2017-07-11

The redundancy effect in human causal learning: no evidence for changes in selective attention

Jones, Peter

<http://hdl.handle.net/10026.1/9616>

10.1080/17470218.2017.1350868

Quarterly Journal of Experimental Psychology

SAGE Publications

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

This is an Accepted Manuscript of an article published by Taylor & Francis in The Quarterly Journal of Experimental Psychology on 11/07/2017, available online:
<http://www.tandfonline.com/10.1080/17470218.2017.1350868>.

The redundancy effect in human causal learning:
no evidence for changes in selective attention

Peter M. Jones and Tara Zaksaitė
Plymouth University

Running head: Redundancy effect

Correspondence concerning this article should be addressed to: Peter M. Jones, School of Psychology, Plymouth University, PL4 8AA, UK. E-mail: peter.m.jones@plymouth.ac.uk

Keywords: redundancy effect, attention, learning, blocking

Abstract

Several recent papers (e.g. Uengoer, Lotz, & Pearce, 2013) have reported a difference in associative learning for two kinds of redundant cues, such that a blocked cue (e.g. X in A+ AX+) apparently forms a stronger association with the outcome than an uncorrelated cue (e.g. Y in BY+ CY-). This difference is referred to as the redundancy effect, and is of interest because it is contrary to the predictions of a number of popular learning models. One way of reconciling these models with the redundancy effect is to assume that the amount of attention paid to redundant cues changes as a result of experience, and that these changes in attention influence subsequent learning. Here we present two experiments designed to evaluate this idea, in which we measured overt attention using an eye tracker while participants completed a learning task that elicited the redundancy effect. In both experiments gaze duration was longer for uncorrelated cues than for blocked cues, but this difference disappeared when we divided gaze durations by trial durations. In Experiment 2, we failed to observe any difference in gaze duration when blocked and uncorrelated cues were subsequently presented together. While the observed difference in gaze duration for the two types of redundant cue may contribute to differences in learning during initial training, we suggest that the principal causes of the redundancy effect are likely to lie elsewhere.

Introduction

Numerous experimental tasks require participants to learn about the relationships between cues and outcomes, and to make predictions about which outcomes will occur when a particular cue or combination of cues is presented. When more than one cue is available, some may be informative about the outcome, and others may be redundant. Participants are typically able to detect this difference in informational value, and learn little about redundant cues. Not all redundant cues are equal, however, and a recent report has highlighted a puzzling difference in how much is learned about two kinds of redundant cues.

Uengoer, Lotz, and Pearce (2013) trained participants with a causal learning task, in which they were required to play the role of a doctor and to diagnose the causes of an allergic reaction in a fictional patient. On each trial, one or two foods (cues, represented here by letters) were presented and the participant had to make a prediction about whether or not the patient would suffer from a stomach ache (the outcome, denoted +; its absence is denoted -). Four types of trial were presented: A+, AX+, BY+, and CY-. Considering only the first two types of trial, we can think of A as an informative cue and X as a redundant cue. Participants needed to learn about A in order to predict the outcome on A+ trials, with the result that A alone was sufficient to predict the outcome on AX+ trials. X therefore provided no additional information about the outcome, and was redundant. This type of redundant cue can be referred to as a blocked cue, because learning about the relationship between X and the outcome is likely to have been ‘blocked’ by A (cf. Kamin, 1969). Considering the latter two trial types, B and C provided reliable information about the presence and the absence of the outcome respectively, and Y was redundant. We refer to Y as an uncorrelated cue, because its occurrence was uncorrelated with the outcome. After a period of training with these four types of trial, Uengoer et al. presented each cue individually and asked participants to rate the likelihood of the outcome occurring using a response scale, with low ratings reflecting a low

perceived probability of the outcome and high ratings reflecting a high probability.

Unsurprisingly, participants rated the likelihood of the outcome occurring as being high for A and B, and low for C. Of particular interest, however, were the ratings for the two redundant cues, X and Y. Ratings for X were higher than for Y, suggesting that participants thought the outcome was more likely given X than Y. This difference between redundant cues is referred to as the redundancy effect (Jones & Pearce, 2015; see also Pearce, Dopson, Haselgrove, & Esber, 2012).

This result is particularly interesting because it is contrary to the predictions of several prominent models of associative learning. These models conceive of learning in situations like that described above as the formation or modification of links between the cues and the outcome. Perhaps the most well-known of these is the model described by Rescorla and Wagner (1972), which assumes that presentation of more than one cue at a time will result in those cues competing to become associated with the outcome, via a common error correction mechanism. Changes in the strength of an association between a cue, S , and an outcome (ΔV_s) on each trial are determined by the following equation:

$$\Delta V_s = \alpha \cdot \beta \cdot (\lambda - \Sigma V)$$

In this equation, α and β are learning rate parameters relating to the cue and the outcome respectively, and represent the salience of those stimuli. The magnitude of the outcome is represented by λ , and ΣV is the total strength of the associations between each cue and the outcome. This model predicts that, as training of the type given by Uengoer et al. (2013) progresses, V_A tends towards λ . There is no competition between cues on A+ trials because A is the only cue present. On AX+ trials, however, learning about both A and X is determined competitively, in proportion to the difference between λ and the combined values of V_A and V_X . Since V_A is high, little learning about X will occur. At asymptote, V_X is predicted to be

close to zero. The prediction regarding Y is more complex, but the eventual result is that Y will become weakly associated with the outcome. Initial increases in the value of V_Y on BY+ trials will lead to an expectation of the outcome on CY- trials, with consequent decreases in both V_C and V_Y . At asymptote, V_Y will be weakly positive and offset on CY- trials by the negative value of V_C . This model therefore predicts that Y should predict the outcome more strongly than X, which is the opposite result to that obtained by Uengoer et al. Similar problems occur for a range of other models (Esber & Haselgrove, 2011; Le Pelley, 2004; Pearce, 1987, 1994; Pearce & Hall, 1980; Pearce & Mackintosh, 2010; for a fuller exploration of the predictions of these models, see Pearce et al., 2012).

