2017

The Redundancy Effect in Human Causal Learning: Attention, Uncertainty, And Inhibition

Zaksaite, Gintare

Copyright statement

THE REDUNDANCY EFFECT IN HUMAN CAUSAL LEARNING:

ATTENTION, UNCERTAINTY, AND INHIBITION

By

GINTARE ZAKSAITE

A thesis submitted to Plymouth University in partial fulfilment for the degree of

DOCTOR OF PHILOSOPHY

School of Psychology

June 2017

Acknowledgements

First of all sincere thanks go to my Doctor of Studies, Dr Peter M. Jones. Thank you for offering me the Ph.D., for your guidance, time, honesty, and comments on the draft of this thesis. Thanks also go to the rest of my supervisory team: Prof Christopher Mitchell, Dr Hannah Drayson, and Dr Martha Blassnigg (1969-2015).

Thanks to Plymouth University for providing the funding and necessary facilities to carry out this research. Thanks also to Angelos Panagiotopoulos for collecting data for Experiment 4 and to Georgios Kavalieros, John Cahill, and Joshua Van Tonder for collecting data for Experiment 6.

I also wish to thank Prof Sue Denham for making CogNovo possible. Time spent as part of CogNovo has been invaluable and unforgettable. Never in my life have I met a group of such kind, caring, fun, funny, brilliant, and dare-I-say-it creative, group of individuals. Special thanks to Thomas, Vaibhav, Katie, Diego, Frank, Christos, and Jack for being kind when I needed it most. Thanks also to members of learning lunch and my friends and colleagues from the Psychology department.

Thanks also go to my grandmothers, Barbara, and Valerija (1935-2016). You may be as different as night and day, but I have learned lessons that made this journey easier from you both, and I am incredibly grateful.

This thesis is dedicated to my parents, Lina and Joe, for being the greatest supporters by valuing my happiness over achievements; it gives me courage every day. Thank you.

Author's Declaration

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award without prior agreement of the Graduate Sub-Committee.

Work submitted for this research degree at the Plymouth University has not formed part of any other degree either at Plymouth University or at another establishment.

This study was financed with the aid of a studentship from Plymouth University.

Relevant scientific seminars and conferences were regularly attended at which work was often presented; several papers prepared for publication.

Publications:

Zaksaite, T., & Jones, P. M. (2017). The Redundancy effect in human causal learning: Evidence against a Comparator Theory explanation. *Proceedings of the 38th Annual Conference of the Cognitive Science Society. Manuscript accepted.*

Jones, P. M., & Zaksaite, T. (2016). The redundancy effect in human causal learning: No evidence for changes in selective attention. *Manuscript submitted.*

Conference presentations:

Zaksaite, T., & Jones, P. M. (2017). *The redundancy effect is partly attributable to a lack of conditioned inhibition.* 21st Associative learning Symposium. Cardiff, Wales.

Zaksaite, T. & Jones, P. M. (2017). *Uncertainty, blocking and the redundancy effect.* EPS meeting, London, England.

Zaksaite, T., & Jones, P. M. (2016). *The redundancy effect and attention for irrelevant cues.* 20th Associative learning Symposium. Cardiff, Wales.

Zaksaite, T. (2015). *Models of associative learning: a brief introduction*. What's UP talk, Plymouth, England.

Word count of the main body of thesis: 36,530

Signed ……..

Date ……..

Abstract

The redundancy effect in human causal learning: Attention, uncertainty, and inhibition

Gintare Zaksaite

Using an allergist task, Uengoer, Lotz and Pearce (2013) found that in a design A+/AX+/BY+/CY-, the blocked cue X was indicated to cause the outcome to a greater extent than the uncorrelated cue Y. This finding has been termed "the redundancy effect" by Pearce and Jones (2015). According to Vogel and Wagner (2017), the redundancy effect "presents a serious challenge for those theories of conditioning that compute learning through a global error-term" (p. 119). One such theory is the Rescorla-Wagner (1972) model, which predicts the opposite result, that Y will have a stronger association with the outcome than X. This thesis explored the basis of the redundancy effect in human causal learning. Evidence from Chapter 2 suggested that the redundancy effect was unlikely to have been due to differences in attention between X and Y. Chapter 3 explored whether differences in participants' certainty about the causal status of X and of Y contributed to the redundancy effect. Manipulations aimed at disambiguating the effects that X had on the outcome, including outcome-additivity training and low outcome rate, resulted in lower ratings for this cue and a smaller redundancy effect. However, the redundancy effect was still significant with both manipulations, suggesting that while participants' uncertainty about the causal status of X contributed to it, there may have been other factors. Chapter 4 investigated whether another factor was a lack of inhibition for cue C. In a scenario where inhibition was more plausible than in an allergist task, a negative correlation between causal ratings for C and for Y, and a positive correlation between ratings for C and the magnitude of the redundancy effect, were found. In addition, establishing C as inhibitory resulted in a smaller redundancy effect than establishing C as neutral. Overall, findings of this thesis suggest that the redundancy effect in human causal learning is the result of participants' uncertainty about the causal status of X, and a lack of

inhibition for C. Further work is recommended to explore whether combining manipulations targeting X and Y would reverse the redundancy effect, whether effects of outcome additivity and outcome rate on X are the result of participants' uncertainty about this cue, and the extent to which participants rely on single versus summed error.

Table of contents

List of figures

List of tables

Chapter 1: Introduction

Causal learning describes the process of acquiring associations between stimuli and their consequences in our environment. These may be between causes and outcomes, events and actions, or actions and their effects. The capacity to acquire such associations in both human and non-human animals is essential for survival. It allows us to appropriately react to changes, anticipate and often manipulate events, predict consequences of our actions, and conserve resources. It is now over a century since the first systematic animal learning experiments were conducted by Thorndike (1898) and Pavlov (1927). However, research into this area remains relevant today. A key issue within causal learning research is what determines the conditions of learning.

Findings of the first associative learning experiments led to the development of some highly successful models of learning (e. g. Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972). Based on experimental findings and the development of theory, these models aimed to predict and explain when learning between potential causes, or cues, and their outcomes takes place and how it proceeds in animals. Following the observation that humans displayed blocking, a learning effect first established in rats by Kamin (1969), Dickinson, Shanks and Evenden (1984) suggested that animal models could be applied to humans as well; that both humans and animals shared a common learning mechanism.

Even though these models remain highly influential in learning research, there is emerging evidence suggesting that learning may be more complex than can be described by these, largely simple, models (e. g. Beckers, De Houwer, Pineno & Miller, 2005; Jones & Pearce, 2015; Livesey & Boakes, 2004; Lovibond, Been, Mitchell, Bouton & Frohardt, 2003; Mitchell & Lovibond, 2002; Mitchell, Lovibond & Condoleon, 2005; Pearce, Dopson, Haselgrove & Esber, 2012; Uengoer, Lotz & Pearce, 2013; Vandorpe & De Houwer, 2006). This thesis focused on the redundancy effect, one result that has posed a particular challenge to some of these models, and aimed to investigate its basis in human causal

learning. The redundancy effect is a finding comparing two types of a redundant cue: a blocked cue and an uncorrelated cue.

An example of a blocked cue is X in A+/AX+ training, in which A predicts the outcome and A and X together predict the same outcome. Kamin (1969) found that X elicited a weaker response in a group of rats which received A+/AX+ training than in a group for which A+ trials were omitted (AX+). In the light of these findings, Kamin proposed that learning about the relationship between X and the outcome was blocked in the A+/AX+ group because X was presented with A, a cue that predicted the outcome very well. Kamin therefore suggested that for redundant cues, learning about the relationship between a cue and an outcome will be blocked unless some form of surprise in the outcome occurrence is present.

Rescorla and Wagner (1972) turned Kamin's (1969) predictions into a formal model. They incorporated an aspect of a connectionist model of human learning first described by Widrow and Hoff (1960; see also Sutton & Barto, 1981). Rescorla and Wagner proposed that learning between a cue and an outcome occurs when the outcome is unexpected and ceases when the outcome is well predicted. On each learning instance, the level of outcome expectation is compared with whether it actually happens; this comparison is referred to as predictive error. If the outcome is not expected and it occurs, the predictive error is large. If the outcome is expected and it occurs, the predictive error is low. Importantly, if several cues are presented at the same time, the level of outcome expectation is summed across these. Therefore cues which are presented at the same time compete for the association with the outcome. If one cue predicts the outcome very well, then learning does not take place for any of the other cues present. The formal Rescorla-Wagner model can be represented by the following equation:

$$\Delta V_A = \alpha_A \beta (\lambda - \Sigma V) \qquad [1]$$

In equation 1, ΔV_A refers to the predicted change in the strength of the association between a cue A and the outcome. The terms α_A and β correspond to the learning rate

associated with the cue and the outcome, respectively. The magnitude of the outcome is represented by λ. The existing expectation of the outcome, summed across all cues present, is signified by ΣV. The expression (λ – ΣV) denotes predictive error; also referred to as the summed error-term. This model predicts that the amount of predictive error for an outcome is calculated by taking away the combined associative strengths of all the cues present at a particular learning instance from the magnitude of the outcome. This model accounts for blocking as follows. When participants are presented with A+/AX+ trials they learn that A causes the outcome. On AX+ trials, the outcome is fully predicted by A and therefore no learning between X and the outcome takes place. This works similarly whether A+ and AX+ trials are presented in different blocks or if they are intermixed. If the trials are intermixed then at the beginning of learning X gains some associative strength, which decreases with repeated presentation of A+/AX+ trials. Blocking has been widely replicated in both humans and non-human animals (e. g. Aitken, Larkin & Dickinson 2001; Arcediano, Matute, & Miller, 1997; Chapman and Robbins, 1990; Dickinson, Shanks & Evenden, 1984; Dwyer, Haselgrove & Jones, 2011; Hinchy, Lovibond, & Ter-Horst, 1995; Kruschke & Blair, 2000; Kruschke, Kappenman & Hetrick, 2005; Shanks, 1985; Waldmann & Walker, 2005, although see Maes et al., 2016). Furthermore, blocking has been described as "a touchstone for theories of learning" (Kruschke, Kappenman & Hetrick, 2005, p. 830).

As stated above, an assumption of the Rescorla-Wagner (1972) model is that learning about X is restricted because the outcome is fully expected on AX+ trials due to the presence of A. This assumption is commonly adopted. One way to evaluate this claim is to compare the blocked cue with a different type of a redundant cue. After all, some redundant cues may still be relatively more informative about the occurrence of the outcome than other redundant cues. To illustrate this, consider the relative validity effect (Wagner, Logan, Haberlandt & Price, 1968; Wasserman, 1990). To simplify the design used in Wasserman's study, I will describe findings from two conditions of his experiment. In these conditions participants received different types of training involving compounds with a shared cue Y (BY/CY). It was found that after being exposed to the contingencies

BY+/CY-, in which the common cue Y was paired with more informative cues B and C, participants judged Y to be less predictive of the outcome than in the Condition BY±/CY±, in which each cue was paired with the outcome 50% of the time. This illustrates that causal judgements about a cue differed based on its relative informational value.

Consider the BY+/CY- condition of the above experiment. The common cue Y in this condition can be referred to as an uncorrelated cue because its presence was not correlated with the presence of the outcome. Since this cue was informationally redundant, it may be expected that after being exposed to BY+/CY- contingencies, participants would not learn that the uncorrelated cue Y caused the outcome. The blocked cue X from A+/AX+ trials is also redundant as it is paired with A which predicts the outcome very well, and it may expected that participants would not learn that X caused the outcome either. However what would happen if the blocked cue X and the uncorrelated cue Y were compared in a single experiment? Uengoer et al. (2013) conducted a study in which they presented participants with A+/AX+/BY+/CY- trials in order to investigate whether the blocked cue X and the uncorrelated cue Y would become associated with the outcome to the same extent. They embedded this design within the allergist task commonly used in human causal learning (e. g. Aitken, Larkin & Dickinson, 2001; Wassermann, 1990). In this task the cues were different foods and the outcome was the presence (+) or absence (-) of stomach ache. After eight blocks of these trials, they asked participants to provide causal ratings for how likely individual cues were to lead to the outcome, using an 11-point, 0 (*Certainly not*) to 10 (*Very certain*) rating scale. These ratings were taken to represent the strength of the association with the outcome for each cue. Their findings are presented in Figure 1.1. As the figure illustrates, A and B had high positive ratings; both of these cues were paired with the outcome during learning. Cue C was paired with no outcome on CY- trials and had the lowest rating. Cue Y also had a low rating, despite having been paired with the outcome 50% of the time on BY+/CY- trials; this was consistent with the findings of Wasserman. The key comparison was between ratings for X and Y. Participants predicted that the blocked cue X caused the outcome to a

greater extent than the uncorrelated cue Y. Jones and Pearce (2015) referred to this difference as "the redundancy effect." This effect has also been observed in rats and pigeons (Jones & Pearce, 2015; Pearce et al., 2012).

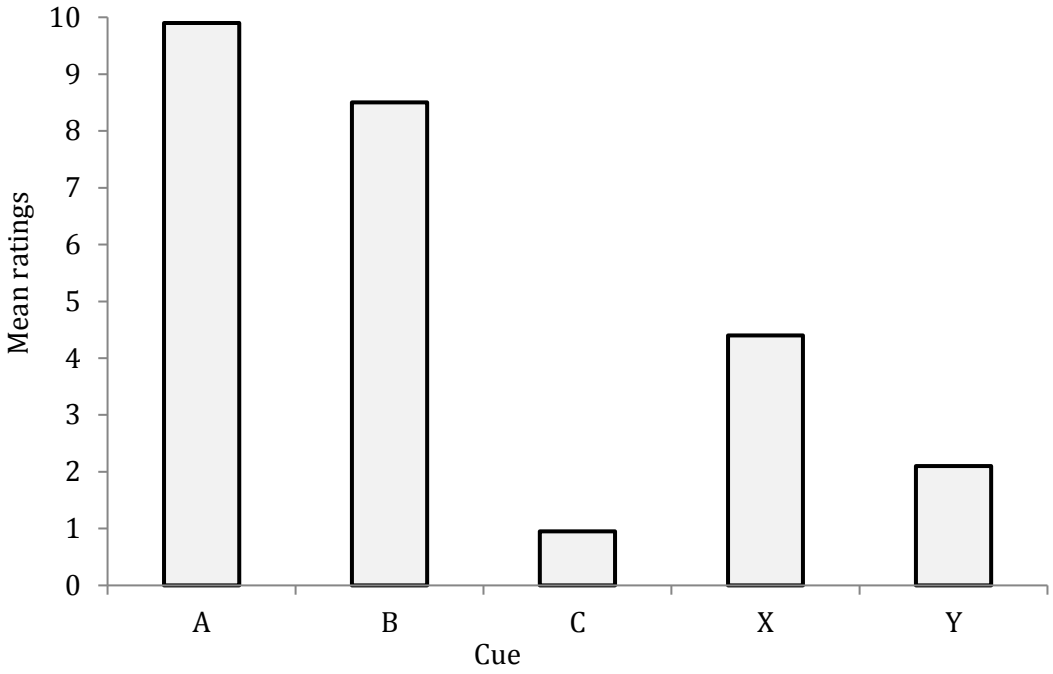

Figure 1.1. The results of Uengoer et al.'s (2013) Experiment 2. Participants received A+/AX+/BY+/CY- training.

The redundancy effect is interesting for two reasons. Firstly, it adds to evidence showing that the way participants treat redundant cues is dependent on their associative history. Secondly, the Rescorla-Wagner (1972) model makes the opposite prediction for the comparison between these cues; the uncorrelated cue Y is predicted to have a stronger association with the outcome than the blocked cue X. The detailed predictions of this model are as follows. Due to its use of a summed error-term, it predicts that learning about X will be restricted by the presence of its accompanying cue A. Therefore, provided learning is at asymptote, X should have a minimal association with the outcome. In Uengoer et al.'s (2013) experiment, Y was paired with B and C which distinguished

between the presence and the absence of the outcome better than Y, however the Rescorla-Wagner model predicts Y to be a weak cause of the outcome at asymptote. This is because of the use of the summed error-term. According to the model, Y will gain some associative strength on BY+ trials. The associative strength for C will therefore decrease on CY- trials, falling below zero, and C will become a weak inhibitor of the outcome as a result. Inhibition for C will allow Y to maintain some positive associative strength. The positive and negative associative strengths for C and Y are predicted to mirror each other at asymptote.

It is important to note however, that the prediction of a stronger association with the outcome for Y than for X made by the Rescorla-Wagner (1972) model hinges on the following assumptions: firstly, that the learning rate for the outcome is higher on the trials on which the outcome occurs (β_E) than on the trials on which the outcome does not occur (β_I) and secondly, that learning is at asymptote. While the former is commonly made by this model, Experiment 1 of Uengoer et al. (2013) aimed to find out whether it could be determined with certainty that β_E was higher than β_I. In this experiment, participants were presented with the design containing the blocked cue X, the blocking-control cue E, and trial types necessary for the relative validity effect (A+/AX+/BY+/CY-/DE+/FG±/HG±). In this design, the uncorrelated cue Y was the informationally-redundant cue paired with more predictive B and C (BY+/CY-; equivalent to the same condition in Wasserman, 1990), and G was as informative as F and H; all three were paired with the outcome 50% of the time (FG±/HG±; equivalent to the BY±/CY± condition in Wasserman, 1990). Firstly, this experiment found that the redundancy effect was observed, as ratings for X were higher than for Y. Secondly, the relative validity effect was also observed, as ratings for G were higher than ratings for Y. This is important because in order for the Rescorla-Wagner model to be able to explain the relative validity effect, an assumption has to be made that $\beta_E > \beta_I$ (Rescorla, 2002). However if this assumption is made, then the redundancy effect is inconsistent with the predictions of this model, even if learning is assumed to be pre-asymptotic.

While in this thesis I focused on the predictions made by the Rescorla-Wagner (1972) model, note that other models that compute learning using summed error-term rules make similar predictions regarding the redundancy effect, that Y will have a stronger association with the outcome than X in a design A+/AX+/BY+/CY- (e. g. Esber & Haselgrove, 2011; Pearce & Mackintosh, 2010).

An alternative to the summed error-term perspective is determining predictive error for each cue individually, or using single error (e. g. Bush & Mosteller, 1955, Denniston, Savastano & Miller, 2001; Mackintosh, 1975; Miller & Matzel, 1988; Stout & Miller, 2007). Given that summed error-term models struggle to explain the redundancy effect, Uengoer et al. (2013) considered whether single error-term models would fare better.

Single error-term models determine the expectation of the outcome for every cue individually, even when several cues are presented on the same trial. One of the earliest single error-term models was proposed by Bush and Mosteller (1955). While this model aimed to predict the probability of responding based on the associative strength between a cue and an outcome, if it is assumed that these are perfectly related, this model can be expressed mathematically as follows:

$$\Delta V_A = \alpha_A \beta (\lambda - V_A) \qquad [2]$$

Equation 2 is very similar to Equation 1. The main difference is that in the Bush and Mosteller (1955) model the expectation of the outcome is determined by single error-term rules, rather than summed error; the "Σ" term is absent. In practice this means that any cue paired with the outcome consistently has a high associative strength. Importantly, this model can predict the redundancy effect because in the Uengeor et al.'s (2013) design A+/AX+/BY+/CY-, X was presented with the outcome 100% of the time while Y was presented with the outcome 50% of the time. Therefore X is predicted to have a stronger association with the outcome than Y. However, this model does not predict cue competition effects. Notably, in Experiment 1 of Uengoer et al., blocking was observed alongside the redundancy effect; ratings for the blocking-control cue E (which received

DE+ training) were higher than for X. The relative validity effect is another example of cue competition that falls outside the scope of the Bush and Mosteller model. However, the relative validity effect was observed together with the redundancy effect in Experiment 1 of Uengoer et al. Cues Y and G in the design BY+/CY-/FG±/HG± were paired with the outcome 50% of the time, however a weaker association between Y and the outcome was observed than between G and the outcome. This suggests that single error-term models may not be best suited to explain the redundancy effect without incorporating an additional process.

One theory that uses a single error-term but incorporates an additional process is the comparator hypothesis (Denniston, Savastano, & Miller, 2001; Miller & Matzel, 1988; Stout & Miller, 2007). This theory proposes that the strength of the association between a cue and the outcome is based on single error-term rules. Therefore any cue which is paired with the outcome consistently is predicted to have a strong association, regardless of any other cues present, including the blocked cue X. At the point of performance however, when participants are asked to estimate the extent to which a cue causes the outcome, a comparator process takes place. This process involves a comparison between the associative strength for the target cue with the associative strength for any companion cues the target was presented with during learning. If the companion cue has a strong association with the outcome, this reduces the strength of the response for the target cue and it will receive a low rating. On the other hand, if the companion cue has a weak association with the outcome then the strength of the response for the target cue is increased and its rating will be high. Applied to the redundancy effect, this theory predicts that even though during learning X has a strong association with the outcome, at test responding for X will be low, because it was presented with A, which had a strong association with the outcome. Importantly however, X is predicted to retain some associative strength. Cue Y on the other hand, was presented with two other cues, one of which had a strong association with the outcome (B), however this cue starts with a lower associative strength because it was paired with the outcome 50% of the time. Therefore

responding for Y is predicted to be lower than for X at test. While the comparator hypothesis has the potential to account for the redundancy effect, Experiment 1 of Jones and Pearce (2015), Experiment 2 of Uengoer et al. (2013), and an experiment reported by Zaksaite and Jones (2017), provide evidence against the redundancy effect being due to its predictions. As such, I will not consider the comparator hypothesis further in this thesis.

Another theory that uses a single error-term and incorporates an additional process is Mackintosh's (1975) theory of selective attention. Based on the observation that a competition for a limited amount of attention was central to attentional theories, Mackintosh proposed that a similar competition for a limited amount of attention occurs between cues in associative learning. He proposed that attention for cues varies with experience. In particular, attention to a cue increases if it is the best predictor of the outcome on a particular trial, and decreases if it is not the best predictor of the outcome. This attention guides learning and cues which receive more attention are learned about to a greater extent than cues which receive less attention. Hence, a person pays more attention to a good predictor of the outcome, and learns about the relationship between this cue and the outcome, than to a poor predictor of an outcome. Mackintosh described the change in attention for cue A ($\Delta\alpha_A$) on an AX+ trial, as follows:

$\Delta\alpha_A$ is positive if $|\lambda - V_A| < |\lambda - V_X|$ [3]

$\Delta\alpha_A$ is negative if $|\lambda - V_A| \geq |\lambda - V_X|$ [4]

The associative strength change is calculated as follows:

$\Delta V_A = \alpha_A \beta (\lambda - V_A)$ [5]

Note that in Equation 5, contrary to the Rescorla-Wagner (1972) model, the associative strength of a cue changes based on single error-term rules, similarly to the Bush and Mosteller (1955) model. Any competition between the cues for an association with the outcome is described by their competition for attention. While some data challenging to this theory have been observed (e. g. Hogarth, Dickinson, Austin, Brown & Duka, 2008;

Pearce & Hall, 1979), there is much evidence to support Mackintosh's (1975) theory of selective attention, especially in humans (e. g. George & Pearce, 1999; Le Pelley & McLaren, 2001, 2003, 2004; Le Pelley, Mitchell, Beesley, George & Wills, 2016; Le Pelley, Vadillo & Luque, 2013; Lochmann & Wills, 2003). This theory also can account for the redundancy effect. Because of its use of a single error-term it predicts a stronger association with the outcome for X than for Y, as X was paired with the outcome 100% of the time, while Y was paired with the outcome 50% of the time. Mackintosh's theory also predicts that learning about cues is determined by the amount of attention they receive on a particular trial. While attention for both X and Y is predicted to decrease as they are presented with other cues more informative about the occurrence of the outcome, it is unclear for which cue this decline will be faster. One possibility is that the decrease in attention is faster for X because its companion, A, was shown to be the perfect predictor of the outcome. An alternative possibility is that the decline in attention for Y is faster, because it was presented twice as often as X, as well as having been paired with other cues that are perfect predictors of both the presence (B), and the absence (C), of the outcome. Due to the lack of specification of this theory it is not possible to accurately predict how attention interacts with previous associative history to determine learning. However, provided that decline in attention for X does not exceed that for Y, this theory can account for the redundancy effect.

This is what I set out to investigate in Chapter 2. In this chapter I used eye-tracking to determine whether there were any differences in attention between blocked and uncorrelated cues. If evidence of differences in attention between these cues was obtained, this would implicate attentional processes as contributing to the redundancy effect and provide further support for Mackintosh's (1975) theory of selective attention.

However, there is also an alternative interpretation of the redundancy effect. As noted previously, Uengoer et al. (2013) used an allergist task in their experiments, in which the cues were different foods resulting in a stomach ache or no stomach ache when consumed

by a fictional patient. At test, participants were asked to provide causal ratings for each cue on an 11-point scale ranging from 0 (*Certainly not*) to 10 (*Very certain*). While X received higher ratings than Y, ratings for X were approximately in the middle of the scale, around 5. This rating has two possible interpretations. It could refer to the fact that on average, participants thought that X was a weak cause of the outcome, consistently with Uengoer et al.'s interpretation. However, it could also mean that participants were uncertain about whether or not X caused the outcome because the rating in the middle of the scale may best reflect the point of uncertainty on this scale. Therefore, rather than indicating a stronger association between X and the outcome than Y and the outcome, this rating could have reflected participants' uncertainty about whether or not X caused the outcome. If this interpretation was correct, then the redundancy effect may be due to differences in participants' certainty about the causal status of X and of Y. Participants may have been uncertain about whether or not X caused the outcome, rating it in the middle of the scale, while they were certain that Y did not cause the outcome, giving it a low rating.

The latter interpretation of the data is consistent with previous evidence indicating that resolving ambiguity about the blocked cue has the potential to reduce its rating (e. g. Beckers, De Houwer, Pineno & Miller, 2005; Livesey & Boakes, 2004; Lovibond, Been, Mitchell, Bouton & Frohardt, 2003; Mitchell & Lovibond, 2002; Mitchell, Lovibond & Condoleon, 2005). Participants' uncertainty about the blocked cue falls outside the predictions made by single and summed error-term models discussed earlier. However, the prediction that people are uncertain about whether or not the blocked cue leads to the outcome has been made by the probabilistic contrast model (Cheng, 1997; Cheng & Holyoak, 1995) and the propositional account of learning proposed by Lovibond (2003; Mitchell & Lovibond, 2002, see also Beckers, Miller, De Houwer & Urushihara, 2006; De Houwer, 2009; Mitchell, De Houwer & Lovibond, 2009). Lovibond found that cues were successfully revalued regardless of whether contingencies were described, experienced, partly-experienced, or partly-described. This revaluation was evidenced by both self-report and skin conductance measures. In the light of these findings, Lovibond suggested

that causal learning is propositional. Propositions refer to beliefs that participants have about how events are related, and can be tested and adjusted as necessary. These propositions can be expressed verbally, specify the nature of the associations encoded, and are formed using complex every-day reasoning processes. The propositional account however, can be seen as answering an orthogonal question to the learning models. Models of learning aim to predict how learning proceeds and under which conditions, and how associative strengths between cues and outcomes change on each trial. The propositional account on the other hand, is concerned with what kind of information is acquired during learning. It predicts that propositions result in knowledge that is conscious and declarative, but does not specify the rules under which learning takes place. The lack of specificity gives this account greater flexibility than the models in explaining a variety of learning effects, but also makes it difficult to evaluate it using my experiments. Furthermore, the conscious status of information acquired during learning still remains a matter of debate (e. g. Livesey & Boakes, 2004; Mitchell et al. 2009; Premack, 2007; Schultz & Helmstetter, 2010) and falls beyond the scope of this thesis. However, because the propositional account and associative models of learning are concerned with answering different types of questions, these approaches could be combined, for example by incorporating a propositional component into the Rescorla-Wagner (1972) model (see Thorwart & Livesey, 2016). This would enable this model to predict uncertainty about the blocked cue therefore potentially enabling this model to account for the redundancy effect.

In Chapter 3 I set out to explore whether participants' uncertainty about the causal status of X contributed to the redundancy effect. Firstly, I investigated whether participants were more uncertain about their causal ratings for X than for Y. In addition, I used manipulations which were predicted to influence ratings of X by resolving the ambiguity about its effects, including outcome-additivity training and outcome-rate variations. If these manipulations successfully reduced the magnitude of the redundancy effect, this would indicate that participants' uncertainty about the causal status of X contributes to the redundancy effect.

Chapter 4 investigated whether the redundancy effect could have been the result of specific properties of the allergist task which has been used to investigate the redundancy effect to date. Recall that the prediction of the Rescorla-Wagner (1972) model is that Y is a weak cause of the outcome because C becomes a weak inhibitor. In order for cues to become inhibitory in the allergist task however, participants would have to learn that a food consumed with a stomach ache-causing food results in no stomach ache. In other words, some foods would have to be seen as preventing a stomach ache. Because this is contrary to most every-day experiences of foods and their effects, inhibition may be difficult to observe in the allergist task. If C does not become an inhibitor of the outcome however, this does not protect Y from extinction and in this case the Rescorla-Wagner model predicts that both X and Y will have a minimal association with the outcome. In Chapter 4 I used a task in which inhibition was more plausible to investigate whether a lack of inhibition for C contributes to the redundancy effect.

Chapter 2: Differences in attention

In the previous chapter I introduced the redundancy effect and why it poses a challenge to the Rescorla-Wagner (1972) model of learning. Previous articles on the redundancy effect (Jones & Pearce, 2015; Pearce et al., 2012; Uengoer et al., 2013) identified that single error-term models (e. g. Bush & Mosteller, 1955) also cannot account for the pattern of observed data without incorporating an additional process. Mackintosh's (1975) theory of selective attention was identified as a strong candidate for explaining the redundancy effect, however. This is because, in addition to a single error-term which predicts a stronger association with the outcome for the blocked cue than the uncorrelated cue, associative strength is also determined by differences in attention. The experiments described in this chapter aimed to explore whether the redundancy effect may be a consequence of differences in attention to blocked and uncorrelated cues.

According to Mackintosh's (1975) theory of selective attention, the amount of attention a cue attracts relative to other cues is related to its predictive history. The amount of attention paid to a cue also determines the speed of learning about this cue, such that more attention results in faster learning. Ample evidence illustrates that more attention is dedicated to cues which are good predictors of outcomes than cues which are poor predictors (e. g. Beesley & Le Pelley, 2011; Kruschke, Kappenman & Hetrick, 2005; Le Pelley, Beesley & Griffiths, 2011; Rehder & Hoffman, 2005). Mackintosh's theory predicts that during A+/AX+/BY+/CY- training, attention decreases for both the blocked cue X and the uncorrelated cue Y relative to their more informative companion cues A, B, and C. When it comes to differentiating how fast this decrease takes place, the predictions are less clear. Because of a lack of specificity of this theory regarding how attention interacts with previous predictive history to determine learning, this theory could predict a stronger association with the outcome for X than for Y, but it could also predict the reverse. Provided that decline in attention for X does not exceed that for Y however, this theory can predict the redundancy effect.

One method used to investigate differences in attention is eye-tracking. Although it is possible to shift attention without corresponding eye-movements (e. g. Posner, 1980), it is generally accepted that eye-gaze reflects attention (e. g. Parkhurst, Law, Niebur, 2002; Rayner, 1998; Wills, Lavric, Croft & Hodgson, 2007). Several studies provided evidence for Mackintosh's (1975) theory of selective attention using eye-tracking (e. g. Beesley & Le Pelley, 2011; Kruschke, Kappenman & Hetrick, 2005; Le Pelley, Beesley & Griffiths, 2011; Rehder & Hoffman, 2005). Therefore in this chapter, I used eye-tracking to investigate whether differences in attention between blocked and uncorrelated cues exist. Given that X had higher ratings than Y in Uengoer et al. (2013), it may be expected that participants would look at X for longer than at Y. However, it is also possible that participants would look at Y for longer than at X. Uengoer et al. assumed that at the beginning of learning, cues were not associated with the outcome (corresponding to a hypothetical rating of zero). The association with the outcome for some cues would have been acquired as learning proceeded, depending on the cue's relationship with the outcome. Therefore higher ratings for X than for Y in Uengoer et al. were interpreted to reflect a stronger association with the outcome for X than for Y. However, an alternative possibility is that at the outset of learning, prior to being presented with any information about the cues, participants were uncertain about whether or not cues led to the outcome. As learning proceeded, participants learned whether cues caused the outcome and uncertainty decreased, resulting in more extreme ratings for cues at test. Cue X in Uengoer et al.'s experiments received ratings of approximately 5 on the 11-point, 0 (*Certainly not*) to 10 (*Very certain*) causal-rating scale. Because ratings in the middle of the scale may best reflect the point of uncertainty on this scale, ratings for X could have indicated that participants were uncertain about whether or not X caused the outcome. There are two reasons why participants may have been uncertain about the effects of X on the outcome. Because X was never presented by itself, its effects could not be verified; participants never saw the effects that X had on the outcome independently of A. In addition, in the allergist task used by Uengoer et al., there was a ceiling on the outcome; two causal cues would have led to

the same outcome as one causal cue. Therefore regardless of whether or not X was a cause of the outcome, on AX+ trials the outcome would have been the same. If participants were uncertain about X, then the redundancy effect in Uengoer et al. may have reflected greater learning about Y than X. At the outset of learning participants may have been uncertain about whether or not Y led to the outcome, but with repeated pairings, they learned that this cue did not cause the outcome and gave it a low rating at test. At the outset of learning participants may have been similarly uncertain about X, however because they did not get an opportunity to verify its effects on the outcome, they remained uncertain about this cue and gave it ratings in the middle of the scale at test. If this was case, more attention for Y than for X may be expected. If the redundancy effect reflected greater learning about X than Y consistently with Uengoer et al.'s interpretation, then more attention for X than for Y may be expected. In Experiment 1 I set out to compare eye-gaze durations for X and for Y using the same allergist task as in Uengoer et al. using the design A+/AX+/BY+/CY-. Longer eye-gaze durations for X than for Y would indicate greater learning about X than about Y, consistently with Uengoer et al.'s interpretation. On the other hand, longer eye-gaze durations for Y than for X would indicate greater learning about Y than about X.

Experiment 1

The design of Experiment 1 is shown in Table 2.1. I used eye-tracking to measure how long participants looked at each cue during learning. Of particular interest was the comparison between eye-gaze durations for X and Y.

Method

Participants. Fifty-one participants took part in this experiment. Twenty-three were undergraduate students receiving course credit and the rest were members of the public receiving payment. They were tested individually.

Table 2.1

The design of Experiment 1.

Stage 1	Test
A+	A
AX+	B
BY+	C
CY-	X
	Y
x 16	x 2

The eye-tracker could not be calibrated for five participants, whose data were excluded from the experiment, reducing the number of participants to 46. The remaining participants had a mean age of 29.52 years.

Materials. The experiment was presented on a 22-inch desktop computer with a 1280 x 1024 screen resolution. Eye-tracking data were collected using an SMI RED remote eye-tracker sampling at 50 Hz (SensoMotoric Instruments, Teltow, Germany). The experiment was designed, presented and responses recorded, using E-prime 2.0 software (Psychology Software Tools, PA, US).

The cues were selected from 11 images of foods on a white background, 300 x 300 pixels. The foods were: apple, banana, broccoli, cabbage, cherries, corn, grapes, orange, pepper, pumpkin, and strawberries. The foods were randomly assigned to each type of cue (A, B, C, X, Y) for each participant. 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 with white text. Participants responded using the mouse. In order to measure gaze

durations for the cues, regions of interest (ROIs) were defined at the boundaries of the cue images.

Procedure. Upon their arrival, participants completed a nine-point calibration procedure for eye-tracking.

The instructions for the learning task were adapted from Uengoer, Lotz and Pearce (2013), and were presented on the screen as follows:

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 Stage 1 participants were presented with 16 blocks of trials, in which each of the different trial types (A+, AX+, BY+, CY-) appeared once per block. The order of the trials within each block was random, with no successive repetitions of the same trial type between blocks. Each trial started with the presentation of either one or two images of foods at 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 once in each sequential pair of blocks. 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, determined randomly, 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". Eye-gaze durations were recorded from when the images appeared on the screen at the start of the trial, until the participant responded. As soon as the participant responded, the response buttons and the sentence above them were replaced by a statement and picture showing the outcome of the trial, while the images of the cues and the sentence "The patient ate the following food(s):", remained. When the outcome was stomach ache, the statement was "The patient has stomach ache" and the picture of the sad face was shown. When the outcome was no stomach ache, the statement was "The patient has no stomach ache", and the picture of the happy face was shown. This feedback display remained on the screen for 3000 ms, followed by a 500 ms blank screen after which the next trial began.

After all of the trials in Stage 1 were completed, participants were then 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 cue, A, B, C, X, and Y, was presented twice, with the order of trials randomly

determined for each participant. For each cue, an average of the two ratings was used in the data analyses.

Data Analyses. To investigate whether there were any significant differences between ratings for the cues, Analysis of Variance (ANOVA) tests and t-tests were performed on the data. When paired comparisons between all levels of a factor were made, Bonferroni corrections were used for all experiments in this thesis. The alpha level of significance was set at .05 for all other comparisons including when using t-tests for the comparisons of interest (e. g. blocking and the redundancy effect) and when using simple main effects to explore interactions. When data violated sphericity, Greenhouse-Geisser corrections were applied to degrees of freedom. For null comparisons of interest, Bayes Factors (BF_{01}) were calculated using a JZS prior with a scaling factor of 0.707. To calculate Bayes Factors, JASP version 0.6 was used (JASP Team, 2015). Bayes Factors greater than three are considered to provide support for the null hypothesis. Bayes Factors less than one-third indicate support for the alternative hypothesis (Jeffreys, 1961). Estimates of effect sizes reported for ANOVA tests were eta squared (η^2) and partial eta squared (η_p^2), and for t-tests, correlation coefficients (r). For figures, error bars showed the standard error of the mean, adjusted to exclude the between-subjects variability as recommended by Cousineau (2005), unless stated otherwise.

Results

Behavioural data. Figure 2.1 shows the proportion of stomach ache predictions throughout the eight epochs of Stage 1, averaged across participants. An epoch was defined as an average of responses on two trials of the same trial type. The figure shows that participants learned the contingencies, with greater proportions of stomach-ache predictions for trial types leading to a stomach ache (A+/AX+/BY+) than the one leading to no stomach ache (CY-). In the final epoch, correct predictions were made on 99.46% (*SD* = 5.2%) of the trials.

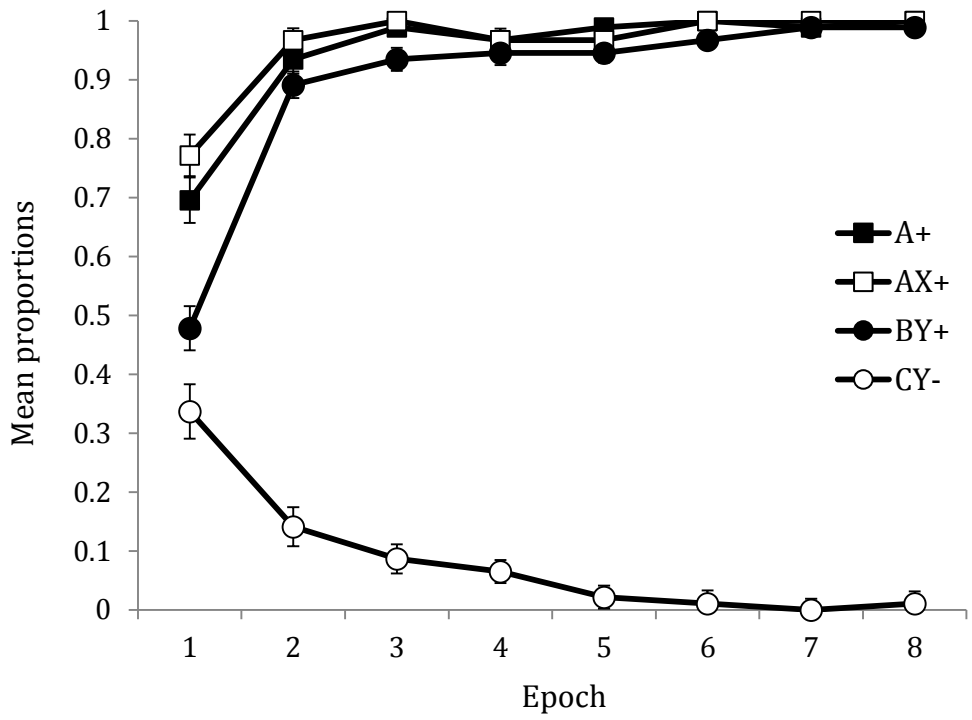

Figure 2.1. Mean proportions of stomach-ache predictions throughout the eight epochs for Stage 1 in Experiment 1 (±SEM).

Causal ratings provided for each cue at test are presented in Figure 2.2. This indicates that ratings for the cues were similar to Uengoer et al.'s (2013) experiments.

A one-way ANOVA (cue [A, B, C, X, Y]) revealed a significant effect of cue, $F(2.36, 106.09) = 195.75$, $p < .001$, $\eta^2 = .81$. Bonferroni-corrected paired comparisons indicated that ratings for all cues differed significantly from each other, ts ≥ 3.59, ps ≤ .008, rs ≥ .47. A t-test confirmed that the redundancy effect was observed; X had significantly higher ratings than Y, $t(45) = 5.02$, $p < .001$, $r = .6$.

