect the rate of subsequent learning involving that outcome. Instead these models have historically been cue-centric, in that they have focused on the degree to which learning about cueing stimuli takes place, and the various mechanisms by which cues might compete for formation of associations with outcomes. There has been substantially less investigation about the dynamic processes involving outcome events, such as how they too might compete for learning (but see Rescorla, 1980). The learned helplessness effect (Maier & Seligman, 1976) is one notable exception to this trend, whereby it has been shown that repeated failure to predict (or control; Burger & Arkin, 1980) an outcome event results in impaired learning *in general*. An important contribution of the present research is to show that this effect can be *stimulus-specific*, whereby the experience of unpredictability with a particular outcome stimulus results in the de-prioritization of that stimulus (but not all stimuli) during subsequent learning. Moreover, Experiment 2 shows that this outcome-specific impairment in learning is itself modulated by the predictiveness of the cue shown in conjunction with that outcome stimulus. In sum, the present data suggest that the learned helplessness effect may be more nuanced than previously thought. More generally, the finding that the training history of an

outcome shapes how readily that outcome enters into future associations demonstrates an important shortcoming of existing models of associative learning. This could potentially be addressed by adding a parameter that tracks the prior predictability of outcome stimuli, similar to the α parameter that indexes cue predictiveness in existing models (Mackintosh, 1976; Le Pelley, 2004). Formal accounts of the outcome predictability bias—perhaps with a quantitative model along these lines—will clearly be necessary at some stage. However, more empirical data are required before we are in a suitable position to develop such a formal model: we need to further probe the empirical parameters of this bias and hence establish more precisely the properties that need to be modelled (see also Thorwart, Livesey, Wilhelm, Liu & Lachnit, 2017, for initial empirical data).

Conclusion

People more readily learned a contingent relationship between a cue stimulus and a target location that was previously predictable, than an otherwise equivalent relationship involving a previously unpredictable location. This finding goes beyond prior causal learning protocols investigating this phenomenon (Griffiths et al, 2015) by providing evidence that this is a learning bias (as distinct from an effect of post-hoc reasoning or biased search), this is evident when learning is incidental to the participant's task. Cumulatively these data show an important new means by which people prioritize learning for some stimuli in their environment over others, that is yet to be properly accommodated in formal theories in the domain.

References

- Apergis, N. (2014). Can gold prices forecast the Australian dollar movements? *International Review of Economics & Finance*, 29, 75-82.
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437-443.
- Baker, A. G. (1976). Learned irrelevance and learned helplessness: Rats learn that stimuli, reinforcers, and responses are uncorrelated. *Journal of Experimental Psychology: Animal Behavior Processes*, 2(2), 130-141.
- Baker, A.G. & Macintosh, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning & Behavior*, 5(3), 315–319.
- Baker, A. G., & Mackintosh, N. J. (1979). Preexposure to the CS alone, US alone, or CS and US uncorrelated: Latent inhibition, blocking by context or learned irrelevance? *Learning and Motivation*, 10(3), 278-294.
- Baker, A. G., Mercier, P., Gabel, J., & Baker, P. A. (1981). Contextual conditioning and the US preexposure effect in conditioned fear. *Journal of Experimental Psychology: Animal Behavior Processes*, 7(2), 109.
- Baker, A. G., & Mercier, P. (1982). Manipulation of the apparatus and response context may reduce the US pre-exposure interference effect. *The Quarterly Journal of Experimental Psychology*, 34(4), 221-234.
- Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(2), 414.

- Beesley, T., & Le Pelley, M. E. (2010). The effect of predictive history on the learning of sub-sequence contingencies. *The Quarterly Journal of Experimental Psychology*, 63(1), 108–35.
- Burger, J. M., & Arkin, R. M. (1980). Prediction, control, and learned helplessness. *Journal of Personality and Social Psychology*, 38(3), 482.
- Cousineau, D. Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, 1, 42-45.
- Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature Neuroscience*, 3, 1218-1223.
- Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.
- De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, 37(1), 1-20.
- Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: a model of attention in associative learning. *Proceedings of the Royal Society of London B: Biological Sciences*, 278(1718), 2553-2561.
- Fugelsang, J. A., & Thompson, V. A. (2001). Belief-based and covariation-based cues affect causal discounting. *Canadian Journal of Experimental Psychology*, 55(1), 70.
- Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3), 15647-15654.
- Griffiths, O. & Le Pelley, M. E. (2009) Attentional changes in blocking are not a consequence of lateral inhibition. *Learning & Behavior*, 37(1), 27-41.