Uengoer et al. (2013) suggest that one way their results might be accommodated is by assuming that the amount of attention paid to each cue varies with experience. Specifically, they suggest that participants might have learned to pay less attention to the redundant cues, X and Y. If participants learned to ignore the uncorrelated cue more quickly than the blocked cue, then they might be expected to have more opportunity to form an association between the blocked cue and the outcome than for the uncorrelated cue. The idea that attention to cues changes with experience is common to a large number of models of learning (e.g. Kruschke, 2001; Le Pelley, 2004; Mackintosh, 1975; Pearce & Hall, 1980; Pearce & Mackintosh, 2010; Sutherland & Mackintosh, 1971), and is consistent with previous evidence. Decrements in the rate of learning about cues have been observed when they have a history of serving as either blocked cues (e.g. Beesley & Le Pelley, 2011) or uncorrelated cues (e.g. Uengoer & Lachnit, 2012). In the absence of any direct comparison of the magnitude of these decrements, however, we have no empirical reason to suggest that one of these two types of cue should be paid more attention than the other. We also have little theoretical reason to assume that the rate at which participants might have learned to ignore the blocked and uncorrelated cues in Uengoer et al.'s experiments should have differed. The theories mentioned above can all

readily predict decrements in attention for these cues, but are largely agnostic on the issue of which will suffer a greater or more rapid decline. The aim of the present work, then, was to provide a first comparison of the attention paid to blocked and uncorrelated cues in a redundancy effect experiment.

Overt attention was measured in the present experiments by monitoring eye gaze while participants performed a redundancy effect task modelled on Uengoer et al. (2013). This eye-tracking method has been used successfully to demonstrate changes in attention in a variety of learning experiments (e.g. Beesley & Le Pelley, 2011; Beesley, Nguyen, Pearson, & Le Pelley, 2015; Hogarth, Dickinson, Austin, Brown, & Duka, 2008; Kruschke, Kappenman, & Hetrick, 2005; Wills, Lavric, Croft, & Hodgson, 2007).

(Table 1 about here)

Experiment 1

The design of Experiment 1 is shown in Table 1. The task was modelled closely on Experiment 2 of Uengoer et al. (2013), which used a variant of the popular allergist task (e.g. Aitken, Larkin, & Dickinson, 2000). On each Stage-1 trial, participants saw either one or two pictures of foods and were asked to predict whether a fictional patient would suffer an allergic reaction after consuming them. Following their choice, they were told whether or not the reaction occurred. Four types of trial were presented: A+ trials on which a single cue was followed by the allergic reaction, AX+ and BY+ trials on which pairs of cues were followed by the reaction, and CY- trials on which a pair of cues was followed by the absence of the reaction. Following Stage-1 training, a test was administered in which participants were shown the individual cues and asked to rate the likelihood of the reaction occurring for each one. We expected to replicate the redundancy effect observed by Uengoer et al. That is, we expected participants to rate the probability of the reaction occurring as higher for the blocked cue X than for the uncorrelated cue Y. The principal aim of the experiment, however, was to compare the attention paid to X and Y during learning. For this purpose, we used an eye tracker to measure the amount of time participants spent looking at each cue during Stage 1. If changes in attention are responsible for the redundancy effect, we should be able to detect a difference in the amount of time participants spend looking at X and Y.

Method

Participants

Fifty-one participants took part in this experiment. Twenty-three participants were undergraduate students taking part in the experiment for course credit, and the remainder were members of the public who received a small payment in return for their participation. The eye-tracker could not be adequately calibrated for five of the participants. As a result, their data were excluded from the experiment, reducing the participant sample to 46 (34 female). The remaining participants had a mean age of 29.52 ($SD = 13.20$).

Apparatus

The experiment was presented using a desktop computer, connected to a 22-inch widescreen monitor with a 1280 x 1024 screen resolution. The eye-tracking data were collected using the SMI RED remote eye-tracker (SensoMotoric Instruments, Teltow, Germany), located immediately below the monitor. The eye tracker sampled the location of participants' gaze at a rate of 50 Hz. Participants sat approximately 50 cm from the monitor, and were instructed to maintain this position as long as it was comfortable. The experiment was designed and administered using E-prime 2.0 software (Psychology Software Tools, PA, US).

Each cue was a square picture of a food on a white background, with a width of 300 pixels. Foods were randomly selected for each participant from the following list: apple, banana, broccoli, cabbage, cherries, corn, grapes, orange, pepper, pumpkin, and strawberry. The foods were then randomly assigned to serve as A, B, C, X, and Y. The outcomes were stomach ache, signified by text and a sad face on a red background, and no stomach ache, indicated by text and a happy face on a green background. The stimuli and outcomes were presented on a

black background using white text. Participants responded using the mouse, by clicking on either binary response buttons or a rating scale.

Procedure

Upon their arrival, the participants completed the 9-point calibration procedure for eye-tracking. They were then asked to read the following on-screen instructions, adapted from Uengoer et al. (2013):

This study is concerned with the question of how people learn about relationships between different events. In the present case, you should learn whether the consumption of certain foods leads to stomach ache or not.

Imagine that you are a medical doctor. One of your patients often suffers from stomach ache after meals. To discover the foods the patient reacts to, your patient eats specific foods and observes whether stomach ache occurs or not.

The results of these tests are shown to you on the screen one after the other. You will always be told what your patient has eaten. Sometimes he has only consumed a single kind of food, and other times he has consumed two different foods. Please look at the foods carefully.

Thereafter you will be asked to predict whether the patient suffers from stomach ache. For this prediction, please click on the appropriate response button. After you have made your prediction, you will be informed whether your patient actually suffered from stomach ache.

Use this feedback to find out what causes the stomach ache your patient is suffering from. Obviously at first you will have to guess because you do not know anything

about your patient, but eventually you will learn which foods lead to stomach ache in this patient and you will be able to make correct predictions.

For all of your answers, accuracy rather than speed is essential. Please do not take any notes during the experiment.