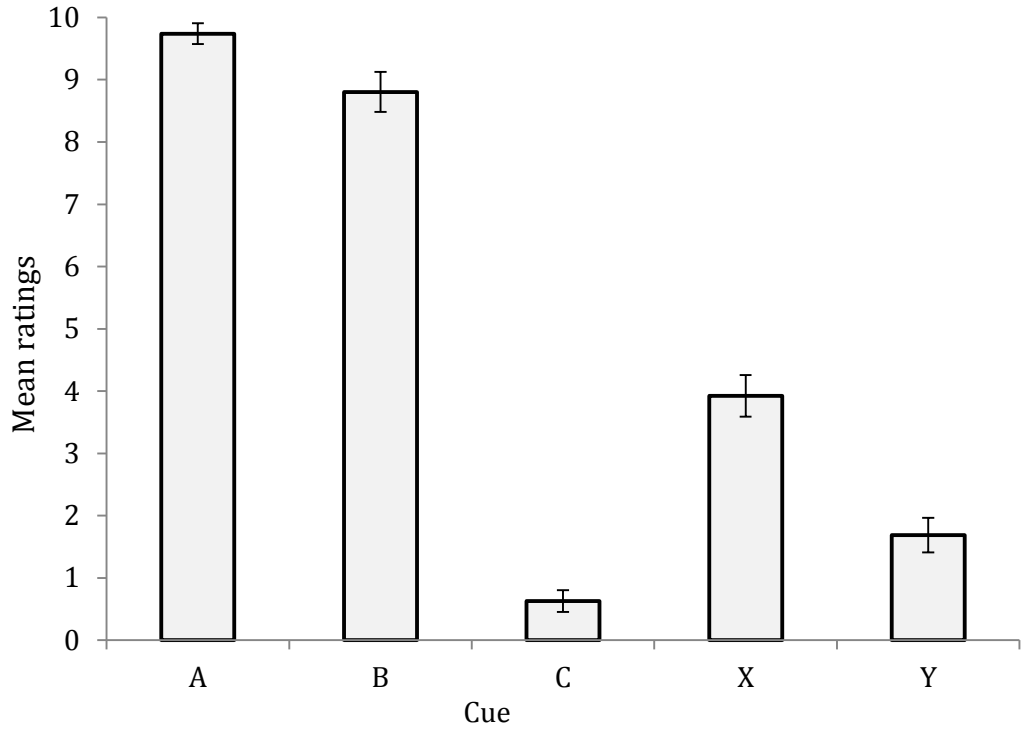

Figure 2.2. Mean causal ratings at test in Experiment 1 (±SEM).

Eye gaze. Eye-gaze durations during Stage 1 are presented in Figure 2.3. A one-way (epoch [1-8]) ANOVA for gaze durations on A+ trials revealed a significant effect of epoch, $F(5.18, 232.94) = 6.43$, $p < .001$, $\eta^2 = .13$ (higher left panel). To compare gaze durations between A and X on AX+ trials a two-way (cue [A, X] by epoch [1-8]) ANOVA was conducted on the data. This revealed a significant effect of cue, $F(1, 45) = 9.71$, $p = .003$, $\eta_p^2 = .18$, indicating that participants looked for longer at A than at X. The effect of epoch was also significant, $F(4.46, 200.66) = 11.93$, $p < .001$, $\eta_p^2 = .21$, but the interaction was not, $F(7, 315) = .64$, $p = .698$, $\eta_p^2 = .01$ (higher right panel). A two-way (cue [B, Y]) by epoch [1-8]) ANOVA on BY+ trials revealed a significant effect of cue, $F(1, 45) = 8.55$, $p = .005$, $\eta_p^2 = .16$, indicating longer gaze durations for B than for Y, a significant effect of epoch, $F(4.21, 189.35) = 5.82$, $p < .001$, $\eta_p^2 = .12$, but no significant interaction, $F(5.26, 246.65) = 1.17$, $p = .325$, $\eta_p^2 = .03$ (lower left panel). Finally, a two-way (cue [C, Y]) by epoch [1-8]) ANOVA

on CY- trials revealed no significant effect of cue, $F(1, 45) = .36$, $p = .551$, $\eta_P^2 = .01$, $BF_{01} = 10.87$, a significant effect of epoch, $F(4.94, 222.05) = 11.34$, $p < .001$, $\eta_P^2 = .2$, but no significant interaction, $F(7, 315) = 1.39$, $p = .21$, $\eta_P^2 = .03$ (lower right panel). Longer gaze durations for the more predictive cue on AX+ and BY+ trials indicated that the eye-tracking procedure used in this experiment was sensitive enough to detect differences gaze durations between cues with different levels of predictiveness.

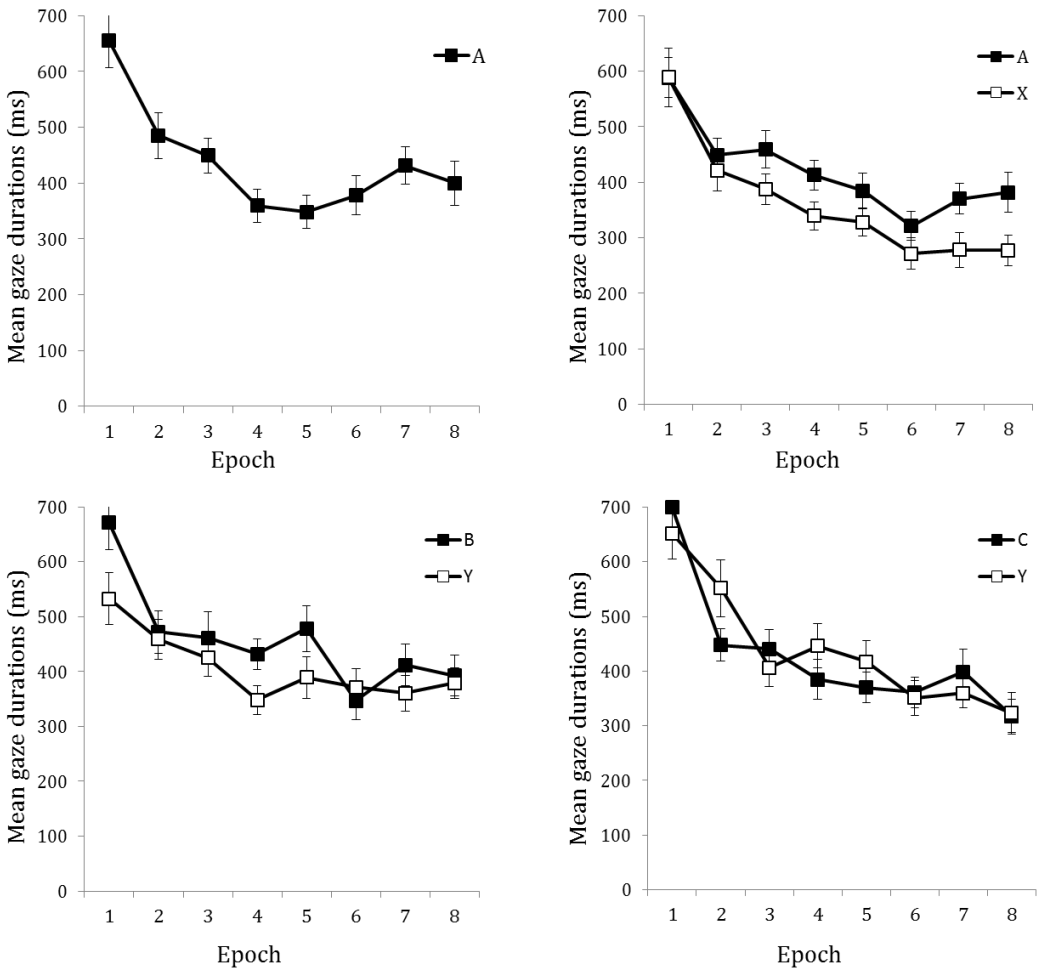

Figure 2.3. Gaze durations for cues within each trial type during Stage 1 in Experiment 1 (ms; ±SEM). Higher left panel indicates gaze durations for A on A+ trials. Higher right panel indicates gaze durations for A and for X on AX+ trials. Lower left panel illustrates gaze durations for B and for Y on BY+ trials. Lower right panel illustrates gaze durations for C and Y on CY- trials.

Figure 2.4 shows gaze durations for X on AX+ trials and for Y averaged across BY+ and CY- trials. In order to determine whether there were any differences between gaze durations for these cues, a two-way (cue [X, Y] by epoch [1-8]) ANOVA was conducted. This revealed a significant effect of cue, $F(1, 45) = 8.91$, $p = .005$, $\eta_p^2 = .17$, indicating longer gaze durations for Y than for X. The effect of epoch was also significant, $F(3.44, 154.89) = 12.97$, $p < .001$, $\eta_p^2 = .22$, while the interaction was not, $F(7, 315) = .64$, $p = .69$, $\eta_p^2 = .01$.

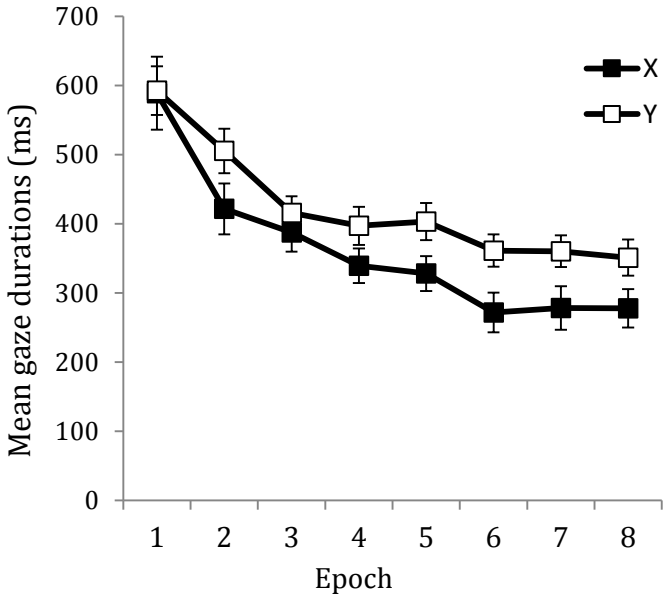

Figure 2.4. Gaze durations for X and for Y across the eight epochs from Stage 1 in Experiment 1 (ms; ±SEM). This graph shows average eye-gaze durations for X on AX+ trials and for Y averaged across BY+ and CY- trials.

Even though the analyses indicated that participants looked for longer at Y than at X, there were several differences between the trials containing X and Y. For example, X was always presented with one other cue, while Y was presented with two other cues, depending on the type of trial. There was also a non-significant trend for longer reaction times on trials containing Y than X, $t(45) = 1.88$, $p = .066$, $r = .27$. While this difference was not significant,

longer trial durations containing Y would also have given participants more opportunity to look at this cue. The reaction times on each type of trial are presented in Figure 2.5.

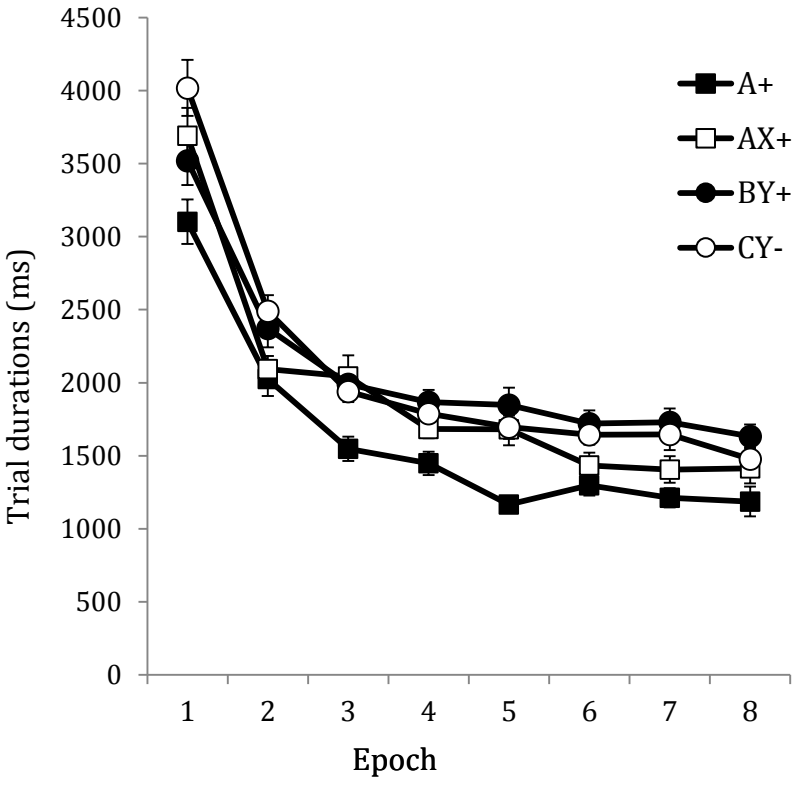

Figure 2.5. Trial durations on each type of trial in Stage 1 of Experiment 1 (±SEM).

Therefore a fairer measure of gaze durations for a comparison between X and Y may have been gaze durations as proportions of trial length. These data are presented in Figure 2.6. A two-way (cue [X, Y] by epoch [1-8]) ANOVA indicated that while the effect of epoch remained significant, $F(4.69, 211.23) = 2.43$, $p = .04$, $\eta_P^2 = .05$, the effect of cue was not, $F(1, 45) = .11$, $p = .746$, $\eta_P^2 < .01$ and there was no significant interaction, $F(5.42, 243.9) = .16$, $p = .982$, $\eta_P^2 < .01$. A Bayesian t-test confirmed that there was no difference between gaze durations as proportions of trial length, averaged across epochs, $t(45) = .33$, $p = .746$, $r = .05$, $BF_{01} = 5.95$.

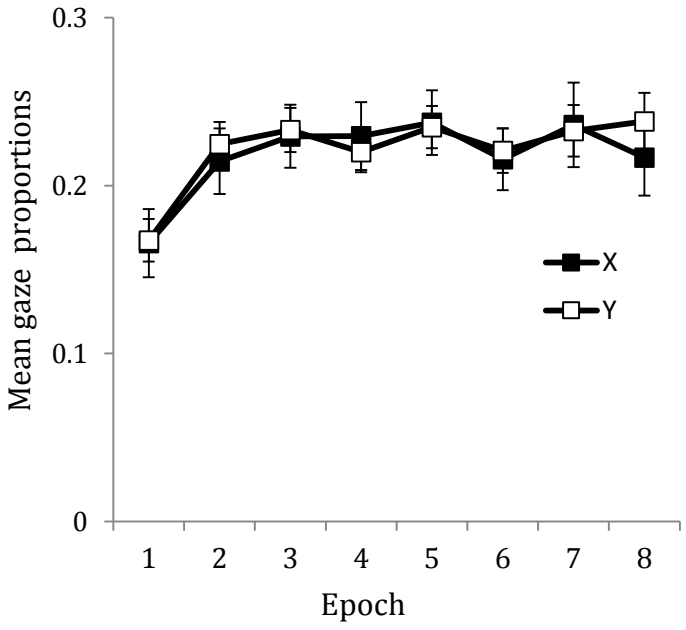

Figure 2.6. Proportional gaze durations for X and Y throughout the eight epochs of Stage 1 in Experiment 1 (±SEM). This figure shows gaze durations for X and Y expressed as proportions of trial durations.

Since Mackintosh's (1975) theory of selective attention predicts that attention drives learning, it is possible that longer uncorrected-gaze durations for Y could still have contributed to the redundancy effect. In other words, the longer participants looked at Y compared to X, the larger the observed redundancy effect. If this was the case, a positive correlation between the gaze-duration bias for Y (calculated as uncorrected gaze durations for Y – uncorrected gaze durations for X) and the magnitude of the redundancy effect (calculated as X ratings – Y ratings) may be expected. However this correlation was negative and at a non-significant trend level, $r(46) = -.249$, $p = .095$. The correlation between gaze-duration bias for Y as proportions of trial length and the magnitude of the redundancy effect was also not significant, $r(46) = -.199$, $p = .185$.

To summarise, in Experiment 1 I set out to explore whether differences in attention between X and Y may be responsible for the redundancy effect. In this experiment, the redundancy effect was significant. While participants looked at Y for longer than at X,

there was a trend for a non-significant difference between the trial durations for these cues; participants took longer to respond on trials containing Y than X. When gaze durations were expressed as proportions of trial durations, no differences between X and Y were found. I also explored the possibility that longer gaze durations for Y than for X contributed to the magnitude of the redundancy effect. If this was the case, a positive correlation between differences in gaze durations between Y and X, and the magnitude of the redundancy effect may have been expected. However, this correlation was not significant.

One reason for longer reaction times on trials containing Y than X may have been that these trials (BY+, CY-) involved a discrimination between the presence and the absence of the outcome. Because Y was not informative about whether or not the outcome occurred, participants would have had to rely on B and C to solve this discrimination. Therefore, if they looked at Y before the other cue on these trials, they had to look at the accompanying cue in order to respond correctly. This would have produced a delay in responding relative to the trials containing X. While A was more informative about the occurrence of the outcome than X on AX+ trials, X was presented on these trials only, always with A, and these trials always led to the outcome. Therefore if participants looked at X at the beginning of the trial, they could have responded correctly, without needing to look at A. If the discrimination for trials containing Y was responsible for reaction-time differences, and participants needed to look at B and C because they were more informative about whether the outcome occurred or not, it may have been expected that participants would have looked at C for longer than at Y on CY- trials. However, participants looked at C and at Y for a similar amount of time on CY- trials. The difference between B and C was that the outcome occurred on trials containing B but not on trials containing C. Therefore, participants could have looked for longer at B than at Y because they thought B caused the outcome on BY+ trials. If participants looked for longer at cues which caused the outcome, this would fit with the pattern of data observed in this experiment. Cue A had longer gaze durations than X on AX+ trials and B had longer gaze durations than Y on BY+ trials. Both

A and B received the highest ratings in this experiment (M_A = 9.74, M_B = 8.8), indicating that participants thought these cues caused the outcome. On the other hand, C, X, and Y received low ratings (M_C = .63, M_X = 3.92, M_Y = 1.68) and therefore participants may have looked at these cues for a similar amount of time because they did not think that these cues caused the outcome. Looking for longer at cues which caused the outcome than cues which did not cause the outcome would have allowed participants to solve the contingencies associated with the task; these cues were more informative about whether Mr X experienced a stomach ache or not. As such, this result falls within the predictions made by Mackintosh's (1975) theory of attention.

To summarise, this experiment failed to evidence for differences in attention between X and Y, once trial-duration differences between these cues were accounted for. However before any strong conclusions about this lack of difference were made, in Experiment 2 I wanted to explore whether any differences in selective attention for these cues were present. In order to test this, blocked and uncorrelated cues were presented on the screen at the same time. This was to eliminate differences in circumstances of presentation between these cues to enable a direct comparison. This experiment also aimed to see whether the results observed in Experiment 1 would be replicated.

Experiment 2

Experiment 2 set out to determine whether there were any differences in selective attention for blocked and uncorrelated cues. In order to do this, blocked and uncorrelated cues were presented on the screen at the same time, therefore they were in direct competition with each other for eye gaze. In Stage 1, the design from the previous experiment was doubled, resulting in two sets of each cue: W and X were blocked cues, and Y and Z were uncorrelated cues. In Stage 2, participants were presented with three key trial types; each consisted of one previously blocked and one previously uncorrelated cue (WY-, WZ+, XY+). I compared gaze durations for blocked and uncorrelated cues using

eye-tracking on these trials. Equivalent gaze durations for these cues in Stage 2 would support the findings from Experiment 1, and would indicate no differences in selective attention between these cues. The trials in Stage 2 were also structured in a way that allowed me to explore differences in further learning about these cues (Pearce, Esber, George & Haselgrove, 2008). If blocked cues were learned about more quickly than uncorrelated cues, then the discrimination WY-/XY+ would be easier to solve than WY-/WZ+, as the former was reliant on blocked cues: participants would learn that the compound XY+ caused the outcome more quickly than the compound WZ+. If uncorrelated cues were learned about more quickly, the discrimination reliant on uncorrelated cues, WY-/WZ+, would be easier to solve than WY-/XY+: participants would learn that the compound WZ+ caused the outcome more quickly than the compound XY+. To slow the speed of acquisition in order for small differences in the speed of learning to be detected, six further trial types were added to Stage 2. The full design of this experiment is presented in Table 2.2.

Table 2.2

The design of Experiment 2. Supplementary trials included to slow learning in Stage 2 are shown in italics.

Stage 1	Test	Stage 2
A+	A	WY-
AW+	B	WZ+
B+	C	XY+
BX+	D	*KM+*
CY+	E	*KN-*
DY-	F	*GM-*
EZ+	W	*PH+*
FZ-	X	*PS-*
	Y	*QH-*
	Z	
x 16	x 2	x 12

Method

Participants. Participants were 60 (three male) Psychology undergraduate students at Plymouth University, aged 18-48 (*M* = 21.47, *SD* = 6.23). They received course credit for participation. Inclusion criteria required participants to be over 18 years old, fluent in English and have normal or corrected-to-normal vision and colour vision. These inclusion criteria were used for all subsequent experiments. The eye-tracker failed to calibrate for one of the participants, therefore their data were excluded.

Materials. The materials and procedure in Experiment 2 were the same as in Experiment 1 unless otherwise stated.

The cues were 18 images of foods on a white background: apple, aubergine, banana, broccoli, cabbage, cherries, coconut, corn, grapes, kiwi, lemon, orange, pear, pepper, pumpkin, strawberries, tomato, and watermelon. The foods were randomly assigned to each type of cue (A, B, C, D, E, F, G, H, K, M, N, P, Q, S, W, X, Y, Z) for each participant.

Procedure. Stage 1 of the experiment consisted of 16 blocks of trials, with each of the eight trial types (A+, AW+, B+, BX+, CY+, DY-, EZ+, FZ-) presented once per block in a random order and with no successive repetitions of the same trial type. For trials with one image, this was presented in the middle of the screen. Following Stage 1, in the test stage, participants were asked to provide causal ratings for each individual cue, twice, with the constraint of no successive repetitions of the same trial type, similarly to Stage 1. Participants responded by clicking on an 11-point rating scale ranging from 0 (*Certainly not*) to 10 (*Very certain*) as in Experiment 1.

In Stage 2, three compounds of interest, WY-, WZ+, and XY+, were created, each combined of one blocked cue (W, X) and one uncorrelated cue (Y, Z). Six types of supplementary trials were also included to slow learning in this stage (KM+, KN-, GM-, PH+, PS-, QH-), to make sure that any differences in further learning for the cues of interest were not masked by participants learning quickly. Since these trial types acted as fillers,

the data for these trials are not reported. Stage 2 consisted of 12 presentations of each of the nine different trial types (WY-, WZ+, XY+, KM+, KN-, GM-, PH+, PS-, QH-), presented in a random order and with the constraint of no successive repetitions of the same trial type.

Results

Behavioural data in Stage 1. Figure 2.7 shows proportions of stomach-ache predictions throughout Stage 1, collapsed for equivalent trial types. An epoch was defined as an average of responses on two trials of the same trial type. As in the previous experiment, participants learned the contingencies. In the final epoch participants responded correctly on 95.13% (*SD* = 15.9%) of the trials.

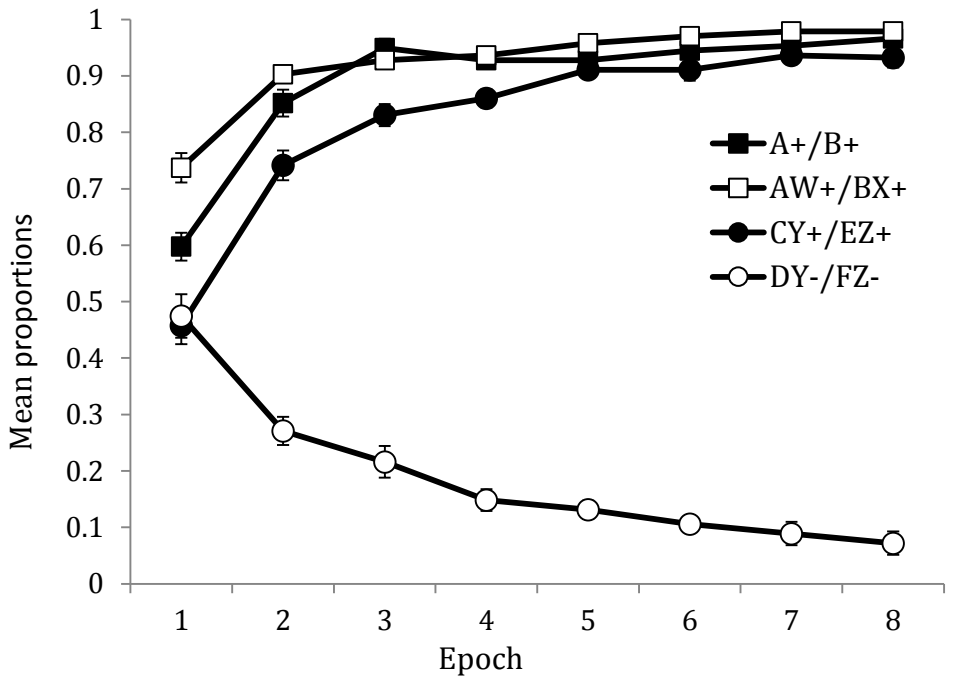

Figure 2.7. Mean proportions of stomach-ache predictions throughout the eight epochs of Stage 1 in Experiment 2 (±SEM). This figure indicates responses collapsed for equivalent trial types.

Causal ratings for each cue at test are presented in Figure 2.8. A one-way (cue [A/B, C/E, D/F, W/X, Y/Z]) ANOVA revealed a significant effect of cue, $F(2.45, 141.8) = 151.6$, p

< .001, η^2 = .72. Bonferroni-corrected paired comparisons indicated that ratings between all cues differed significantly, *t*s ≥ 5.93, *p*s ≤ .001, *r*s ≥ .61. A t-test confirmed that the redundancy effect was observed; W/X received significantly higher ratings than Y/Z, $t(58) = 7$, $p < .001$, $r = .68$.

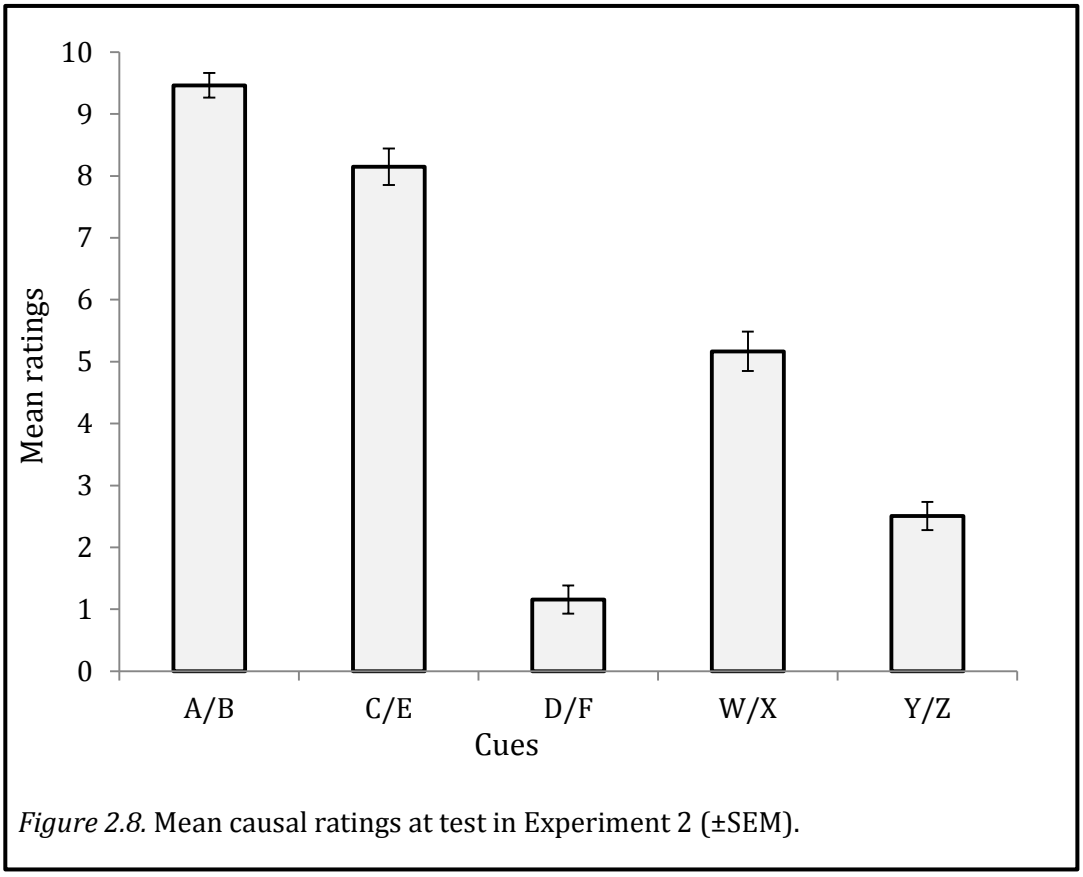

Figure 2.8. Mean causal ratings at test in Experiment 2 (±SEM).

Eye gaze in Stage 1. Figure 2.9 shows gaze durations during Stage 1, collapsed for equivalent cues. A one-way (epoch [1-8]) ANOVA for gaze durations on A+/B+ trials revealed a significant effect of epoch, $F(4.76, 275.81) = 5.52$, $p < .001$, $\eta^2 = .09$ (higher left panel). A two-way (cue [A/B, W/X] by epoch [1-8]) ANOVA on AW+/BX+ trials revealed a significant effect of cue, $F(1, 58) = 7.55$, $p = .008$, $\eta_p^2 = .12$, indicating longer eye-gaze durations for A/B than for W/X. The effect of epoch was also significant, $F(3.77, 218.87) =$

17.2, $p < .001$, $\eta_p^2 = .23$, and there was a non-significant trend for the interaction, $F(5.49, 318.16) = 1.86$, $p = .095$, $\eta_p^2 = .03$ (higher right panel). Because there was a trend for the interaction, simple main effects analyses were conducted on these data. These revealed significantly longer eye-gaze durations for A/B than for W/X on epoch five, $t(58) = 3.13$, $p = .003$, $r = .38$, and non-significant trends for longer durations for A/B than for W/X on epochs four and seven, $ts \geq 1.87$, $ps \leq .067$, $rs \geq .24$; differences between the cues were not significant on any of the other epochs, $ts \leq .77$, $ps \geq .447$, $rs \leq .1$.

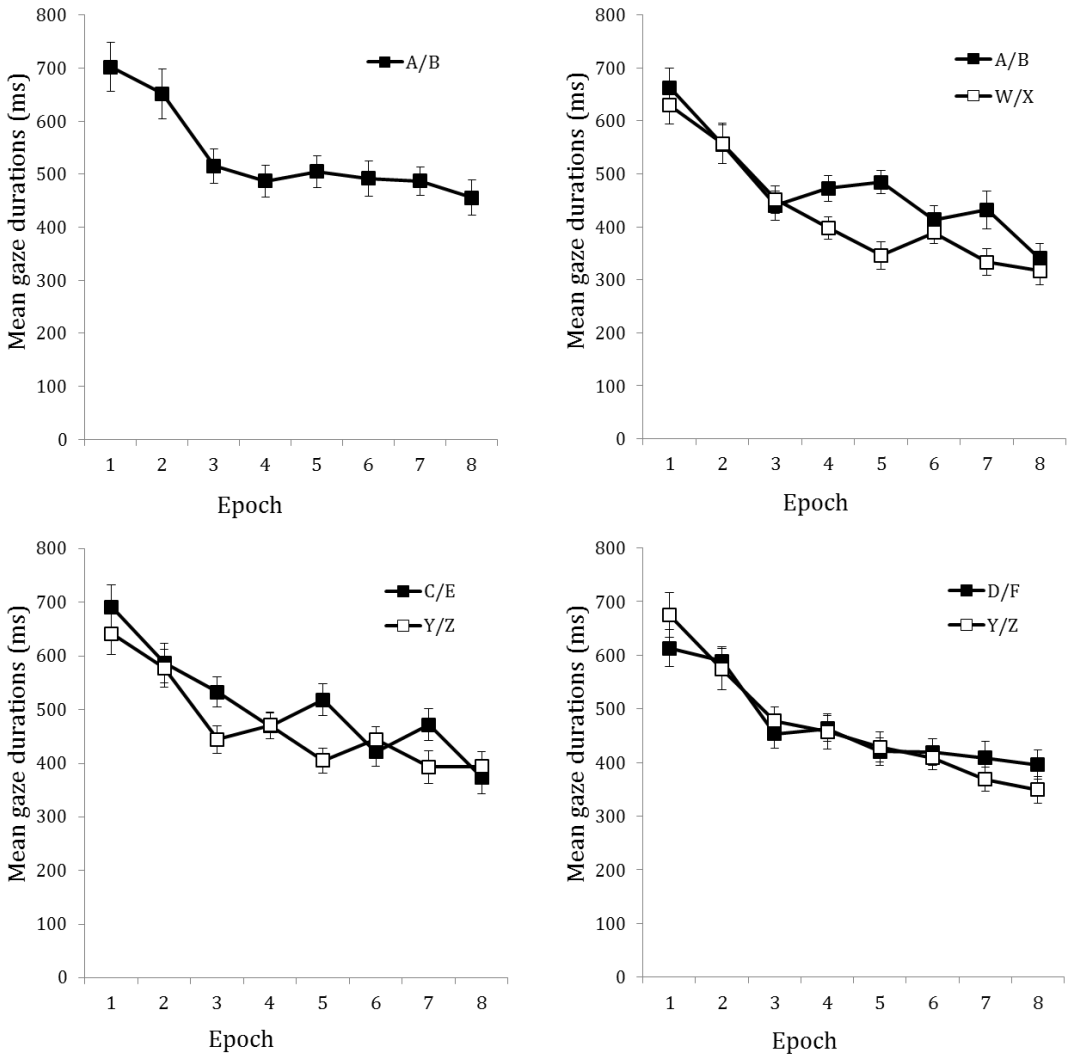

Figure 2.9. Gaze durations for cues within each trial type during Stage 1 in Experiment 2, averaged for equivalent cues (ms; ±SEM). Higher left panel shows gaze durations for A and for B on A+/B+ trials. Higher right panel shows gaze durations for A/B and for W/X on AW+/BX+ trials. Lower left panel shows gaze durations for C/E and for Y/Z on CY+/EZ+ trials. Lower right panel shows gaze durations for D/F and for Y/Z on CY-/DZ trials.

A two-way (cue [C/E, Y/Z] by epoch [1-8]) ANOVA on CY+/EZ+ trials revealed a significant effect of cue, $F(1, 58) = 4.21$, $p = .045$, $\eta_P^2 = .07$, indicating longer gaze durations for C/E than for Y/Z. The effect of epoch was also significant, $F(4.27, 247.47) = 11.92$, $p < .001$, $\eta_P^2 = .17$, and there was a non-significant trend for the interaction, $F(5.73, 332.19) = 2.04$, $p = .063$, $\eta_P^2 = .03$ (lower left panel). Simple main effects analyses indicated significantly longer gaze durations for C/E than for Y/Z on epochs three and five, *t*s ≥ 2.52, *p*s ≤ .014, *r*s ≥ .31, and a non-significant trend for longer gaze durations for C/E than for Y/Z on epoch seven, $t(58) = 1.88$, $p = .065$, $r = .24$; gaze-duration differences between the cues were not significant on the other epochs, *t*s ≤ 1.29, *p*s ≥ .202, *r*s ≤ .17. A two-way (cue [D/F, Y/Z] by epoch [1-8]) ANOVA on DY-/FZ- trials revealed no significant effect of cue, $F(1, 58) = .04$, $p = .852$, $\eta_P^2 \leq .01$, $BF_{01} = 13.7$, a significant effect of epoch, $F(4.57, 264.76) = 14.59$, $p < .001$, $\eta_P^2 = .2$, and no significant interaction, $F(7, 406) = .91$, $p = .485$, $\eta_P^2 = .02$ (lower right panel).

The left panel of Figure 2.10 shows uncorrected gaze durations for W/X on AW+/BX+ trials and Y/Z averaged across CY+/EZ+/DY-/EZ- trials. The right panel of Figure 2.10 shows gaze durations for these cues expressed as proportions of trial length.

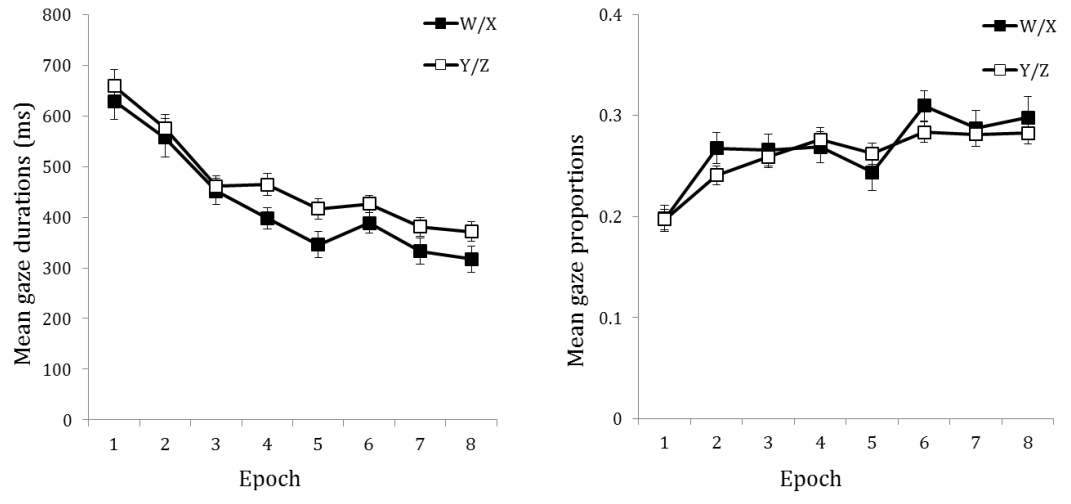

Figure 2.10. Gaze durations for W/X and for Y/Z across the eight epochs of Stage 1 in Experiment 2 (±SEM). The left panel shows the uncorrected-gaze durations for W/X and for Y/Z. The right panel shows gaze durations for W/X and for Y/Z expressed as proportions of trial length.

To compare gaze durations between W/X and Y/Z in Stage 1, a two-way (cue [W/X, Y/Z] by epoch [1-8]) ANOVA was conducted. This revealed a significant effect of cue, $F(1, 58) = 5.76$, $p = .02$, $\eta_P^2 = .09$, indicating longer gaze durations for Y/Z than W/X. The effect of epoch was also significant, $F(3.53, 204.67) = 25.08$, $p < .001$, $\eta_P^2 = .3$. The interaction was not significant, $F(4.58, 265.56) = .4$, $p = .831$, $\eta_P^2 = .01$. Therefore longer uncorrected-gaze durations for uncorrelated cues than blocked cues found in Experiment 1 were replicated.

In this experiment, longer reaction times on trials containing Y/Z than W/X reached significance, $t(58) = 3.36$, $p = .001$, $r = .4$. The reaction times on each type of trial in Stage 1 are presented in Figure 2.11.

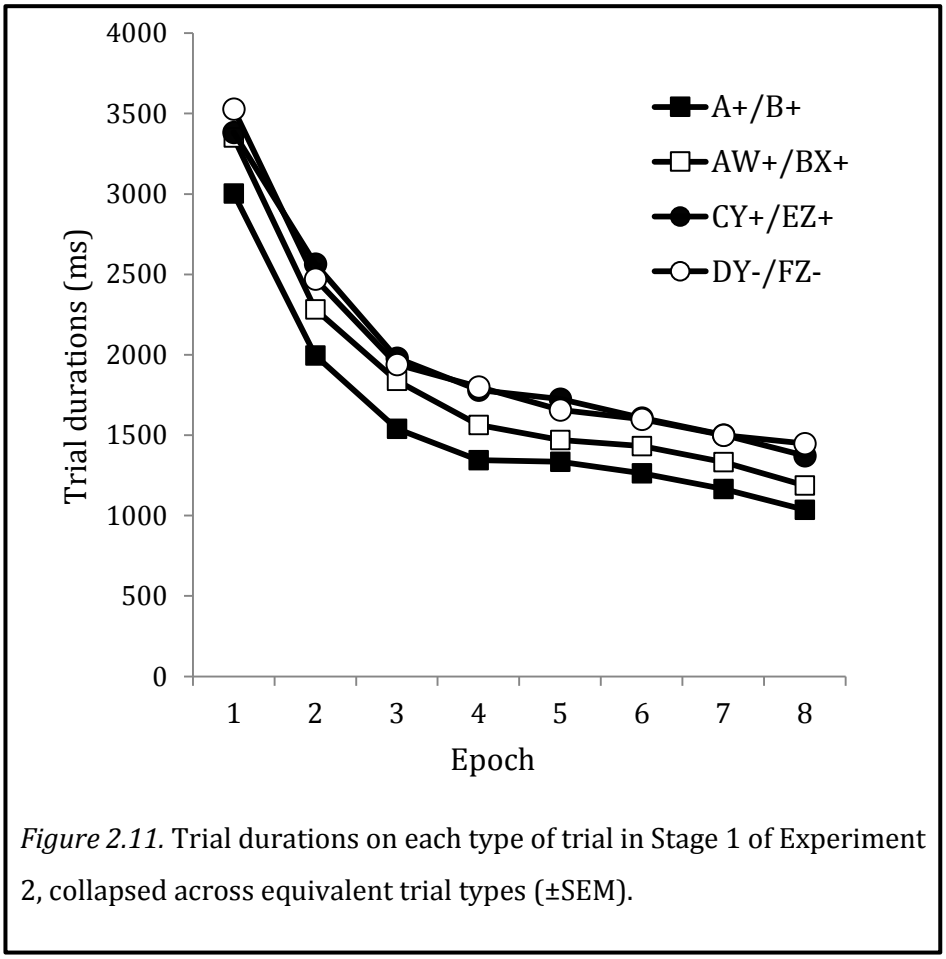

Figure 2.11. Trial durations on each type of trial in Stage 1 of Experiment 2, collapsed across equivalent trial types (±SEM).