- Griffiths, O. & Thorwart, A. (2017) Effects of outcome predictability on human learning. *Frontiers in Psychology*, 8, 511.
- Griffiths, O., Mitchell, C. J., Bethmont, A., & Lovibond, P. F. (2015). Outcome predictability biases learning. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41(1), 1.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51(4), 334-384.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, 116, 661–716
- Johansen, J. P., & Fields, H. L. (2004). Glutamatergic activation of anterior cingulate cortex produces an aversive teaching signal. *Nature Neuroscience*, 7(4), 398-403.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.
- Kruschke, J. K. (2006) Locally Bayesian learning with applications to retrospective revaluation and highlighting. *Psychological Review*, 113 (4), 677.
- Larkin, M. J., Aitken, M. R., & Dickinson, A. (1998). Retrospective revaluation of causal judgments under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(6), 1331.
- Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology: Section B*, 56(1), 68-79.
- Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology*, 57B(3), 193-243.

- Le Pelley, M. E., Beesley, T. & Griffiths, O. (2014) Relative salience versus relative validity: Cue salience influences blocking in human associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition*, 40 (1), 116.
- Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychological Bulletin*, 142(10), 1111.
- Lochmann, T. & Wills, A.J. (2003) Predictive history in an allergy prediction task. *Proceedings of European Cognitive Science Society*, 3, 217-222.
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological Review*, 115(4), 955.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82(4), 276-298.
- Mackintosh, N.J. (1973). Stimulus selection: Learning to ignore stimuli that predict no change in reinforcement. In R.A. Hinde & J.S. Hinde (Eds.), *Constraints on learning* (pp. 75-96). London: Academic Press.
- The MathWorks, Inc. (2012) *MATLAB and Statistics Toolbox Release 2012b*. Natick, Massachusetts, United States.
- Maier, S. F., & Seligman, M. E. (1976). Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General*, 105(1), 3.
- Mitchell, C. J., Griffiths, O., Seetoo, J., & Lovibond, P. F. (2012). Attentional mechanisms in learned predictiveness. *Journal of Experimental Psychology: Animal Behavior Processes*, 38(2), 191.
- Pearce, J. M., & Mackintosh, N. J. (2010). Two theories of attention: A review and a possible integration. In C. Mitchell and M. Le Pelley's (Eds.) *Attention and associative learning: From brain to behaviour*, 11-39.

- Pearce, J. M. (2013). *Animal learning and cognition: An introduction*. Psychology Press.
- Randich, A., & LoLordo, V. M. (1979). Associative and nonassociative theories of the UCS preexposure phenomenon: Implications for Pavlovian conditioning. *Psychological Bulletin*, 86(3), 523–548.
- Rescorla, R. A. (1980). *Pavlovian second-order conditioning*. Hillsdale, NJ: Erlbaum
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2, 64-99.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2), 174.
- Thorwart, A., Livesey, E., Wilhelm, F., Liu, W. & Lachnit, H. (2017) Learned predictiveness and outcome predictability effects are not simply two sides of the same coin. *Journal of Experimental Psychology: Animal Behavior Processes*, 43, 341-365.
- Waldmann, M. R. (1996). Knowledge-based causal induction. *Psychology of Learning and Motivation*, 34, 47-88.

Acknowledgments.

Dr Griffiths was supported by an Australian Research Council (ARC) Discovery Early Career Research Award (DE150100667). Dr Beesley was supported by ARC Discovery Project (DP140103268). A/Prof Le Pelley was supported by National Health and Medical Research (APP1069487).

Figure captions.

Figure 1. Example sequence of trial events in Experiment 1. After fixation, a cue stimulus appeared, followed by a search display. Participants were required to respond according to the direction of a single-headed target arrow in this search display.

Figure 2. The design of Experiment 1, for a participant for whom upper and lower locations were predictable, and left and right locations were unpredictable (this assignment was counterbalanced across participants). Each line drawing represents a cue (labelled A – J in each cell). The next column depicts the possible locations of the target arrow on each trial of that trial-type. The third column notes the conditional probability (in percentage) of the target appearing in one of the possible locations, given the presence of the cue shown in that row.

Figure 3. Mean response latencies in Experiment 1, separated by the predictability of the location in which the target icon appeared. Data are shown for trials in the 1st and 2nd half of each training stage. In Stage 1, some locations were predictable on the basis of the preceding cue (“Predictable”) whereas other locations were only 50% predictable (“Unpredictable”). In Stage 2 the “Predictable” and “Unpredictable” locations were both perfectly predictable on the basis of the preceding cue. Error bars indicate the within-subjects standard error of the mean (Cousineau, 2005).