If you have any questions, please ask them now. If you do not have any questions, please start the experiment by clicking the mouse.

In the first stage of the experiment the participants were presented with sixteen blocks of trials, with each of the different trial types (A+, AX+, BX+, CY-) occurring once per block. The order of the trials within each block was random, except for the constraint that the first trial of a block could not be of the same type as the last trial of the preceding block. Each trial started with presentation of either one or two images of foods in the top half of the screen, below the phrase “The patient ate the following food(s)”. For trials with two images, one was located on the left and one on the right. The left-right allocation of positions for pairs of images was balanced, with each of the two possible arrangements occurring an equal number of times. In order to standardise the locations at which pictures were presented for trials containing one image, this appeared either on the left or on the right side of the screen, with half of the images presented on the left and half on the right. The sentence “Which reaction do you expect?” was presented below the images. Participants responded by clicking one of two response buttons placed at the bottom of the screen. The left-hand button was labelled “No stomach ache”, and the right-hand button was labelled “Stomach ache”. As soon as the participant responded, a screen was displayed that showed the relevant cues together with the outcome of the trial. When the outcome was stomach ache, the statement “The patient has stomach ache” and the picture of the sad facial expression was shown. When the outcome

was no stomach ache, the statement “The patient has no stomach ache” and the picture of the happy expression was shown. This feedback display remained on the screen for 3 s, followed by a 500 ms blank screen after which the next trial began.

After all of the blocks in Stage 1 were completed, participants were shown the following instructions:

Now, your task is to judge the probability with which specific foods cause stomach ache in your patient. For this purpose, single foods will be shown to you on the screen.

In this part, you will receive no feedback about the actual reaction of the patient. Use all the information that you have collected up to this time.

The test stage then began. On each trial, the sentence “What is the probability that the food causes stomach ache?” was shown above a single food image. Participants responded by clicking on an 11-point rating scale ranging from 0 (*certainly not*) to 10 (*very certain*). The rating scale was located in the lower half of the screen, oriented horizontally. After each response, a blank screen was shown for 500 ms and was followed by the next trial. Each of A, B, C, X and Y was presented twice, with the order of trials randomly determined for each participant. For each cue, the average of the two ratings was used in subsequent data analysis.

Eye-tracking data were recorded for each training trial from stimulus onset to the time at which the participant made a response. Regions of interest were defined along the edges of each cue, such that each region of interest formed a square with a width of 300 pixels. For each trial, we calculated the total amount of time during which the participant’s gaze was located in the region of interest corresponding to each cue (referred to below as gaze duration).

Statistical analysis

Two features of our statistical analysis of the results merit comment. Firstly, estimates of effect size for each Analysis of Variance (ANOVA) are given as partial eta squared, and estimates of effect size for paired t -tests are given as Cohen's d_{av} (as recommended by Lakens, 2013). Secondly, Bayesian t -tests were used to evaluate the strength of support for the null hypothesis where appropriate. The guidelines proposed by Rouder, Speckman, Sun, Morey, and Iverson (2009) were adhered to by using the JZS prior and a scaling factor of 1. Each of these tests yielded a Bayes factor (B_{01}), the value of which indicates the level of support for the null and alternative hypotheses. Values higher than 3 can be regarded as support for the null hypothesis, whereas values lower than 1/3 can be regarded as support for the alternative hypothesis. Although a scaling factor of 1 was used, outcomes were not dependent on this choice. This was confirmed by additional, unreported analyses with a scaling factor of 0.5. In no case did these analyses differ substantively from those reported here.

Results

(Figure 1 about here)

The left-hand panel of Figure 1 shows the proportion of trials on which participants predicted stomach ache during Stage 1, divided into epochs that each contained an average response from two trials of each type. Participants learned this discrimination very well, making errors on fewer than 1% of trials during the final epoch. The right-hand panel of Figure 1 shows the mean ratings for each cue (A, B, C, X, Y) during the test stage. These results closely resemble those reported by Uengoer et al. (2013). A one-way ANOVA confirmed that ratings differed among cues, $F(4, 180) = 195.75, p < .001, \eta_p^2 = .813$. All pairwise differences between cue ratings were significant, smallest $t(45) = 3.59, p = .001, d_{av} = 0.69$. Importantly, the redundancy effect was observed; ratings for X were higher than those for Y, $t(45) = 5.02, p < .001, d_{av} = 0.84$.

(Figure 2 about here)

Figure 2 shows eye-tracking data for Stage 1. Mean gaze duration for A during A+ trials is shown in the top-left panel; a one-way ANOVA demonstrated a significant effect of epoch, $F(7, 315) = 6.43, p < .001, \eta_p^2 = .125$. The top-right panel of Figure 2 shows mean gaze duration for A and X during AX+ trials, with longer durations for A than for X. A two-way ANOVA using epoch and cue variables demonstrated that this difference between cues was significant, $F(1, 45) = 9.71, p = .003, \eta_p^2 = .177$. The analysis also revealed a significant effect of epoch, $F(7, 315) = 11.93, p < .001, \eta_p^2 = .210$, and no interaction between the two variables, $F < 1$. Similarly, as shown in the middle-left panel, participants spent more time looking at B than Y during BY+ trials. A two-way ANOVA revealed a significant effect of epoch, $F(7, 315) = 5.82, p < .001, \eta_p^2 = .115$, a significant effect of cue, $F(1, 45) = 8.55, p = .005, \eta_p^2 = .160$, and no interaction, $F(7, 315) = 1.17, p = .320, \eta_p^2 = .025$. The middle-right

panel shows gaze durations for C and Y during CY- trials. In this case, participants spent equivalent amounts of time looking at the two types of cue. A two-way ANOVA demonstrated a significant effect of epoch, $F(7, 315) = 11.34, p < .001, \eta_p^2 = .201$, but no effect of cue, $F < 1$, and no interaction between the variables, $F(7, 315) = 1.39, p = .210, \eta_p^2 = .030$.