A comparison for gaze durations between W/X and Y/Z expressed as proportions of trial durations was assessed by a two-way (cue [W/X, Y/Z] by epoch [1-8]) ANOVA. This indicated that the effect of epoch was significant, $F(4.87, 282.65) = 7.46$, $p < .001$, $\eta_P^2 = .11$, but the effect of cue was not, $F(1, 58) = .62$, $p = .436$, $\eta_P^2 = .01$, nor was the interaction, $F(5.4, 313) = .69$, $p = .643$, $\eta_P^2 = .01$. Once again, a Bayesian t-test indicated strong support for no differences between proportions of gaze durations for W/X and Y/Z averaged across Stage-1 epochs, $t(58) = .79$, $p = .436$, $r = .1$, $BF_{01} = 7.24$.

Similarly to Experiment 1, the correlation between gaze-duration bias for uncorrelated cues (gaze durations for Y/Z – gaze durations for W/X) and the magnitude of the redundancy effect (W/X ratings – Y/Z ratings) was not significant, $r(59) = -.166$, $p = .208$, nor was the corresponding correlation using gaze-duration values corrected for trial length, $r(59) = -.058$, $p = .66$.

Stage 2. In Stage 2, I compared gaze durations between W/X and Y/Z directly, while they were presented on the screen at the same time. These were approximately equivalent throughout the six epochs of Stage 2, as shown in Figure 2.12. A two-way (cue [W/X, Y/Z] by epoch [1-6]) ANOVA revealed a significant effect of epoch, $F(3.85, 223.14) = 5.25$, $p = .001$, $\eta_P^2 = .08$, no significant effect of cue, $F(1, 58) = .99$, $p = .323$, $\eta_P^2 = .02$, and no significant interaction, $F(5, 290) = .37$, $p = .871$, $\eta_P^2 = .01$. A Bayesian t-test for the difference between gaze durations for W/X and Y/Z averaged across Stage-2 epochs indicated support for the null result, $t(58) = 1$, $p = .323$, $r = .13$, $BF_{01} = 6.02$.

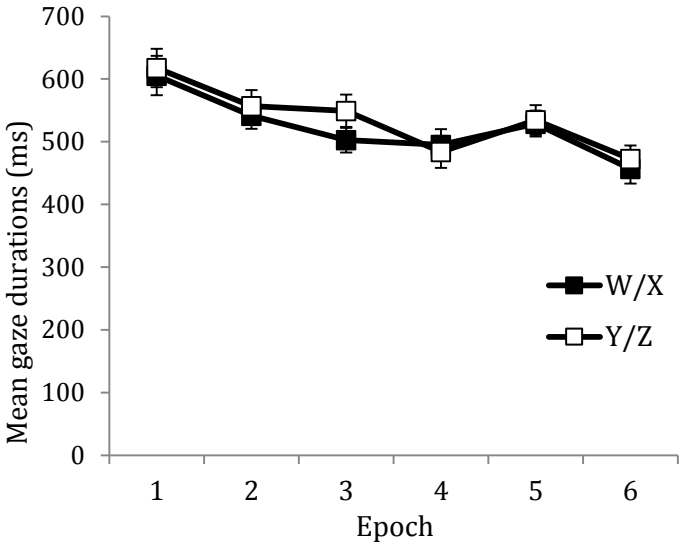

Figure 2.12. Mean gaze durations for W/X and Y/Z across the six epochs for Stage 2 in Experiment 2 (±SEM).

Proportions of stomach-ache predictions in Stage 2 on WY-, WZ+, and XY+ trials are presented in Figure 2.13. As can be seen from the figure, learning on WZ+ and on XY+ trials proceeded at a similar rate. A two-way (trial type [WZ+, XY+] by epoch [1-6]) ANOVA revealed a significant effect of epoch, $F(3.42, 198.53) = 29.24$, $p < .001$, $\eta_P^2 = .34$, no significant effect of trial type, $F(1, 58) = .01$, $p = .905$, $\eta_P^2 < .001$, and no significant interaction, $F(4.23, 245.14) = 1.16$, $p = .331$, $\eta_P^2 = .02$. A Bayesian t-test for the difference between responses on WZ+ and XY+ trials averaged across epochs indicated support for the null hypothesis, $t(58) = .12$, $p = .905$, $r = .02$, $BF_{01} = 6.98$.

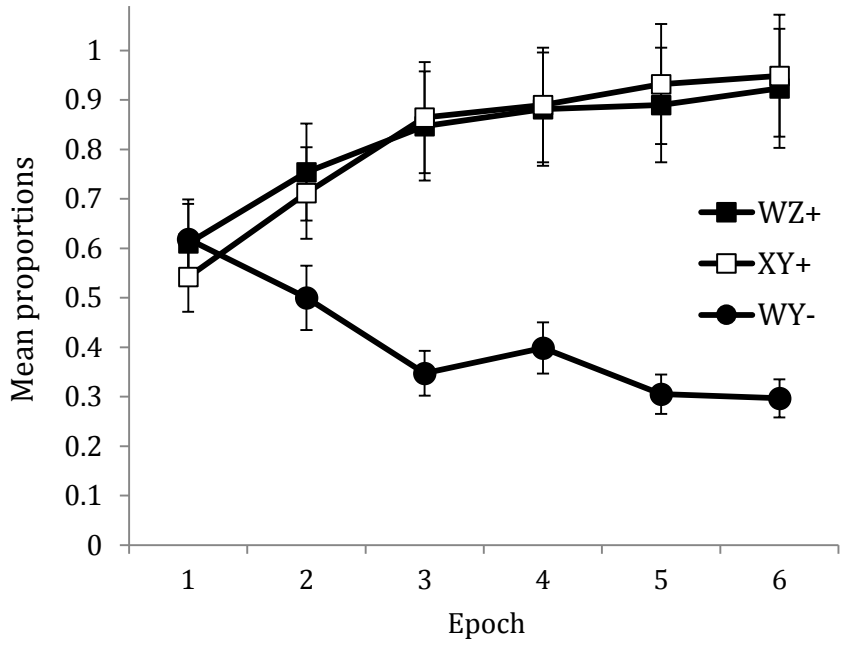

Figure 2.13. Proportions of stomach-ache predictions across the six epochs of Stage 2 in Experiment 2 (±SEM). This figure shows mean proportions of stomach-ache predictions for the three key trial types each consisting of one blocked (W, X) and one uncorrelated cue (Y, Z).

Discussion

In this chapter I set out to investigate whether the redundancy effect could be a consequence of differences in attention for blocked and uncorrelated cues. In two experiments I used eye-tracking to measure gaze durations for these cues during learning. I found that although participants spent more time looking at uncorrelated cues than at blocked cues, when these gaze durations were corrected for trial length, this difference disappeared. This was found in both experiments. I also considered the possibility that longer uncorrected-gaze durations for uncorrelated cues than for blocked cues could have contributed to the redundancy effect. This was indicated to be unlikely for two reasons. Firstly, I failed to obtain any differences in gaze durations for these cues when they were presented on the screen at the same time in Stage 2 of Experiment 2. This shows that

when conditions of presentation were matched, differences in gaze durations for blocked and uncorrelated cues were not obtained. This suggested no differences in selective attention for these cues in this experiment. Secondly, if longer gaze durations for uncorrelated cues than blocked cues contributed to the redundancy effect, then the magnitude of the redundancy effect would be related to this gaze-duration bias. To see whether this was the case I computed correlations between gaze duration bias for uncorrelated cues (gaze durations for uncorrelated cues – gaze durations for blocked cues) and the magnitude of the redundancy effect (ratings for blocked cues – ratings for uncorrelated cues). These correlations were not significant.

In addition to this, I did not obtain any differences in further learning about these cues in Stage 2 of Experiment 2. Participants learned the discrimination reliant on blocked cues (XY+, WY-) at a similar rate as the discrimination reliant on uncorrelated cues (WZ+, WY-). However, it is possible that this procedure was not sensitive enough to differences in learning rates between cues with different levels of predictiveness, including between blocked and uncorrelated cues. No suitable control condition was present in this experiment to show that this was not the case (e. g. with trials AY-, AZ+, BY+ presented in Stage 2. Therefore evidence from this part of Experiment 2 is inconclusive.

Based on this evidence, it is unlikely that differences in attention between blocked and uncorrelated cues constitute the basis of the redundancy effect. While it may have been possible that these differences were too slight to reveal using eye-tracking in these experiments, differences in gaze durations were detected between informative and redundant cues. In particular A was looked at for longer than X and B was looked at for longer than Y, indicating that this was not the case.

It is also interesting to consider why there were no significant differences in gaze durations between C and Y on CY- trials (and equivalent cues in Experiment 2). One possibility is that gaze durations were longer for the cues that participants thought caused the outcome. Participants looked at A for longer than at X on AX+ trials and at B for longer

than at Y on BY+ trials. Cues A and B received higher ratings than C, X, and Y, indicating that participants thought A and B caused the outcome. The only difference between B and C was that the outcome occurred on trials containing B but not C. Therefore no differences in gaze durations between C and Y on CY- trials may have been due to participants thinking that neither C nor Y caused the outcome. This may also suggest that participants treated stomach-ache and no stomach-ache trials asymmetrically in these tasks. I will return to this issue later.

Chapter 3: Uncertainty about the blocked cue

Findings in Chapter 2 indicated that differences in attention between blocked and uncorrelated cues were unlikely to have been responsible for the redundancy effect. In this chapter, I explored whether participants' uncertainty about the relationship between the blocked cue and the outcome contributed to the redundancy effect. In Uengoer et al.'s (2013) studies and in the experiments from the previous chapter, the blocked cue X received ratings at approximately in the middle of the 11-point, 0-10 causal-rating scale (M = 5.33 in Uengoer et al.'s experiments and M = 4.55 in Experiments 1 and 2 of this thesis). It is ambiguous however, what this rating represented. It could have reflected the fact that on average, participants thought X was a weak cause of the outcome. Alternatively, ratings in the middle of the scale may have best reflected the point of uncertainty on this scale. Therefore ratings for X could have indicated that participants were uncertain about whether or not X caused the outcome. Participants' uncertainty about whether or not X caused the outcome falls outside the predictions of single and summed error-term models. However this prediction has been made previously, including by causal models (e. g. Cheng, 1997; Cheng & Holyoak, 1995; Waldmann & Holyoak, 1992) and the propositional account of learning proposed by Lovibond (2003; see also Beckers, Miller, De Houwer & Urushihara, 2006; De Houwer, 2009; Mitchell et al., 2009). Waldmann and Holyoak noted that because the blocked cue X on A+/AX+ trials is never presented by itself, participants may be unsure about its effects on the outcome. This particularly applies to tasks in which the outcome is at maximal intensity as participants are unable to verify whether on A+/AX+ trials X has any causal effects on the outcome in addition to A. Therefore X may or may not cause the outcome. Evidence of participants' uncertainty about the blocked cue can be illustrated by Vandorpe and De Houwer (2006). In their study participants were presented with A+/AX+/B-/BY+ trials. In this design X was the blocked cue and Y was a release-from-overshadowing cue. The main difference between X and Y in this design was that the companion cue for Y did not cause the outcome by itself (B-), while the companion cue for X did (A+). Vandorpe and De Houwer argued that

because there was a ceiling on the outcome in their study, participants should have been uncertain about whether or not X caused the outcome. On the other hand, participants could have inferred that Y was a cause even with the ceiling on the outcome, because B was shown to have not caused the outcome. Under the assumption that participants would seek information to reduce uncertainty, Vandorpe and De Houwer gave participants a choice of revealing the outcome associated with either X or Y. They found that a greater number of participants chose X than Y to be revealed, consistently with predictions. However when the ceiling on the outcome was removed by showing the outcome was at sub-maximal intensity (allergist reaction: 10/20), a greater number of participants chose Y than X to be revealed.

Strong evidence that participants are uncertain about the causal status of the blocked cue also comes from experiments using outcome-additivity training (Lovibond et al., 2003; Mitchell & Lovibond, 2002). Mitchell and Lovibond suggested that participants' uncertainty about the blocked cue is mediated by their beliefs about how causal cues interact. Similarly to Waldmann and Holyoak (1992), they argued that when the outcome is at ceiling, participants will be uncertain about whether or not X caused the outcome. However if participants believed that two causes can add to produce a stronger outcome (e. g. G causes a stomach ache, H causes a stomach ache, and G and H together cause a strong stomach ache [G+/H+/GH++]), their uncertainty about the causal status of X would be reduced. Since the outcome that occurred on A+/AX+ trials is of the same magnitude, they can conclude that X was not a cause. If it was, a stronger outcome would have occurred on AX+ trials relative to A+ trials. Lovibond et al. experimentally manipulated assumptions about how two causal cues interacted in two groups of participants. In Group Additive, participants were shown that two individual causes (G+/H+) led to a stronger outcome when they were presented together (GH++). In Group Non-additive participants were shown that two individual causes (G+/H+) led to an outcome of the same magnitude when presented together (GH+). Lovibond et al. found that participants were able to learn and generalise the outcome-additivity rule to the blocked cue X; participants in Group

Additive had lower ratings for X and higher blocking than participants in Group Non-additive. Results showing that outcome-additivity training reduced ratings for X and increased blocking have been widely replicated (e. g. Beckers et al., 2005; Livesey & Boakes, 2004; Lovibond et al., 2003; Mitchell & Lovibond, 2002; Mitchell et al., 2005).

There was a ceiling on the outcome in Uengoer et al.'s (2013) experiments and in Experiments 1 and 2 of this thesis. Therefore, it is possible that average ratings for X at approximately in the middle of the rating scale reflected participants' uncertainty about the relationship between X and the outcome. Therefore, rather than indicating a stronger relationship between X and the outcome than Y and the outcome as noted by Uengoer et al., the redundancy effect could have been due to differences in certainty about the effects that X and Y had on the outcome. Participants may have been uncertain about whether or not X caused the outcome, rating it in the middle of the scale, but they were certain that Y did not cause the outcome, rating it low on the scale. If this was the case, asking participants to estimate how confident they are in their causal ratings for X and Y should reveal greater confidence for Y than for X. This is what I investigated in Experiment 3. In this experiment, in addition to providing causal ratings, participants were asked to indicate how confident they were about them. Given that Y was presented with the outcome 50% of the time, and, like X, it was never presented on its own, it is possible that participants would have been uncertain about both. If participants were less certain about their causal ratings for X than for Y however, this would indicate that uncertainty about the causal status of X could have contributed to the redundancy effect. In this case, I would expect that using manipulations to disambiguate the effects of X on the outcome would reduce the redundancy effect. To pre-emptively disclose the findings of Experiment 3, participants were less certain about their causal ratings for X than for Y. Therefore in Experiment 4, I used outcome-additivity training, as used by Lovibond et al. (2003) to explore whether the redundancy effect would be smaller in Group Additive than Group Non-additive.

If participants were uncertain about the causal status of X, it is possible that they used other task-specific information when judging the likelihood with which this cue led to the outcome. One suggestion provided by Livesey, Lee and Shone (2013) is that in the absence of disambiguating information, the probability that a cue causes an outcome is determined by outcome rate, or the proportion of trials that lead to the outcome in a particular task. Informally, this prediction can be illustrated as follows. In a task in which many trial types lead to the outcome, participants may expect that cues which they are uncertain about would also lead to the outcome. In contrast, in a task in which very few trial types lead to the outcome, participants may expect cues which they are uncertain about would not lead to the outcome either. Livesey et al. proposed that because participants are particularly uncertain about the blocked cue, the probability of the outcome for this cue will correspond to outcome rate. As outcome rate approaches 1, the probability of the outcome for the blocked cue will also approach 1. As outcome rate approaches zero, the probability of the outcome for the blocked cue will also approach zero. I am not aware of any prior experimental demonstrations of this prediction. Notably, in Uengoer et al.'s (2013) experiments, the outcome rate was high (71.43% in Experiment 1, 75% in Experiment 2, and 66.67% in Experiment 3). This was also the case for experiments in Chapter 2 of this thesis (75% in Experiments 1 and 2). Therefore, high outcome rate may have contributed to higher ratings for X than for Y in these experiments. In Experiment 5, I manipulated outcome rate to test Livesey et al.'s (2013) predictions that variations in outcome rate would affect causal ratings for X. In Experiment 6, I varied outcome rate in order to see whether it would affect the redundancy effect.

Experiment 3

In order to establish whether participants' uncertainty about the causal status of X contributed to the redundancy effect, it was important to first determine whether participants were less certain about their causal ratings for X than for Y. The design of

Experiment 3 is presented in Table 3.1. Since there has been only one prior demonstration of the redundancy effect together with blocking (Experiment 1 of Uengoer et al., 2013), in this experiment I included the blocking-control cues P/Q which received PQ+ training, in order to see whether blocking could be observed in tandem with the redundancy effect. This would constitute an independent replication of these effects in the same experiment and provide further evidence that the redundancy effect was not due to a failure of blocking. I also included a compound of cues which did not lead to the outcome (EF-) to mirror the PQ+ trials. This allowed me to determine whether participants were more or less certain about their causal ratings for a compound of cues which lead to the outcome (PQ+) than for a compound of cues which did not lead to the outcome (EF-). In this experiment, subsequent to Stage-1 trials, participants were asked to provide ratings for how likely each cue was to cause the outcome individually (causal ratings), similarly to Experiments 1 and 2. Following this, participants were asked to indicate how confident they were that these ratings were accurate (confidence ratings). I hypothesised that causal ratings for X would be higher than for Y (the redundancy effect), but lower than for blocking-control cues P/Q (blocking). If participants were less certain about X than about Y, confidence ratings for X would be lower than for Y.

Method

Participants. Twenty-one Psychology undergraduate students at Plymouth University took part in this experiment. They were 18-39 years old (M = 20.71, SD = 4.37) and seven were male. They were tested in individual cubicles and received course credit for their participation.

Materials. The materials and procedure for this experiment were identical to Experiment 2 unless otherwise stated.

The stimuli were 10 images of foods. The foods were: apple, banana, cherries, grapes, kiwi, mango, orange, pineapple, strawberries, and watermelon. The foods were randomly assigned to each type of cue (A, B, C, D, E, F, P, Q, X, Y) for each participant.

Table 3.1

The design of Experiment 3.

Stage 1	Test
A+	A
AX+	B
BY+	C
CY-	D
D-	E
EF-	F
PQ+	P
	Q
	X
	Y
x 12	x 2

Procedure. In Stage 1 participants were presented with 12 blocks of trials, in each of which the seven trial types (A+, AX+, BY+, CY+, D-, EF-, PQ+) appeared once in a random order with no repetitions. Each trial started with the presentation of either one or two images of foods in the top half of the screen, below the phrase "The patient ate the following foods:" (or "The patient ate the following food:" for trials with single images). For trials with two images, one was located on the left and one on the right, while images of single foods were located in the middle.

After all blocks in Stage 1 were completed, the test stage began. On each trial, the sentence "What is the probability that the food causes stomachache?" 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*). Below the rating scale the sentence "How confident are you that this rating is accurate?" was presented, together with another 11-point rating scale from 0 (*Not at all*) to 10 (*Extremely*). Both rating scales were oriented

horizontally. Each cue that appeared in Stage 1 was presented twice, in a random order, again with the constraint of no successive repetitions of the same trial type. For each participant, an average of the two causal ratings was calculated and used in the analyses. This was also the case for the confidence ratings. Ratings were also averaged between equivalently-treated cues P/Q and E/F.

Results

Stage 1. For Stage-1 data, an epoch was defined as an average response of two trials of the same trial type calculated for each participant. The proportions of stomach ache predictions for the six epochs of Stage 1 are shown in Figure 3.1. This figure indicates that participants were able to learn the contingencies. They responded accurately on 99.64% (*SD* = 8.16%) of the trials in the final epoch.

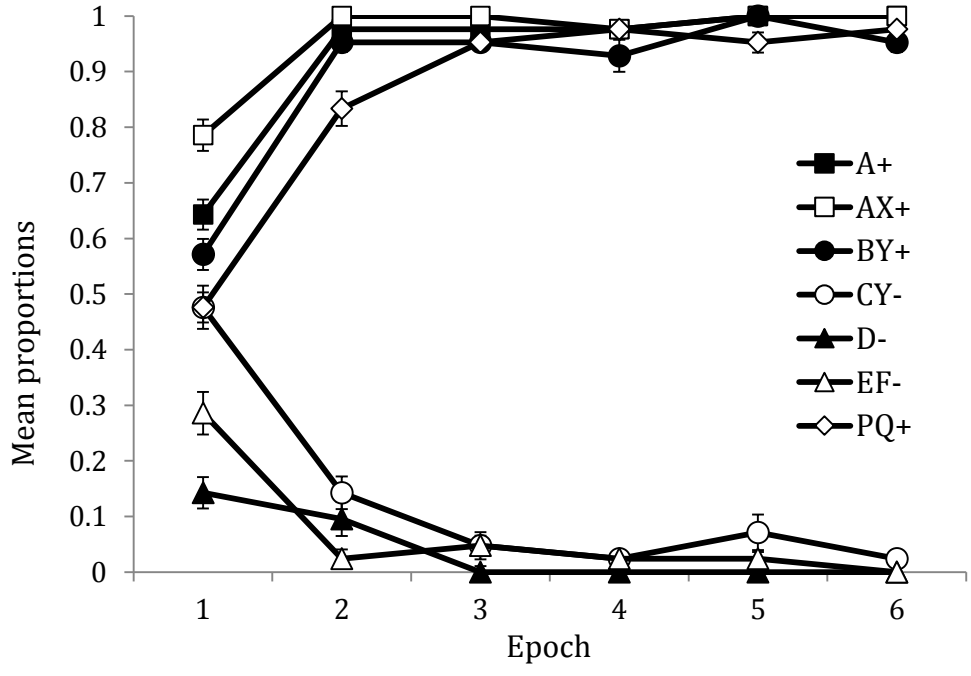

Figure 3.1. Mean proportions of stomach-ache predictions for Stage 1 in Experiment 3 (±SEM).

Causal ratings. Causal and confidence ratings for the different cues at test are presented in Figure 3.2. A one-way (cue [A, B, C, D, E/F, P/Q, X, Y]) ANOVA on the causal ratings revealed a significant effect of cue, $F(2.73, 54.65) = 132.09$, $p < .001$, $\eta^2 = .87$. Bonferroni-corrected paired comparisons indicated that A and B had higher ratings than the other cues, $ts \geq 11.04$, $ps < .001$, $rs \geq .93$, but did not differ from each other, $t(20) = 2.36$, $p = .801$, $r = .47$. Cues C, D, E/F, and Y had significantly lower ratings than the other cues, $ts \geq 4.23$, $ps \leq .012$, $rs \geq .69$, but did not differ from each other, $ts \leq 1.27$, $ps \geq .772$, $rs \leq .27$. Cues P/Q and X had intermediate ratings that were significantly different from all other cues, $ts \geq 4.23$, $ps \leq .012$, $rs \geq .69$.

Importantly, X received higher causal ratings than Y, as indicated by a t-test $t(20) = 4.61$, $p < .001$, $r = .72$, confirming that the redundancy effect was obtained. Cues P/Q had significantly higher ratings than X, confirming that blocking was also demonstrated, $t(20) = 2.67$, $p = .015$, $r = .51$.

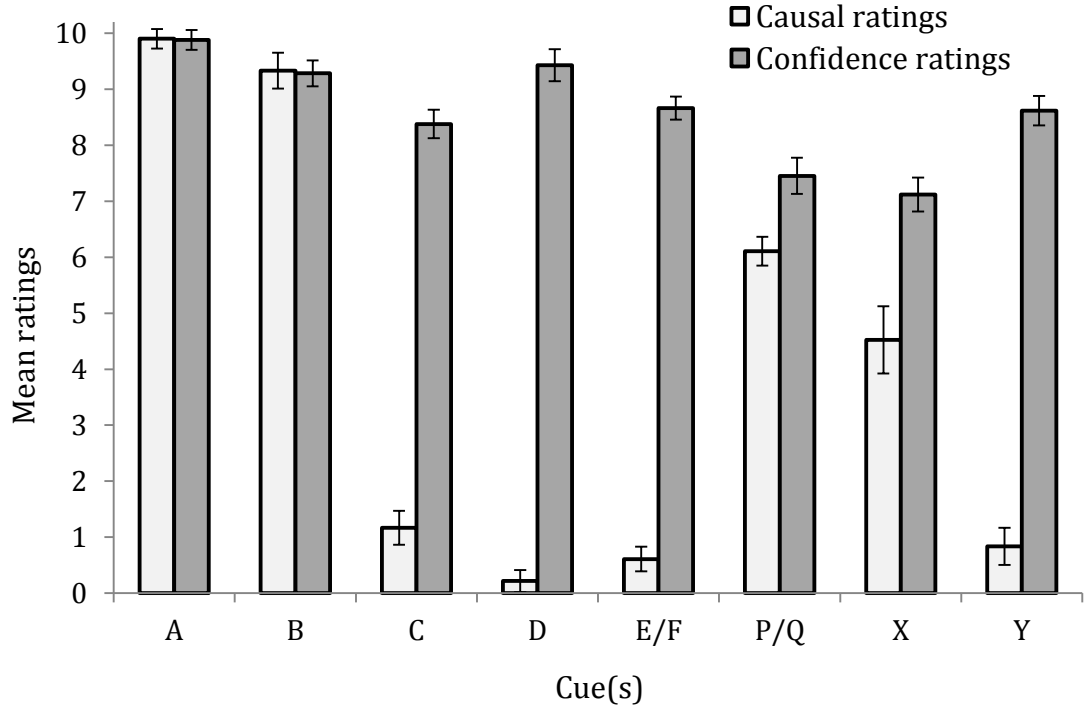

Figure 3.2. Mean causal and confidence ratings for the cues at test in Experiment 3 (±SEM).

Confidence ratings. A one-way (cue [A, B, C, D, E, F, P, Q, X, Y]) ANOVA on the confidence ratings revealed a significant effect of cue, $F(4.51, 90.22) = 11.78$, $p < .001$, $\eta^2 = .37$. Bonferroni-corrected paired comparisons indicated that participants were less confident in their ratings for P/Q and X than A, B, and D, ts ≥ 3.99, ps ≤ .02, rs ≥ .67. While A had higher confidence ratings than C, confidence ratings for C, E/F, and Y were intermediate and did not differ significantly from the other cues, ts ≤ 3.45, ps ≥ .07, rs ≤ .61. Importantly, participants were less confident about their causal ratings for X than for Y, $t(20) = 3.18$, $p = .005$, $r = .58$. There were no differences between confidence ratings for P/Q and for X, $t(20) = .82$, $p = .423$, $r = .18$, $BF_{01} = 3.26$. Participants were also more confident in their causal ratings for E/F than for P/Q, $t(20) = 2.84$, $p = .01$, $r = .54$.

To summarise, the redundancy effect was obtained in Experiment 3: ratings for X were higher than for Y. Blocking was also demonstrated, as ratings for X were lower than for P/Q. In addition, participants indicated less confidence in their causal ratings for X than for Y. Therefore the redundancy effect could have been due to participants' higher uncertainty about the causal status of X than of Y. I failed to observe any differences between confidence ratings for X and for P/Q however. This difference may have been expected because, arguably, more information was available for P/Q than for X. Given a compound of two cues leading to the outcome (PQ+) at least one of these cues would have been a cause. On the other hand, X may or may not have been a cause of the outcome. However, it is also worth considering that while more information was available for the cue compound PQ+ collectively, individually each cue may or may not have caused the outcome, similarly to X. This could have resulted in this lack of difference.

Confidence ratings were higher for E/F, which did not lead to the outcome than for P/Q, which led to the outcome in Stage 1. This indicated asymmetry in the way that participants treated trials on which the outcome occurred and trials on which the outcome did not occur, rather than viewing them as equal and opposite; these cues were otherwise equivalent.

Experiment 4

Given that participants were less certain about their causal ratings for X than for Y in Experiment 3, uncertainty about the effects of X may have contributed to the redundancy effect. If this was the case, then manipulations aimed at reducing participants' uncertainty about the causal status of X may influence the redundancy effect. One manipulation that has been used to reduce participants' uncertainty about whether or not X caused the outcome is outcome-additivity training (Lovibond et al., 2003; Mitchell & Lovibond, 2002). As discussed in the introduction to this chapter, in Group Additive, Lovibond et al., showed participants that two causal cues led to the outcome of a greater magnitude than each individual cue (G+/H+/GH++). Participants in this group were able to extract and apply the additivity rule to X; ratings for this cue were lower than in Group Non-additive, in which two causal cues led to the same outcome (G+/H+/GH+). Therefore if outcome-additivity training was used with the design A+/AX+/BY+/CY- it would be expected that ratings for X would be lower in Group Additive than in Group Non-additive. As a result, the redundancy effect should also be smaller in Group Additive than Group Non-additive. Since participants indicated that they were more certain about their causal ratings for Y than for X in Experiment 3, outcome-additivity training should not affect ratings for Y.

Therefore, in Experiment 4 I used outcome-additivity training to see whether it would reduce ratings for X but not for Y, and result in a smaller redundancy effect in Group Additive than in Group Non-additive. This would constitute good evidence that participants' uncertainty about the causal status of X contributed to the redundancy effect. The design of Experiment 4 is presented in Table 3.2. In Stage 1 two cues were established as causes of the outcome (G+/H+) and when presented together led to a stronger outcome (GH++) in Group Additive. Participants in Group Non-additive saw two causes (G+/H+) which led to the same outcome when presented together (GH+). To make sure that participants in Group Non-additive experienced a stronger outcome, they were presented with cue I++; the same cue led to a weaker outcome in Group Additive (I+). In Stage 2

participants were presented with a design containing X, Y, and blocking-control cues P/Q (A+/AX+/BY+/CY-/PQ+). To determine whether additive training successfully reduced causal ratings for X, resulted in larger blocking and a smaller redundancy effect in Group Additive than Group Non-additive, ratings for X were compared with ratings for P/Q and Y between the groups.

Table 3.2

The design of Experiment 4.

Stage 1	Stage 2	Test 2
Group Additive	A+	A
G+	AX+	B
H+	BY+	C
GH++	CY-	P
I+	PQ+	Q
J-		X
		Y
Group Non-additive		
G+		
H+		
GH+		
I++		
J-		
x 8	x 8	x 2

Method

Participants. Participants were 64 Plymouth University students studying Psychology (12 male). There were 31 participants in Group Additive and 33 participants in Group Non-additive.

Materials. The materials and procedure for Experiment 4 were the same as in the previous experiments, unless otherwise stated.

The cues were 11 images of foods: apple, banana, broccoli, cabbage, cherries, coconut, corn, grapes, kiwi, lemon, and orange. The foods were randomly assigned to each type of cue (A, B, C, D, G, H, I, J, P, Q, X, Y) for each participant.

Procedure. In this experiment, during the learning stages, instead of a binary outcome, participants were asked to predict the strength of a stomach ache that the patient experienced. The initial instructions were the same as in Experiment 3, incorporating changes in the type of response.

In Stage 1 participants were presented with eight blocks of trials, in each of which the five trial types (Group Additive: G+/H+/GH++/I+/J-; Group Non-additive: G+/H+/GH+/I++/J-) appeared once. Participants were presented with either one or two foods on the screen and asked to predict the strength of the stomach ache that the patient experienced. On each trial they were asked to make a rating on an 11-point scale ranging from 0 (*No stomach ache*) to 10 (*Severe stomach ache*). After making the response, they were presented with feedback showing the strength of stomach ache experienced. The outcome consisted of three levels: no stomach ache (-), a mild stomach ache (+), and a strong stomach ache (++). The image of the outcome was an arrow pointing left, middle or right within a semi-circle dial. The left side of the dial was coloured green, the right side was coloured red and the middle was a blend of green and red. The arrow representing no stomach ache pointed left within the dial, the arrow representing a mild stomach ache pointed in the middle of the dial and the arrow representing a strong stomach ache pointed right within the dial.

After Stage-1 training participants were presented with a screen summarising the contingencies from Stage 1 in order to ensure adequate learning for all of the participants. On the top of the screen text “These are the relationships you should have learned:” was presented. Images of the cues were located on the left side of the screen (G and H at the

top, I in the middle, and H at the bottom). To the right of the images the relationship with the outcome was stated. Participants in Group Additive were shown that G and H "Each caused mild stomach ache, and when eaten together they caused severe stomachache", I "Caused mild stomach ache", and H "Caused no stomach ache". Participants in Group Non-additive were shown that G and H "Each caused mild stomach ache, and when eaten together they caused mild stomachache", I "Caused a severe stomach ache", and H "Caused no stomach ache". This information remained on the screen until participants clicked a button to continue.

Instructions to indicate the commencement of Stage-2 trials were as follows:

You will now see a further series of meals, featuring different foods. Just as before, you will make predictions about the strength of stomach ache that you expect, and you will be given feedback on each trial. Please click the mouse to continue.

Stage 2 consisted of eight blocks of trials with each of the five trial types (A+/AX+/BY+/CY-/PQ+) presented once per block. The outcomes in this stage were no stomach ache (-) and mild stomach ache (+). Participants made their predictions on the same rating scale as in Stage 1.

At test participants were asked "What strength stomach ache do you expect?" for each individual cue, twice. They made their ratings on the same 11-point scale used in the learning stages. For each participant, the average of the two stomach ache-strength ratings was calculated and used in the analyses. Ratings were also averaged between equivalently-treated cues P/Q.

Results

Stage 1. Participants learned the responses in Stage 1 well, responding correctly on 98.19% (*SD* = 9.33%) of the trials in the final epoch (Group Additive: *M* = 97.62%, *SD* = 10.45%; Group Non-additive: *M* = 98.72%, *SD* = 8.15%). These data are presented in the higher panels of Figure 3.3. A two-way (trial type [G+,H+,GH++/+, I+/++, J-] by group [Group Additive vs Group Non-additive]) ANOVA on the responses of the final epoch of

Stage 1 revealed a significant effect of trial type, $F(1.57, 97.1) = 1129.39$, $p < .001$, $\eta_p^2 = .95$, no significant effect of group, $F(1, 62) = .64$, $p = .425$, $\eta_p^2 = .01$, and a significant interaction, $F(1.57, 97.1) = 375.3$, $p < .001$, $\eta_p^2 = .86$. Bonferroni-corrected paired comparisons indicated that while ratings on G+, H+, and J- trials did not differ between the groups, *t*s ≤ 1.03, *p*s ≥ .306, *r*s ≤ .13, Group Additive had higher ratings on GH trials, $t(32.1) = 27.53$, $p < .001$, $r = .96$, and lower ratings on I trials, $t(62) = 28.53$, $p < .001$, $r = .96$, than Group Non-Additive, consistently with the manipulation.

Stage 2. Participants learned the contingencies in Stage 2, with 97.98% (*SD* = 10.49%) of participants responding correctly in the final epoch (Group Additive: *M* = 96.68%, *SD* = 13.77%; Group Non-additive: *M* = 99.21%, *SD* = 5.74%). Lower panels of Figure 3.3 show learning responses during Stage 2 for Group Additive (left panel) and for Group Non-additive (right panel).

A two-way (trial type [A+/AX+/BY+/CY-/PQ+] by group [Group Additive vs Group Non-additive]) ANOVA on the responses of the final epoch revealed a significant effect of trial type, $F(1.45, 89.64) = 883.24$, $p < .001$, $\eta_p^2 = .93$, a trend for the effect of group, $F(1, 62) = 3.5$, $p = .066$, $\eta_p^2 = .05$, and no significant interaction, $F(1.45, 89.64) = .15$, $p = .794$, $\eta_p^2 = .002$.

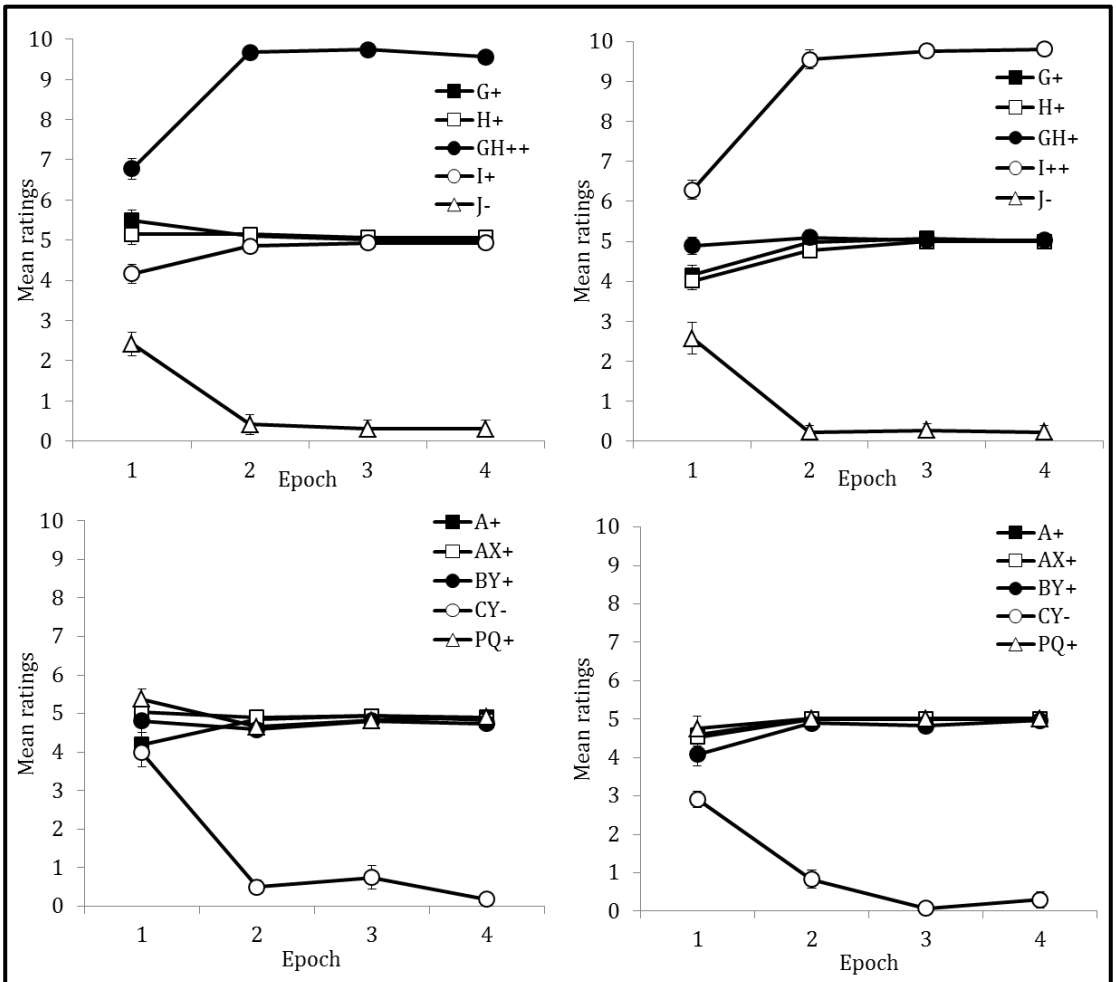

Figure 3.3. Mean strength of stomach-ache ratings for Stage 1 and Stage 2 in Experiment 4 (±between-subjects SEM). Higher left panel: Stage-1 responses in Group Additive. Higher right panel: Stage-1 responses in Group Non-additive. Lower left panel: Stage-2 responses in Group Additive. Lower right panel: Stage-2 responses in Group Non-additive.

Test. Causal ratings for each cue at test are presented in Figure 3.4. A two-way (cue [A, B, C, P/Q, X, Y] by group [Group Additive vs Group Non-additive]) ANOVA revealed a significant effect of cue, $F(3.46, 214.33) = 153.52$, $p < .001$, $\eta_p^2 = .71$. Bonferroni-corrected paired comparisons revealed that while ratings between A and B and between X and P/Q did not differ significantly, $ts \leq 1.9$, $ps \geq .64$, $rs \leq .23$, ratings between the other cues differed significantly from each other, $ts \geq 3.28$, $ps \leq .025$, $rs \geq .38$. There was no significant effect of group, $F(1, 62) = 2.1$, $p = .152$, $\eta_p^2 = .03$, but a significant interaction, $F(3.46,$

214.33) = 4.37, p = .001, η_p^2 = .07. Simple main effects analyses indicated that the groups did not differ in their ratings for A, B, C, and Y, ts ≤ 1.57, ps ≥ .126, rs ≤ .2, however Group Additive had significantly lower ratings for X than Group Non-additive, t(39.14) = 3.38, p = .002, r = .39. There was also a non-significant trend for lower ratings for P/Q in Group Additive than Group Non-additive, t(54.21) = 1.84, p = .072, r = .23.

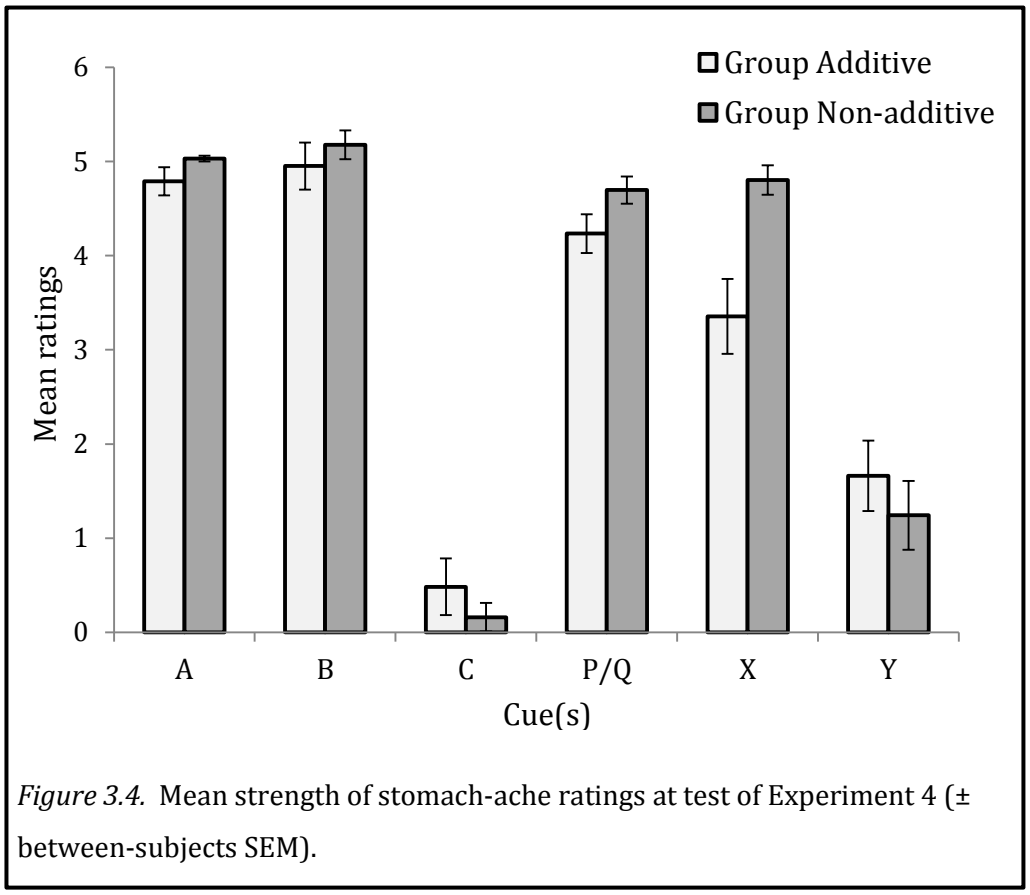

Figure 3.4. Mean strength of stomach-ache ratings at test of Experiment 4 (± between-subjects SEM).