Figure 4. Mean response latencies in Stages 1 and 2 of Experiment 2. Data are shown separately for trials in which the target appeared at a location that was predictable (“Pred”) during Stage 1, versus trials in which the target appeared at a location that was unpredictable during Stage 1 (“UnPred”), and for trials in which a valid cue preceded the search array (cues A-D in Stage 1, cues F-I in Stage 2), versus trials in which an invalid cue preceded the search array. Data are shown for trials in the 1st and 2nd half of each training stage. Error bars indicate the within-subjects standard error of the mean (Cousineau, 2005).

Figures

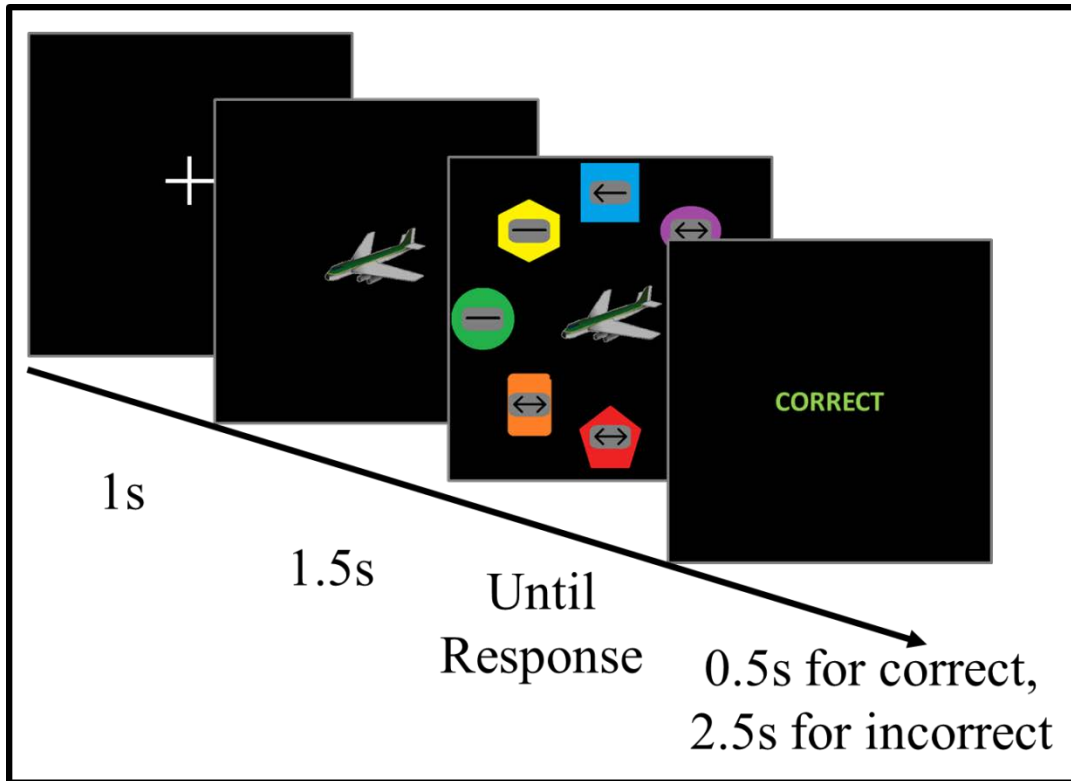


Figure 1.


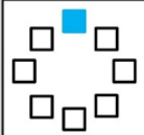

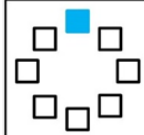

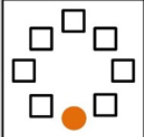

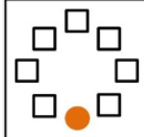

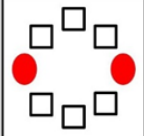

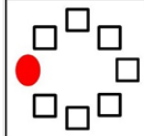

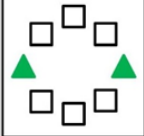

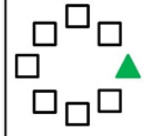
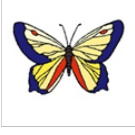
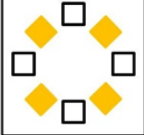
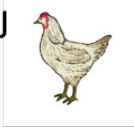
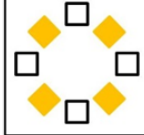
Stage 1			Stage 2		
Cue	Location	Probability	Cue	Location	Probability
A 		100%	F 		100%
B 		100%	G 		100%
C 		50%	H 		100%
D 		50%	I 		100%
E 		25%	J 		25%

Figure 2.