The data of primary interest are those shown in the lower panels of Figure 2, which provide a comparison of the blocked and uncorrelated cues. As shown in the bottom-left panel, mean gaze duration for Y (averaged across BY+ and CY- trials) was longer than for X (on AX+ trials). A two-way ANOVA using stimulus and epoch variables confirmed this difference between cues, $F(1, 45) = 8.91, p = .005, \eta_p^2 = .165$, and found an effect of epoch, $F(7, 315) = 12.97, p < .001, \eta_p^2 = .224$, but no interaction between these variables, $F < 1$.

Uengoer et al. (2013) suggested that participants might learn more about the blocked cue than the uncorrelated cue because they pay more attention to the blocked cue, but we observed the opposite result. It is worth noting that Uengoer et al.'s suggestion depends on a conception of learning in this kind of task as the acquisition of an association between each cue and the stomach ache outcome. If participants' expectation of this outcome is low at the beginning of the task, higher ratings presumably indicate superior learning. However, it is not clear that this interpretation is correct. The rating scale used in the present experiments and by Uengoer et al. conflates associative strength with certainty. In other words, moderate ratings might reflect either an association of moderate strength or simply uncertainty about whether the cue and outcome are associated. It is possible that participants' initial expectation in this kind of experiment would be that each cue may or may not be associated with the outcome. In this case, the moderate ratings assigned to the blocked cue may in fact be an indication of poor learning, whereas the lower ratings given to the uncorrelated cue may indicate that participants have learned that it is not a cause of the outcome. The question of interest then

becomes one of why participants learn more about uncorrelated cues than blocked cues, rather than the opposite. The findings reported here provide a possible explanation. Since participants looked more at Y than at X, they presumably had more opportunity to learn that Y was not predictive of stomach ache.

It should be emphasised that any difference in gaze duration for X and Y may be a consequence of the manner in which they were presented. This is because X and Y were presented on different trials, alongside different cues. We must consider, for instance, whether the duration of the trials on which X and Y were presented was the same. If trials featuring Y were longer than those featuring X, this would provide a ready explanation for the difference in gaze duration between these cues. Trial durations in this experiment were determined by the participant, as the cues remained on screen until the participant had clicked on one of the two response buttons. The mean reaction time (RT) was marginally longer for trials containing Y (Mean, $M = 2.09s$, standard error of the mean, $SEM = 0.09s$) than for trials containing X ($M = 1.93s$, $SEM = 0.12s$); This difference did not reach significance, $t(45) = 1.88$, $p = .066$, $d_{av} = 0.22$. Since no significant difference in RT for trials containing X and Y was observed, it is tempting to conclude that the significant difference in eye gaze we observed was independent of RT. We think this conclusion would be premature, however. Although the difference in RT for trials containing X and Y was not statistically significant, it was large enough to warrant further examination. We sought to find out whether the difference in eye gaze we observed would persist if we controlled for RT. To do this, we expressed gaze duration on each trial as a proportion of the RT; these data are shown in the bottom-right panel of Figure 2. No difference in gaze duration for X and Y is apparent when the data are expressed in this way, and statistical analyses confirmed this result. A two-way ANOVA with cue and epoch variables revealed an effect of epoch, $F(7, 315) = 2.51$, $p = .016$, $\eta_p^2 = .053$, but no effect of cue and no interaction, $F_s < 1$. We also compared mean

gaze as a proportion of RT for X and Y across all Stage 1 trials, using a Bayesian t -test. This analysis revealed strong support for the null hypothesis, $B_{01} = 8.40$. As a result, we conclude that the difference in gaze duration for X and Y is likely to have been primarily a consequence of a difference in RT, even though that difference in RT did not itself reach statistical significance.

Finally, we tested for a correlation between gaze bias and the redundancy effect. If the difference in gaze for X and Y contributes to the redundancy effect, then we might expect individuals with a large gaze bias to give higher ratings for X than for Y. To obtain single scores for both gaze bias and the redundancy effect, we simply calculated difference scores for each participant. Positive difference scores for gaze bias indicated longer dwell times for Y than for X, while positive redundancy effect scores indicated higher ratings for X than for Y. We found no correlation between these variables, Pearson $r = -.249$, $p = .095$. A similar analysis was conducted using gaze durations that were corrected for RT as above, which also revealed no correlation between gaze bias and redundancy effect score, $r = -.199$, $p = .185$. This suggests that the influence of gaze bias on the redundancy effect in Experiment 1 was, at best, weak.

Experiment 1 was designed to test the idea that the amount of attention paid to blocked and uncorrelated cues might differ, following changes in attention of the kind described by various models of learning (e.g. Kruschke, 2001; Le Pelley, 2004; Mackintosh, 1975; Pearce & Hall, 1980; Pearce & Mackintosh, 2010; Sutherland & Mackintosh, 1971). While we observed a difference in gaze duration for X and Y, the results of Experiment 1 appear to be subtly different from the predictions of those models. Each model predicts that the properties of the cues that elicit attention (which, for the sake of clarity, we will refer to as salience) can change as a result of experience. By contrast, no changes in the salience of X and Y are needed to explain the difference in attention observed in Experiment 1, because that

difference may have been simply a consequence of the circumstances in which they were presented. When we corrected for RT, the difference in gaze duration for X and Y disappeared. However, it is possible that dividing gaze duration by RT was unnecessarily conservative. Some smaller changes in salience may have occurred that were obscured by this transformation. A more sensitive alternative might be to eliminate RT as a confounding variable by presenting blocked and uncorrelated cues simultaneously during a subsequent training stage, thus ensuring that they appear for the same length of time regardless of RT. This was one of the primary aims of Experiment 2, which contained an additional training stage following the test stage which utilised compounds consisting of one previously blocked cue and one previously uncorrelated cue. Blocked and uncorrelated cues were therefore compared directly on these trials. If the salience of blocked cues differs from that of uncorrelated cues as result of their initial training, then a difference in gaze might be apparent when they are presented together. This additional training stage also provided an opportunity to assess the ease with which novel learning about these cues took place. It is a common assumption that changes in salience are accompanied by consequent changes in the speed of learning (e.g. Esber & Haselgrove, 2011; Le Pelley, 2004; Mackintosh, 1975; Pearce & Hall, 1980; Pearce & Mackintosh, 2010). Indeed, differences in learning rate are frequently used as an indirect measure of salience in both humans and non-human animals (for a review, see Le Pelley, 2004), and the present investigation into the attention paid to blocked and uncorrelated cues was motivated by differences in learning. We were therefore interested to see whether participants' ability to use previously blocked or uncorrelated cues as discriminative stimuli would differ.