To check whether blocking differed between the groups, a two-way (cue [P/Q, X] by group [Group Additive vs Group Non-additive]) ANOVA was conducted. It revealed a significant effect of cue, F(1, 62) = 4.28, p = .043, η_p^2 = .07, a significant effect of group, F(1, 62) = 10.89, p = .002, η_p^2 = .15, and a significant interaction, F(1, 62) = 6.96, p = .011, η_p^2 = .1. Simple main effects analyses indicated that ratings for X were lower in Group Additive than Group Non-additive, t(39.14) = 3.38, p = .002, r = .39, and there was a trend for lower

ratings for P/Q in Group Additive than Group Non-additive, $t(54.21) = 1.84$, $p = .072$, $r = .23$, as stated previously. Ratings for P/Q were higher than for X in Group Additive, $t(30) = 2.7$, $p = .011$, $r = .44$, (P/Q ratings – X ratings: $M = .88$, $SEM = .33$), but not in Group Non-additive, $t(32) = -.55$, $p = .587$, $r = .1$, $BF_{01} = 4.67$, (P/Q ratings – X ratings: $M = -.11$, $SEM = .19$).

Similar analysis was performed for the redundancy effect. A two-way (cue [X, Y] by group [Group Additive vs Group Non-additive]) ANOVA revealed a significant effect of cue, $F(1, 62) = 75.43$, $p < .001$, $\eta_p^2 = .55$, no significant effect of group, $F(1, 62) = 1.99$, $p = .163$, $\eta_p^2 = .03$, and a significant interaction, $F(1, 62) = 9.52$, $p = .003$, $\eta_p^2 = .13$. Simple main effects analyses indicated that the groups did not differ in their ratings for Y, $t(62) = .8$, $p = .427$, $r = .1$, but as mentioned previously, significantly lower ratings for X were observed in Group Additive than Group Non-additive, $t(39.14) = 3.38$, $p = .002$, $r = .39$. Consequently, the redundancy effect was smaller in Group Additive (X ratings – Y ratings: $M = 1.69$, $SEM = .42$) than in Group Non-additive (X ratings – Y ratings: $M = 3.56$, $SEM = .44$), however it was significant in both groups (Group Additive: $t(30) = 4.07$, $p < .001$, $r = .6$; Group Non-additive: $t(32) = 8.14$, $p < .001$, $r = .82$).

To summarise, in Experiment 4 I used outcome-additivity training, which aimed to resolve participants' uncertainty about the causal status of X. In Group Additive, participants were shown that two causal cues presented together led to a stronger outcome than either individual cue. In Group Non-additive, participants were shown that two causal cues presented together led to the outcome of the same strength as either individual cue. Participants were able to learn and generalise the additivity rule to X; Group Additive had lower causal ratings for X than Group Non-additive, consistently with Lovibond et al.'s (2003) findings. Although there was a trend for lower ratings for P/Q in Group Additive than in Group Non-additive, ratings for Y did not differ between the groups. As a consequence, blocking and the redundancy effect were reduced in Group Additive. This suggests that participants' uncertainty about the causal status of X

contributed to the redundancy effect. However, the redundancy effect was significant in both groups. Hence while outcome-additivity training was successful at reducing ratings for X and the redundancy effect, it was not enough to abolish it. Two possibilities for these data follow: either this manipulation was not strong enough to fully resolve uncertainty about the causal status of X, or there were other factors contributing to the redundancy effect independently of this manipulation.

In order to see whether the former could be the case, I considered whether a different manipulation would influence ratings for X. In Experiments 5 and 6, I varied outcome rate to explore whether this would affect ratings for X, blocking (Experiment 5), and the redundancy effect (Experiment 6).

Experiment 5

In Experiment 5, I aimed to test Livesey et al.'s (2013) predictions regarding the effects of outcome rate on ratings for X. They predicted that for uncertain cues outcome rate will be used to judge the probability with which they lead to the outcome. They argued that because the causal status of the blocked cue X is particularly uncertain, this cue should be more susceptible to variations in outcome rate than other cues. Livesey et al. predicted that in a task where many trial types led to the outcome, X would receive higher ratings than in a task in which only a few trial types led to the outcome. Evidence that participants were uncertain about whether or not X caused the outcome relative to Y was obtained in the previous experiments. Participants indicated less confidence about their causal ratings for X than for Y in Experiment 3. In addition, outcome-additivity training reduced the ratings for X but not for Y in Experiment 4. The prediction that causal ratings for X would correspond to outcome rate however, remained to be tested.

While Livesey et al. (2013) predicted that outcome rate would influence ratings for all uncertain cues, blocking control cues P/Q should be affected to a lesser extent than X. This

is because more information is available for these cues: if a compound of two cues leads to the outcome, at least one of them must be a cause. They predicted that when outcome rate approaches one, the probability of the outcome for P/Q should also approach one, and when outcome rate approaches zero, the probability of the outcome will approximate 0.5. As a result, blocking will be inversely related to outcome rate: when outcome rate is low, blocking will be large and when outcome rate is high, blocking will be small. If it was found that low outcome rate resulted in lower ratings for X than high outcome rate, this would establish variation in outcome rate as a viable technique for manipulating ratings of X. In addition, if low outcome rate reduced ratings for X to a greater extent than for P/Q, resulting in lower blocking, this would confirm Livesey et al.'s predictions regarding blocking. However it is also worth noting that in Experiment 3 participants were similarly uncertain about their causal ratings for X and for P/Q. As such, if uncertainty alone determines the extent to which these cues are affected by outcome rate, both X and P/Q should be affected to the same extent.

This is what I set out to determine in this experiment. Outcome rate was varied in two groups of participants. In one group the outcome occurred on 25% of the trials (Group 25%) and in the other group the outcome occurred on 75% of the trials (Group 75%). Both groups were presented with trial types necessary for blocking, a blocking control, and a compound of two cues which did not lead to the outcome (A+/AX+/PQ+/BC-). In addition to these, I added eight trial types to manipulate outcome rate. These included single and compound cues. In Group 25%, none of the added trial types led to the outcome, while in Group 75%, six out of eight added trial types led to the outcome. The full design of Experiment 5 is presented in Table 3.3.

Method

Participants. Participants were 40 (10 male) Psychology undergraduate students at Plymouth University, aged 18-53 (M = 22.28, SD = 6.71), who received course credit for participation. There were 20 participants in each group.

Materials. The materials and procedure in Experiment 5 were the same as in Experiment 3 unless otherwise stated. The stimuli were 19 images of foods. The foods were: apple, banana, broccoli, cabbage, cherries, corn, grapefruit, grapes, kiwi, mango, orange, peach, pear, pepper, pineapple, pomegranate, pumpkin, strawberries, and watermelon. They were randomly assigned to each cue (A, B, C, D, E, F, G, H, I, J, K, L, M, N, O, R, P, Q, X) for each participant.

Table 3.3

The design of Experiment 5. Supplementary trials used to manipulate outcome rate are shown in italics.

Stage 1 – Group 25%	Stage 1 – Group 75%	Test
A+	A+	A
AX+	AX+	B
BC-	BC-	C
PQ+	PQ+	P
D-	*D+*	Q
E-	*E+*	X
F-	*F-*	
GH-	*GH+*	
IJ-	*IJ+*	
KL-	*KL+*	
MN-	*MN+*	
OR-	*OR-*	
x 24	x 24	x 2

Procedure. Stage 1 consisted of 24 blocks of trials. Each of the 12 trial types was presented once per block in a random order. In Stage 1 participants made predictions about whether a stomach ache or no stomach ache occurred. After 12 blocks of trials participants were given the option to have a short break.

At test, participants were asked to provide causal ratings for the cues of interest (A, B, C, P, Q, X). Other details were the same as for Experiment 3.

Results

Stage 1. Figure 3.5 illustrates the mean proportions of stomach-ache predictions in Stage 1 averaged across participants in Group 25% (higher panel) and Group 75% (lower panel).

Participants learned the contingencies in both groups. The acquisition was more rapid for trial types not associated with the outcome in Group 25% than in Group 75%, however this difference was negligible by the final epoch. In the final epoch, participants responded correctly on 97.61% (*SD* = 13.36%) of the trials (Group 25%: *M* = 97%, *SD* = 13.89%; Group 75%: *M* = 98.13%, *SD* = 12.94%).

To make sure that any effects of outcome rate on X were not due to differences in Stage-1 learning between the groups, I conducted a two-way (trial type [A+/AX+/BC-/PQ+] by group [Group 25% vs Group 75%]) ANOVA on the responses from the final epoch of Stage 1. This revealed a significant effect of trial type, $F(1.75, 66.45) = 674.97$, $p < .001$, $\eta_p^2 = .95$, a non-significant trend for the effect of group, $F(1, 38) = 3.04$, $p = .089$, $\eta_p^2 = .07$, and no significant interaction, $F(1.75, 66.45) = .311$, $p = .704$, $\eta_p^2 = .01$. Because there was a trend for the effect of group, I explored this further. Simple main effects analyses indicated a non-significant trend for higher proportions of stomach-ache predictions in Group 75% than Group 25% on A+ trials, $t(19) = 1.83$, $p = .083$, $r = .28$. Responses did not differ between groups on any of the other trial types, $ts \leq 1$, $ps \geq .33$, $rs \leq .16$. This difference was unexpected, however, given its direction, it did not prevent me from exploring the results at test. Lower ratings on A+ trials in Group 25%, indicated poorer learning that this cue led to the outcome. This could have resulted in higher ratings for X in Group 25% while I predicted lower ratings for X in this group relative to Group 75%. Therefore this would have exerted an influence in the opposite direction to the predicted result.

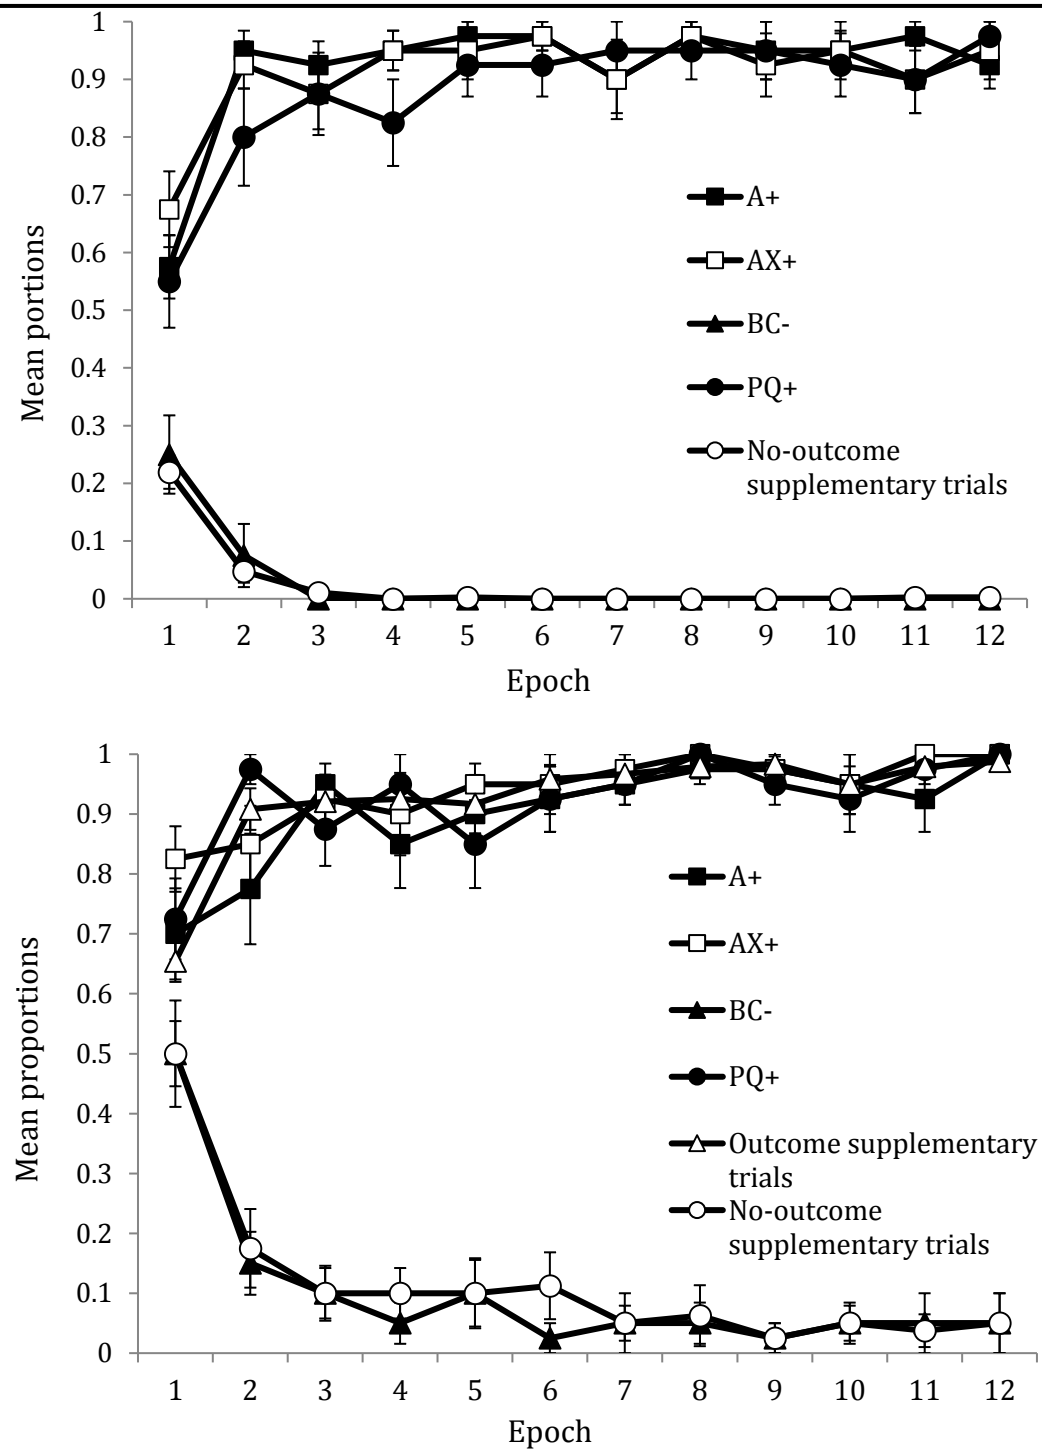

Figure 3.5. Mean proportions of stomach-ache predictions for Stage 1 in Experiment 5 (± between-subjects SEM). Higher panel: Responses in Group 25%. "No-outcome supplementary trials" show responses averaged across supplementary trials leading to no outcome (BC-, D-, E-, F-, GH-, IJ-, KL-, MN-, OR-). Lower panel: Responses in Group 75%. "Outcome supplementary trials" show responses averaged across supplementary trials leading to the outcome (D+, E+, GH+, IJ+, KL+, MN+), "No-outcome supplementary trials" show responses averaged across supplementary trials leading to no outcome (F-, OR-).

Test. Test ratings are presented in Figure 3.6.

A two-way (cue [A, B/C, P/Q, X] by group [Group 25% vs Group 75%]) ANOVA revealed a significant effect of cue, $F(2.19, 83.03) = 226.77$, $p < .001$, $\eta_P^2 = .86$; ratings for all the cues significantly differed from each other, ts ≥ 3.07, ps ≤ .017, rs ≥ .44. There was no significant effect of group, $F(1, 38) = 1.4$, $p = .245$, $\eta_P^2 = .04$, but the interaction just reached significance, $F(3, 114) = 3.08$, $p = .047$, $\eta_P^2 = .08$. The groups did not differ in their ratings for A, B/C, and P/Q, ts ≤ 1.04, ps ≥ .307, rs ≤ .17, but ratings for X were higher in Group 75% than in Group 25%, $t(38) = 2.09$, $p = .044$, $r = .32$.

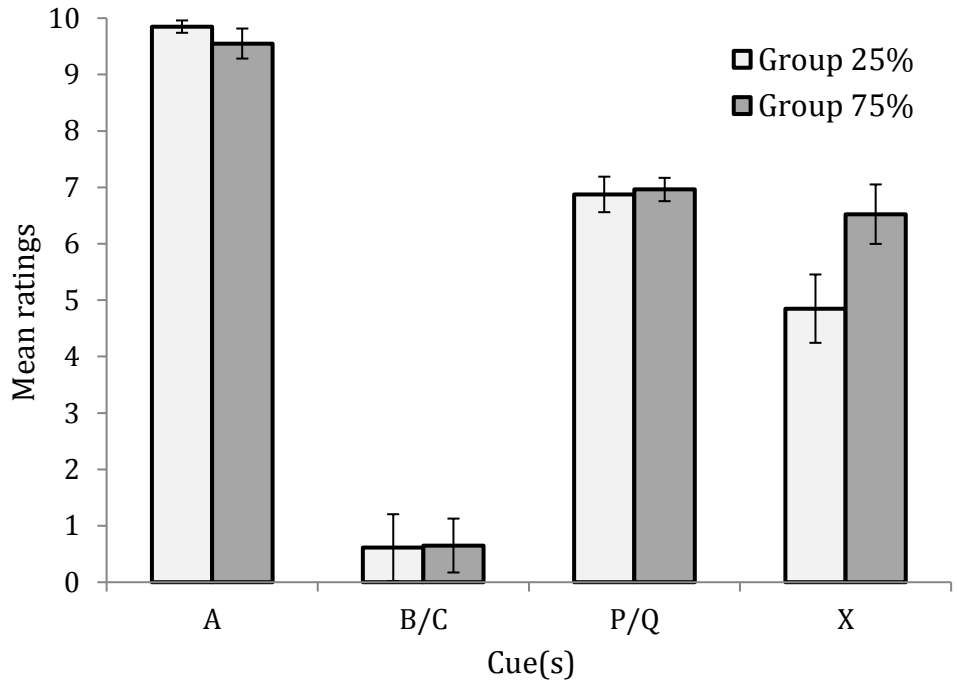

Figure 3.6. Mean causal ratings at test in Experiment 5 (± between-subjects SEM).

In order to explore differences between the groups in ratings for X and for P/Q, a two-way (cue [P/Q vs X] by group [Group 25% vs Group 75%]) ANOVA was conducted on these data. This revealed a significant effect of cue, $F(1, 38) = 10.23$, $p = .003$, $\eta_P^2 = .21$; X

had lower ratings than P/Q. The effect of group was not significant, $F(1, 38) = 2.35$, $p = .134$, $\eta_p^2 = .06$, but the interaction was, $F(1, 38) = 4.25$, $p = .046$, $\eta_p^2 = .1$. Importantly, X had significantly higher ratings in Group 75% than in Group 25%, $t(38) = 2.09$, $p = .044$, $r = .32$, as stated previously, while there were no differences between the groups in ratings for P/Q, $t(38) = .16$, $p = .877$, $r = .03$, $BF_{01} = 3.21$. As a consequence, blocking was significant in Group 25% $t(19)=3.18$, $p = .005$, r = .59 (P/Q ratings – X ratings: $M = 2.03$, $SEM = .64$), but not in Group 75%, $t(19) = 1.01$, $p = .324$, $r = .23$, $BF_{01} = 2.74$ (P/Q ratings – X ratings: $M = .44$, $SEM = .43$).

To summarise, in Experiment 5 I set out to test Livesey et al.'s (2013) predictions regarding the effects of outcome rate on ratings for X and blocking. They predicted that ratings for X would be higher with high outcome rate than with low outcome rate, and blocking would vary accordingly; it would be large with low outcome rate and small with high outcome rate. Findings of this experiment confirmed these predictions: ratings for X were higher, and consequently blocking was smaller, in Group 75% than in Group 25%. Outcome-rate variations did not affect ratings for the other cues, including P/Q. These results established that variation in outcome rate was a viable manipulation to influence ratings for X and blocking. Since the experiments in this chapter and elsewhere indicated that participants were at least somewhat uncertain about the causal status of X, it is possible that uncertainty may have been the mediating factor for the effects of outcome rate on X, in line with Livesey et al.'s (2013) predictions. However, since this experiment did not measure participants' certainty in their causal ratings, this remains to be determined.

In Experiment 6 I aimed to see whether the effects of outcome-rate variation on ratings for X could be replicated and whether this would affect the redundancy effect as a result.

Experiment 6

In Experiment 6, I manipulated outcome rate to determine whether this would affect the redundancy effect. Two groups of participants were presented with the trial types necessary for the redundancy effect (A+/AX+/BY+/CY-). In addition, eight trial types were added to manipulate outcome rate. These were identical to the previous experiment. I retained the same percentage of trial types leading to the outcome in each group as in the previous experiment (Group 25% and Group 75%). The full design of Experiment 6 is presented in Table 3.4.

Table 3.4.

The design of Experiment 6. Supplementary trials used to manipulate outcome rate are shown in italics.

Stage 1 – Group 25%	Stage 1 – Group 75%	Test
A+	A+	A
AX+	AX+	B
BY+	BY+	C
CY-	CY-	X
D-	*D+*	Y
E-	*E+*	
F-	*F-*	
GH-	*GH+*	
IJ-	*IJ+*	
KL-	*KL+*	
MN-	*MN+*	
OP-	*OP-*	
x 24	x 24	x 2

Method

Participants. Fifty-eight Psychology undergraduate students (29 per group) at Plymouth University took part in this experiment. They were aged 18-50 (*M* = 22.07, *SD* = 7) and seven were male. They were tested individually and received course credit for their participation.

Materials. The materials and procedure in Experiment 6 were the same as in Experiment 5 unless otherwise stated.

The stimuli were 18 images of foods. The foods were: apple, banana, broccoli, cabbage, cherries, corn, grapefruit, grapes, kiwi, mango, orange, peach, pepper, pineapple, pomegranate, pumpkin, strawberries, and watermelon. They were randomly assigned to each type of cue (A, B, C, D, E, F, G, H, I, J, K, L, M, N, O, P, X, Y) for each participant.

Procedure. In Stage 1 participants were presented with 24 blocks of trials, with each of the 12 trial types appearing once per block. After 12 blocks of trials participants were given the option to have a break. At test, participants were asked to provide causal ratings for the cues of interest (A, B, C, X, Y) with the same details as in the previous experiment.

Results

Stage 1. Figure 3.7 shows proportions of stomach-ache predictions across the 12 epochs of Stage 1 for each trial type in Group 25% (higher panel) and Group 75% (lower panel). This figure indicates that participants were able to learn the contingencies. In the final epoch, correct responses were made on 96.15% (*SD* = 13.59%) of the trials (Group 25%: *M* = 95.09%, *SD* = 17%; Group 75%: *M* = 97.03%, *SD* = 13.59%).

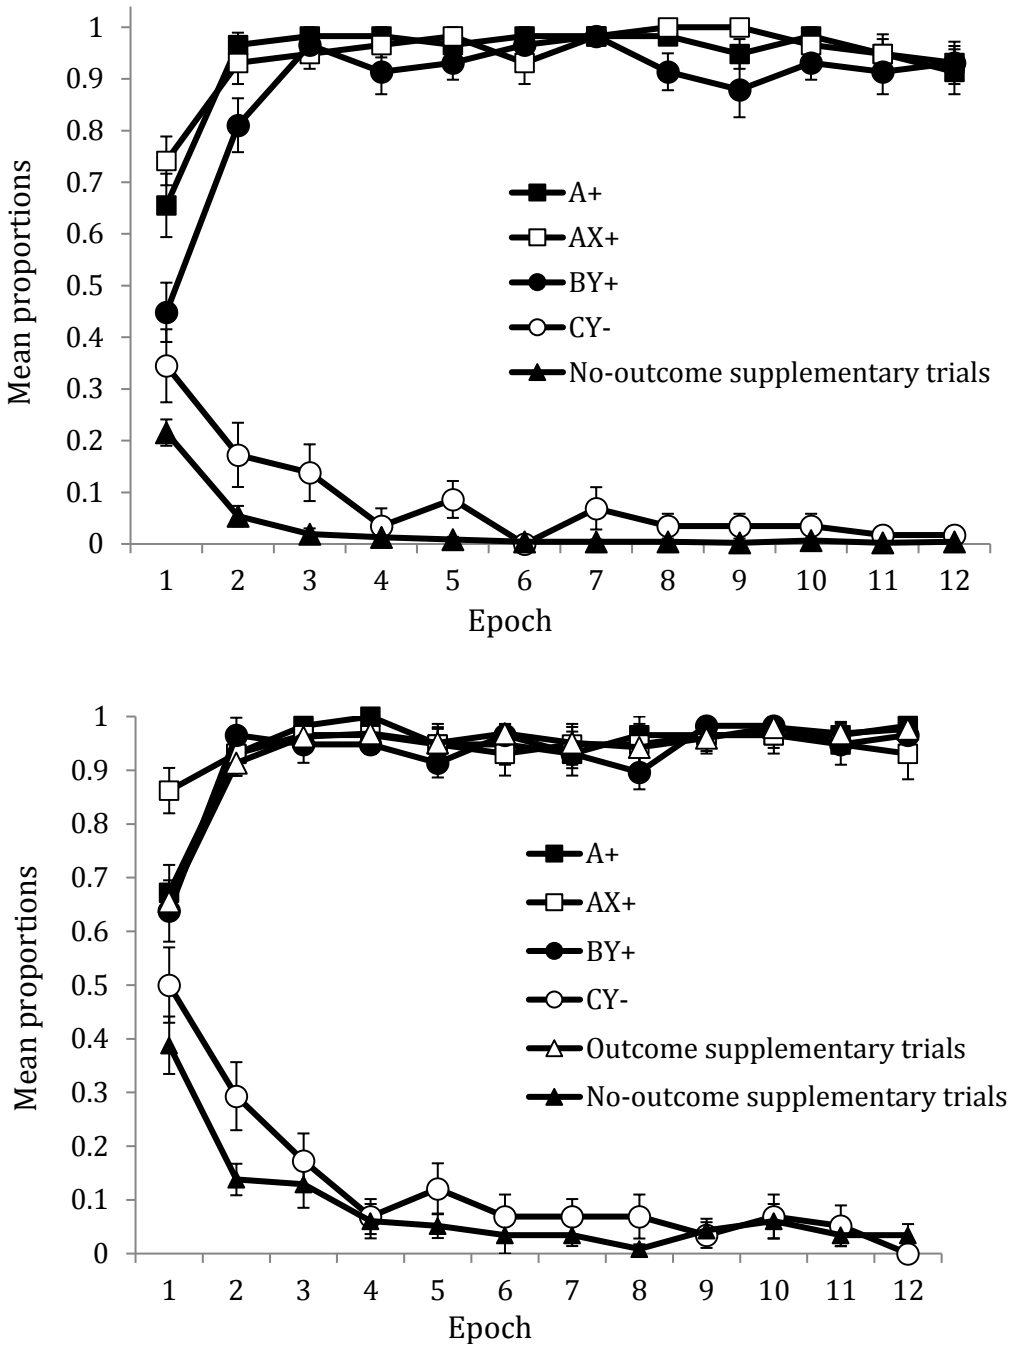

Figure 3.7. Mean proportions of stomach-ache predictions for Stage 1 in Experiment 6 (± between-subjects SEM). Higher panel: Group 25%. "No-outcome supplementary trials" show responses averaged across supplementary trials leading to no outcome (D-, E-, F-, GH-, IJ-, KL-, MN-, OP-). Lower panel: Group 75%. Outcome supplementary trials" show responses averaged across supplementary trials leading to the outcome (D+, E+, GH+, IJ+, KL+, MN+) "No-outcome supplementary trials" show responses averaged across supplementary trials leading to no outcome (F-, OP-).

To make sure that any effects of outcome rate on X and the redundancy effect were not due to differences in Stage-1 learning between the groups, I conducted a two-way (trial type [A+/AX+/BY+/CY-] by group [Group 25% vs Group 75%]) ANOVA on the responses from the final epoch. This revealed a significant effect of trial type, $F(3, 168) = 599.12$, $p < .001$, $\eta_P^2 = .92$. The effect of group was not significant, $F(1, 56) = .5$, $p = .483$, $\eta_P^2 = .01$, nor was the interaction, $F(3, 168) = 1$, $p = .393$, $\eta_P^2 = .02$.

Test. Causal ratings at test, averaged across participants in Group 25% and Group 75%, are presented in Figure 3.8. A two-way (cue [A, B, C, X, Y] by group [Group 25% vs Group 75%]) ANOVA conducted on these data revealed a significant effect of cue, $F(2.86, 160.22) = 182.06$, $p < .001$, $\eta_P^2 = .77$; ratings for all the cues differed significantly from each other, $ts \geq 3.54$, $ps \leq .007$, $rs \geq .42$. There was no significant effect of group, $F(1, 56) = .33$, $p = .568$, $\eta_P^2 = .01$, but a significant interaction, $F(2.86, 160.22) = 4.72$, $p = .004$, $\eta_P^2 = .08$. The groups did not differ in their ratings for A, C, and Y, $ts \leq .75$, $ps \geq .455$, $rs \leq .1$, but ratings for X were higher in Group 75% than in Group 25%, $t(56) = 2.85$, $p = .006$, $r = .36$. Ratings for B were higher in Group 25%, than in Group 75%, $t(56) = 2.19$, $p = .033$, $r = .28$; this was not expected.

Next, to compare ratings for X and Y between the groups, I conducted a two-way (cue [X vs Y] by group [Group 25% vs Group 75%]) ANOVA. This revealed a significant effect of cue, $F(4, 224) = 73.95$, $p < .001$, $\eta_P^2 = .57$, indicating that ratings for X were higher than for Y. The main effect of group did not reach significance, $F(1, 56) = 3.77$, $p = .057$, $\eta_P^2 = .06$, but the interaction was significant, $F(1, 56) = 5.43$, $p = .023$, $\eta_P^2 = .09$. Simple main effects analyses indicated that Group 75% had higher ratings for X than Group 25%, $t(56) = 2.85$, $p = .006$, $r = .36$, as stated previously, while the groups did not differ significantly in their ratings for Y, $t(56) = .26$, $p = .795$, $r = .03$, $BF_{01}=3.07$. Therefore the effects of outcome rate on ratings for X from the previous experiment were replicated. As a consequence, the redundancy effect (X ratings – Y ratings) was larger in Group 75% ($M = 3.28$, $SEM = .67$)

than Group 25% (M = 2.69, SEM = .69), however it was significant in both groups (Group 25%: $t(28) = 3.92$, $p = .001$, $r = .6$; Group 75%: $t(28) = 9.12$, $p < .001$, $r = .86$). This indicated that participants' uncertainty about the causal status of X contributed to the redundancy effect.

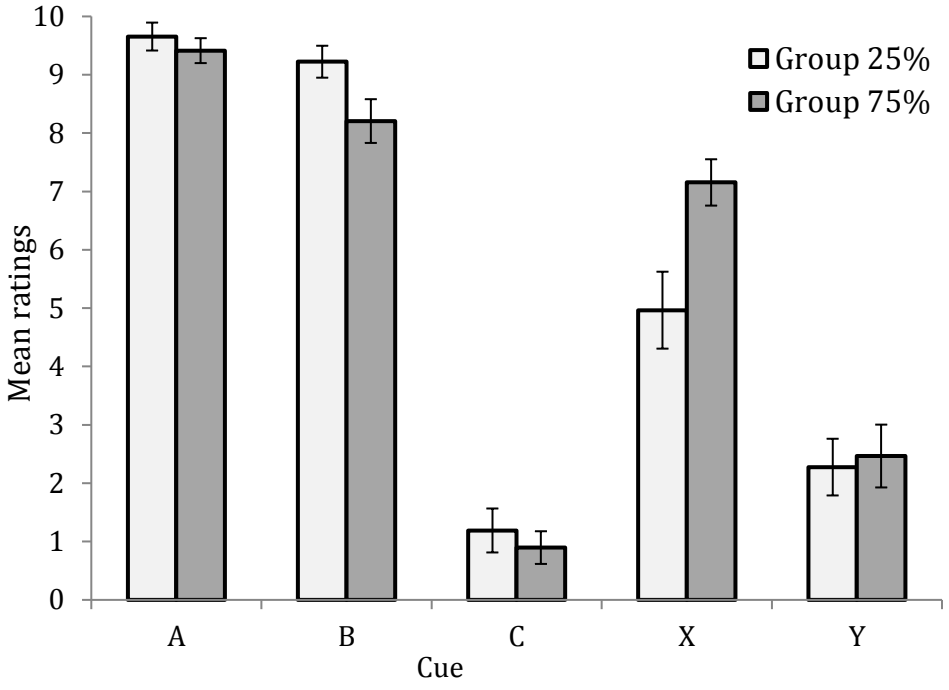

Figure 3.8. Mean causal ratings at test in Experiment 6 (± between-subjects SEM).

Common stimulus elements. Findings of this experiment fall outside the scope of predictions of the Rescorla-Wagner (1972) model, which predicts that Y should have a stronger relationship with the outcome than X, and no effects of outcome rate on the ratings for X compared with other cues. However, recently Vogel and Wagner (2017) suggested one way that the redundancy effect could be reconciled with this model. They proposed that some stimulus elements are common to all trial types presented during learning. Therefore cues presented on the same trial compete with each other, and also with the common elements, for the association with the outcome. Since common elements

are present on every trial, to incorporate these, the design A+/AX+/BY+/CY- can be rewritten as AK+/AXK+/BYK+/CYK-; common elements are denoted by K. In this design K is predicted to partly restrict learning of the association between A and the outcome; this will enable X to gain some associative strength. On CY- trials, the expectation of the outcome will be high due to the presence of both Y and K, resulting in a greater predictive error. This will reduce the associative strength of Y to a greater extent than would be assumed under the circumstances in which no common stimulus elements are present. This interpretation of the Rescorla-Wagner model can predict the redundancy effect as shown by simulations in Vogel and Wagner's paper. However, there is an additional prediction made by this account. Vogel and Wagner stated that the associative strength of K will be partly determined by outcome rate. If outcome rate is high, this should increase the associative strength of K. Alternatively, if outcome rate is low, the associative strength of K should be reduced and K will exert a smaller influence on learning. This is equivalent to the outcome rate variations I carried out in in this experiment. This account predicts that with high outcome rate the redundancy effect should be observed, but with low outcome rate it should be reversed; Y will have a higher associative strength than X. I simulated the design used in Experiment 6 using the Rescorla-Wagner model with the addition of common stimulus elements. The resulting associative strengths for the cues which participants rated at test are displayed in Figure 3.9.

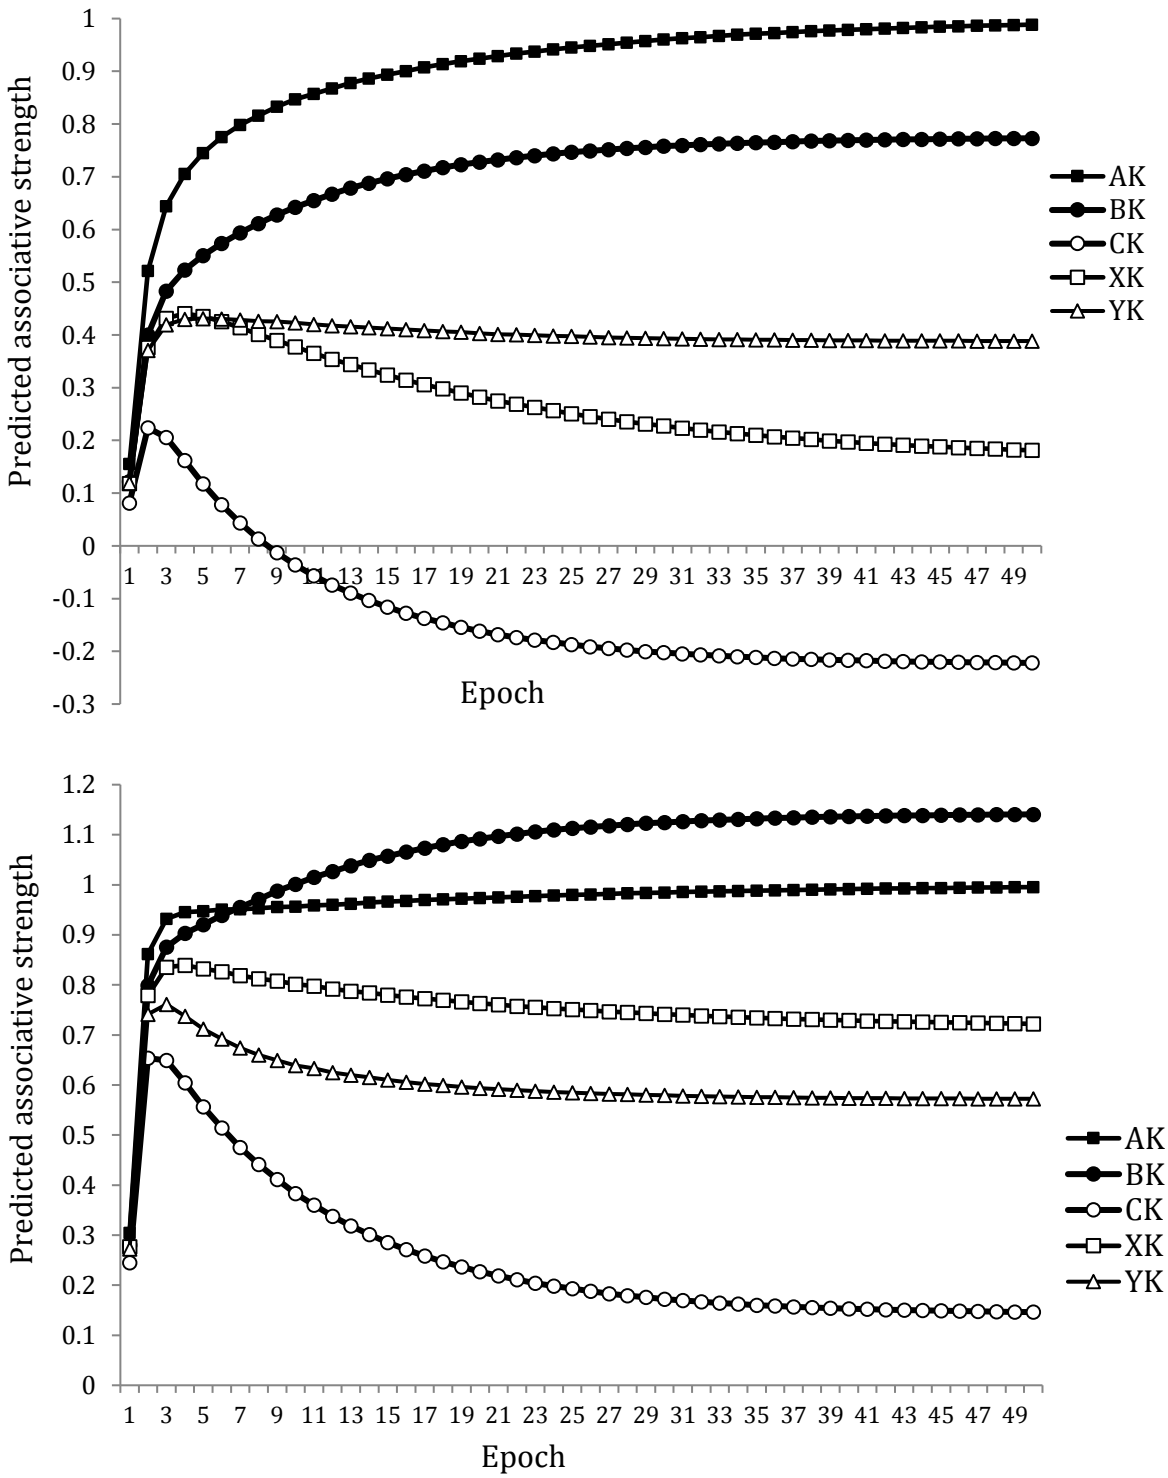

Figure 3.9. Simulations of the Rescorla-Wagner (1972) model using common stimulus elements as proposed by Vogel and Wagner (2017). This figure shows associative strengths for A, B, C, X, and Y using the design in Experiment 6. Higher panel: Group 25% (design AK+/AXK+/BYK+/CYK-/DK-/EK-/FK-/GHK-/IJK-/LMK-/NOK-/PQK-). Lower panel: Group 75% (design AK+/AXK+/BYK+/CYK-/DK+/EK+/FK-/GHK+/IJK+/LMK+/NOK+/PQK-). Consistently with the parameters chosen by Vogel and Wagner (2017), α was set at 0.4, $Beta_E$ was 0.2, and $Beta_I$ was 0.1.