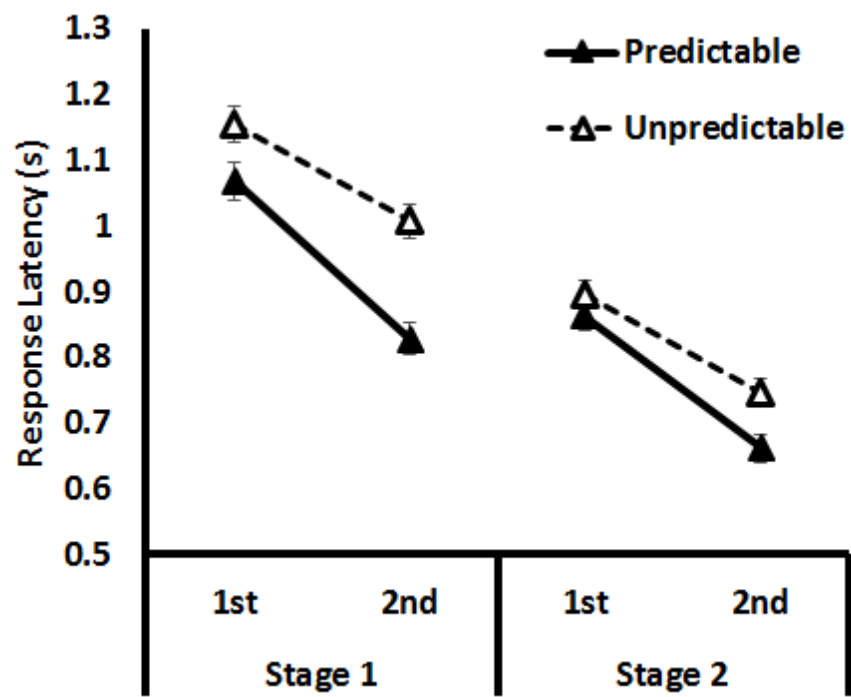


Figure 3.

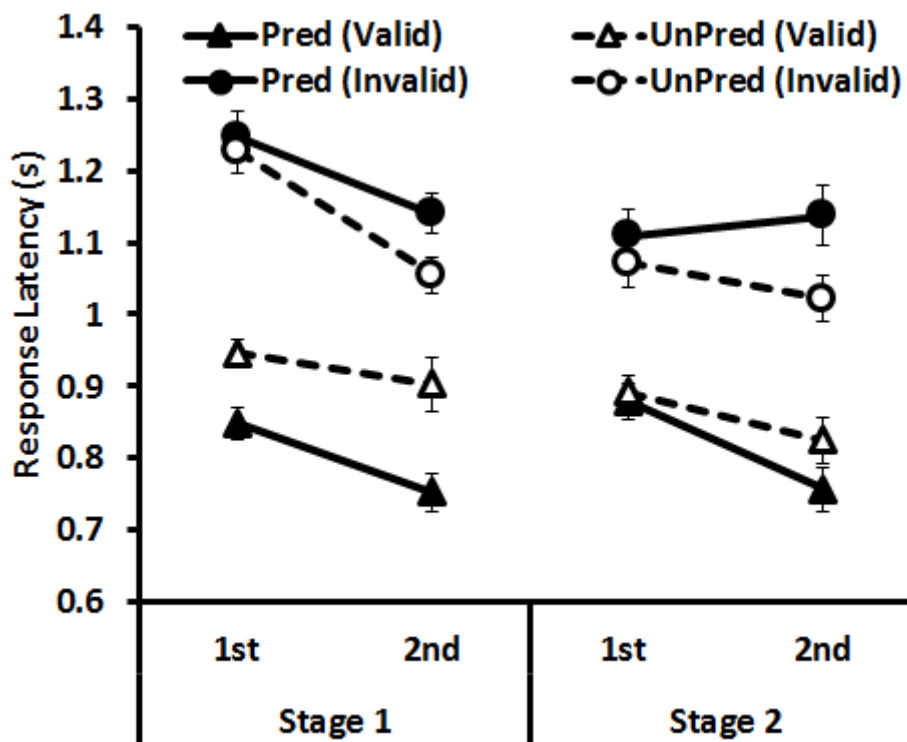


Figure 4.