Experiment 2

Experiment 2 aimed to replicate the difference in gaze duration observed in Experiment 1, and to test whether any difference in either gaze or learning would be observed during a discrimination involving compounds of previously blocked and uncorrelated cues. The design of Experiment 2 is shown in Table 1. Stage 1 consisted of a similar discrimination to that used in Experiment 1, except that additional trials were included so that there were two blocked cues (W and X) and two uncorrelated cues (Y and Z). Participants' eye gaze during this stage was monitored; Following Experiment 1, we expected gaze duration to be longer for Y/Z than for W/X, but for this difference to be eliminated when a correction for RT was applied. Following this training stage, a test stage was administered during which participants rated the likelihood of the outcome occurring for each cue independently. In a similar manner to Uengoer et al. (2013) and Experiment 1 here, we expected ratings to be higher for W/X than for Y/Z. Finally, participants were trained with a WY-/WZ+/XY+ discrimination. The main purpose of this training was to present blocked and uncorrelated cues together, allowing us to compare eye gaze for these cues without the confounding variable of RT. However, this discrimination also provided an opportunity to compare the ease with which learning occurs for cues with differing associative histories (see Rescorla, 2000, 2002; Pearce, Esber, George, & Haselgrove, 2008). The discrimination can be thought of as consisting of two component sub-discriminations: WY-/WZ+ and WY-/XY+. In order for participants to solve the first of these, they must use the previously uncorrelated cues, Y and Z, as discriminative stimuli. Likewise, the solution of the second sub-discrimination depends on the use of the previously blocked cues, W and X, as discriminative stimuli. If there is any difference in the salience of these two sets of cues, and this difference affects the rate of novel learning about them, then the two sub-discriminations might be solved at differing rates. If participants learn to predict the outcome on WZ+ trials more readily than on XY+ trials, we may infer that the salience of

the previously uncorrelated cues was higher than that of the previously blocked cues.

Conversely, higher salience for the previously blocked cues would be indicated by faster learning about XY than WZ.

Method

Participants

Sixty participants were recruited, all of whom were undergraduate students taking part in exchange for course credit. The eye tracker could not be calibrated for one participant. The remaining 59 participants (56 female) had a mean age of 21.51 (SD = 6.27)..

Stimuli and apparatus

Eighteen food cues were used. These were: apple, aubergine, banana, broccoli, cabbage, cherries, coconut, corn, grapes, kiwi, lemon, orange, pear, pepper, pumpkin, strawberry, tomato, and watermelon. These were randomly assigned for each participant to serve as A, B, C, D, E, F, G, H, I, J, K, L, M, N, W, X, Y, and Z. All other details were the same as for Experiment 1.

Procedure

The procedure for calibrating the eye tracker and the instructions given to participants at the start of the experiment were the same as for Experiment 1. The first stage of the experiment consisted of 16 blocks of trials, with each block containing one of each of the following trial types: A+, AW+, B+, BX+, CY+, DY-, EZ+, FZ-. The test stage then followed, during which each individual cue was rated twice. Finally, participants completed a test discrimination. The instructions for this stage were:

You will now see a further series of meals, each made up of two foods. You will be asked to make predictions about whether stomach ache will occur, and you will be given feedback as before.

There were twelve blocks of trials, each containing one of each of WY-, WZ+, and XY+. We were concerned that this discrimination would be solved too quickly for any difference in the rate of learning about the different components of the discrimination to be observed; to combat this, each trial block also contained one of each of GH-, IJ-, IH+, KL-, MN-, and KN+. These trials were included only to reduce the speed at which the test discrimination was solved, and no data are reported here. Other details of the test discrimination were the same as for Stage 1. Experimental details that are omitted here were the same as for Experiment 1.

Results

(Figure 3 about here)

For ease of exposition, the results reported here for Stage 1 and the test stage are averages of equivalent cues and compounds. For example, W and X are treated as equivalent blocked cues, and Y and Z as equivalent uncorrelated cues. The proportions of Stage-1 trials on which participants predicted stomach ache are shown in the left-hand panel of Figure 3, divided into epochs in the same way as for Experiment 1. This discrimination was acquired readily, with fewer than 5% of trials resulting in an error during the final epoch. Ratings from the test stage are shown in the right-hand panel of Figure 3; these are consistent with the results of Experiment 1. A one-way ANOVA confirmed that these ratings differed among cue types, $F(4, 232) = 151.60, p < .001, \eta_p^2 = .723$. All pairwise comparisons of cue types yielded significant differences, smallest $t(58) = 5.93, p < .001, d_{av} = .79$. Crucially, ratings for blocked cues were higher than for uncorrelated cues, $t(58) = 7.00, p < .001, d_{av} = .94$.