As the figure illustrates, the redundancy effect is predicted in Group 75% (lower panel) and is predicted to reverse, with a higher associative strength for Y than X, in Group 25% (higher panel). These predictions are contrary to the findings observed in Experiment 6; the redundancy effect was significant in Group 75% and Group 25%. There are also other discrepancies between the simulations and the data observed in this experiment. In Group 75% all cues but A are predicted to have higher associative strengths than in Group 25%, however the only cue which had higher ratings in Group 75% was X. In addition, in Group 75%, B is predicted to have a greater associative strength than A, however ratings were greater for A than B in this group, $t(28) = 3.51$, $p = .002$, $r = .55$ (A: $M = 9.93$, $SEM = .05$; B: $M = 8.84$, $SEM = .34$). Therefore, these data appear to challenge the account proposed by Vogel and Wagner (2017).

Discussion

Experiments in this chapter aimed to investigate whether participants' uncertainty about the relationship between X and the outcome contributed to the redundancy effect. Participants' uncertainty about the causal status of X is predicted by several accounts of learning, including the probabilistic contrast model (Cheng, 1997; Cheng & Holyoak, 1995) and the propositional account of learning (Beckers et al., 2006; De Houwer, 2009; Lovibond, 2003, Mitchell et al., 2009). In Experiment 3 I found that participants were less confident in their causal ratings for X than for Y. This indicated that differences in certainty between X and Y could have contributed to the redundancy effect. In subsequent experiments I employed different manipulations aimed to resolve ambiguity about the causal status of X. Resolving ambiguity about X should have reduced ratings for X but not for Y and therefore reduced the magnitude of the redundancy effect. In Experiment 4 I used outcome-additivity training. In Group Additive, participants were presented with two cues which individually caused the outcome and resulted in a stronger outcome when presented together. Participants in the comparison group, Group Non-additive were

presented with two cues which individually caused the outcome and when presented together led to same outcome as either individual cue. I found that ratings for X were lower in Group Additive than in Group Non-additive. Therefore, outcome-additivity training successfully reduced ratings for X and increased blocking. This was consistent with findings from other studies (e. g. Beckers et al., 2005; Livesey & Boakes, 2004; Lovibond et al., 2003; Mitchell & Lovibond, 2002; Mitchell et al., 2005). As a consequence, the redundancy effect was also smaller in Group Additive. However, while additivity training successfully reduced the redundancy effect, the redundancy effect was still significant in both groups.

In subsequent experiments I employed a different manipulation to influence ratings for the blocked cue X, variations in outcome rate. This tested a prediction made by Livesey et al. (2013). They proposed that due to uncertainty about the blocked cue, participants will use outcome rate when judging the probability with which this cue leads to the outcome. Because outcome rate was high in Uengoer et al.'s (2013) experiments and in the experiments in Chapter 2 of this thesis, high outcome rate could have contributed to the redundancy effect observed in these studies. In order to establish whether outcome-rate variations influenced ratings for the blocked cue X, in Experiment 5, I varied outcome rate and observed how this affected ratings for X and blocking. I found that ratings for X varied with outcome rate; they were higher in the group in which the outcome occurred on 75% of the trials (Group 75%) than in a group in which the outcome occurred on 25% of the trials (Group 25%). The effects of outcome rate were specific to X; ratings for the other cues, including P/Q, did not differ between the groups. As a result, blocking was larger in Group 25%, than in Group 75%. This confirmed Livesey et al.'s predictions. Having established that this manipulation was effective at influencing ratings for X and blocking, in Experiment 6 I aimed to replicate the effects of outcome-rate variation on ratings for X. In addition, I investigated whether the redundancy effect would be smaller in Group 25%, relative to Group 75%. Once again X had higher ratings in Group 75% than Group 25%. Ratings for Y did not differ between groups. Consequently, the redundancy effect was

smaller in Group 25% than in Group 75%. However, redundancy effect was significant in both groups.

Overall, findings in this chapter indicated that outcome-additivity training and low outcome rate successfully reduced ratings for X and increased blocking. The redundancy effect was also reduced, but was still significant. However, neither manipulation influenced X directly, rather they relied on participants extracting information from other cues and applying it to X. Therefore, there is no reason to expect for these manipulations to have outweighed direct experience and X may have remained primarily uncertain in these experiments. A future experiment could verify this by asking participants to indicate how certain they are about their causal ratings in different groups in which outcome-additivity/non-additivity training and variations in outcome rate are applied.

Alternatively however, it is possible that other factors contributed to the redundancy effect; the redundancy effect may be multiply-determined. In Experiment 3 I found that participants were less certain about their causal ratings for X than for Y. Perhaps instead of focusing on lower certainty about the causal status of X, the focus should be shifted to why participants were more certain about the causal status of Y.

Chapter 4: Inhibition

In the previous chapter it was found that participants' uncertainty about the causal status of X contributed to the redundancy effect, but did not eliminate it. As such, I suggested that the redundancy effect may have been multiply determined. In this chapter, I explored whether another factor contributing to the redundancy effect was a lack of inhibition for one of the companion cues to Y, cue C.

In the introduction I stated that the redundancy effect poses a challenge to the Rescorla-Wagner (1972) model. This model predicts that in a design A+/AX+/BY+/CY-, Y will have a stronger association with the outcome than X at asymptote. Because of its use of a summed error-term, it predicts that on AX+ trials learning about X will be restricted due to the presence of A, which is a good predictor of the outcome. Cue C is predicted to become a weak inhibitor of the outcome, enabling Y to retain some positive associative strength from BY+ trials.

However, in the previous experiments the allergist task was used to investigate the redundancy effect. In this task, participants were asked to learn which foods led to a stomach ache or no stomach ache in a fictional patient. While it may appear as though the trials on which the outcome did and did not occur were symmetrical, I observed evidence suggesting that participants treated these trials differently. In Experiment 3, participants indicated more confidence in their ratings for two foods which led to no stomach ache (EF-) than two foods which led to a stomach ache (PQ+). This may not be surprising however, because of the particularities of the allergist task. If two foods led to a stomach ache, then the possibilities were that one or both of them were causes. If two foods led to no stomach ache, then the possibilities were that neither food was a cause or one food caused a stomachache while the other prevented it. However the assumption that one food caused a stomach ache and the other prevented it, and they combined to result in no stomach ache, is contrary to most every-day experiences of foods and their effects. Therefore, participants could have assumed that neither food caused a stomach ache on EF- trials,

resulting in higher confidence for E/F than for P/Q in Experiment 3. If this was the case, it may be difficult to demonstrate inhibition using the allergist task.

Because participants experienced trials which did not lead to a stomach ache in the context of Y (CY-), these concerns are relevant to the redundancy effect. If participants assumed that both C and Y were neutral, this could have contributed to lower causal ratings for Y than for X found in the previous experiments. Evidence in Experiments 1 and 2 indicated that participants may have treated these cues similarly, as they looked at C and Y for a similar amount of time on CY- trials, while B was looked at for longer than Y on BY+ trials. If C failed to become an inhibitor, the Rescorla-Wagner (1975) model predicts that Y would have the same associative strength as X. However, even though the redundancy effect was observed in experiments in the previous chapter, participants' uncertainty about the causal status of X, or pre-asymptotic learning for X would have been enough to observe the redundancy effect. Therefore, I explored whether a lack of inhibition for C contributed to the redundancy effect in this chapter. If it was found that inhibition for C resulted in a stronger association with the outcome for Y than for X, the redundancy effect could be partly reconciled with the predictions of the Rescorla-Wagner model.

Experiment 7

The primary aim of Experiment 7 was to explore whether C gains inhibitory associative strength in a design A+/AX+/BY+/CY-, using the allergist task. If inhibition for C and the redundancy effect were observed in the same experiment, this would indicate that a lack of inhibition for C was not likely to have contributed to the redundancy effect in this task. If C was not inhibitory and the redundancy effect was observed, then a lack of inhibition for C could have contributed to the redundancy effect.

Another aim of this experiment was to see whether inhibition could be demonstrated for a cue trained as an inhibitor. Because inhibition is contrary to most every-day experiences

of foods and their effects, inhibition may be difficult to obtain using this task. Cues G and H were established as neutral (GH-) while E was inhibitory (D+/DE-). This type of inhibitory cue manipulation has been used to establish inhibition in the past (e. g. Savastano, Cole, Barnet & Miller, 1999; Urcelay & Miller, 2006). To assess whether inhibition was obtained, a summation test was used (Rescorla, 1969). This involved presenting the cues of interest together with a causal cue (transfer cue; F+) and asking participants to provide causal ratings for three compounds at test: CF, EF, and GF. This enabled me to explore the extent to which C, E, and G were inhibitory in this experiment. If I found that EF received lower ratings than GF, this would constitute evidence of general inhibition. If CF had lower ratings than GF, this would indicate that C became inhibitory. Because the Rescorla-Wagner (1972) model predicts C to be a weak inhibitor, I compared ratings for EF with CF to investigate whether C was as weaker inhibitor than E. The design of this experiment is presented in Table 4.1.

Table 4.1

The design of Experiment 7.

Stage 1	Test
A+	A
AX+	B
BY+	C
CY-	D
D+	E
DE-	F
F+	G
GH-	H
	X
	Y
	CF
	EF
	GF
x 16	x 2

Method

Participants. Participants were 33 Plymouth University students studying Psychology. They received course credit for their participation in this experiment. They were aged 18-26 years (M = 19.67, SD = 1.8) and four were male.

Materials. The materials and procedure of this experiment were the same as Experiment 3, unless otherwise stated.

The cues were 10 images of foods on a white background. The foods were: apple, banana, broccoli, cabbage, cherries, grapes, orange, pumpkin, strawberries, and watermelon. The foods were randomly assigned to each type of cue (A, B, C, D, E, F, G, H, X, Y) for each participant.

Procedure. In Stage 1 participants were presented with 16 blocks of trials, with the eight trial types (A+, AX+, BY+, CY-, D+, DE-, F+, GH-) appearing once per block in a random order, with no successive repetitions of the same trial type.

At test participants were asked to rate the single foods as well as three compounds of two cues (CF, EF, GF). For each participant, an average of the two causal ratings for each type of trial was calculated and used in the analyses.

Results

Stage 1. Figure 4.1 shows proportions of stomach-ache predictions throughout the eight epochs of Stage 1. This figure indicates that participants learned the contingencies. In the final epoch they responded correctly on 99.22% (SD = 6.21%) of the trials.

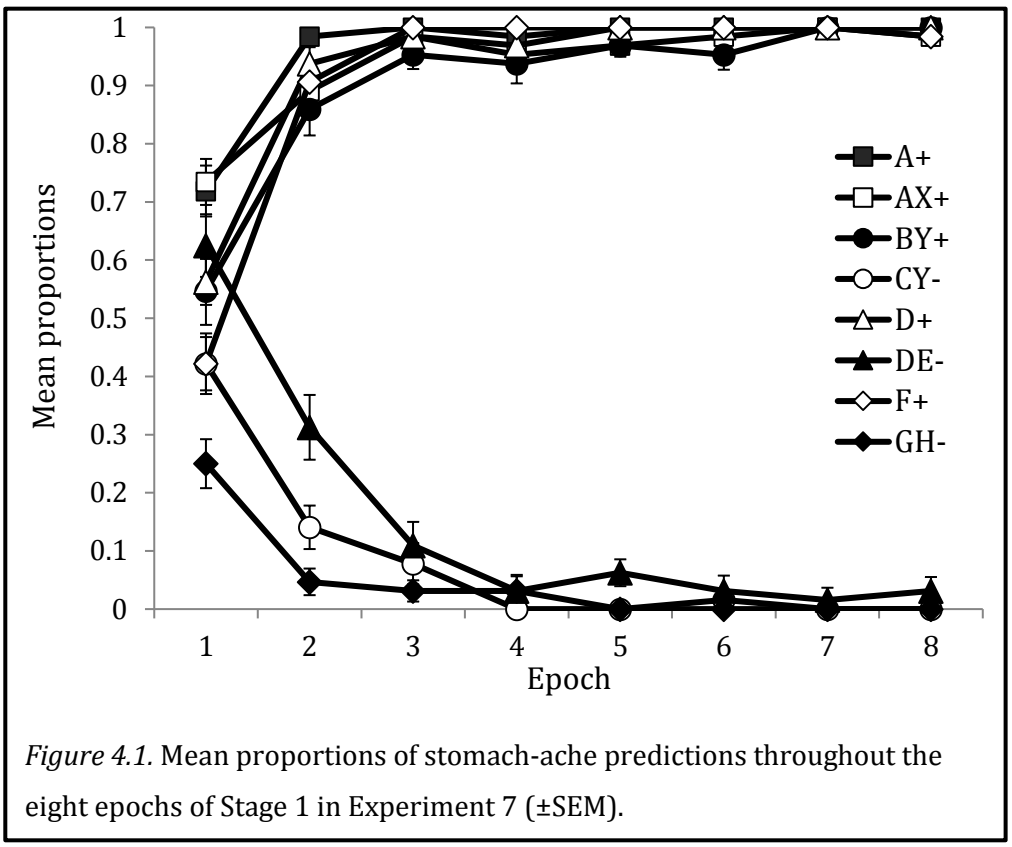

Figure 4.1. Mean proportions of stomach-ache predictions throughout the eight epochs of Stage 1 in Experiment 7 (±SEM).

Test. Figure 4.2 shows the causal ratings for each trial type at test. Ratings for G and H were collapsed as these cues were treated equivalently. Differences between ratings for the single cues were assessed using a one-way (cue [A, B, C, D, E, F, G/H, X, Y]) ANOVA. It revealed a significant effect of cue, $F(3.81, 110.42) = 112.16$, $p < .001$, $\eta^2 = .78$. Bonferroni-corrected paired comparisons indicated that A, B, D, and F had higher ratings than all other cues, $ts \geq 3.96$, $ps \leq .014$, $rs \geq .57$, and while there was a non-significant trend for higher ratings for A than for D, $t(32) = 3.33$, $p = .08$, $r = .51$, ratings for these cues did not differ from each other, $ts \leq 2.73$, $ps \geq .37$, $rs \leq .43$. Cues C, E, and G/H had lower ratings than all other cues, $ts \geq 3.66$, $ps \leq .032$, $rs \geq .54$, but did not differ from each other, $ts \leq .86$, $ps \geq .99$, $rs \leq .15$. Cues X and Y had intermediate ratings, which were different from all other cues, $ts \geq 3.66$, $ps \leq .032$, $rs \geq .54$. A t-test indicated that the redundancy effect was observed: X had significantly higher ratings than Y, $t(32) = 3.94$, $p < .001$, $r = .57$.

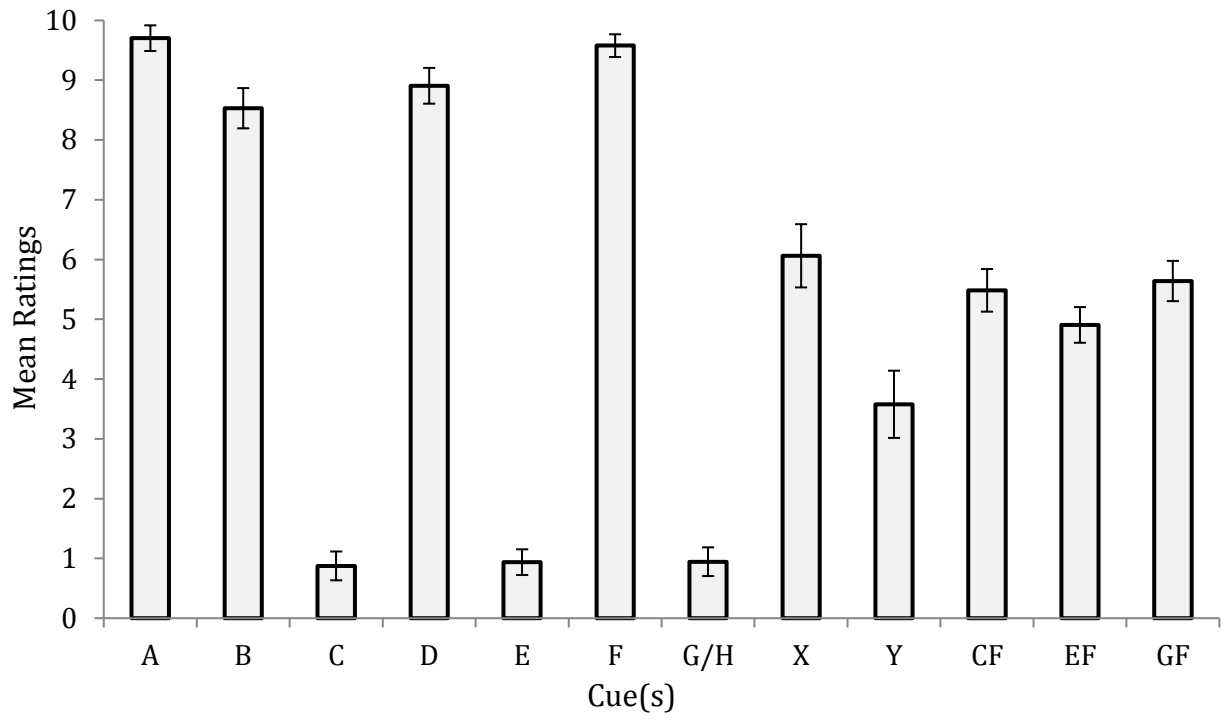

Figure 4.2. Mean causal ratings at test in Experiment 7 (±SEM).

Regarding inhibition comparisons, a one-way (trial type [CF, EF, HF]) ANOVA indicated that ratings between CF, EF, and GF did not differ significantly, $F(2, 64) = 1.13$, $p = .33$, $\eta^2 = .03$, $BF_{01} = 4.46$.

Paired t-tests indicated that C did not become inhibitory as ratings for CF did not differ significantly from ratings for GF, $t(32) = .49$, $p = .63$, $r = .09$, $BF_{01} = 4.81$. One might expect that the best evidence for inhibition, if present, would be demonstrated by the comparison between EF and GF. A paired t-test indicated no significant differences between ratings for EF and GF, $t(32) = 1.4$, $p = .172$, $r = .24$, $BF_{01} = 2.22$, suggesting a failure to obtain inhibition in this experiment. However the Bayes Factor was less than 3 for this comparison and thus did not provide evidence for the null result.

To summarise, in Experiment 7, I explored whether a lack of inhibition for C could have contributed to the redundancy effect in the allergist task. While the redundancy effect was observed, I failed to obtain evidence of inhibition for C, despite showing that participants

learned the contingencies correctly. Therefore, it was possible that a lack of inhibition for C could have contributed to the redundancy effect in this experiment. In addition to this, I failed to obtain evidence of inhibition for E, which was trained as an inhibitor. This has implications for research investigating cues which are predicted to become inhibitory in an allergist task. It suggests that obtaining inhibition may be difficult, particularly for cues which are predicted to be weak inhibitors, such as C. Therefore the next experiments investigated whether evidence of inhibition for C would be obtained in a task in which inhibition was more plausible. If, using this task, inhibition for C was obtained and the redundancy effect was not observed, this would support that inhibition for C may have contributed to the redundancy effect. If inhibition for C was obtained and the redundancy effect was reversed, with higher causal ratings for Y than for X, this would have the potential to partly reconcile the redundancy effect with the predictions of the Rescorla-Wagner (1972) model.

Experiment 8

In the next experiments I used a task in which inhibition was more plausible than in the allergist task. This task was similar to others used to establish inhibition previously (e. g. Melchers, Wolff & Lachnit, 2006). In this task, participants were presented with a fictional patient who consumed medicines which led to a decrease, an increase, or no change in the levels of a fictional hormone. An increase in hormone levels represented excitatory effects on the outcome, a decrease in hormone levels represented inhibitory effects on the outcome, and no change in hormone levels represented neutral effects on the outcome.

Firstly, I aimed to see whether evidence of inhibition could be obtained using this task, and in particular, inhibition for C. Secondly, I aimed to see whether the redundancy effect would be observed.

In order to be consistent with the allergist paradigm in which no-allergy outcome indicated an absence of an allergy as opposed to its prevention, the notation for CY- trials was changed to CY0 (indicating no change in the outcome). Therefore the trial types necessary for the redundancy effect in Experiment 8 appeared as follows: A+/AX+/BY+/CY0. The design of Experiment 8 is presented in Table 4.2.

Table 4.2

The design of Experiment 8. Letters refer to different medicines and "+" represents an increase, "0" no change, and "-" a decrease in hormone levels.

Stage 1	Test
A+	A
AX+	B
BY+	C
CY0	D
D+	E
DE0	F
F+	G
G-	H
H0	X
	Y
	CF
	EF
	HF
x 20	x 2

To check whether inhibition was obtained, E was established as an inhibitory cue, inhibiting an increase in hormone levels (D+/DE0). Cue H was shown to be neutral and led

to no change in hormone levels (H0). Participants were also presented with a single cue which led to a decrease in hormone levels (G-) to make sure they saw evidence that single cues could lead to a decrease. In order to test whether inhibition was obtained, once again a summation test was used with a causal transfer cue F (F+). At test, F was paired with C, the inhibitory cue E, and the neutral cue H (CF, EF, HF). To test whether inhibition occurred, ratings for EF and for HF were compared. To test whether C became an inhibitor, ratings for CF and for HF were compared.

Method

Participants. Participants were 31 Plymouth University students aged 18-62 years (*M* = 27.06, *SD* = 12.63) and 12 were male.

Materials. The materials and procedure in Experiment 8 were the same as in Experiment 7 unless otherwise stated.

The stimuli were 10 images of different colour medicines on a white background, 300 x 300 pixels. Images of the medicines were: brown, green, magenta, orange, pink, purple, red, turquoise, and white. The medicines were randomly assigned to each type of cue (A, B, C, D, E, F, G, H, X, Y) for each participant. The levels of the outcome were: an increase, signified by text "The level of hormone increased" and an image of a yellow arrow pointing upwards on a white background; a decrease, indicated by text "The level of hormone decreased" and an image of a blue arrow pointing downwards on a white background; no change, indicated by text "The level of hormone did not change" and an image of a grey-horizontal arrow pointing left and right on a white background.

Procedure. The instructions for the learning task were adapted from Experiment 7, and were presented on the screen as follows:

Imagine that you are a medical researcher, interested in the effects of different medicines on hormone levels. Your task is to figure out whether the consumption of different medicines will result in an increase, no change, or a decrease in hormone levels. Sometimes one medicine will be consumed and sometimes two medicines will be consumed together.

In the cases where consuming two medicines leads to no change in hormone levels, one medicine may cause an increase and the other a decrease in hormone level, cancelling each other's effects out. However it is also possible that both of the medicines lead to no change.

On the following screens you will see the medicines that participants consume and will be asked to predict whether hormone level will increase, decrease or will not change by clicking the corresponding button. Then, you will be informed of the resulting hormone level change, if any.

At the beginning you will have to guess but by using the feedback provided your guesses should become more accurate. Accuracy is more important than speed for your answers; you may take as long as you like on each trial.

If you have any questions, please ask the experimenter now. Alternatively please click the mouse to start the experiment.

In Stage 1 participants were presented with 20 blocks of trials, with each of the eight trial types (A+, AX+, BY+, CY0, D+, DE0, F+, G-, H0) occurring once per block. As in the previous experiments, the order of the trial types was random with no successive repetitions of the same trial type. On each trial the screen displayed either one or two medicines with the caption "The following medicines were administered" (or "The following medicine was administered" for single medicines). The sentence "What effect on hormone level do you expect?" was presented below the images. Participants responded by clicking one of three response buttons placed at the bottom of the screen. The left-hand button was labelled "Decrease", the right-hand button was labelled "Increase", and the middle button "No change". As soon as participants responded, the screen was replaced by a statement and an image showing the outcome of the trial. When the outcome was an increase, the statement was "The level of hormone increased" and the image was a yellow arrow pointing upwards. When the outcome was a decrease, the statement was "The level of hormone decreased" and the image was a blue arrow pointing downwards. When the outcome was no change, the statement was "The level of hormone did not change" and the image was a grey-horizontal arrow.

The test instructions were as follows:

Next, you are asked to provide your final report. Medicines will be presented on the screen and your task is to use all of the information you have collected up to this time to judge the probability to which specific medicines will change hormone level. Please rate them on a scale from "Decrease" to "Increase" by clicking on the corresponding button.

You will receive no feedback about the resulting hormone level in this stage.

Please click the mouse to begin.

On each test trial either one or two medicines were presented on the screen. Above the image(s) was the sentence "Following the consumption of this medicine the level of hormone will:" (or "Following the consumption of these medicines the level of hormone will:" for trials with two medicines). Participants responded by clicking on a 21-point horizontal rating scale ranging from -10 (*Decrease*) through 0 (*Not change*) to 10 (*Increase*). Similarly to the previous experiment, each trial type at test was presented twice with no successive repetitions, and an average of the two ratings was used in the analyses.

Results

Stage 1. Predictions of hormone-level changes in Stage 1 are displayed in Figure 4.3. This figure shows that participants were able to learn the pairings. On the final epoch they responded correctly on 96.42% (*SD* = 16.28%) of the trials.

Test. Figure 4.4 shows mean ratings at test. A one-way (cue [A, B, C, D, E, F, G, H, X, Y]) ANOVA revealed a significant effect of cue, $F(3.54, 106.05) = 97.43$, $p < .001$, $\eta^2 = .77$. Bonferroni-corrected paired comparisons indicated that ratings for A, D, and F did not significantly differ from each other, $ts \leq .12$, $ps \geq .99$, $rs \leq .02$, and while there was a non-significant trend for higher ratings for D than for X, $t(30) = 3.52$, $p = .062$, $r = .54$, A, D, and F had significantly higher ratings than all other cues, $ts \geq 4.41$, $ps \leq .005$, $rs \geq .63$. While ratings between E and C did not differ, $t(30) = 3.1$, $p = .189$, $r = .49$, ratings for E and for G were lower than for all other cues, $ts \geq 5.4$, $ps < .001$, $rs \geq .7$. Ratings for Y were not significantly different from ratings for B, C, and H, $ts \leq 1.88$, $ps \geq .085$, $rs \leq .32$, but were significantly different from all other cues, $ts \geq 4.87$, $ps \leq .002$, $rs \geq .66$. Ratings for B were not significantly different from X and Y, $ts \leq 2.47$, $ps \geq .872$, $rs \leq .41$, but differed significantly from all other cues, $ts \geq 4.08$, $ps \leq .014$, $rs \geq .6$. Ratings for X differed from all other cues other than B, $ts \geq 4.46$, $ps \leq .005$, $rs \geq .63$.

A paired t-test indicated that the redundancy effect was obtained; ratings for X were higher than for Y, $t(30) = 2.42$, $p = .022$, $r = .4$.

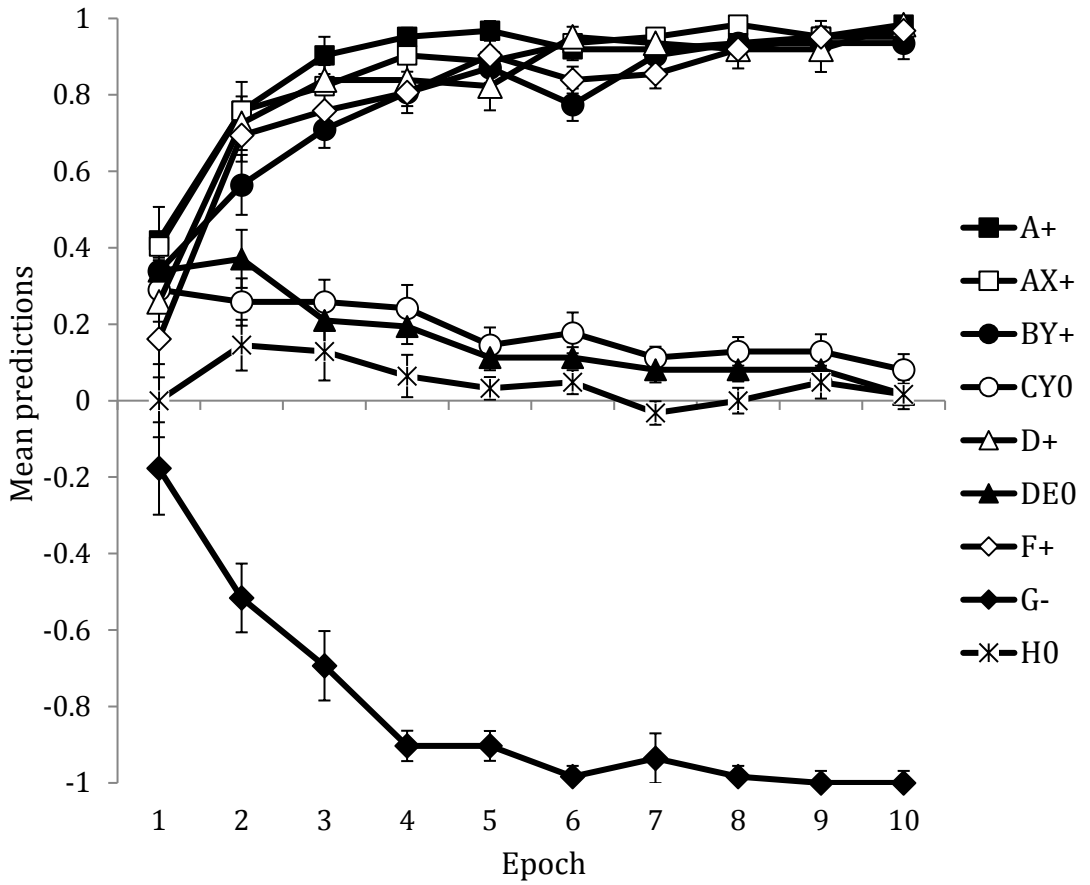

Figure 4.3. Mean hormone-change predictions throughout the ten epochs of Stage 1 in Experiment 8 (±SEM). Positive values refer to a prediction of an increase, zero values refer to a prediction of no change, and negative values refer to a prediction of a decrease in hormone levels.

To investigate the comparisons between CF, EF, and HF, a one-way (trial type [CF, EF, HF]) ANOVA was conducted on the data. It revealed a significant effect of trial type, $F(2, 60) = 13.08$, $p < .001$, $\eta^2 = .3$. A t-test showed that ratings for EF were lower than for HF, $t(30) = 5.4$, $p < .001$, $r = .7$, indicating that inhibition was obtained in this experiment. In addition, CF had significantly lower ratings than HF, $t(30) = 2.22$, $p = .034$, $r = .38$,

suggesting that overall C acquired some inhibitory associative strength. However, EF had lower ratings than CF, $t(30) = 2.82$, $p = .008$, $r = .46$, indicating that inhibition was greater for E than for C.

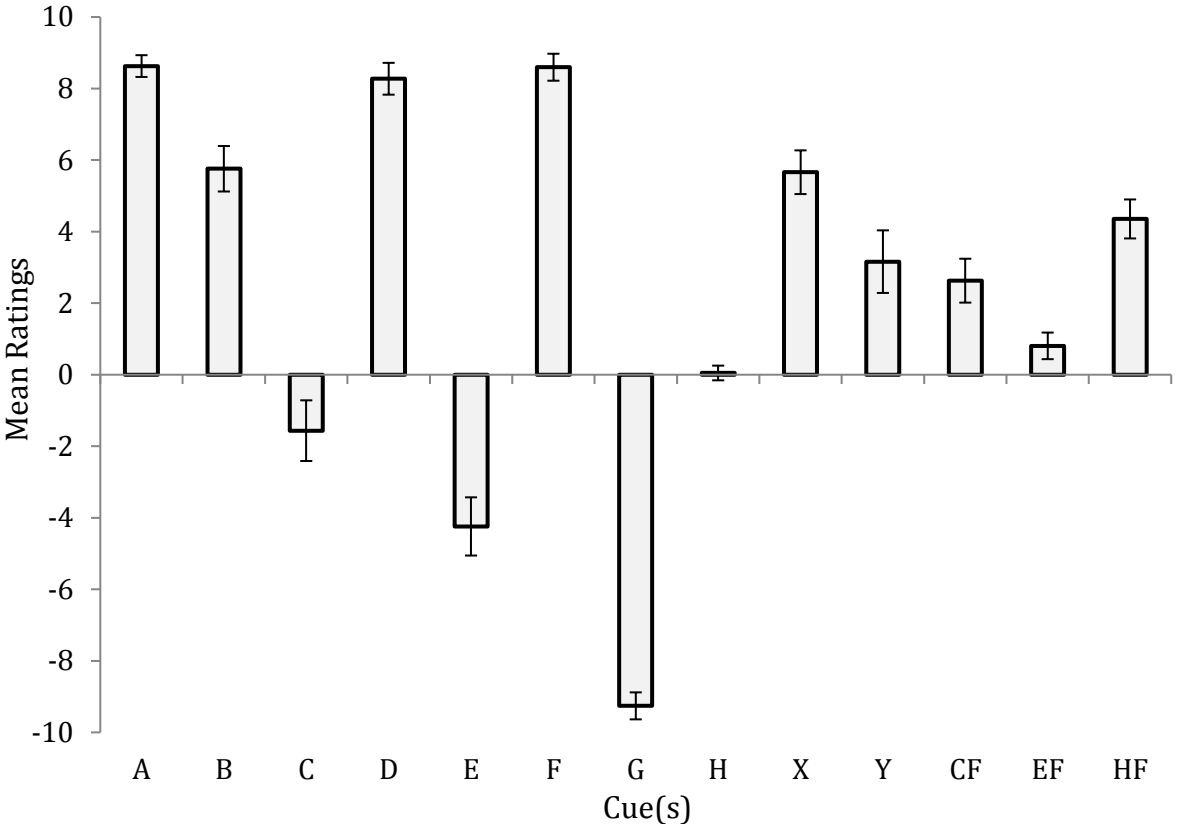

Figure 4.4. Mean hormone-change ratings at test in Experiment 8 (±SEM). Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

Even though C gained some inhibitory associative strength in this experiment, the redundancy effect was significant. Therefore it appeared that a lack of inhibition for C was not responsible for the redundancy effect. However, there were individual differences in ratings for C; some participants gave C negative ratings, indicating inhibition for this cue. It was possible that these participants may have had higher ratings for Y than the participants who did not. To see whether there was a relationship between ratings for C

and for Y, correlations on these data were performed. These revealed a significant negative correlation between ratings for C and for Y, $r(31) = -.494$, $p = .005$ (Figure 4.5, higher panel). There was also a non-significant trend for a negative correlation between ratings for CF and for Y, $r(31) = -.318$, $p = .082$ (Figure 4.5, lower panel). To make sure that higher ratings for Y were not due to general inhibition, I also performed a correlation between ratings for Y and ratings for the inhibitory cue E ; this correlation was not significant, $r(31) = .152$, $p = .413$. Given the negative relationship between ratings for C and for Y, it may have been expected that participants who had lower ratings for C would also have had a smaller redundancy effect. In other words, ratings for C would be positively correlated with the magnitude of the redundancy effect (calculated as X ratings – Y ratings). However, while this correlation was positive, it did not reach significance, $r(31) = .242$, $p = .19$ (Figure 4.6).

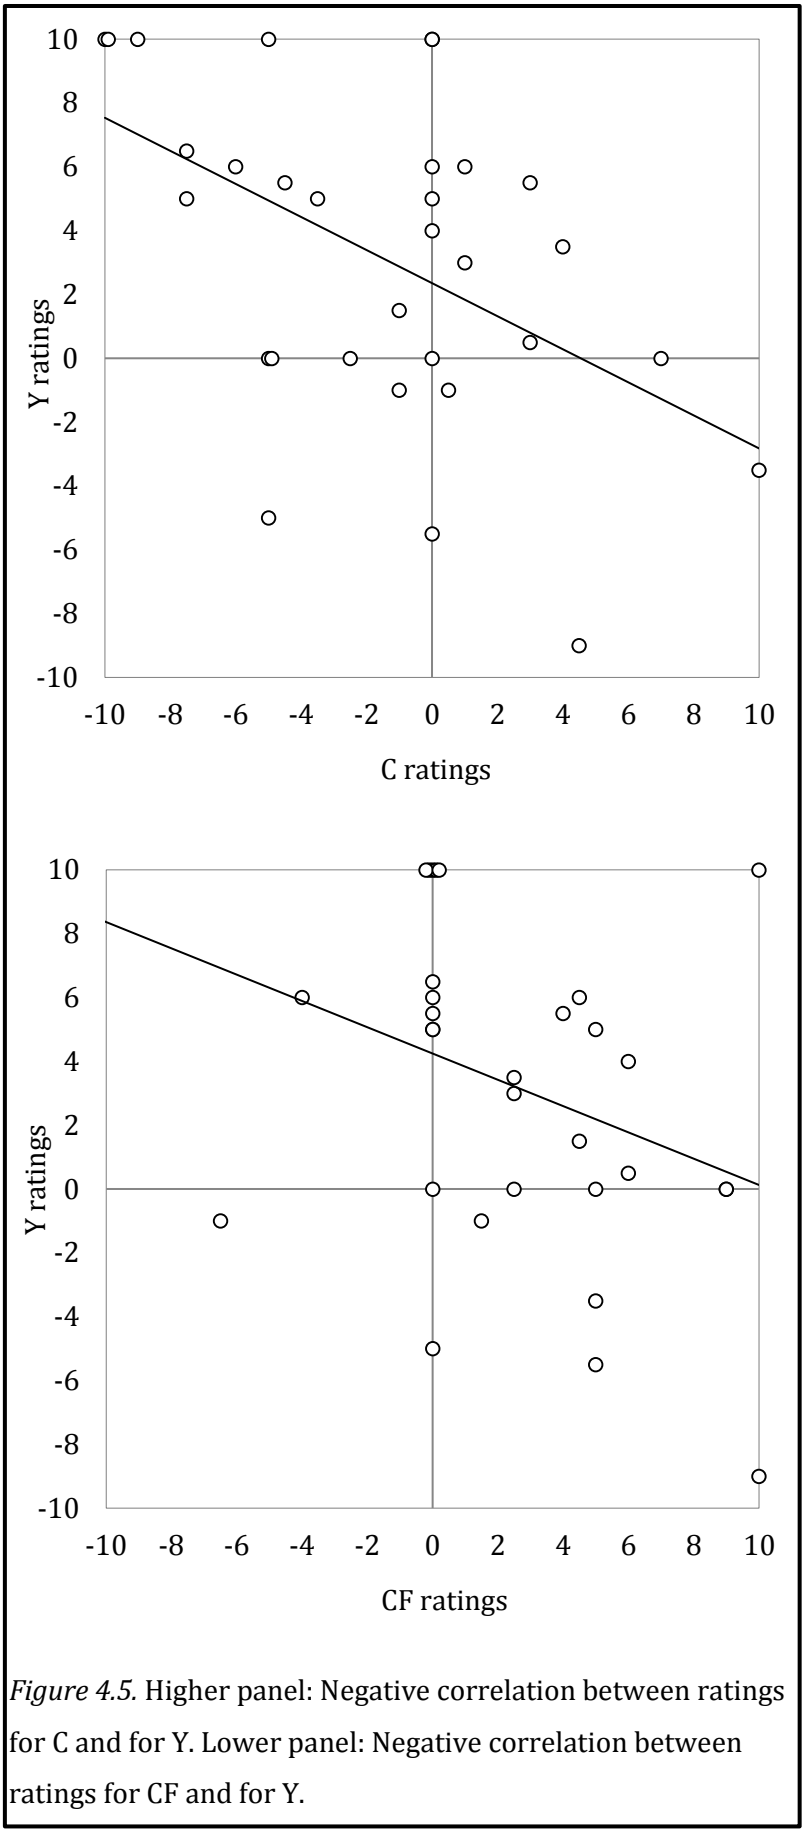

Figure 4.5. Higher panel: Negative correlation between ratings for C and for Y. Lower panel: Negative correlation between ratings for CF and for Y.

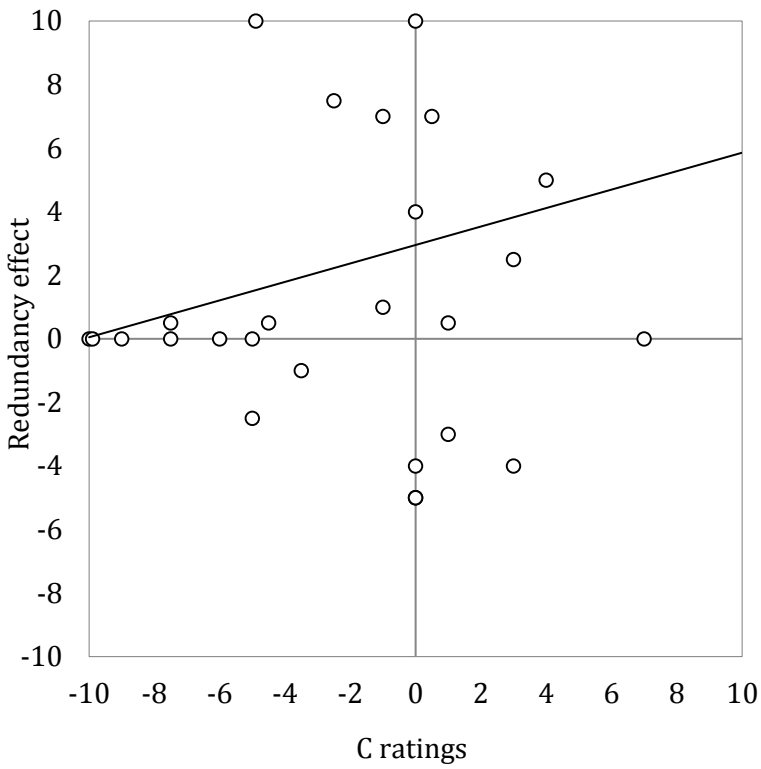

Figure 4.6. Positive correlation between ratings for C and the magnitude of the redundancy effect (X ratings – Y ratings).

To summarise, Experiment 8 used a task in which the consumption of different medicines led to an increase, a decrease, or no change in hormone levels. Firstly, the redundancy effect was significant. Secondly, evidence for inhibition was obtained in this study, as a cue compound which included an inhibitory cue E (EF), had lower ratings than the cue compound which included a neutral cue H (HF). Thirdly, results indicated that C gained some inhibitory associative strength: CF had lower ratings than HF. There was also a significant negative correlation between ratings for C and for Y. While a relationship between ratings for C and the magnitude of the redundancy effect may have been expected, I found that this correlation did not reach significance. It is possible that even though ratings for C and for Y were related, this did not affect the magnitude of the redundancy effect. However, it is also possible that the limited sample size in this

experiment was the reason that this correlation failed to reach significance. I explored whether the latter could have been the case in the next experiment.

Experiment 9

Given the negative relationship between ratings for C and for Y observed in Experiment 8, Experiment 9 aimed to determine whether an experimental manipulation could encourage participants to think of C as inhibitory and therefore would result in higher ratings for Y than in a comparison group. In addition, I aimed to see whether the negative correlation between ratings for C and for Y could be replicated, and whether the correlation between ratings for C and the magnitude of the redundancy effect would reach significance with a larger sample of participants. This experiment had two groups of participants. In Stage 1, Group Inhibitory were shown that one cue led to an increase (G+), one cue lead to a decrease (H-), and when presented together, these cues led to no change (GH0) in hormone levels. Group Neutral were presented with no information about how the cues interacted. They were presented with a cue which led to an increase (G+), a cue which led to a decrease (H-), and two different cues which led to no change (IJ0) in hormone levels.

In Stage 2, all participants were presented with A+/AX+/BY+/CY0/D- trials. It was expected that Stage-1 training in Group Inhibitory would generalise to CY0 trials and would encourage participants to consider that C could have led to a decrease and Y to an increase in hormone levels, since Y was paired with an increase on BY+ trials. Therefore, if Stage-1 manipulation was successful, participants in this group would have higher ratings for Y than participants in Group Neutral. The full design of Experiment 9 is displayed in Table 4.3. In addition to causal ratings for the single cues at test, participants were also asked to rate the cue compound AD. This served as a check for whether participants understood that a cue paired with an increase (A), and a cue paired with a decrease (D), would lead to no change in hormone levels when presented together. I also used

instructions to encourage inhibition in Group Inhibitory. Participants in this group were presented with the instruction "*In the cases where consuming two medicines leads to no change in hormone levels, one medicine may cause an increase and the other a decrease in hormone level, cancelling each other's effects out*" to encourage participants to note that inhibition may be possible on trials which led to no change. This instruction was omitted in Group Neutral.

Table 4.3

The design of Experiment 9.

Stage 1	Stage 2	Test
Group Inhibitory	A+	A
G+	AX+	B
H-	BY+	C
GH0	CY0	D
	D-	X
Group Neutral		Y
G+		AD
H-		
IJ0		
x 8	x 12	x 2

Method

Participants. Participants were 50 Plymouth University students aged 18-46 years (M = 21.56, SD = 6.4) and six were male. There were 25 participants in each group.

Materials. The materials and procedure in Experiment 9 were the same as in Experiment 8 unless otherwise stated.

Images of the medicines were: blue, dark yellow, green, light yellow, orange, pink, purple, red, turquoise, and white. The medicines were randomly assigned to each type of cue (A, B, C, D, G, H, I, J, X, Y) for each participant.

Procedure. The initial instructions were the same as in Experiment 8, with the following changes. In Group Neutral, the instruction "*In the cases where consuming two medicines leads to no change in hormone levels, one medicine may cause an increase and the other a decrease in hormone level, cancelling each other's effects out. However it is also possible that both of the medicines lead to no change*" was omitted. Participants in group Inhibitory were presented with the first part of this instruction "*In the cases where consuming two medicines leads to no change in hormone levels, one medicine may cause an increase and the other a decrease in hormone level, cancelling each other's effects out*" in order to alert participants to this possibility.

In Stage 1 participants were presented with eight blocks of trials, in each of which the three trial types (Group Inhibitory: G+/H-/GH0; Group Neutral: G+/H-/IJ0) appeared once. After Stage 1 participants were presented with instructions to indicate the commencement of Stage-2 trials, as follows:

You will now see a further series of medicines. Just as before, you will make predictions about whether hormone levels will decrease, not change or increase, and you will be given feedback as before.
Please click the mouse to continue.

Stage 2 consisted of 12 blocks of trials with each of the five trial types (A+/AX+/BY+/CY0/D-) presented once per block in a random order.

At test, participants were shown single medicines which appeared in Stage 2 below the statement "Following the consumption of this medicine the level of hormone will:" and were asked to rate the effects of the medicines on hormone level on a 21-point scale ranging from -10 (*Decrease*) through 0 (*Not change*) to 10 (*Increase*). Each cue was rated twice, as in the previous experiment. The last two trials of this stage were designed to act as a manipulation check to make sure that participants learned that inhibition was

possible and involved a presentation of the cue compound AD. Since A had been paired with an increase and D with a decrease in Stage 1, I expected that, provided participants learned that a cue leading to a decrease and a cue leading to an increase summed to no change, this trial type would receive a rating of zero.

Results

Stage 1. Participants learned the contingencies in Stage 1, and in the final epoch responded correctly on 98.32% (*SD* = 14.67%) of the trials (Group Inhibitory: *M* = 96.67%, *SD* = 20.69%; Group Neutral: *M* = 100%, *SD* = 0%). These data are shown in the higher panels of Figure 4.7. A two-way (trial type [G+, H-, GH0/IJ0] by group [Group Inhibitory vs Group Neutral]) ANOVA on these responses revealed a significant effect of trial type, $F(2, 96) = 4201$, $p < .001$, $\eta_P^2 = .99$. The effect of group was not significant, $F(1, 48) = 1$, $p = .322$, $\eta_P^2 = .02$, nor was the interaction, $F(2, 96) = 1$, $p = .372$, $\eta_P^2 = .02$.

Stage 2. Lower panels of Figure 4.7 show the predicted hormone levels, averaged across participants throughout the six epochs of Stage 2 in Group Inhibitory (lower left panel) and Group Neutral (lower right panel). Once again, participants learned these pairings, with correct responses on 97.8% (*SD* = 12.88%) of the trials in the final epoch (Group Inhibitory: *M* = 98%, *SD* = 11.71%; Group Neutral: *M* = 97.58%, *SD* = 14%). A two-way (trial type [A+, AX+, BY+, C0/-, CY0, D-] by group [Group Inhibitory vs Group Neutral]) ANOVA on the final epoch of Stage-2 predictions revealed a significant effect of trial type, $F(4, 192) = 2306.87$, $p < .001$, $\eta_P^2 = .98$, no significant effect of group, $F(1, 48) = .73$, $p = 397$, $\eta_P^2 = .02$, and no significant interaction, $F(4, 192) = .24$, $p = .916$, $\eta_P^2 = .01$.

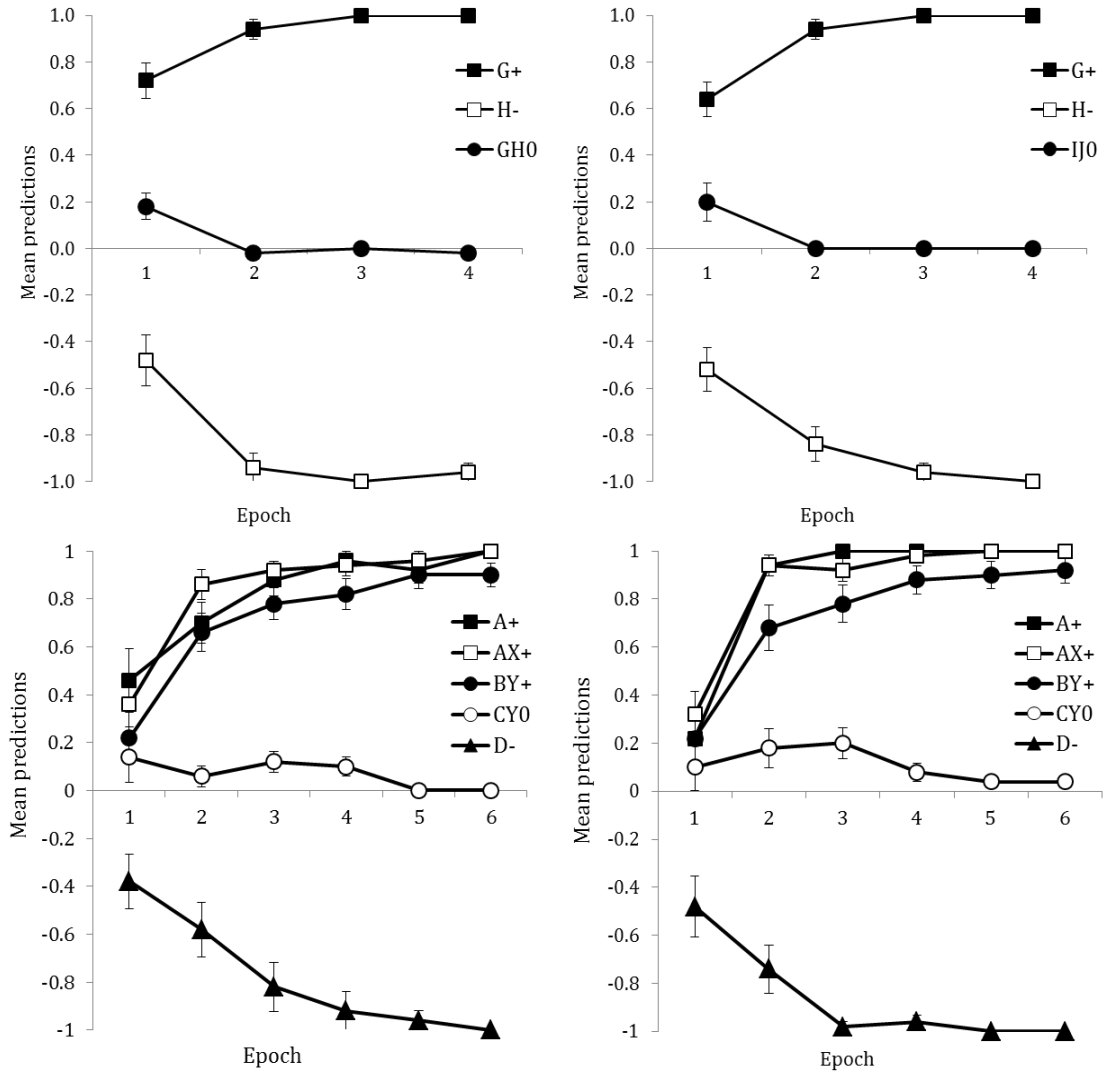

Figure 4.7. Mean hormone-change predictions from Stage 1 and Stage 2 in Experiment 9 (± between-subjects SEM). Higher left panel: Stage-1 responses in Group Inhibitory. Higher right panel: Stage-1 responses in Group Neutral. Lower left panel: Stage-2 responses in Group Inhibitory. Lower right panel: Stage-2 responses in Group Neutral. Positive values refer to a prediction of an increase, zero values refer to a prediction of no change, and negative values refer to a prediction of a decrease in hormone levels.

Test. Causal ratings at test within each group are presented in Figure 4.8.

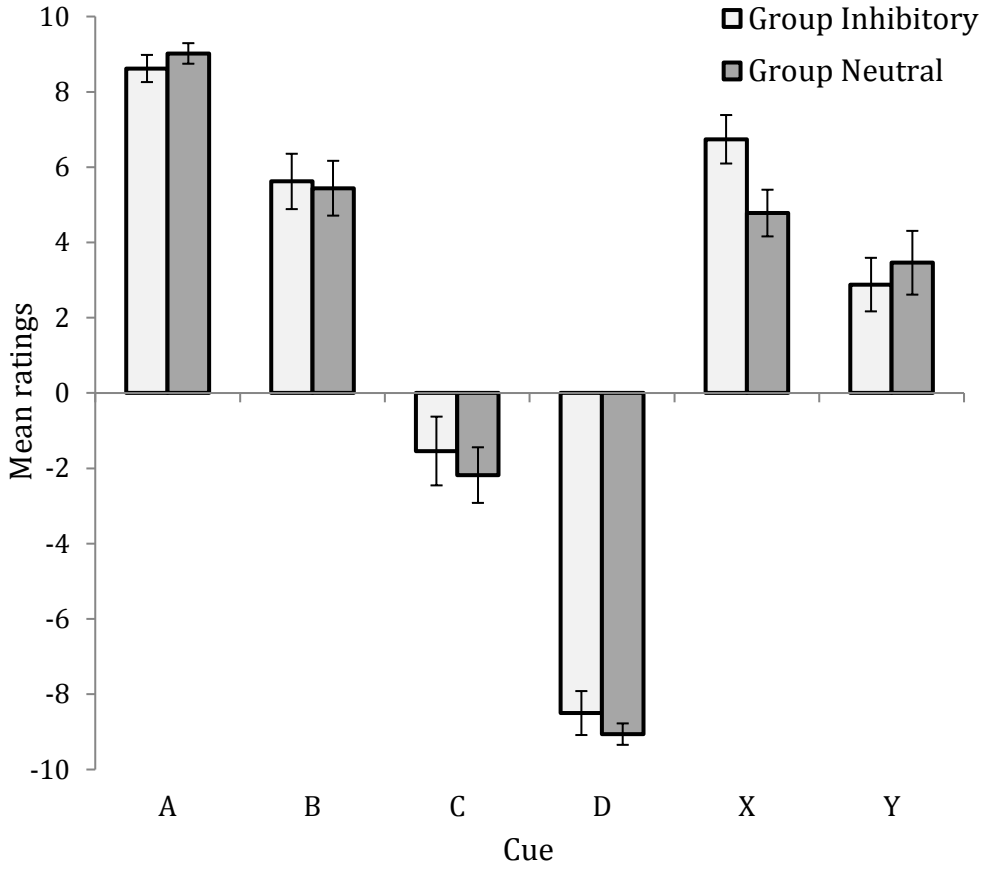

Figure 4.8. Mean hormone-change ratings at test in Experiment 9 (± between-subjects SEM). Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

Manipulation check. In order to check whether the manipulation was successful, a two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA was conducted on the data at test. It revealed a significant effect of cue, $F(1, 48) = 185.96$, $p < .001$, $\eta_p^2 = .8$. Bonferroni-corrected paired comparisons revealed that while ratings between B and X did not differ from each other, $t(49) = .46$, $p > .99$, $r = .07$, ratings between all other cues differed significantly, $ts \geq 3.65$, $ps \leq .01$, $rs \geq .46$. A t-test confirmed that the redundancy effect was obtained, $t(49) = 3.76$, $p < .001$, $r = .47$. The effect of group was not significant, $F(1, 48) = 1.37$, $p = .247$, $\eta_p^2 = .03$, nor was the interaction, $F(3, 144.2) = .94$, $p = .422$, $\eta_p^2 = .02$. In order to check whether C, the cue targeted by the manipulation, yielded different ratings between the groups, a t-test was conducted. It indicated no significant

differences between the groups, $t(48) = .55$, $p = .588$, $r = .08$, $BF_{01} = 3.13$. Therefore I concluded that the manipulation to influence ratings for C was not successful. Interestingly however, ratings for X were higher in Group Inhibitory than in Group Neutral, $t(48) = 2.19$, $p = .033$, $r = .3$.

Regarding AD trials, all but 6 participants rated AD as zero. Mean ratings for AD did not significantly differ from zero, as shown by a one-sample t-test, $t(50) = 1.14$, $p = .259$, $r = .16$. This indicated that overall participants understood that a cue leading to an increase and a cue leading to a decrease would result in no change when presented together. Out of the 6 participants who gave AD a non-zero rating, 1 had a negative rating (-7) while the others had positive ratings (0.5, 3, 4, 6, 10). Two of the latter participants were in Group Inhibitory while the rest were in Group Neutral.

Correlations. These data were then collapsed across groups and used for correlational analyses. A one-way (cue [A, B, C, D, X, Y]) ANOVA revealed a significant effect of cue, $F(3.1, 149.91) = 186.18$, $p < .001$, $\eta^2 = .79$.

Based on the correlations identified in Experiment 8, I investigated the relationships of interest on the collapsed data. Firstly, there was a significant negative correlation between ratings for C and for Y, $r(50) = -.576$, $p < .001$ (Figure 4.9, higher panel). In this experiment, the positive correlation between ratings for C and the magnitude of the redundancy effect (X ratings – Y ratings) was significant, $r(50) = .423$, $p = .002$ (Figure 4.9, lower panel).

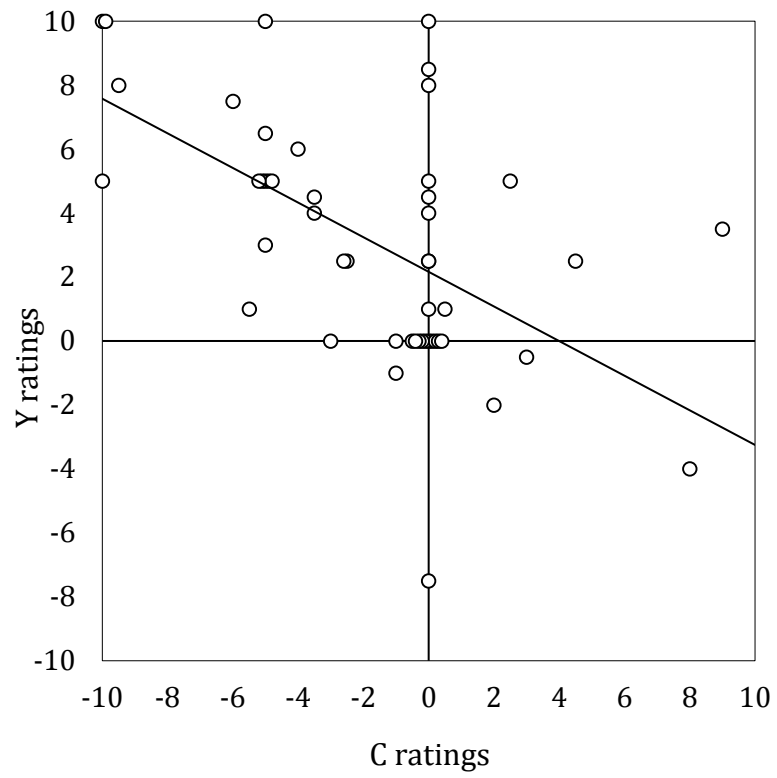

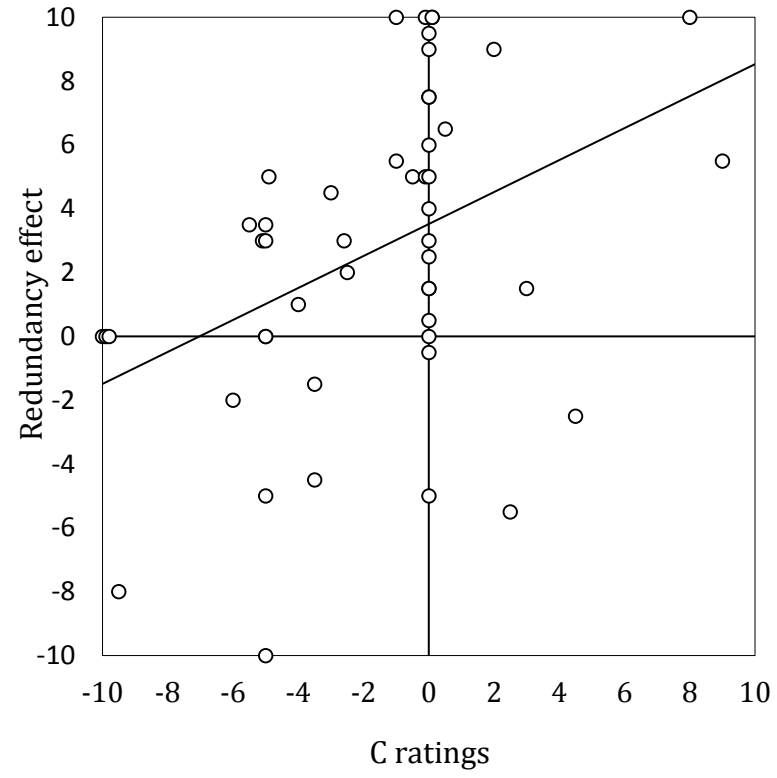

Figure 4.9. Correlations in Experiment 9. Higher panel: Negative correlation between ratings for C and for Y. Lower panel: Positive correlation between ratings for C and the magnitude of the redundancy effect (X ratings – Y ratings).

Data-split analyses. Correlational analyses indicated that lower ratings for C were associated with higher ratings for Y, and a smaller redundancy effect. Therefore it was possible that the magnitude of the redundancy effect would differ depending on whether ratings for C were negative or neutral. To investigate this, I performed data-split analyses on these data, based on participants' ratings for C. Given that the Rescorla-Wagner (1972) model predicts that C will be a weak inhibitor of the outcome, I compared ratings for participants who gave negative ratings for C (C-negative group, *N*=24) with those who rated C as zero or positive (C-neutral/positive group, *N*=26). Mean causal ratings at test within each group are presented in Figure 4.10.

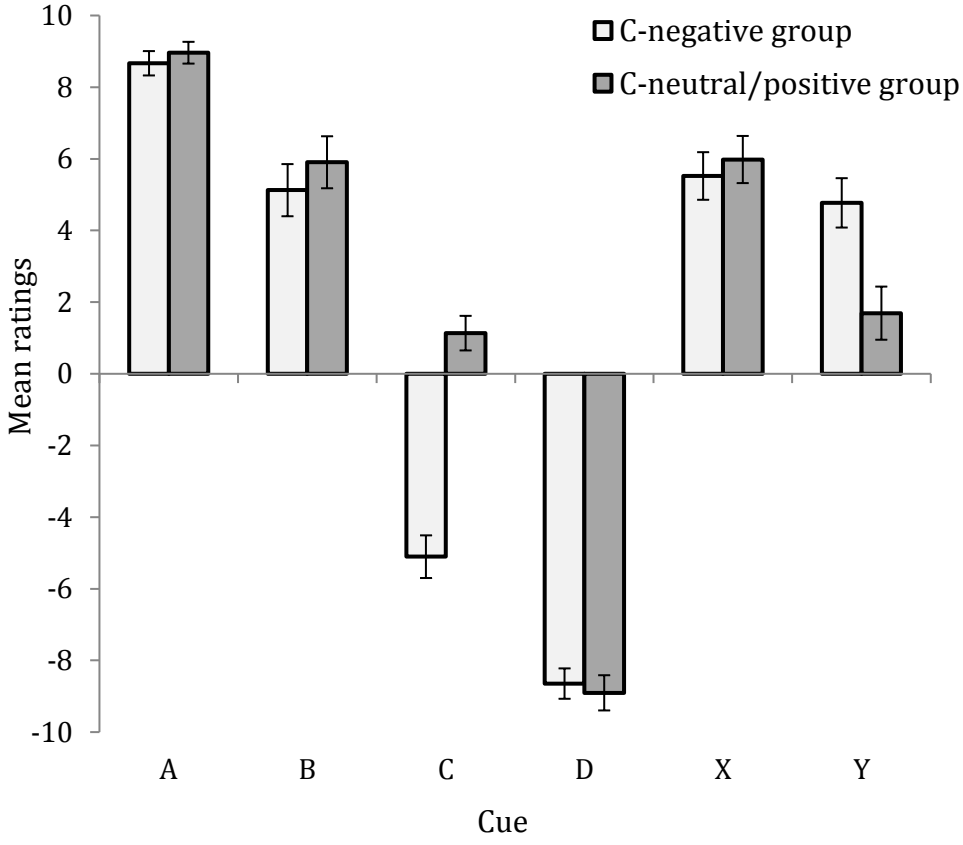

Figure 4.10. Mean hormone-change ratings at test in Experiment 9 split by group (± between-subjects SEM). Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

A two-way (cue [A, B, C, D, X, Y] by group [C-negative group vs C-neutral/positive group]) ANOVA revealed a significant effect of cue, $F(3.33, 159.73) = 232.16$, $p < .001$, $\eta_P^2 = .83$, a significant effect of group, $F(1, 48) = 5.22$, $p = .027$, $\eta_P^2 = .1$, and a significant interaction, $F(3.33, 159.73) = 12.97$, $p < .001$, $\eta_P^2 = .21$. Simple main effects analyses indicated that in the C-neutral/positive group ratings for C were higher, $t(48) = 8.2$, $p < .001$, $r = .76$, and ratings for Y were lower, $t(48) = 3.03$, $p = .004$, $r = .4$, than in the C-negative group. The groups did not differ in their ratings for the other cues, ts ≤ .76, ps ≥ .453, rs ≤ .11.

A two-way (cue [X vs Y] by group [C-negative group vs C-neutral/positive group]) ANOVA revealed a significant effect of cue, $F(1, 48) = 15.15$, $p < .001$, $\eta_P^2 = .24$, with greater ratings for X than for Y, and a non-significant trend for the effect of group, $F(1, 48) = 3.2$, $p = .08$, $\eta_P^2 = .06$. There was also a significant interaction, $F(1, 48) = 7.47$, $p = .009$, $\eta_P^2 = .24$. Simple main effects analyses indicated that the redundancy effect was significant in the C-neutral/positive group, $t(25) = 4.683$, $p < .001$, $r = .68$ (X ratings – Y ratings: $M = 4.29$, $SEM = .92$), but not in the C-negative group, $t(23) = .822$, $p = .419$, $r = .17$, $BF_{01} = 3.42$ (X ratings – Y ratings: $M = .75$, $SEM = .91$). This confirmed that the redundancy effect was related to whether or not C was indicated to be inhibitory.

To summarise, in Experiment 9 I set out to explore whether the redundancy effect could be influenced by encouraging inhibition for C in a between-groups manipulation. While participants learned the contingencies, this manipulation was not successful. Therefore, the data were collapsed between groups and used for correlational analyses. Similarly to Experiment 8, I found a significant negative correlation between ratings for C and for Y. The positive correlation between ratings for C and the magnitude of the redundancy effect was also significant in this experiment, indicating that lower ratings for C were related to a smaller redundancy effect. Data-split analyses based on participants' ratings for C revealed that the redundancy effect was larger for participants who rated C as neutral, than for participants who rated C as inhibitory. Importantly, no significant redundancy effect for

participants who rated C as inhibitory refers to the first failure to observe the redundancy effect in this thesis. In Experiment 10, I aimed to replicate these results, but without the between-subjects manipulation in Stage 1 of Experiment 9.

Experiment 10

Results which Experiment 10 aimed to replicate included the negative relationship between ratings for C and for Y, the positive relationship between ratings for C and the magnitude of the redundancy effect, and data-split analyses showing that negative ratings for C resulted in a smaller redundancy effect than neutral or positive ratings. In this experiment, I omitted the Stage-1 manipulation, but otherwise the design was identical to the previous experiment (Table 4.4).

Table 4.4

The design of Experiment 10.

Stage 1	Test
A+	A
AX+	B
BY+	C
CY0	D
D-	X
	Y
	AD
x 12	x 2

Method

Participants. Participants were 50 Plymouth University students aged 18-34 years (*M* = 19.68, *SD* = 3) and nine were male.

Materials. The materials and procedure in Experiment 10 were the same as in Experiment 9 unless otherwise stated.

The cues were six images of different colour medicines: blue, green, orange, pink, purple, and yellow. These medicines were randomly assigned to each type of cue (A, B, C, D, X, Y) for each participant.

Procedure. The instruction from the previous experiment "*In the cases where consuming two medicines leads to no change in hormone levels, one medicine may cause an increase and the other a decrease in hormone level, cancelling each other's effects out*" was retained in this experiment in order to alert participants to the possibility of inhibition on CY0 trials.

In Stage 1 participants were presented with 12 blocks of trials with the five trial types (A+/AX+/BY+/CY0/D-) appearing once per block in a random order, with no successive repetitions.

Once again participants rated all of the individual medicines at test, followed by AD trials.

Results

Stage 1. Participants learned the pairings in Stage 1 (Figure 4.11). In the final epoch they responded correctly on 98.6% (*SD* = 8.27%) of the trials.

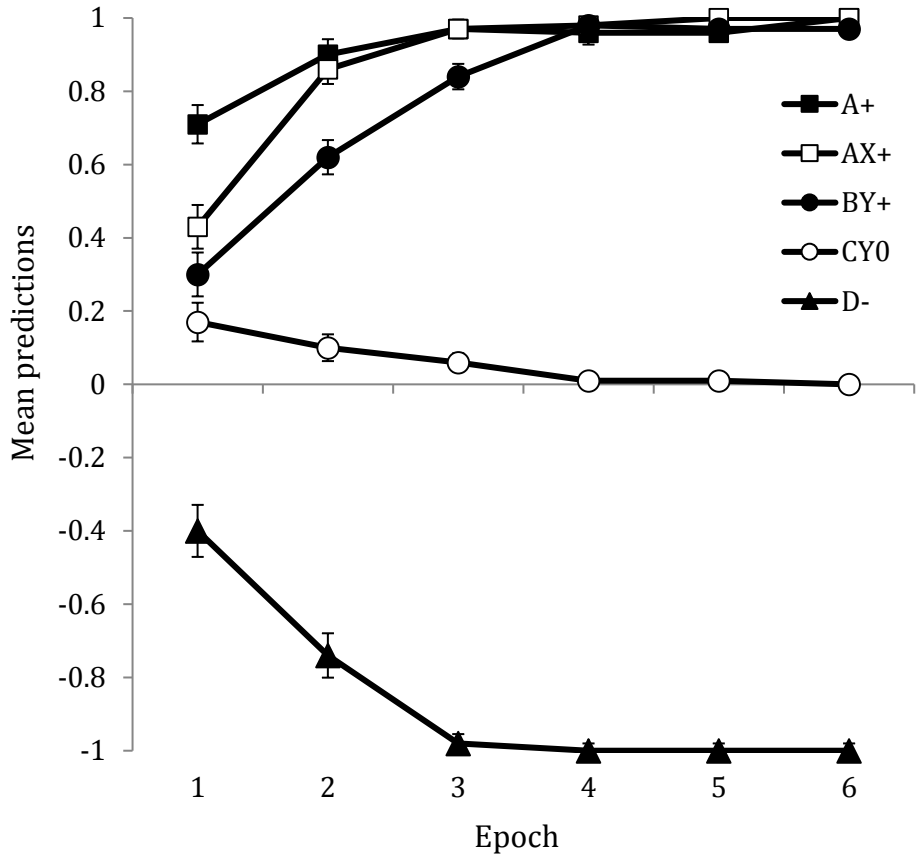

Figure 4.11. Mean hormone-change predictions throughout the six epochs of Stage 1 in Experiment 10 (±SEM). Positive values refer to a prediction of an increase, zero values refer to a prediction of no change, and negative values refer to a prediction of a decrease in hormone levels.

Test. Figure 4.12 shows ratings for cues at test, averaged across participants. A one-way (cue [A, B, C, D, X, Y]) ANOVA revealed a significant effect of cue, $F(2.42, 118.67) = 185.94$, $p < .001$, $\eta^2 = .79$. Bonferroni-corrected paired comparisons indicated that ratings for B, X, and Y did not differ from each other, $ts \leq 2.66$, $ps \geq .159$, $rs \leq .35$, but all other cues did, $ts \geq 7.09$, $ps < .001$, $rs \geq .71$. A t-test showed that the redundancy effect was not obtained, $t(49) = .96$, $p = .343$, $r = .14$, $BF_{01} = 4.22$. Since in this experiment I was interested in replicating the relationships observed in Experiment 9, the redundancy effect failing to reach significance was surprising, but did not prevent me from exploring these.

Average ratings for AD did not significantly differ from zero, as shown by a one-sample t-test, $t(49) = 1.57$, $p = .123$, $r = .22$. All but four participants rated AD as zero and all four had positive ratings (1, 2, 4, 10).

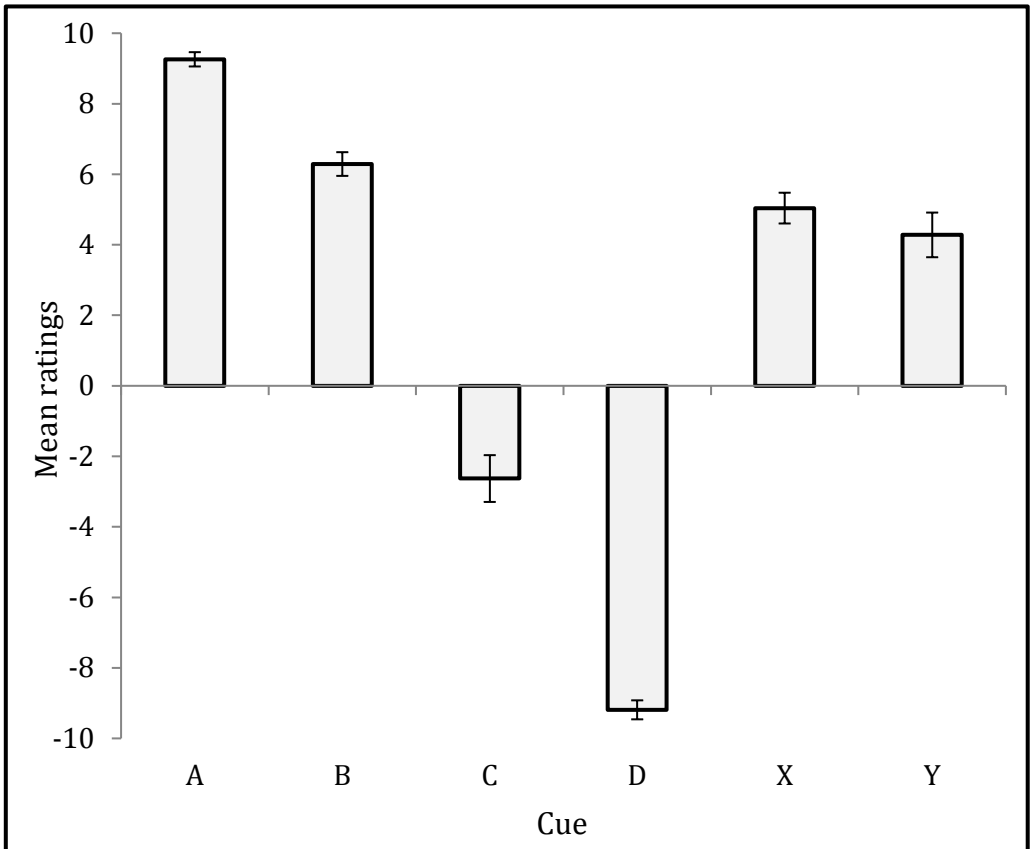

Figure 4.12. Mean hormone-change ratings at test in Experiment 10 (±SEM). Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

Correlations and data-split analyses. Once again there was a significant negative correlation between ratings for C and for Y, $r(50) = -.645$, $p < .001$ (Figure 4.13, higher panel). There was also a significant positive correlation between ratings for C and the magnitude of the redundancy effect, $r(50) = .518$, $p < .001$ (Figure 4.13, lower panel).

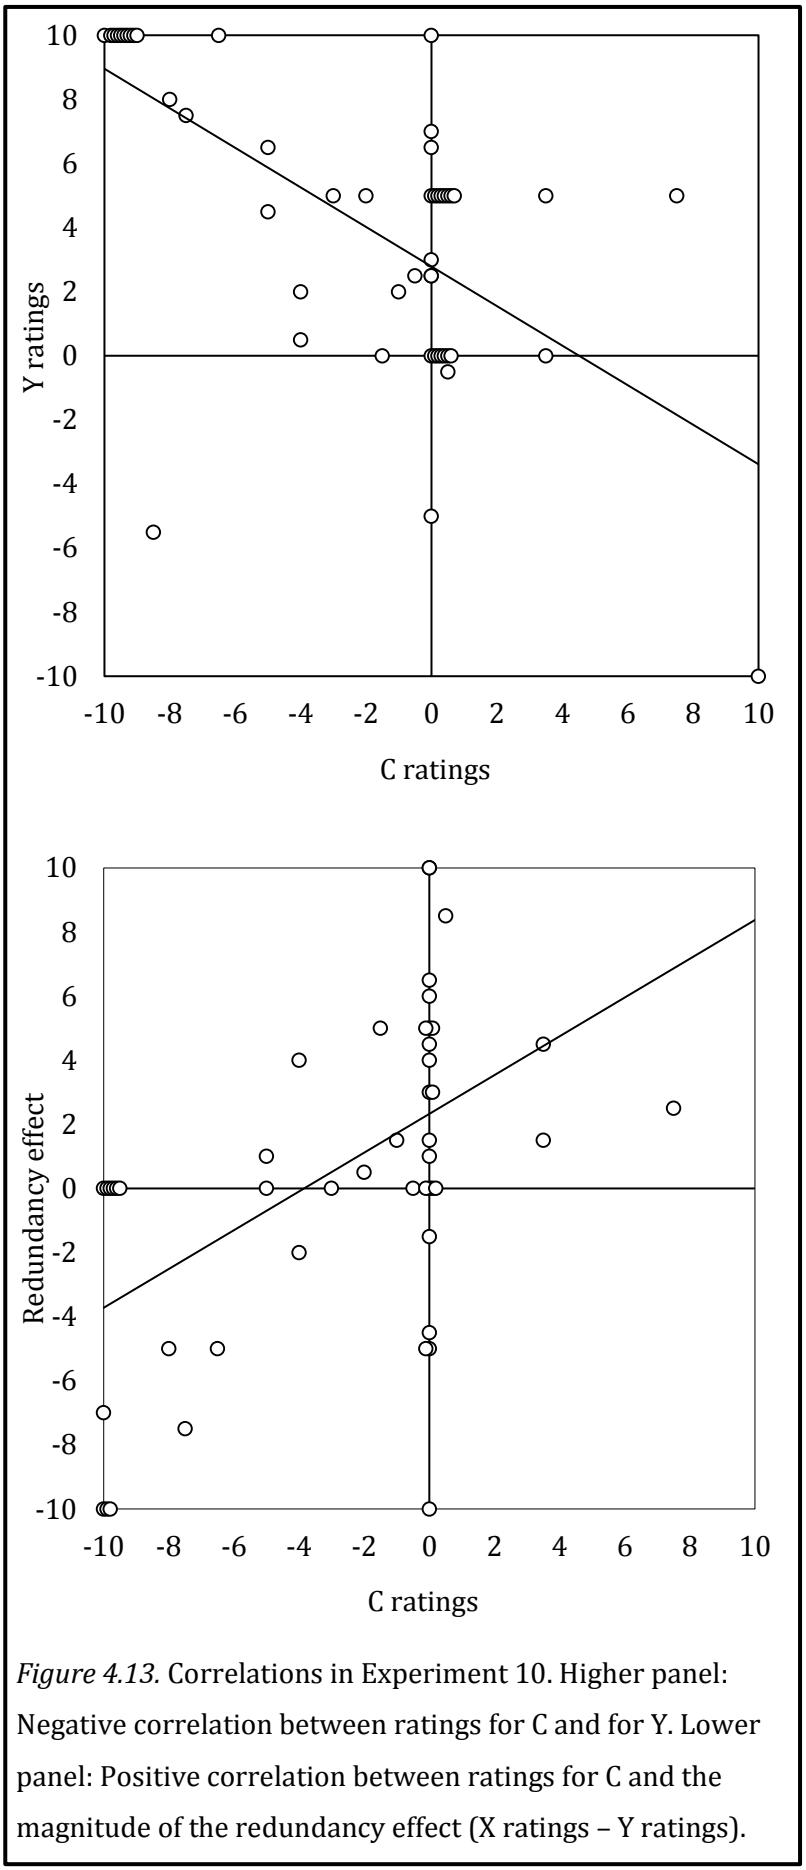

Figure 4.13. Correlations in Experiment 10. Higher panel: Negative correlation between ratings for C and for Y. Lower panel: Positive correlation between ratings for C and the magnitude of the redundancy effect (X ratings – Y ratings).

Because the correlation between ratings for C and the magnitude of the redundancy effect was significant, I once again performed data-split analyses. Participants were split into two groups based on their ratings for C in the manner of the previous experiment. Participants who rated C as negative were included in the C-negative group (N=23) and those who rated C as zero and above were included in the C-neutral/positive group (N=27). Causal ratings in each group are presented in Figure 4.14. A two-way (cue [A, B, C, D, X, Y] by group [C-negative group vs C-neutral/positive group] ANOVA revealed a significant effect of cue, $F(3.16, 151.29) = 260.05$, $p < .001$, $\eta_p^2 = .84$, a significant effect of group, $F(1, 48) = 5.96$, $p = .018$, $\eta_p^2 = .11$, and a significant interaction, $F(3.16, 151.29) = 20.04$, $p < .001$, $\eta_p^2 = .3$. Simple main effects analyses indicated that in the C-neutral/positive group ratings for C were higher, $t(38.72) = 8.9$, $p < .001$, $r = .79$, and ratings for Y were lower, $t(49) = 3.34$, $p = .002$, $r = .43$, than in the C-negative group. There was also a non-significant trend for lower ratings for A in the C-negative group than in the C-neutral/positive group, $t(36.57) = 1.75$, $p = .088$, $r = .24$, while ratings of the other cues did not differ between the groups, $ts \leq .27$, $ps \geq .786$, $rs \leq .04$. Differences between the groups in ratings for A were not expected, however they did not prevent me from exploring differences in the magnitude of the redundancy effect. Lower ratings for A in the C-negative group than in the C-positive/neutral group may have resulted in higher ratings for X in this group. However I predicted a smaller redundancy effect in this group than the C-neutral/positive group. Therefore this difference would have exerted the opposite influence to the anticipated result.