(Figure 4 about here)

Eye gaze data are shown in Figure 4, in a similar manner to Experiment 1. Mean gaze durations for A and B during A+ and B+ trials are shown in the top-left panel; a one-way ANOVA revealed a significant effect of epoch, $F(7, 406) = 5.52, p < .001, \eta_p^2 = .087$. The upper-right panel shows mean gaze durations for each type of cue during AW+ and BX+ trials, with durations being longer for blocking cues (A and B) than for blocked cues (W and X). A two-way ANOVA revealed an effect of epoch, $F(7, 406) = 17.20, p < .001, \eta_p^2 = .229$, a difference between the two cue types, $F(1, 58) = 7.55, p = .008, \eta_p^2 = .115$, and no interaction, $F(7, 406) = 1.86, p = .076, \eta_p^2 = .031$. Gaze durations for CY+ and EZ+ trials are shown in the middle-left panel, with longer durations evident for predictive cues (C and E) than for uncorrelated cues (Y and Z) during some epochs. A two-way ANOVA revealed a

significant effect of epoch, $F(7, 406) = 11.92, p < .001, \eta_p^2 = .170$, a significant difference between cue types, $F(1, 58) = 4.21, p = .045, \eta_p^2 = .068$, and a significant interaction between the two variables, $F(7, 406) = 2.04, p = .049, \eta_p^2 = .034$. Simple effects analysis of this interaction indicated that the difference between cue types was significant for the third epoch, $F(1, 58) = 6.36, p = .014, \eta_p^2 = .099$, and the fifth epoch, $F(1, 58) = 8.52, p = .005, \eta_p^2 = .128$, and approached significance on the seventh epoch, $F(1, 58) = 3.53, p = .065, \eta_p^2 = .057$. The difference between cue types did not approach significance for any of the five remaining epochs, largest $F(1, 58) = 1.67, p = .202, \eta_p^2 = .028$. Finally, mean gaze durations for DY- and FZ- trials are shown in the middle-right panel of Figure 4. Gaze was equivalent for predictive (D and F) and uncorrelated (Y and Z) cues. A two-way ANOVA revealed a significant effect of epoch, $F(7, 406) = 14.59, p < .001, \eta_p^2 = .201$, but no effect of cue type and no interaction between the two variables, $F_s < 1$.

Overall mean durations of eye gaze for blocked and uncorrelated cues are shown in the bottom-left panel of Figure 4. A two-way ANOVA confirmed that, as for Experiment 1, gaze durations were longer for uncorrelated cues than for blocked cues, $F(1, 58) = 5.76, p = .02, \eta_p^2 = .090$. There was also a significant effect of epoch, $F(7, 406) = 25.08, p < .001, \eta_p^2 = .302$, but no interaction between these variables, $F < 1$. In a similar manner to Experiment 1, reaction times (RTs) on trials containing blocked cues ($M = 1.81s, SEM = .11s$) and uncorrelated cues ($M = 1.99s, SEM = .08s$) were compared. The difference between them was significant, $t(58) = 3.36, p = .001, d_{av} = .24$. Consequently, a correction for RT was applied by dividing gaze durations for each trial by the relevant RT; these data are shown in Figure 2. A two-way ANOVA revealed a significant effect of epoch, $F(7, 406) = 6.37, p < .001, \eta_p^2 = .099$, but no difference between scores for blocked and uncorrelated cues, and no interaction, $F_s < 1$. A Bayesian t -test comparing mean corrected gaze durations for blocked and uncorrelated cues demonstrated strong support for the null hypothesis, $B_{01} =$

7.39. As in Experiment 1, then, participants spent longer looking at uncorrelated cues than at blocked cues, but this is likely to have been a result of differences in RT.

(Figure 5 about here)

Eye gaze during Stage 2 was also monitored; the left-hand panel of Figure 5 shows mean gaze duration for blocked and uncorrelated cues during this final stage. These two types of cue were presented on the screen at the same time during WY-, WZ+, and XY+ trials. Mean gaze duration for blocked and uncorrelated cues was equivalent here. A two-way ANOVA revealed an overall effect of trial block, $F(5, 290) = 5.25, p < .001, \eta_p^2 = .083$, but no effect of cue type and no interaction, $F_s < 1$. Gaze duration for blocked and uncorrelated cues across the whole of Stage 2 was also compared using a Bayesian t -test, which supported the null hypothesis, $B_{01} = 6.02$. It is of course possible that this result reflects a lack of sensitivity in our eye-tracking measure. Although we observed a difference in gaze duration for blocked and uncorrelated cues during Stage 1, that difference was seemingly driven by a difference in RT. It is possible that our method was not able to detect differences in gaze for cues that were presented together. However, several differences between simultaneously-presented cues were evident during Stage 1 of both experiments. For instance, mean gaze durations were higher for blocking cues than for blocked cues in both experiments. It seems more likely, therefore, that this was not the reason why no difference between cue types was observed for Stage 2 of Experiment 2. Both experiments can therefore be summarised by saying that gaze durations were consistently longer for uncorrelated cues than for blocked cues, but that these differences were not present when we controlled for RT by either expressing gaze duration as a proportion of RT, or by presenting blocked and uncorrelated cues simultaneously. In the same manner as for Experiment 1, additional analyses tested for a correlation between gaze bias and the redundancy effect. No correlation between the two was found, either when using

uncorrected gaze durations, $r = -.166$, $p = .208$, or gaze durations that were corrected for RT, $r = -.057$, $p = .594$.

The subsequent associability of blocked and uncorrelated cues was compared, by testing how readily participants were able to use them as discriminative cues in Stage 2. The right-hand panel of Figure 5 shows the proportion of trials on which participants predicted stomach ache on WY-, WZ+, and XY+ trials. If the associability of the previously blocked and uncorrelated cues were different here, we might expect to see a difference in the rate of learning about WZ and XY. No such difference was observed. A two-way ANOVA comparing learning about these two compounds revealed a significant effect of epoch, $F(5, 290) = 29.24$, $p < .001$, $\eta_p^2 = .335$, but no difference between the two compounds, $F < 1$, and no interaction, $F(5, 290) = 1.16$, $p = .330$, $\eta_p^2 = .020$. A Bayesian t -test comparing mean ratings for WZ and XY during Stage 2 found strong evidence for the null hypothesis, $B_{01} = 9.72$.

General Discussion

In two experiments, participants' eye gaze was monitored during training that was designed to elicit the redundancy effect. If the redundancy effect is a consequence of differences in attention for blocked and uncorrelated cues, as Uengoer et al. (2013) suggest, then gaze duration for these cues might be expected to differ. Both experiments demonstrated the redundancy effect, and a difference in gaze duration such that participants looked at uncorrelated cues for longer than at blocked cues. We suggest that this difference in gaze duration was a consequence of the conditions in which the two kinds of cues were presented rather than any change in the salience of the cues, for two reasons. Firstly, in neither experiment did the difference in gaze duration for blocked and uncorrelated cues survive a correction for the duration of the trials. Secondly, when the two types of cue were presented together simultaneously in Experiment 2, gaze duration was equivalent for each.