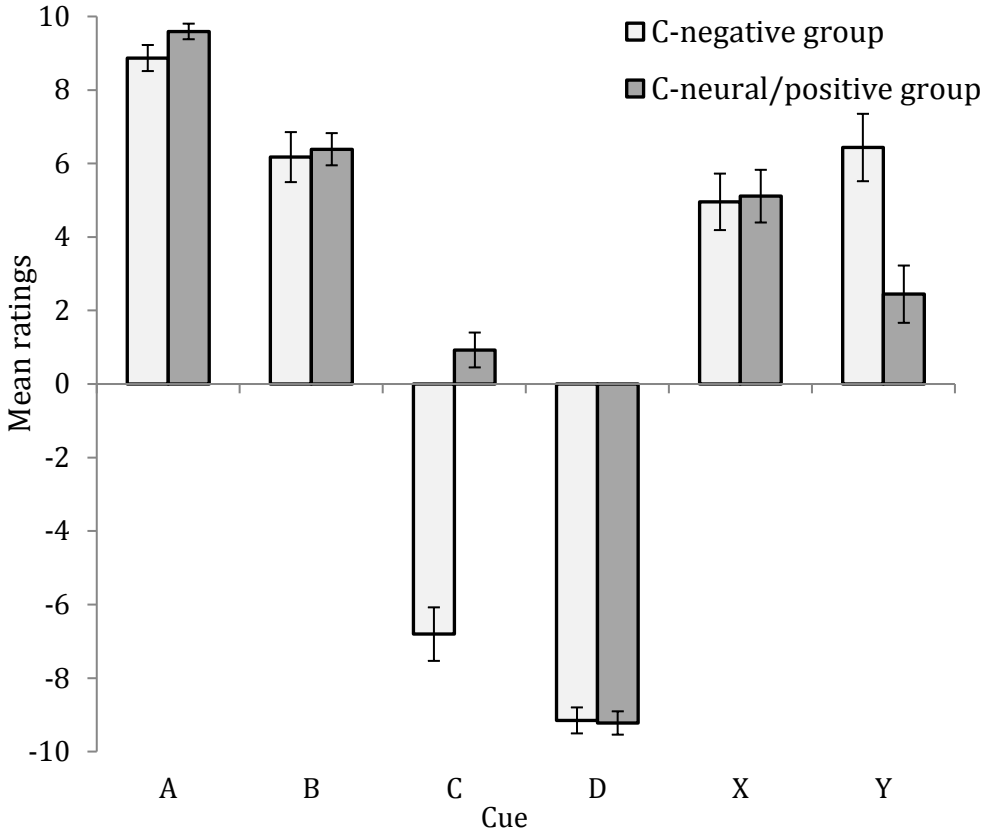

Figure 4.14. Mean hormone-change ratings at test in Experiment 10 split by group (± between-subjects SEM). Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

In order to explore whether the redundancy effect differed between the groups, a two-way (cue [X vs Y] by group [C-negative group vs C-neutral/positive group] ANOVA was conducted on the data. This revealed no significant effect of cue, $F(1, 48) = .63$, $p = .43$, $\eta_P^2 = .01$, a significant effect of group, $F(1, 48) = 5.18$, $p = .027$, $\eta_P^2 = .1$, and a significant interaction, $F(1, 48) = 7.71$, $p = .008$, $\eta_P^2 = .14$. Simple main effects analyses indicated that the redundancy effect was significant in the C-neutral/positive group, $t(26) = 2.56$, $p = .017$, $r = .48$ (X ratings – Y ratings: $M = 2.67$, $SEM = 1.04$) but not in the C-negative group, $t(22) = 1.4$, $p = .175$, $r = .29$, $BF_{01} = 9.91$ (one-tailed Bayesian t-test; X ratings – Y ratings: $M = -1.47$, $SEM = 1.06$), similarly to the previous experiment.

To summarise, in Experiment 10 I set out to replicate a number of findings from Experiment 9, which I was able to do. Firstly, I replicated the negative correlation between ratings for C and for Y, and a positive correlation between ratings for C and the magnitude of the redundancy effect. Since these correlations were significant, I split the data into two groups, based on participants' ratings for C. I found that, similarly to Experiment 9, the redundancy effect was larger for participants who had neutral or positive ratings for C than for participants who had negative ratings for C. Therefore, findings of this and the previous experiment indicated that the redundancy effect was related to the extent to which C was rated as inhibitory.

The redundancy effect was not observed in this experiment. It is unclear why it was not significant in this experiment but significant in Experiments 8 and 9; all three used the same task and the design of Experiment 9 was very similar to Experiment 10. However, in this experiment mean ratings for C were lower (M = -2.63) than in Experiment 8 (M = -1.56) and in Experiment 9 (M = -1.86). Therefore, individual variation resulting in lower ratings for C would have been related to higher ratings for Y, and a smaller redundancy effect which failed to reach significance.

In Experiment 11 I aimed to see whether it was possible to manipulate the inhibitory/neutral assumptions about C directly, and observe the corresponding differences in the magnitude of the redundancy effect.

Experiment 11

Experiments 8, 9, and 10 found a negative correlation between ratings for C and for Y, and Experiments 9 and 10, a positive correlation between ratings for C and the magnitude of the redundancy effect. In these experiments data-split analyses indicated that negative ratings for C resulted in a smaller redundancy effect than positive or neutral ratings. In Experiment 11 I set out to investigate whether ratings for Y and subsequently the

redundancy effect could be manipulated by overtly changing the causal nature of C. I expected that demonstrating that C had a negative effect on the outcome would result in higher ratings for Y and a smaller redundancy effect than demonstrating that C was neutral and had no effect on the outcome. Experiment 11 proceeded as follows. In Stage 1 participants were presented with A+/AX+/BY+/CY0/D- trials similarly to the previous experiment. Following Stage 1, they were asked to provide ratings for each cue (Test 1). In Stage 2, participants were shown the same trial types as in Stage 1, but with additional C-alone trials. For participants in Group Inhibitory, C was shown to lead to a decrease (C-), while for participants in Group Neutral, C was shown to lead to no change (C0) in hormone levels. Subsequently to this manipulation participants were asked to provide ratings for each cue once again (Test 2). The full design of Experiment 11 is shown in Table 4.5. I expected that the redundancy effect would be smaller in Group Inhibitory than in Group Neutral at Test 2.

Table 4.5

The design of Experiment 11. In Stage 2, participants in Group Inhibitory were presented with C- trials while participants in Group Neutral were presented with C0 trials.

Stage 1	Test 1	Stage 2	Test 2
A+	A	A+	A
AX+	B	AX+	B
BY+	C	BY+	C
CY0	D	C-/C0	D
D-	X	CY0	X
	Y	D-	Y
x 12	x 2	x 8	x 2

Method

Participants. Participants were 50 Plymouth University students aged 18-39 years (*M* = 20.88, *SD* = 4.26) and five were male. There were 25 participants in each group.

Materials. The materials and procedure in Experiment 11 were the same as in the previous experiment unless otherwise stated.

Images of the medicines were: blue, green, orange, pink, purple, and yellow. The medicines were randomly assigned to each type of cue (A, B, C, D, X, Y) for each participant.

Procedure. In Stage 1 participants were presented with 12 blocks of trials with the five trial types (A+/AX+/BY+/CY0/D-) appearing once per block, in a random order. Subsequently participants rated all of the individual medicines twice at Test 1. In Stage 2 participants were presented with eight blocks of trials, including C-alone trials. Group Inhibitory were presented with A+/AX+/BY+/C-/CY0/D-, while Group Neutral were presented with A+/AX+/BY+/C0/CY0/D-. At Test 2 participants were asked to rate each cue again, twice.

Results

Stage 1. Participants learned the contingencies in Stage 1. In the final epoch they responded correctly on 97.8% (*SD* = 12.88%) of the trials (Group Inhibitory *M* = 98.8%, *SD* = 7.68%; Group Neutral *M* = 96.8%, *SD* = 16.47%). These data are presented in the higher panels of Figure 4.15. A two-way (trial type [A+, AX+, BY+, C0/-, CY0, D-] by group [Group Inhibitory vs Group Neutral]) ANOVA on the final epoch of Stage-1 responses revealed a significant effect of trial type, $F(4, 192) = 2151.85$, $p < .001$, $\eta_P^2 = .98$, no significant effect of group, $F(1, 48) = .08$, $p = .783$, $\eta_P^2 < .01$, and no significant interaction, $F(4, 192) = .62$, $p = .65$, $\eta_P^2 = .01$.

Stage 2. In the final epoch of Stage 2 participants responded correctly on 99% (SD = 9.09%) of the trials (Group Inhibitory: *M* = 99%, *SD* = 9.1%; Group Neutral: *M* = 99%, *SD* = 9.1%). These data are shown in the lower left (Group Inhibitory) and lower right (Group Neutral) panels of Figure 4.15.

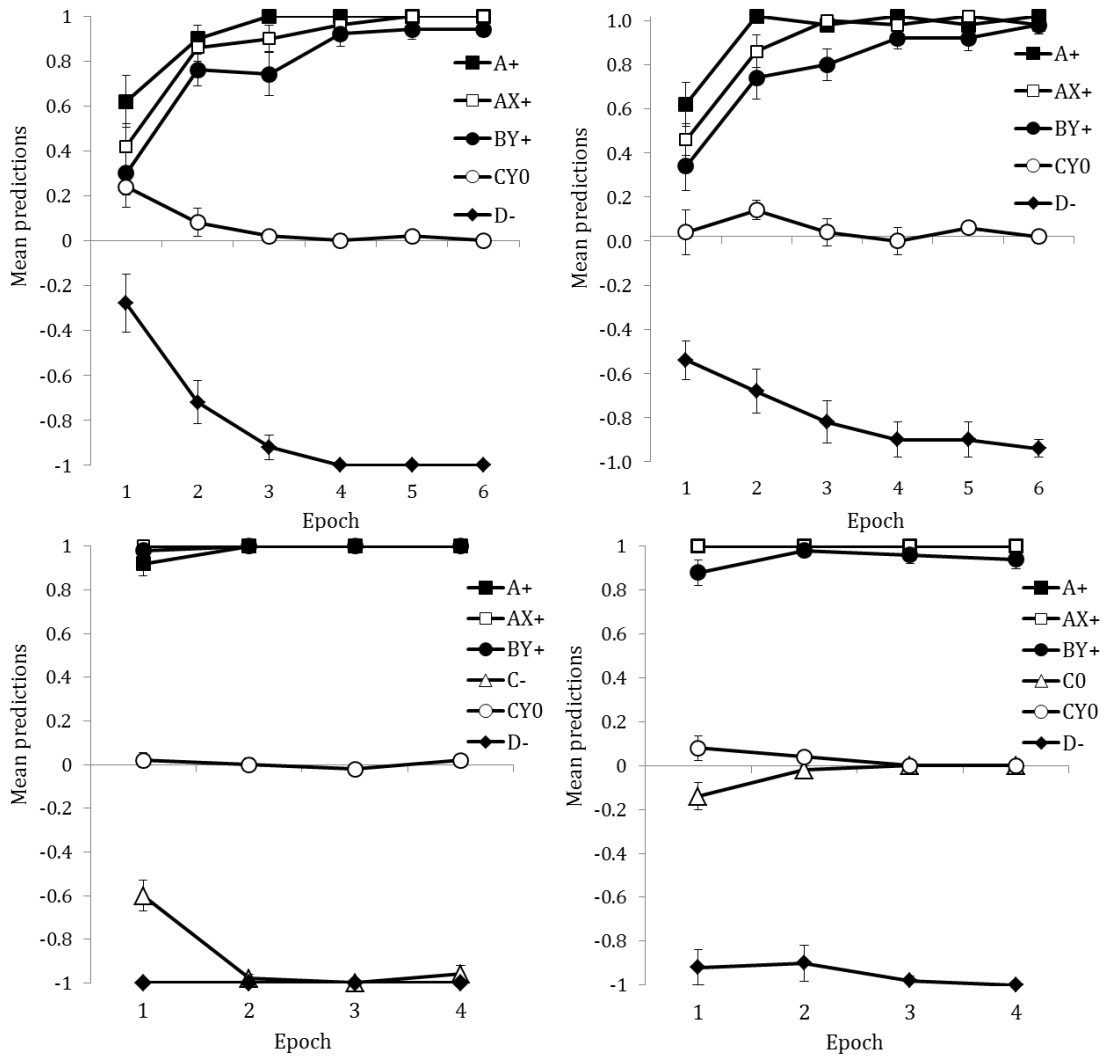

Figure 4.15. Mean hormone-change predictions in Stage 1 and Stage 2 in Experiment 11 (± between-subjects SEM). Higher left panel: Stage-1 responses in Group Inhibitory. Higher right panel: Stage-1 responses in Group Neutral. Lower left panel: Stage-2 responses in Group Inhibitory. Lower right panel: Stage-2 responses in Group Neutral. Positive values refer to a prediction of an increase, zero values refer to a prediction of no change, and negative values refer to a prediction of a decrease in hormone levels.

A two-way (trial type [A+, AX+, BY+, C0/-, CY0, D-] by group [Group Inhibitory vs Group Neutral]) ANOVA on the final epoch of Stage-2 responses indicated significant effects of trial type, $F(5, 240) = 4617.24$, $p < .001$, $\eta_p^2 = .99$, and group, $F(1, 48) = 200.28$, $p < .001$, $\eta_p^2 = .81$, and a significant interaction, $F(5, 240) = 242.19$, $p < .001$, $\eta_p^2 = .84$. Importantly, the only significant differences between the groups were for responses on C-alone trials, with participants in Group Inhibitory predicting a greater decrease in hormone levels than participants in Group Neutral, $t(24) = 24$, $p < .001$, $r = .98$, consistently with the manipulation. Responses for the other trial types did not differ significantly between the groups, ts ≤ 1.37, ps ≥ .185, rs ≤ .19.

Test. Mean hormone-change ratings in both groups at Test 1 and Test 2 are shown in the higher and lower panels of Figure 4.16. Due to lengthy analyses for this experiment, this section includes a shortened version. For complete analyses see Appendix 4A.

Exploring data with a three-way ANOVA. A three-way (cue [A, B, C, D, X, Y] by test [1 vs 2] by group [Group Inhibitory vs Group Neutral]) ANOVA revealed a significant three-way interaction, $F(2.8, 134.55) = 11.47$, $p < .001$, $\eta_p^2 = .19$.

Differences between the groups at Test 1 and at Test 2. Firstly, to establish that ratings for the cues did not differ between the groups at Test 1, a two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-1 data was conducted. It revealed a significant effect of cue, $F(3.09, 148.24) = 157.29$, $p < .001$, $\eta_p^2 = .77$. Bonferroni-corrected paired comparisons indicated that while ratings between B, X, and Y did not differ significantly, ts ≤ 2.75, ps ≥ .134, rs ≤ .37, ratings between all other cues did, ts ≥ 6.01, ps < .001, rs ≥ .65. There was no significant effect of group, $F(1, 48) = .85$, $p = .361$, $\eta_p^2 = .02$, and no significant interaction, $F(3.09, 148.24) = .8$, $p = .502$, $\eta_p^2 = .02$.

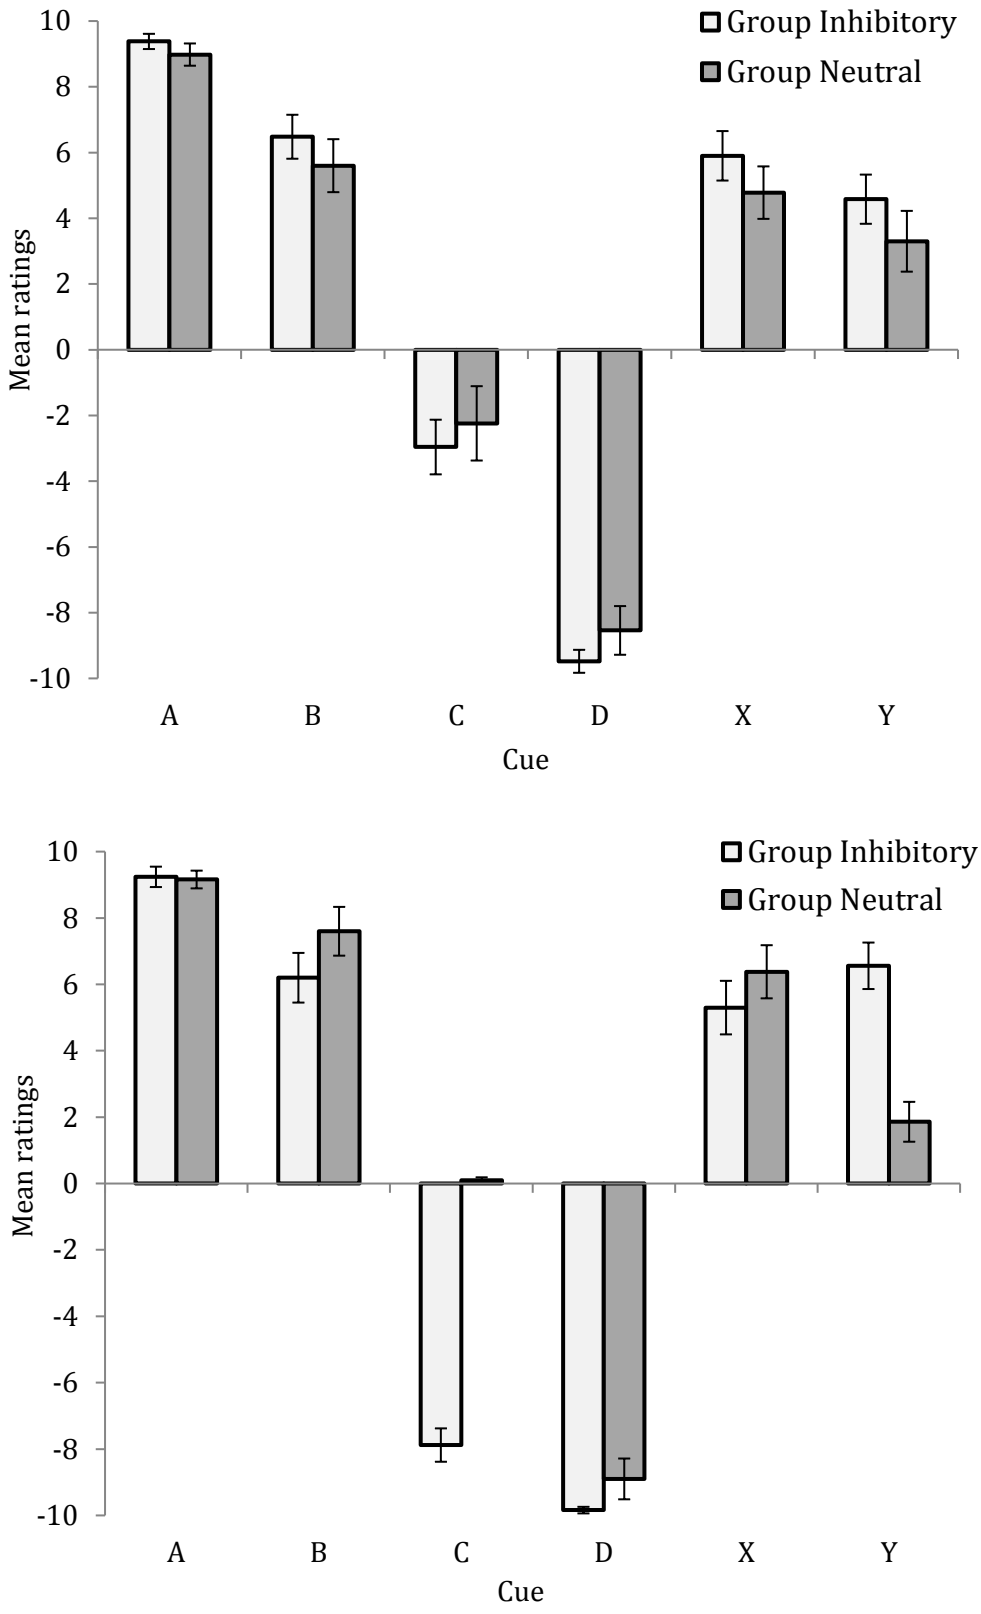

Figure 4.16. Mean hormone-change ratings in Experiment 11 (± between-subjects SEM). Higher panel: Responses from Test 1. Lower panel: Responses from Test 2. Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

To investigate whether establishing C as inhibitory or neutral affected ratings for cues at Test 2, a two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data was conducted. It revealed a significant effect of cue, $F(3.17, 152.27) = 299.27$, $p < .001$, $\eta_P^2 = .86$, a significant effect of group, $F(1, 48) = 12.38$, $p = .001$, $\eta_P^2 = .21$, and a significant interaction, $F(3.17, 152.27) = 23.8$, $p < .001$, $\eta_P^2 = .33$. Simple main effects analyses indicated that the groups did not differ in their ratings for A, B, D, and X, ts ≤ 1.51, ps ≥ .137, rs ≤ .21, however Group Inhibitory had lower ratings for C, $t(25.42) = 15.64$, $p < .001$, $r = .91$, and higher ratings for Y, $t(48) = 5.09$, $p < .001$, $r = .59$, than Group Neutral.

Therefore while ratings at Test 1 did not differ between the groups, at Test 2 Group Inhibitory had lower ratings for C and higher ratings for Y than Group Neutral, consistently with the manipulation.

Exploring the redundancy effect with a three-way ANOVA. A three-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral] by test [1 vs 2]) ANOVA revealed a significant three-way interaction, $F(1, 48) = 11.95$, $p = .001$, $\eta_P^2 = .2$.

Differences in X and Y ratings between the tests in each group. To test how the manipulation affected ratings for X and for Y between the tests, 2 two-way ANOVA tests were conducted, to compare ratings for these cues at Test 1 with ratings at Test 2 in each group.

In Group Neutral, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed a significant effect of cue, $F(1, 24) = 14.74$, $p = .001$, $\eta_P^2 = .38$, with higher ratings for X than for Y, no significant effect of test, $F(1, 24) = .02$, $p = .899$, $\eta_P^2 < .01$, and a significant interaction, $F(1, 24) = 6.8$, $p = .015$, $\eta_P^2 = .22$. Simple main effects analyses indicated that ratings for X were higher at Test 2 than at Test 1, $t(24) = 3.03$, $p = .006$, $r = .53$, while ratings for Y did not differ between the tests, $t(24) = 1.32$, $p = .199$, $r = .26$, $BF_{01} = 2.18$.

As a result, the redundancy effect was larger at Test 2, $t(24) = 4.87$, $p < .001$, $r = .7$ (X ratings – Y ratings: $M = 4.52$, $SEM = .93$), than at Test 1, $t(24) = 1.45$, $p = .16$, $r = .28$ (X ratings – Y ratings: $M = 1.48$, $SEM = 1.02$), in Group Neutral. However, it appeared that higher ratings for X at Test 2 than at Test 1 were responsible for this result.

In Group Inhibitory, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed no significant effect of cue, $F(1, 24) < .01$, $p = .97$, $\eta_p^2 < .01$, no significant effect of test, $F(1, 24) = 2.63$, $p = .118$, $\eta_p^2 = .1$, but a significant interaction, $F(1, 24) = 5.19$, $p = .032$, $\eta_p^2 = .18$. Simple main effects analyses indicated ratings for X did not differ between the tests, $t(24) = 1.34$, $p = .191$, $r = .26$, while ratings for Y were higher at Test 2 than at Test 1, $t(24) = 2.21$ $p = .037$, $r = .41$.

As a result, the redundancy effect was smaller at Test 2, $t(24) = -1.27$, $p = .217$, $r = -.25$, $BF_{01} = 9.79$ (one-tailed Bayesian t-test; X ratings – Y ratings: $M = -1.26$, $SEM = .99$), than at Test 1, $t(24) = 1.42$, $p = .17$, $r = .28$ (X ratings – Y ratings: $M = 1.32$, $SEM = .93$), in Group Inhibitory.

In order to compare the magnitude of the redundancy effect (X ratings – Y ratings) at Test 2 between the groups, a t-test was performed on the data. This confirmed that the magnitude of the redundancy effect was larger in Group Neutral than Group Inhibitory, $t(48) = 4.25$, $p < .001$, $r = .52$.

Therefore, while the redundancy effect did not differ between the groups Test 1, at Test 2, the redundancy effect was larger in Group Neutral than in Group Inhibitory. However, it appeared that higher ratings for X at Test 2 than at Test 1 in Group Neutral contributed to this result.

To summarise, in Experiment 11 I set out to determine whether ratings for Y and the magnitude of the redundancy effect could be influenced by establishing C as inhibitory or neutral. At Test 1 ratings for X and for Y did not differ between the groups. At Test 2 however, higher ratings for Y were observed in Group Inhibitory than in Group Neutral. As

a result, the redundancy effect was smaller in Group Inhibitory than in Group Neutral. However, while results indicated that the causal status of C influenced the redundancy effect at Test 2, there were several factors that contributed to this change outside of the main manipulation.

Firstly, a larger redundancy effect in Group Neutral than in Group Inhibitory at Test 2 appeared to have been partly due to higher ratings for X at Test 2 than at Test 1 in Group Neutral. It is unclear why this increase in ratings for X was observed.

Secondly, in Group Inhibitory, the redundancy effect was not significant at Test 1. While ratings for Y were higher at Test 2 than at Test 1, a smaller redundancy effect at Test 2 could have been due partly to the non-significant redundancy effect observed at Test 1 in this group.

Overall, manipulation of the causal status of C was partly successful at influencing the redundancy effect. Establishing C as inhibitory resulted in increased ratings for Y and a smaller redundancy effect than establishing it as neutral. However, because of the problems detailed above, I decided to replicate this study. In the next experiment I aimed to see whether lower ratings for Y in Group Inhibitory at Test 2 than at Test 1 would be replicated. In addition I investigated whether in Group Neutral higher ratings for X would persist, and whether lower ratings for Y would be observed at Test 2 than at Test 1, with a greater sample of participants.

Experiment 12

The design of Experiment 12 was identical to Experiment 11 (Table 4.5).

Method

Participants. Participants were 70 Plymouth University students aged 18-41 (*M* = 20.64, *SD* = 4.09) and 11 were male. There were 34 participants in Group Inhibitory and 36 participants in Group Neutral.

The materials and the procedure in this experiment were identical to Experiment 11.

Results

Stage 1. Participants learned the contingencies in Stage 1. They responded correctly in the final epoch on 99.57% (*SD* = 4.62%) of the trials (Group Inhibitory: *M* = 99.71%, *SD* = 3.83%; Group Neutral: *M* = 99.44%, *SD* = 5.26%). These data are shown in the higher panels of Figure 4.17. A two-way (trial type [A+, AX+, BY+, CY0, D-] by group [Group Inhibitory vs Group Neutral]) ANOVA on the final epoch of Stage-1 predictions indicated a significant effect of trial type, $F(4, 272) = 26046.09$, $p < .001$, $\eta_P^2 > .99$, no significant effect of group, $F(1, 68) = .35$, $p = .558$, $\eta_P^2 < .01$, nor a significant interaction, $F(4, 272) = .31$, $p = .872$, $\eta_P^2 < .01$.

Stage 2. Lower panels of Figure 4.17 show predicted changes in hormone levels in Stage 2 for participants in Group Inhibitory (lower left panel) and Group Neutral (lower right panel). In the final epoch of Stage 2 correct responses were made on 99.93% (*SD* = 8.74%) of the trials (Group Inhibitory: *M* = 98.04%, *SD* = 12%; Group Neutral: *M* = 99.77%, *SD* = 3.4%). A two-way (trial type [A+, AX+, BY+, C-/C0, CY0, D-] by group [Group Inhibitory vs Group Neutral]) ANOVA on the final epoch of Stage-2 predictions revealed a significant effect of trial type, $F(5, 340) = 6454.23$, $p < .001$, $\eta_P^2 = .99$, a significant effect of group, $F(1, 68) = 643.75$, $p < .001$, $\eta_P^2 = .9$, and a significant interaction, $F(5, 340) = 330.09$, $p < .001$, $\eta_P^2 = .83$. Importantly, the only significant differences between the groups were observed on C-alone trials; as expected, Group Inhibitory predicted a greater decrease in hormone levels than Group Neutral, $t(33) = 67$, $p < .001$, $r = .99$. Ratings for the other cues did not differ significantly between the groups, $ts \leq 1$, $ps \geq .325$, $rs \leq .12$.

Test. Figure 4.18 shows the mean hormone-change ratings in both groups at Test 1 (higher panel) and at Test 2 (lower panel).

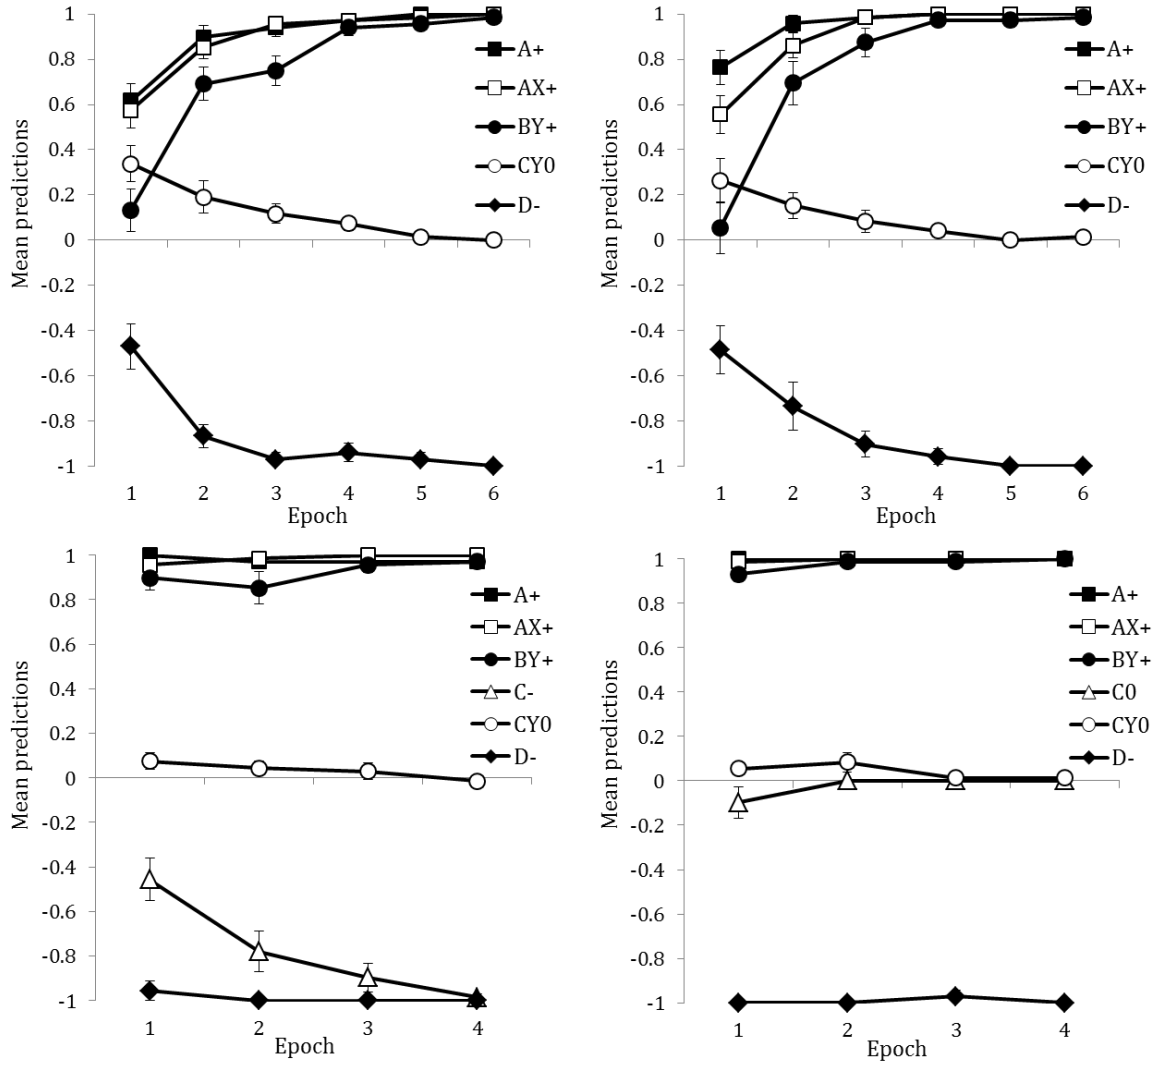

Figure 4.17. Mean hormone-change predictions in Stage 1 and in Stage 2 in Experiment 12 (± between-subjects SEM). Higher left panel: Stage-1 responses in Group Inhibitory. Higher right panel: Stage-1 responses in Group Neutral. Lower left panel: Stage-1 responses in Group Inhibitory. Lower right panel: Stage-2 responses in Group Neutral. Positive values refer to a prediction of an increase, zero values refer to a prediction of no change, and negative values refer to a prediction of a decrease in hormone levels.

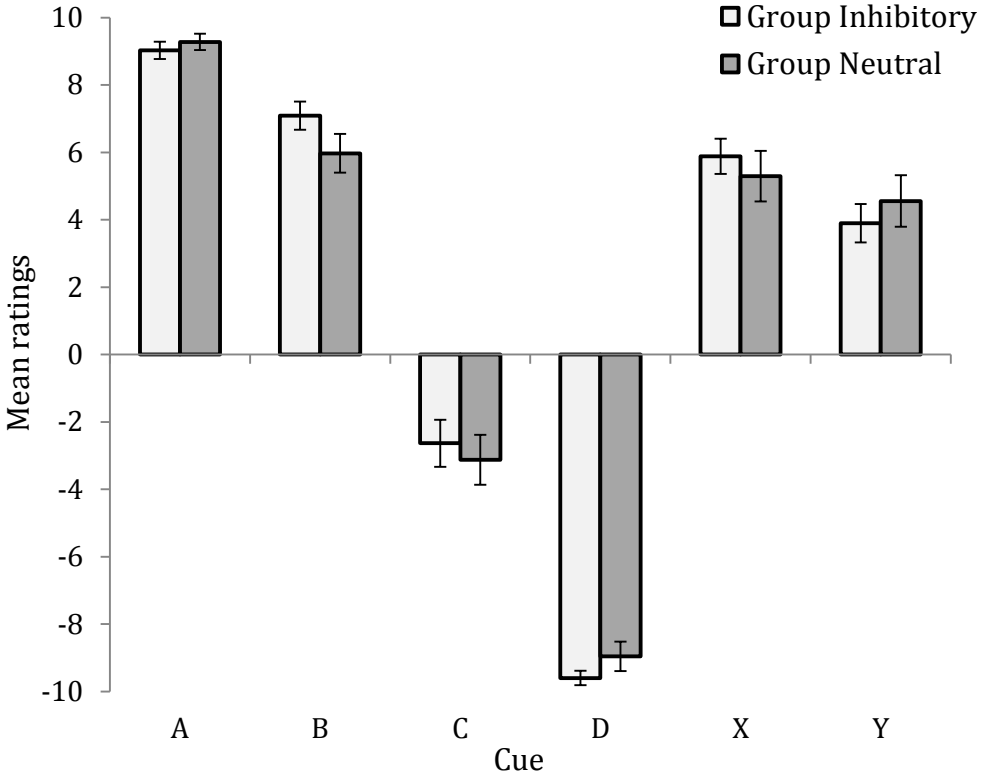

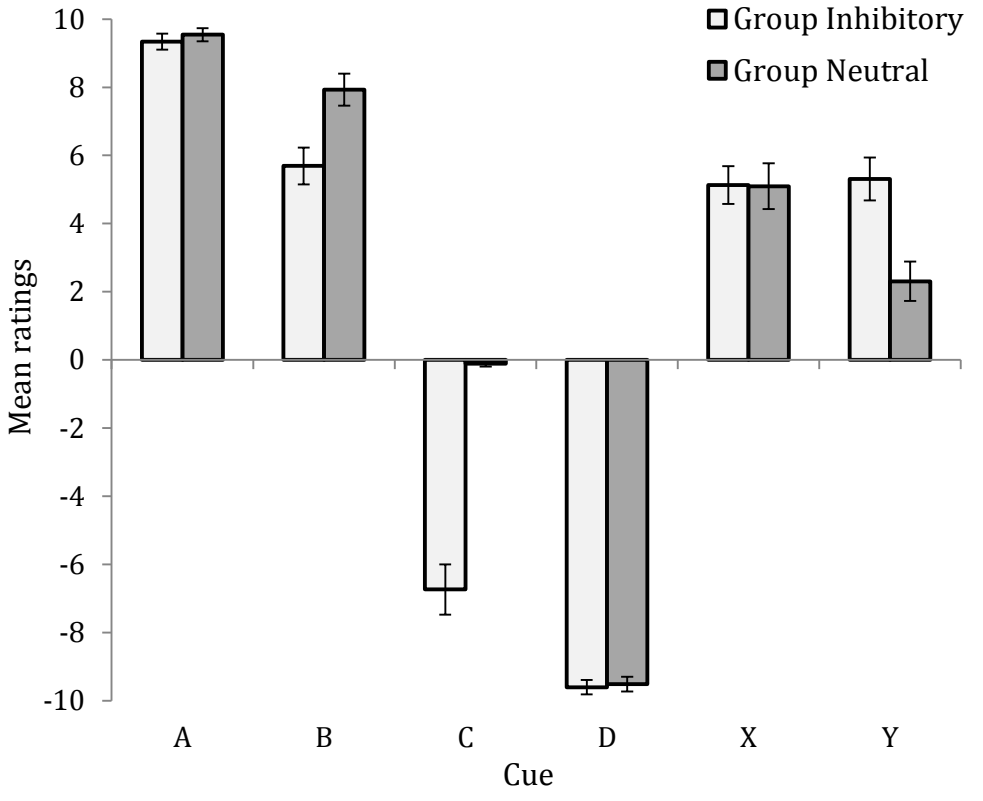

Figure 4.18. Mean hormone-change ratings in Experiment 12 (± between-subjects SEM). Higher panel: Responses from Test 1. Lower panel: Responses from Test 2. Positive values refer to an increase, negative values to a decrease, and zero values to no change in hormone levels.

Because of lengthy analyses, this section includes a shortened version. The same analyses were performed in this experiment as in the previous experiment. For full details of the analyses see Appendix 4B.

Exploring data with a three-way ANOVA. A three-way (cue [A, B, C, D, X, Y] by test [1 vs 2] by group [Group Inhibitory vs Group Neutral]) ANOVA revealed a significant three-way interaction, $F(2.76, 187.64) = 19.68$, $p < .001$, $\eta_P^2 = .22$.

Differences between the groups at Test 1 and at Test 2. In order to make sure that there were no differences between the groups at Test 1, a two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA was conducted on Test-1 data. It revealed a significant effect of cue, $F(3.04, 206.66) = 314.93$, $p < .001$, $\eta_P^2 = .82$. Bonferroni-corrected paired comparisons indicated that ratings between B and X, and between X and Y, did not differ significantly, ts ≤ 2.57, ps ≥ .167, rs ≤ .3, while ratings between all other cues did, ts ≥ 3.98, ps ≤ .002, rs ≥ .43. There was no significant effect of group, $F(1, 68) = .14$, $p = .714$, $\eta_P^2 < .01$, and no significant interaction, $F(5, 340) = .88$, $p = .496$, $\eta_P^2 = .02$.

In order to check whether establishing C as either inhibitory or neutral affected ratings for the cues at Test 2, a two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data was conducted. It revealed a significant effect of cue, $F(3.83, 260.47) = 464.98$, $p < .001$, $\eta_P^2 = .87$, a significant effect of group, $F(1, 68) = 13.78$, $p < .001$, $\eta_P^2 = .17$, and a significant interaction, $F(3.83, 260.47) = 23.63$, $p < .001$, $\eta_P^2 = .26$. Paired comparisons indicated that the groups did not differ in their ratings for A, D, and X, ts ≤ .68, ps ≥ .5, rs ≤ .08, but Group Inhibitory had lower ratings for B, $t(68) = 3.18$, $p = .002$, $r = .36$, lower ratings for C, $t(33.83) = 8.95$, $p < .001$, $r = .74$, and higher ratings for Y, $t(68) = 3.57$, $p = .001$, $r = .4$, than Group Neutral.

Therefore, ratings for cues at Test 1 did not differ between the groups, while at Test 2 Group Inhibitory had lower ratings for C and higher ratings for Y than Group Neutral,

consistently with the manipulation. Ratings for B were also lower in Group Inhibitory than in Group Neutral at Test 2.

Exploring the redundancy effect with a three-way ANOVA. A three-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral] by test [1 vs 2]) ANOVA revealed a significant three-way interaction, $F(1, 68) = 14.09$, $p < .001$, $\eta_P^2 = .17$.

Differences in X and Y ratings between the tests in each group. To test how the manipulation affected ratings for X and for Y between the tests, 2 two-way ANOVA tests were conducted, to compare ratings for these cues at Test 1 with ratings at Test 2 in each group.

In Group Neutral, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed a significant effect of cue, $F(1, 35) = 6.37$, $p = .016$, $\eta_P^2 = .15$, with higher ratings for X than for Y, a significant effect of test, $F(1, 35) = 4.48$, $p = .041$, $\eta_P^2 = .11$, and a significant interaction, $F(1, 35) = 8.72$, $p = .006$, $\eta_P^2 = .2$. Simple main effects analyses indicated that ratings for X did not differ between the tests, $t(35) = .35$, $p = .732$, $r = .06$, while ratings for Y were lower at Test 2, $t(35) = 2.92$, $p = .006$, $r = .44$, than at Test 1.

As a result, the redundancy effect was larger at Test 2, $t(35) = 3.8$, $p = .001$, $r = .54$ (X ratings – Y ratings: $M = 2.79$, $SEM = .76$), than at Test 1, $t(35) = .89$, $p = .378$, $r = .15$, $BF_{01} = 3.86$ (X ratings – Y ratings: $M = .74$, $SEM = .85$), in Group Neutral.

In Group Inhibitory, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed no significant effect of cue, $F(1, 33) = 2.57$, $p = .119$, $\eta_P^2 = .07$, no significant effect of test, $F(1, 33) = .63$, $p = .433$, $\eta_P^2 = .02$, but a significant interaction, $F(1, 33) = 5.89$, $p = .021$, $\eta_P^2 = .15$. Simple main effects analyses indicated that ratings for X did not differ between the tests, $t(33) = 1.64$, $p = .112$, $r = .27$, while there was a non-significant trend for higher ratings for Y at Test 2 than at Test 1, $t(33) = 1.93$, $p = .062$, $r = .32$.

As a result, the redundancy effect was smaller at Test 2, $t(33) = -.22$, $p = .827$, $r = .04$, $BF_{01} = 5.32$ (X ratings – Y ratings: $M = -.18$, $SEM = .8$), than at Test 1, $t(33) = 3.18$, $p = .003$, $r = .48$ (X ratings – Y ratings: $M = 1.99$, $SEM = .62$), in Group Inhibitory.