This does not mean that the difference in gaze duration is trivial, or unrelated to the redundancy effect itself. It is not necessary to suppose that the salience of blocked and uncorrelated cues differs in order to propose a role for attention in producing the redundancy effect. If participants pay more attention to one type of cue than another because of the conditions of presentation, this itself may be sufficient to produce a difference in learning about them. Following our results, we might want to conclude that participants are better able to learn about uncorrelated cues than blocked cues because they spend more time looking at them. While this is plausible, this conclusion should be treated with caution. One likely consequence of such a relationship between the observed difference in eye gaze and the redundancy effect is that participants whose gaze shows the largest bias towards uncorrelated cues should also demonstrate the largest difference in learning. In other words, eye gaze bias and the difference in test ratings for blocked and uncorrelated cues should be correlated. In neither experiment was such a correlation evident. The simplest interpretation, therefore, is

that the observed difference in gaze duration does not contribute meaningfully to the redundancy effect.

As summarised in the Introduction, the redundancy effect poses a challenge for associative models of learning. Although we have focussed mainly on the model described by Rescorla and Wagner (1972), similar problems are evident for a host of other models (e.g. Esber & Haselgrove, 2011; Le Pelley, 2004; Pearce, 1987, 1994; Pearce & Hall, 1980; Pearce & Mackintosh, 2010), even if the details of the models are quite different. For instance, Uengoer et al. (2013) pointed out that the shortcomings of the Rescorla-Wagner model in predicting the redundancy effect stem in large part from its prediction that blocked cues will have approximately zero associative strength at asymptote. This prediction arises because of its use of a summed error term. Consequently, Uengoer et al. suggested that a model which does not incorporate this summed error term, and which therefore predicts weaker blocking, might be a better candidate. As an example they suggested Pearce's (1987, 1994) configural model, but noted that this model still makes the prediction that uncorrelated cues will become better associated with the outcome than blocked cues.

The authors of all previous demonstrations of the redundancy effect, whether in humans (Uengoer et al., 2013) or in non-human animals (Jones & Pearce, 2015; Pearce et al., 2012) have suggested that their results might be accounted for if we assume that the amount of attention paid to redundant cues changes as a result of experience. This suggestion has been made despite any strong theoretical reason to expect a difference in attention for blocked and uncorrelated cues. The experiments reported here represent a first attempt to find experimental evidence to support this idea, but the resultant data are largely inconsistent with this account. While further empirical work would be welcomed, we suggest that the explanation for the redundancy effect is likely to lie elsewhere.

Pearce et al. (2012) provided simulations of Rescorla and Wagner's (1972) model which show that it incorrectly predicts that the associative strength of uncorrelated cues will be higher than that of blocked cues. However, as Vogel and Wagner (2016) have highlighted, this prediction is dependent on each cue consisting of a non-overlapping set of elements. Vogel and Wagner considered the alternative possibility that each cue contains a set of common elements, *K*. The essence of this idea is that each cue has some properties in common with the others, and that these properties are learned about in the same way as any other aspect of the cues. Vogel and Wagner presented a series of simulations showing that, with certain parametric assumptions, the Rescorla-Wagner model is able to predict the redundancy effect in this way. It does this because the common element *K* gains a considerable amount of associative strength, which restricts the accumulation of associative strength more for the unique *Y* elements than for the unique *X* elements. Vogel and Wagner note that this modification also allows the Pearce (1987, 1994) configural model to make a similar prediction. Because this account depends on the acquisition of associative strength by *K*, it makes testable predictions about the consequences of changing either the salience or the associative history of *K*. Firstly, decreasing the salience of *K* should result in it gaining less associative strength and having less influence on learning about other cues. In the limiting case where the salience of *K* is zero, the model will revert to the prediction given here and by Uengoer et al. (2013). Secondly, the influence of *K* on learning about other cues should be modified by the addition of other trials including *K*. Vogel and Wagner provided simulations to demonstrate that additional trials on which *K* is followed by the outcome will result in a larger redundancy effect, whereas additional trials on which *K* is followed by the absence of the outcome will result in a smaller redundancy effect or the opposite result. Further experimental work should enable the evaluation of this account.

We have discussed the possibility that higher causal ratings for blocked cues than for uncorrelated cues might reflect superior learning about uncorrelated cues. An alternative is that participants learn adequately about both kinds of cue, but that higher causal ratings for blocked cues are a consequence of their ambiguous causal status. This ambiguity was highlighted by Lovibond, Been, Mitchell, Bouton, and Frohardt (2003), who noted that blocking effects in human causal learning are often modest in comparison with those seen in experiments with non-human animals. Lovibond et al. suggested that blocking effects are weak because there is insufficient information to deduce the causal status of blocked cues in many instances. In the present experiments, for instance, consumption of X may or may not cause stomach ache. This is because participants learn on A+ trials that consumption of A is by itself sufficient to cause stomach ache, and consequently the occurrence of the stomach ache on AX+ trials does not tell us whether or not X is also a cause. Lovibond et al. pointed out that this problem only occurs if the magnitude of the outcome is the same when it is preceded by either a single cause or two causes. Accordingly, they demonstrated that blocking could be enhanced if participants were trained to expect a larger outcome following two causes than just one, presumably because ambiguity about the causal status of the blocked cue was resolved. The redundancy effect, then, might be at least partly a consequence of this ambiguity. Future experiments could address this issue by providing information that allows participants to infer the causal status of the blocked cue.

In summary, the experiments reported here provide no support for an attentional account of the redundancy effect. This account was proposed as a result of the proposition that learning about blocked and uncorrelated cues progresses at different rates, but another possibility is that the redundancy effect occurs despite adequate learning about both. Further experiments are needed to distinguish these two possibilities.

References

- Aitken, M. R. F., Larkin, M. J. W., & Dickinson, A. (2000). Super-learning of causal judgements. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, 53, 59-81. doi: 10.1080/027249900392995
- Beesley, T., & Le Pelley, M. E. (2011). The influence of blocking on overt attention and associability in human learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 114-120. doi: 10.1037/a0019526
- Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *Quarterly Journal of Experimental Psychology*, 68, 2175-2199. doi: 10.1080/17470218.2015.1009919
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, 1, 2-5. doi: 10.20982/tqmp.01.1.p042
- 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 B*, 278, 2553-2561. doi: 10.1098/rspb.2011.0836
- Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: the role of uncertainty. *Quarterly Journal of Experimental Psychology*, 61, 1658-1668. doi: 10.1080/17470210701643439
- Jones, P. M., Pearce, J. M. (2015). The fate of redundant cues: Further analysis of the redundancy effect. *Learning and Behavior*, 43, 72-82. doi: 10.3758/s13420-014-0162-x

- Kamin, L. J. (1969). Selective attention and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42-64). Halifax, Nova Scotia: Dalhousie University Press.
- Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, 45, 812-863. doi: 10.1006/jmps.2000.1354
- Kruschke, J. K., Kappenman, E. S., Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 830-845. doi: 10.1037/0278-7393.31.5.830
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4. doi: 10.3389/fpsyg.2013.00863
- 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 B*, 57, 193-243. doi: 10.1080/02724990344000141
- Lovibond, P. E., Been, S. L., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory and Cognition*, 31, 133-142. doi: 10.3758/BF03196088
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298. doi: 10.1037/h0076778
- Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, 94, 61-73. doi: 10.1037/0033-295X.94.1.61

Pearce, J. M. (1994). Similarity and discrimination: a selective review and a connectionist model. *Psychological Review*, 101, 587-607. doi: 10.1037/0033-295X.101.4.587

Pearce, J. M., Dopson, J. C., Haselgrove, M., & Esber, G. R. (2012). The fate of redundant cues during blocking and a simple discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 167-179. doi: 10.1037/a0027662

Pearce, J. M., Esber, G. R., George, D. N., & Haselgrove, M. (2008). The nature of discrimination learning in pigeons. *Learning and Behavior*, 36, 188-199.
doi:10.3758/LB.36.3.188

Pearce, J. M., and Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, 87, 532-552. doi: 10.1037/0033-295X.87.6.532

Pearce, J. M., and Mackintosh, N. J. (2010). Two theories of attention: a review and a possible integration. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and Associative Learning: From Brain to Behaviour* (pp. 11-39). Oxford, UK: Oxford University Press.

Rescorla, R. A. (2000). Associative changes in excitors and inhibitors differ when they are conditioned in compound. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 428-438

Rescorla, R. A. (2002). Savings Tests: Separating differences in rate of learning from differences in initial levels. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 369-377.

Rescorla, R. A., and 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. 64-99). New York, NY: Appleton-Century-Crofts.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review*, *16*, 225-237. doi: 10.3758/PBR.16.2.225

Sutherland, N. S., and Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York, NY: Academic Press.

Uengoer, M., & Lachnit, H. (2012). Modulation of attention in discrimination learning: the roles of stimulus relevance and stimulus-outcome correlation. *Learning and Behavior*, *40*, 117-127. doi: 10.3758/s13420-011-0049-z

Uengoer, M., Lotz, A., & Pearce, J. M. (2013). The fate of redundant cues in human predictive learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*, 323-333. doi: 10.1037/a0034073

Vogel, E. H., & Wagner, A. R. (2016). A theoretical note on the interpretation of the “redundancy effect” in associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition*. Advance online publication. doi: 10.1037/xan0000123

Wills, A. J., Lavric, A., Croft, G. S., & Hodgson, T. L. (2007). Predictive learning, prediction errors, and attention: evidence from event-related potentials and eye tracking. *Journal of Cognitive Neuroscience*, *19*, 843-54.

Table 1

The Designs of the Experiments

Experiment	Stage 1	Test	Stage 2
1	A+ AX+ BY+ CY-	A B C X Y	
2	A+ AW+ B+ BX+ CY+ DY- EZ+ FZ-	A B C D E F W X Y Z	WY- WZ+ XY+

Figure captions

Figure 1. Response data from Experiment 1. The left panel shows the mean proportion of trials on which participants predicted stomach ache during Stage 1 for each trial type, and the right panel shows mean ratings for individual cues at test. Error bars for this and subsequent figures show the standard error of the mean, following an adjustment for within-subjects designs described by Cousineau (2005).

Figure 2. Eye gaze data for Experiment 1. Mean dwell times are shown for individual cues on A+ trials (top left), AX+ trials (top right), BY+ trials (middle left), and CY- trials (middle right). The bottom-left panel shows mean dwell times for X and Y, irrespective of trial type, and the bottom-right panel shows mean dwell times for X and Y as a proportion of reaction time (RT).

Figure 3. Response data from Stage 1 of Experiment 2. The left panel shows the mean proportion of trials on which participants predicted stomach ache during Stage 1 for each trial type, with equivalent trials combined. The right panel shows mean test ratings for each cue type, with equivalent cues combined.

Figure 4. Eye gaze data for Stage 1 of Experiment 2. Mean dwell times are shown for each cue type on A+ and B+ trials (top left), AW+ and BX+ trials (top right), CY+ and EZ+ trials (middle left), and DY- and FZ- trials (middle right). The bottom-left panel shows mean dwell times for W/X and Y/Z, irrespective of trial type, and the bottom-right panel shows mean dwell times for W/X and Y/Z as a proportion of reaction time (RT).

Figure 5. Response and gaze data for Stage 2 of Experiment 2. The left panel shows mean dwell times for W/X and Y/Z during Stage 2, and the right panel shows the mean proportion of trials on which participants predicted stomach ache during Stage 2, for each trial type.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5