In order to compare the magnitude of the redundancy effect (X ratings – Y ratings) at Test 2 between groups, a t-test was performed on the data. This confirmed that the magnitude of the redundancy effect was larger in Group Neutral than Group Inhibitory, $t(68) = 2.73$, $p = .008$, $r = .31$.

Therefore, while the groups did not differ in the magnitude of the redundancy effect at Test 1, at Test 2, the redundancy effect was larger in Group Neutral than in Group Inhibitory as a result of changes in ratings for Y.

To summarise, in Experiment 12 I set out to replicate findings from Experiment 11. I found that manipulating the causal status of C had an effect on the ratings for Y; they were higher when C was established as inhibitory than when C was established as neutral. As a result, establishing C as neutral resulted in a larger redundancy effect than establishing C as inhibitory.

Higher ratings for X at Test 2 than at Test 1 in Group Neutral, observed in Experiment 11, were not present in these data. Therefore, a larger redundancy effect in Group Neutral than in Group Inhibitory at Test 2 was due to changes in ratings for Y only. Ratings for B were also higher in Group Neutral than in Group Inhibitory at Test 2. It may not be surprising that manipulating the causal status of C affected ratings for B, however. When C was established as neutral, resulting in lower ratings for Y, participants may have assumed that B caused the outcome on BY+ trials. When C was established as inhibitory, resulting in higher ratings for Y, participants may have assumed that Y caused the outcome on BY+ trials.

Discussion

In this chapter I set out to investigate whether a lack of inhibition for C in the design A+/AX+/BY+/CY-, contributed to the redundancy effect. The Rescorla-Wagner (1972) model predicts a stronger association with the outcome for Y than for X. However the predicted positive associative strength for Y is contingent on C becoming inhibitory.

In Experiment 7 I failed to find evidence of inhibition for C in an allergist task and the redundancy effect was significant, suggesting that a lack of inhibition could have contributed to the redundancy effect in this experiment. In the following experiments I used a task in which inhibition was more plausible. The outcome in this task was levels of a fictional hormone which could decrease, not change, or increase, representing inhibitory, neutral, and excitatory effects on the outcome, respectively.

In Experiments 8, 9, and 10 I found a negative correlation between ratings for C and for Y, indicating that a prediction of a decrease for C was related to a prediction of an increase for Y. In Experiments 9 and 10 I also observed a significant positive correlation between ratings for C and the magnitude of the redundancy effect, indicating that lower ratings for C were related to a smaller redundancy effect. In these experiments, using data-split analyses, I demonstrated that the redundancy effect was smaller in the group which had negative ratings for C, than in the group which had zero or positive ratings for C. In Experiment 12, I directly manipulated the causal status of C and found the corresponding change in the redundancy effect, consistently with predictions derived from Experiments 9 and 10. The redundancy effect was larger in the group in which C was established as neutral (Group Neutral) than in the group in which C was established as inhibitory (Group Inhibitory).

Taken together, results from these experiments indicated that causal ratings for Y, and the magnitude of the redundancy effect, were related to whether C was neutral or inhibitory. Therefore the previous demonstrations of the redundancy effect in this thesis and elsewhere using the allergist task (Uengoer et al., 2013) may have been at least partly

attributable to a lack of inhibition for C. If C did not become inhibitory, then Y may have been seen as neutral, contributing to the redundancy effect.

These findings are not easy to reconcile with predictions of the Rescorla-Wagner (1972) model. This model predicts that Y will gain a small amount of positive associative strength while X will have zero associative strength at asymptote. However, the association between Y and the outcome relies on C becoming inhibitory. If, in the allergist task, C does not gain inhibitory associative strength but is seen as neutral as indicated in Experiment 7, this does not protect Y from extinction. In this case, the Rescorla-Wagner model predicts that both X and Y will have zero associative strength. In Experiment 12, when C was established as neutral, the redundancy effect was still observed, in an apparent contradiction to these predictions. When C is inhibitory, the model predicts a stronger relationship between Y and the outcome than X and the outcome. Establishing C as inhibitory in Experiment 12 did not enable me to obtain the predicted result either; causal ratings for X and Y did not differ significantly. However, there is one simple way to reconcile these findings with the predictions of the Rescorla-Wagner model. If learning for A is assumed to have been pre-asymptotic, then X would have gained some associative strength as a result. Therefore, when C was established as neutral, resulting in lower ratings for Y, the positive associative strength for X would have resulted in the redundancy effect being observed in Experiment 12. When C was established as inhibitory, resulting in higher ratings for Y, ratings for X would not have differed from Y in Experiments 11 and 12. However, learning for A was high in experiments within this chapter (99.68% correct responding within the final epoch on A+ trials), making it difficult to claim that learning was incomplete for this cue. If learning is assumed to have been at asymptote however, it is difficult to see how these findings could be reconciled with the predictions of the Rescorla-Wagner model fully.

However, it is possible that the redundancy effect was not significant in these experiments for an alternative reason. Findings in Chapter 3 showed that participants'

uncertainty about the causal status of X contributed to the redundancy effect. The manipulation in Experiment 12 focused on Y only, without aiming to resolve ambiguity about the causal status of X, therefore it may not be surprising that ratings for Y were not higher than for X when C was established as inhibitory. It is possible that if two manipulations were used, one targeting X and one targeting Y, this would result in a stronger relationship with the outcome for Y than for X. This would be an important future direction to pursue relating to the redundancy effect research.

It is also worth noting that while mean ratings for C in Experiments 8, 9, 10, and Test 1 of Experiments 11 and 12 indicated weak inhibition, participants varied in their ratings for C: some rated it as highly negative and others as neutral. Since both of these interpretations are consistent with the contingencies in the task then perhaps it is not surprising that participants were divided in their assumptions of the causal status of C. However, I suggest that one possible reason for these differences could have been individual variation in the extent to which participants relied on single and summed error. I will return to this issue in the next chapter.

In Experiment 7, mean ratings for C indicated that inhibition was not obtained for this cue using the allergist task. In Experiment 8 I found that in a task in which inhibition was more plausible, mean ratings for C indicated some inhibitory associative strength for this cue. It may be interesting to consider how employing a task in which inhibition is more plausible may affect other associative-learning effects. For example extinction of inhibitory cues is rarely observed under ordinary circumstances (e. g. Yarlas, Cheng & Holyoak, 1995). However, Melchers, Wolff and Lachnit (2006) found that, contrary to previous results, extinction of a conditioned inhibitor was observed in a task similar to the one used in this chapter, in which cues could lead to an increase or a decrease in hormone levels.

Interestingly, it appeared that simply using a task which promoted inhibition was not enough to eliminate the redundancy effect, however the data for this were inconsistent. In four out of five experiments using this task (Experiments 8, 9, 10, 11: Test 1, 12: Test 1),

the redundancy effect was obtained for the overall data. However, it was not significant in Experiment 10, in each individual group at Test 1 of Experiment 11, and in one of the groups at Test 1 of Experiment 12. In the discussion of Experiment 10 I suggested that one reason why the redundancy effect was not significant in this experiment was lower ratings for C (M = -2.63) than in Experiments 8 (M = -1.56) and 9 (M = -1.86), which could have resulted in higher ratings for Y and therefore a smaller redundancy effect. This may also apply to Test-1 data from Experiments 11 and 12. In both experiments, the redundancy effect was obtained overall at Test 1 and mean ratings for C were as low as in Experiment 10 (Experiment 11: M = -2.6, Experiment 12: M= -2.88). However, the magnitude of the redundancy effect (X ratings – Y ratings) was also smaller in these experiments (Experiment 11: M = 1.4, Experiment 12: M = 1.37) than in Experiments 8 (M = 2.5) and 9 (M = 2.59). Therefore, it is possible that low ratings for C contributed to the smaller redundancy effect in these experiments. As such, it is possible that using the hormone task encouraged more people to consider that C was likely to be inhibitory and Y was likely to be causal than in the allergist task. However, this is a speculative claim and should be investigated experimentally. It is also ambiguous whether ratings for C reflected inhibition in Experiments 9, 10, 11, and 12. Unfortunately, I did not have appropriate control cues to test for this. It is possible that low ratings for C could have reflected, for example, a general tendency for participants to have lower ratings for all cues.

To summarise, this chapter found the following. Firstly, inhibition for C was not obtained using the allergist task in Experiment 7. Secondly, in a task in which inhibition was more plausible, participants varied in their ratings for C. These ratings were negatively related to ratings for Y, as shown in Experiments 8, 9, and 10. Ratings for C were also positively related to the magnitude of the redundancy effect as found in Experiments 9 and 10. Experiment 12 showed that establishing C as an inhibitory cue which led to a decrease in the outcome resulted in a smaller redundancy effect than establishing C as a neutral cue which led to no change in the outcome. It follows that the

previous demonstrations of the redundancy effect within experiments in Chapters 2 and 3 and in Uengoer et al. (2013) could have been partly due to a lack of inhibition for C.

Chapter 5: Discussion

In a human causal learning task, Uengoer et al. (2013) found that the blocked cue X was indicated to cause the outcome to a greater extent than the uncorrelated cue Y in a design A+/AX+/BY+/CY-. This finding has been referred to as the redundancy effect (Jones & Pearce, 2015) and it has also been observed in rats and pigeons (Jones & Pearce, 2015; Pearce et al., 2012). This effect is puzzling because the Rescorla-Wagner (1972) model predicts the opposite result; that Y will have a stronger association with the outcome than X. Because of its use of a summed error-term, this model predicts that firstly, learning the association between X and the outcome will be restricted by the presence of A, which predicts the outcome very well. Secondly, C is predicted to become a weak inhibitor, allowing Y to retain some positive associative strength from BY+ trials.

In contrast, Mackintosh's (1975) theory of selective attention could account for the redundancy effect. This model uses a single error-term, and therefore predicts that the blocked cue X will have a stronger relationship with the outcome than the uncorrelated cue Y. Associative strength will also be determined by attention, which will decline for both X and Y as they are worse predictors of the outcome than their companion cues A, B, and C. Provided that the decline in attention for X does not exceed that for Y, Mackintosh's theory could account for the redundancy effect. I explored whether the redundancy effect could be a consequence of differences in attention between blocked and uncorrelated cues in Chapter 2. In this chapter I used eye-tracking to compare eye-gaze durations between these cues during learning. In Experiment 1 I failed to find any differences in gaze durations for X and for Y, once trial-duration differences between these cues had been accounted for. These findings were replicated in Experiment 2. In this experiment I also failed to find evidence for differences in gaze durations when blocked and uncorrelated cues were presented on the screen at the same time, suggesting no differences in selective attention between these cues. Therefore findings from Chapter 2 indicated that an attentional explanation for the redundancy effect was unlikely.

In Chapter 3, I explored whether differences in participants' certainty about the causal status of X and of Y contributed to the redundancy effect. In Uengoer et al.'s (2013) experiments and in Experiments 1 and 2 of this thesis, X received ratings at approximately in the middle of the 11-point, 0-10 causal-rating scale while Y received low ratings. Uengoer et al. interpreted this difference to reflect a stronger association between X and the outcome than between Y and the outcome. Alternatively, if ratings in the middle of the scale best reflected the point of uncertainty on this scale, ratings for X could have indicated that participants were uncertain about whether or not X caused the outcome. Participants' uncertainty about the causal status of the blocked cue has been suggested previously (e. g. Cheng, 1997; Cheng & Holyoak, 1995; De Houwer, Beckers & Glautier, 2002; Lovibond et al., 2003; Mitchell & Lovibond, 2002; Vandorpe & De Houwer, 2006; Waldmann & Holyoak, 1992). Waldman and Holyoak argued that participants will be particularly uncertain about the blocked cue when the outcome in a given task is of maximal intensity, in other words, when there is a ceiling on the outcome. In this case, in a design A+/AX+, participants cannot verify whether X has any causal effects on the outcome in addition to A on AX+ trials. Because there was a ceiling on the outcome in Uengoer et al.'s studies and in Experiments 1 and 2 of this thesis, participants' uncertainty about the causal status of X could have contributed to the redundancy effect observed in these experiments.

This is what I set out to explore in Chapter 3. Findings from Experiment 3 showed that participants were less certain about their causal ratings for X than for Y; therefore it was possible that this difference contributed to the redundancy effect. In Experiment 4 I used outcome-additivity training to disambiguate the effects that X had on the outcome. Participants in Group Additive were presented with two causal cues (G+, H+) which when presented together, led to a stronger outcome (GH++). Participants in Group Non-additive were presented with two causal cues (G+, H+) which when presented together, led to the same outcome as either cue (GH+). If participants in Group Additive extracted and applied the additive rule to X, they could have deduced that X was not a cause of the outcome. I found that ratings for X were lower and blocking was larger in Group Additive than Group

Non-additive, consistently with other research which used outcome-additivity training (e. g. Beckers et al., 2005; Livesey & Boakes, 2004; Lovibond et al., 2003; Mitchell & Lovibond, 2002; Mitchel et al., 2005). Consequently, the redundancy effect was also smaller in Group Additive than in Group Non-additive. Findings of this experiment indicated that participants' uncertainty about X contributed to the redundancy effect. However, the redundancy effect was significant in both Group Additive and Group Non-additive.

In Experiments 5 and 6 I varied the rate of the outcome to investigate whether this would influence ratings for X and the redundancy effect. This was based on suggestions by Livesey et al. (2013), who predicted that in the absence of other information, participants use outcome rate to determine the extent with which cues cause the outcome. Because participants are particularly uncertain about the blocked cue, causal ratings for this cue should vary with outcome rate. Consequently, blocking was predicted to be inversely related to outcome rate; it would be large when outcome rate is low and small when outcome rate is high. In Experiments 5 and 6 I varied outcome rate in two groups of participants. In one group the outcome occurred on 25% of the trials (Group 25%) and in the other group the outcome occurred on 75% of the trials (Group 75%). I found that ratings for X differed between the groups: they were lower in Group 25% than in Group 75% in both experiments. In Experiment 5, outcome rate also influenced blocking; blocking was larger in Group 25% than Group 75%. This confirmed that outcome rate successfully affected ratings for the blocked cue and blocking, in line with Livesey et al.'s predictions. In Experiment 6 I found that outcome rate also affected the magnitude of the redundancy effect; it was larger in Group 75% than Group 25%. However, the redundancy effect was significant in both groups.

Overall, findings from Chapter 3 suggested that participants' uncertainty about the causal status of X contributed to the redundancy effect. The redundancy effect was reduced with outcome-additivity training and low outcome rate, however it remained significant with both. It is possible that since these manipulations did not manipulate X

directly, uncertainty continued to influence ratings for X. However, the redundancy effect may also have been multiply-determined.

In Chapter 4 I explored whether another contributing factor to the redundancy effect in a design A+/AX+/BY+/CY- was a lack of inhibition for C. The Rescorla-Wagner (1972) model predicts that at asymptote Y will be a weak cause of the outcome because C will become a weak inhibitor. However, Uengoer et al.'s (2013) experiments and the experiments in Chapters 2 and 3 of this thesis used the allergist task to investigate the redundancy effect. In order to demonstrate inhibition in this task, participants would have had to learn that a stomach ache-causing food, when eaten with a stomach ache-preventing food, resulted in no stomach ache. However this is contrary to most every-day experiences of foods and their effects. Furthermore, Experiment 7 failed to demonstrate inhibition for C in an allergist task. If C failed to become inhibitory, this may not have protected Y from extinction. In this case the Rescorla-Wagner model would predict that both X and Y will have a minimal association with the outcome. While the redundancy effect was observed in Experiment 7, this could have been due to participants' uncertainty about the causal status of X, or pre-asymptotic learning. If it was shown that when C was inhibitory the redundancy effect was reversed, with Y having higher ratings than X, this could have partly reconciled the redundancy effect with the predictions of the Rescorla-Wagner model.

Therefore, in the subsequent experiments, I used a task in which inhibition was more plausible. In this task participants were asked to predict whether different medicines would lead to a decrease, an increase, or no change in fictional levels of a hormone. In Experiment 8 I found that even when C gained some inhibitory associative strength, the redundancy effect was still observed. However, participants varied in their ratings for C: some rated it as inhibitory and others as neutral. I also found a negative correlation between ratings for C and ratings for Y in this experiment. This correlation was also observed in Experiments 9 and 10, in which I also found a positive correlation between

ratings for C and the magnitude of the redundancy effect. In these experiments I performed data-split analyses based on participants' ratings for C. I found that the redundancy effect was larger for participants who rated C as neutral or positive, than for participants who rated C as negative. Findings of these experiments indicated that inhibition for C had the capacity to reduce the redundancy effect. Therefore in Experiments 11 and 12 I experimentally manipulated the causal status of C. In one group C was established as a neutral cue that led to no change in the outcome and in the other group as an inhibitory cue which led to a decrease in the outcome. In these experiments I found that establishing C as neutral resulted in a larger redundancy effect than establishing C as inhibitory. Taken together, these findings indicated that inhibition for C reduced the magnitude of the redundancy effect. Therefore it is likely that previous demonstrations of the redundancy effect in Uengoer et al.'s (2013) studies and the experiments in Chapters 2 and 3 of this thesis were at least partly due to a lack of inhibition for C.

Overall findings of this thesis suggest that the redundancy effect is multiply-determined. Two factors which contribute to the redundancy effect are participants' uncertainty about the causal status of X, and a lack of inhibition for C. It may be that manipulations aimed at reducing participants' uncertainty about X, if employed together with manipulations to establish C as an inhibitor, may reverse the redundancy effect. This would be one avenue for future research to explore.

The blocked cue

Experiments in Chapter 3 provided strong evidence that ratings for the blocked cue X incorporated uncertainty. Firstly, participants indicated that they were less certain about their causal ratings for X than for other cues, including Y. Secondly, I showed that outcome-additivity training and outcome-rate variations significantly influenced ratings

for X relative to other cues; ratings for X were lower with outcome-additivity training and with low outcome rate compared with the comparison groups.

However, there were also some indications that X may have gained some positive association with the outcome in Chapter 4. In Experiments 1, 2, 3, and 7, I used an allergist task with no manipulations. In these experiments, X received ratings at approximately +5, in the middle of the 11-point, 0-10 rating scale (M = 4.92). I previously argued that these ratings coincided with the point of uncertainty within this scale. In experiments within Chapter 4, I extended the rating scale to 21-points, ranging from -10 to +10. In this scale, the rating of +5 and the middle point of the scale (0) had different values. However ratings for X continued to average at approximately +5 (M = 5.47 in Experiments 8, 9, 10, Test 1 of Experiment 11, and Test 1 of Experiment 12). It is worth considering whether the positive mean ratings for X could have reflected some positive associative strength for this cue.

Livesey et al. (2013) argued that in an allergist task the possibilities for the causal status of the blocked cue X were that it either did or did not cause the outcome. In the hormone task the outcome consisted of three levels; cues led to a decrease, an increase, or no change in the outcome. However, not all three of these possibilities were equally likely for X. The results in Experiments 9 and 10 revealed that after being presented with A+/AX+/BY+/CY0/D- trials, all but a small number of participants rated the compound AD as zero. This indicated that participants understood that a cue which led to an increase and a cue which led to a decrease would result in no change when presented together. Because of this, X could have led to an increase or to no change, similarly to the allergist task. If X had led to a decrease, then AX trials would have resulted in no change in hormone levels (AX0). Therefore, while it is possible that positive ratings for X in experiments within Chapter 4 reflected some positive association with the outcome, uncertainty could have still primarily guided ratings for X in these experiments.

Some models of learning, such as Pearce's configural model (Pearce, 1987, 1994, 2002) do predict that the blocked cue X will have a positive association with the outcome, therefore exploring this idea further may be beneficial.

One possible way to test whether X gained some positive association with the outcome may be as follows. Following training, participants in a hormone task could be asked which of two choices was more likely for X: an increase or no change in hormone levels. If X was guided by uncertainty then a similar number of participants should choose each outcome. If ratings for X were guided by some positive association with the outcome, then a greater number of participants may indicate that X was more likely to lead to an increase than to no change in hormone levels.

Another possibility may be to compare ratings for X with ratings for Z in a design A+/AX+/AXZ+. Cue Z in this design could be referred to as a "double-blocked cue". If participants were completely uncertain about the causal status of X, causal ratings for X should not differ from ratings for Z. Alternatively, if X gained some positive association with the outcome, ratings for X should be greater than ratings for Z.

Since I showed that outcome rate affected ratings for X in Experiments 5 and 6, it is worth considering whether this could have resulted in higher ratings for X in experiments within Chapter 4 which used the hormone task. While outcome rate was greater than 50% in these experiments (an average of 60% of the trials led to an increase), this was also the case for Experiments 1, 2, 3, and 7 (an average of 59% of the trials led to the outcome in Experiments 8, 10, 11, and 12). However, it is difficult to estimate how having three levels of the outcome influenced outcome rate; outcome rate could have been considered higher in these experiments, for example. If positive ratings for X were due to higher outcome rate however, blocking may have been smaller in this task than for an equivalent design in an allergist task. Unfortunately I did not include a blocking-control cue in the experiments within Chapter 4; this could be explored by a future experiment.

Uncertainty at the beginning of learning

The Rescorla-Wagner (1972) model predicts that at the beginning of learning all cues start with zero associative strength. While this assumption is commonly made, an associative strength of zero corresponds to no association with the outcome. In an allergist task it would indicate that foods are assumed to not cause a stomach ache, or in a hormone task, that medicines are assumed to lead to no change in hormone levels. However, it may be more likely that at the beginning of learning, before having been presented with any information, participants are uncertain about the effects that cues will have on the outcome. While this may depend to some extent on whether participants experienced relationships between these cues and outcomes in the past, this should particularly apply to novel relationships between cues and outcomes. If this was the case, then if participants were asked to provide ratings for cues prior to being presented with any contingencies, these ratings should average at around the middle of the causal-rating scale. An experiment not reported in this thesis provides support for this view. In this experiment participants were asked to rate single cues at the beginning of learning, and after every block of trials in a design A+/AX+/BY+/CY-/D-/DE-/PQ+. This design was embedded in a task in which participants were asked to predict which medicines caused side effects. Specifically, some medicines led to a headache and some did not. I found that the initial ratings for cues averaged around the middle of the 11-point scale ranging from 0 (*Certainly not*) to 10 (*Very certain; M* = 4.98). At the end of learning X received ratings that did not differ from the ratings at the beginning of learning (*M* = 5), while ratings for Y were lower at the end (*M* = 1.5) than at the beginning of learning. In order to reconcile uncertainty about novel cues at the beginning of learning with predictions of the learning models, novel cues may be best represented as having an associative strength of 0.5λ, rather than 0 prior to learning.

The propositional account

My findings could fit with the propositional account of learning (e. g. De Houwer, 2009; De Houwer et al., 2005; Lovibond, 2003; Mitchell et al., 2009). This account predicts that participants use propositional reasoning processes when learning causal associations. In the case of my experiments, participants may have made propositions such as "I could not verify whether this cue [X] was a cause or not, therefore I am uncertain whether it caused the outcome", "Many cues in this task led to the outcome therefore it is likely that any cues I am uncertain about did too", or "This medicine [C] led to a decrease in hormone levels and when presented with this medicine [Y] it led to no change, therefore this medicine [Y] must have caused an increase in hormone levels". This account is concerned with specifying the nature of the information acquired, contrary to the learning models which specify conditions under which learning takes place. Therefore the propositional account has greater flexibility in explaining certain learning phenomena. For example, it can predict contributions of uncertainty to ratings for X observed in Chapter 3. In addition, there were individual differences in ratings for C in Experiments 8, 9, and 10: some participants gave C a negative rating, while others rated C as neutral, even though all participants were exposed to the same contingencies. The propositional account can incorporate such individual differences because it predicts that participants act based on their beliefs about the contingencies rather than the contingencies themselves. As I mentioned in the introduction, even though my findings could fit with the propositional account, the predictions of this account at this time are too general to be evaluated using my evidence, and would also include a retrospective explanation. Therefore I am tentative to draw any strong conclusions about whether my findings support the propositional account or not.

Individual differences

In Chapter 4 I found that participants varied in their ratings for C. When participants indicated that C was inhibitory, the redundancy effect was not obtained. This was partly consistent with the predictions of summed error-term models. When participants indicated that C was neutral, the redundancy effect was observed, consistently with the predictions of single error-term models. Because of these findings it may be of benefit to incorporate individual differences when considering the extent to which people rely on single or summed error in associative learning tasks, as suggested in Chapter 4. Hybrid models of learning, which include both single and summed error-terms, do exist (e. g. Le Pelley, 2004), however these assume that properties of the cues determine the extent to which each will be utilised, with little scope for the inclusion of systematic individual differences. Even though it has been argued that an approach incorporating individual differences in learning may be needed to understand the full complexity of learning and behaviour (e. g. Byrom, 2013; Sauce & Matzel, 2013), to my knowledge variation in the extent to which participants utilise processes consistent with single and summed error has not been considered in published work to date. There is some indication that this could be a promising approach to explore, however. Several researchers have investigated individual differences for effects related to irrelevant cues, such as blocking, latent inhibition, and highlighting. Firstly, the magnitude of these effects has been shown to be related to certain personality variables, for example reduced blocking with high schizotypy (e. g. Haselgrove & Evans, 2010; Moran et al., 2003) and reduced latent inhibition with high schizotypy (e. g. Allan et al., 1995; Gray, Fernandez, Williams, Ruddle & Snowden, 2002; Wuthrich & Bates, 2001). Secondly, the magnitude of these effects has been found to be related within participants, for example larger blocking was related to larger highlighting (e. g. Kruschke, Kappenman & Hetrick, 2005). Given that the redundancy effect involves a comparison between two redundant cues, a further exploration of whether the magnitude of the redundancy effect would be related to the magnitude of blocking, latent inhibition, and highlighting, may be of benefit.

One possibility for what may mediate the extent to which participants utilise single or summed-error may be the amount of working-memory resources spent on a particular task. Evidence indicates that working-memory resources affect learning (e. g. Lewandowsky, 2011; Wills, Barrasin & MacLaren, 2011). Relying on contingencies between cues and outcomes only (such that in a hormone task C would be neutral as the compound CY led to no change) may reflect a strategy that relies on memory retrieval. This strategy is more consistent with the predictions of single error-term models. The alternative cue competition strategy, considering how two cues interact to cause the outcome, may be more effortful. This strategy is more consistent with the predictions of summed error-term models. Therefore, higher working-memory resources may produce results more in line with summed error-term models while lower working-memory resources may produce results more consistent with single error-term models. In particular, higher working-memory capacity may be related to greater inhibition for C, higher ratings for Y and a smaller redundancy effect. On the other hand, lower working-memory capacity may be associated with less inhibition for C, lower ratings for Y, and a larger redundancy effect. It is important to consider that the redundancy effect also depends on ratings for X. Consistently with the predictions of summed error-term models, higher working memory may be associated with lower ratings for X, more substantial blocking, and a smaller redundancy effect. Lower working memory on the other hand, may be related to higher ratings for X, smaller blocking, and a larger redundancy effect. In line with these predictions, De Houwer and Beckers (2003) found that reducing the amount of working-memory resources by the use of a secondary task reduced blocking. What remains to be tested is whether working-memory resources would be related to the magnitude of the redundancy effect.

Implications and future directions

The implications of this thesis regarding the redundancy effect in a design A+/AX+/BY+/CY- are threefold. Firstly, I obtained evidence showing that the redundancy effect was unlikely to be due to differences in attention between X and Y. Secondly, participants' uncertainty about the causal status of X contributed to the redundancy effect. Thirdly, a lack of inhibition for C also contributed to the redundancy effect. Therefore it appears that the basis of the redundancy effect in human causal learning is the result of participants' uncertainty about the causal status of X and a lack of inhibition for C.

One future direction to investigate the redundancy effect further, would be to use several manipulations within one experiment, one targeting X and the other targeting Y. For example, a manipulation used to disambiguate the effects that X has on the outcome, such as outcome-additivity training, together with a manipulation to establish C as an inhibitor, could be used. These manipulations may reverse the redundancy effect, strengthening the findings in this thesis. This would be an interesting direction for future research to explore.

In addition to the redundancy effect, findings of this thesis have implications for blocking.

Firstly, I confirmed that variations in outcome rate contributed to causal ratings of the blocked cue and blocking as predicted by Livesey et al. (2013). This was consistent with other evidence indicating that participants were at least partly uncertain about the causal status of the blocked cue. What remains to be shown however, is whether the effects of outcome rate on the blocked cue are specific to this type of cue, or mediated by uncertainty. If outcome rate also influenced ratings for other uncertain cues, such as novel cues which were not presented during training, this would indicate support for the latter.

Secondly, my findings raised an interesting question about whether the blocked cue gained some positive association with the outcome. Average ratings for the blocked cue

were around the middle of the scale in an allergist task, in which the middle point of the scale corresponded to a medium positive value. In an alternative task in which the medium-positive value and the middle point of the scale were separated, the blocked cue continued to have positive ratings. While uncertainty could have continued to have been responsible for these ratings, this could be explored by future research. For example, experiments could use force choice methods or the design A+/AX+/AXZ+ to investigate whether X gains some positive association with the outcome. Establishing that the blocked cue reflects some positive association with the outcome, while intuitive, would have important theoretical implications, and would provide support for the models of learning which predict positive associative strength for this cue.

My findings also underline the importance of considering individual differences when investigating learning effects. Participants varied greatly in the ratings they gave for C in Chapter 4. In the previous section I suggested that ratings for C may have been related to the amount of working-memory resources. An additional speculation was whether participants may differ in the extent to which they rely on single and summed-error processes. This could be another avenue for future research to explore.

Finally, as noted previously, the redundancy effect has been observed in rats and pigeons (Jones & Pearce, 2015; Pearce et al., 2012). While my findings apply to humans, it may still be applicable to non-human animal data. The allergist task is similar to the experiments used in animal studies in which cues led to the delivery of food (e. g. Jones & Pearce, 2015; Pearce et al., 2012). A future experiment could explore whether inhibition for C could be demonstrated with the redundancy effect in the same experiment using this task in animals. My findings also pose an additional question. I showed that participants were less certain about the causal status of X than of Y, and ratings for X were influenced by manipulations such as outcome rate. Would outcome-rate variations affect the redundancy effect in animals? Arguably animals do not have the same concept of uncertainty as do humans. If outcome rate affected responding for X in animals in a similar

manner to my experiments however, it would raise an interesting question as to the basis of these processes. For example, Beckers, Miller, De Houwer and Urushihara (2006) found that outcome-additivity training influenced blocking in rats. Showing whether or not the redundancy effect was related to outcome rate in animals would further contribute to the debates regarding whether or not animals and humans possess a common-learning mechanism (e. g. Mitchell et al., 2009).

Lastly, perhaps the broadest implication of my findings is that they cannot be easily accommodated within a single model of learning. They also highlight the limitations of the Rescorla-Wagner (1972) model, which has been very successful in the past. Perhaps, as we move towards exploring increasingly complex learning and behaviour effects, there is a greater need to bring together different perspectives. An attempt of this has been made with researchers proposing hybrid models of learning which incorporate processes which historically have been considered as opposing one another (e. g. Le Pelley, 2004; Pearce & Mackintosh, 2010) and investigating the circumstances under which these effects could be obtained (e. g. Beesley, Nguyen, Pearson & Le Pelley, 2015). My findings underline the importance of such work.

Conclusion

In this thesis I explored the basis of the redundancy effect in human causal learning. I found that participants' uncertainty about the blocked cue X and a lack of inhibition for C, constituted the basis of the redundancy effect in a design A+/AX+/BY+/CY-. My findings also indicated that an attentional explanation for the redundancy effect was unlikely. These findings pose a challenge to the Rescorla-Wagner (1972) model of learning and the common-stimulus elements account suggested by Vogel and Wagner (2017). In addition, my findings indicated that individual differences in ratings for C contributed to the redundancy effect. Future directions regarding this work could use manipulations targeting X and Y within one experiment, investigate the nature of individual differences in

ratings for C, their relationship to working-memory resources and the extent to which participants rely on single and summed error. In addition, I suggested that ratings for the blocked cue X may have been determined by both uncertainty and a weak positive association with the outcome.

Appendices

Appendix 4A: Analyses from Experiment 11.

Analyses not included in the chapter are underlined.

Exploring the data with a three-way ANOVA.

A three-way (cue [A, B, C, D, X, Y] by test [1 vs 2] by group [Group Inhibitory vs Group Neutral]) ANOVA revealed:

- a significant effect of cue, $F(3.21, 154.24) = 298.03$, $p < .001$, $\eta_P^2 = .86$,
- no significant effect of group, $F(1, 48) = 1.7$, $p = .199$, $\eta_P^2 = .03$,
- no significant effect of test, $F(1, 48) < .001$, $p > .99$, $\eta_P^2 < .001$,
- a significant interaction between test and group, $F(1, 48) = 17.61$, $p < .001$, $\eta_P^2 = .27$,
- a significant interaction between cue and group, $F(3.21, 154.24) = 8.64$, $p < .001$, $\eta_P^2 = .15$,
- no significant interaction between test and cue, $F(2.8, 134.55) = 2.11$, $p = .106$, $\eta_P^2 = .042$,
- a significant three-way interaction, $F(2.8, 134.55) = 11.47$, $p < .001$, $\eta_P^2 = .19$.

Differences between the groups at Test 1 and at Test 2.

A two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-1 data revealed:

- a significant effect of cue, $F(3.09, 148.24) = 157.29$, $p < .001$, $\eta_P^2 = .77$,
    - while ratings for B, X, and Y did not differ significantly, ts ≤ 2.75, ps ≥ .134, rs ≤ .37, ratings between all other cues did, ts ≥ 6.01, ps < .001, rs ≥ .65.

- no significant effect of group, $F(1, 48) = .85$, $p = .361$, $\eta_P^2 = .02$,
- no significant interaction, $F(3.09, 148.24) = .8$, $p = .502$, $\eta_P^2 = .02$.

A two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data revealed:

- a significant effect of cue, $F(3.17, 152.27) = 299.27$, $p < .001$, $\eta_P^2 = .86$,
    - while ratings between B and X, and X and Y, did not differ significantly, $ts \leq 2.07$, $ps \geq .306$, $rs \leq .28$, ratings between all other cues did, $ts \geq 3.06$, $ps \leq .016$, $rs \geq .4$.
- a significant effect of group, $F(1, 48) = 12.38$, $p = .001$, $\eta_P^2 = .21$,
- a significant interaction, $F(3.17, 152.27) = 23.8$, $p < .001$, $\eta_P^2 = .33$,
    - simple main effects analyses indicated that the groups did not differ in their ratings for A, B, D, and X, $ts \leq 1.51$, $ps \geq .143$, $rs \leq .21$, however Group Inhibitory had lower ratings for C, $t(25.42) = 15.64$, $p < .001$, $r = .91$, and higher ratings for Y, $t(48) = 5.09$, $p < .001$, $r = .59$, than Group Neutral.

Differences between the tests in each group

To explore how establishing C as neutral or inhibitory affected the ratings for all cues at Test 2 relative to Test 1, two ANOVA tests were conducted comparing data at Test 1 and Test 2 in each group.

In Group Neutral, a two-way (cue [A, B, C, D, X, Y] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(3, 71.19) = 113.58$, $p < .001$, $\eta_P^2 = .83$,
- a significant effect of test, $F(1, 24) = 7.84$, $p = .010$, $\eta_P^2 = .25$,
- a significant interaction, $F(2.71, 64.92) = 3.21$, $p = .033$, $\eta_P^2 = .12$,

- paired comparisons indicated that ratings for B, C, and X were higher at Test 2 than Test 1 (cue B: $t(24) = 2.49$, $p = .02$, $r = .45$; C: $t(24) = 2.12$, $p = .045$, $r = .4$; cue X: $t(24) = 3.03$, $p = .006$, $r = .53$) and ratings for the other cues did not differ between tests, $ts \leq 1.32$, $ps \geq .199$, $rs \leq .26$.

In Group Inhibitory, a two-way (cue [A, B, C, D, X, Y] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(2.75, 65.89) = 204.95$, $p < .001$, $\eta_p^2 = .9$,
- a significant effect of test, $F(1, 24) = 10.04$, $p = .004$, $\eta_p^2 = .3$,
- a significant interaction, $F(2.62, 62.9) = 13.18$, $p < .001$, $\eta_p^2 = .35$,
  - paired comparisons indicated that at Test 2 ratings for C were lower, $t(24) = 5.84$, $p < .001$, $r = .77$, and ratings for Y were higher, $t(24) = 2.21$, $p = .037$, $r = .41$, than at Test 1, while ratings for the other cues did not differ, $ts \leq 1.34$, $ps \geq .191$, $rs \leq .26$.

Exploring the redundancy effect with a three-way ANOVA

A three-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(1, 48) = 7.53$, $p = .008$, $\eta_p^2 = .14$,
- a non-significant trend for the main effect of group, $F(1, 48) = 3.92$, $p = .053$, $\eta_p^2 = .08$,
- a significant interaction between cue and group, $F(1, 48) = 7.24$, $p = .01$, $\eta_p^2 = .13$,
- no significant effects of test, no significant interaction between test and group, no significant interaction between test and cue, $Fs \leq 1.03$, $ps \geq .314$, $\eta_p^2 \leq .02$,
- a significant three-way interaction, $F(1, 48) = 11.95$, $p = .001$, $\eta_p^2 = .2$.

Differences in X and Y ratings between tests in each group

In Group Neutral, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(1, 24) = 14.74$, $p = .001$, $\eta_P^2 = .38$,
    - ratings for X were higher than for Y,
- no significant effect of test, $F(1, 24) = .02$, $p = .899$, $\eta_P^2 < .01$,
- a significant interaction, $F(1, 24) = 6.8$, $p = .015$, $\eta_P^2 = .22$,
    - simple main effects analyses indicated that ratings for X were higher at Test 2 than Test 1, $t(24) = 3.03$, $p = .006$, $r = .53$, while ratings for Y did not differ between the tests, $t(24) = 1.32$, $p = .199$, $r = .26$.
- As a result, the redundancy effect was larger at Test 2, $t(24) = 4.87$, $p < .001$, $r = .7$ (X ratings – Y ratings: $M = 4.52$, $SEM = .93$), than at Test 1, $t(24) = 1.45$, $p = .16$, $r = .28$ (X ratings – Y ratings: $M = 1.48$, $SEM = 1.02$), in Group Neutral.

In Group Inhibitory, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed:

- no significant effect of cue, $F(1, 24) < .01$, $p = .097$, $\eta_P^2 < .01$,
- no significant effect of test, $F(1, 24) = 2.63$, $p = .118$, $\eta_P^2 = .1$,
- a significant interaction, $F(1, 24) = 5.19$, $p = .032$, $\eta_P^2 = .18$,
    - simple main effects analyses indicated ratings for X did not differ between tests, $t(24) = 1.34$, $p = .191$, $r = .26$, while ratings for Y were higher at Test 2, $t(24) = 2.21$ $p = .037$, $r = .41$, than at Test 1.
- As a result, the redundancy effect was smaller at Test 2, $t(24) = -1.27$, $p = .217$, $r = -.25$, $BF_{01} = 9.79$ (one-tailed Bayesian t-test; X ratings – Y ratings: $M = -1.26$, $SEM = .99$), than at Test 1, $t(24) = 1.42$, $p = .17$, $r = .28$ (X ratings – Y ratings: $M = 1.32$, $SEM = .93$), in Group Inhibitory.

Magnitude of the redundancy effect (X ratings – Y ratings) at Test 2 was larger in Group Neutral than Group Inhibitory, $t(48) = 4.25$, $p < .001$, $r = .52$.

Differences in X and Y ratings between the groups at Test 1 and at Test 2

In order to investigate differences in the redundancy effect between groups at each test, 2 two-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA tests were conducted on the data at Test 1 and Test 2.

A two-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test 1 data revealed:

- a significant effect of cue, $F(1, 48) = 4.1$, $p = .048$, $\eta_P^2 = .08$,
    - ratings for X were higher than for Y,
- no significant effect of group, $F(1, 48) = 1.73$, $p = .2$, $\eta_P^2 = .04$,
- no significant interaction, $F(1, 48) = .01$, $p = .91$, $\eta_P^2 < .001$.

A two-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data revealed:

- a significant effect of cue, $F(1, 48) = 5.75$, $p = .02$, $\eta_P^2 = .12$,
    - ratings for X were higher than for Y
- a significant effect of group, $F(1, 48) = 5.38$, $p = .025$, $\eta_P^2 = .1$,
- a significant interaction, $F(1, 48) = 18.08$, $p < .001$, $\eta_P^2 = .27$,
    - simple main effects analyses indicated that X had higher ratings than Y in Group Neutral, $t(24) = 4.87$, $p < .001$, $r = .71$, but ratings for these cues did not differ in Group Inhibitory, $t(24) = 1.27$, $p = .217$, $r = .25$, $BF_{01} = 9.79$.

Appendix 4B. Analyses from Experiment 12

Analyses not included in the chapter are underlined.

Exploring ratings with a three-way ANOVA

A three-way (cue [A, B, C, D, X, Y] by test [1 vs 2] by group [Group Inhibitory vs Group Neutral]) ANOVA revealed:

- a significant effect of cue, $F(3.58, 243.58) = 564.2$, $p < .001$, $\eta_P^2 = .89$,
- a non-significant trend for the effect of group, $F(1, 68) = 3.32$, $p = .073$, $\eta_P^2 = .05$,
- no significant effect of test, $F(1, 68) = 2.01$, $p = .16$, $\eta_P^2 = .03$,
- a significant interaction between test and group, $F(1, 68) = 17.47$, $p < .001$, $\eta_P^2 = .2$,
- a significant interaction between cue and group, $F(3.58, 243.58) = 5.78$, $p < .001$, $\eta_P^2 = .08$,
- no significant interaction between cue and test, $F(2.76, 187.64) = .82$, $p = .474$, $\eta_P^2 = .01$,
- a significant three-way interaction, $F(2.76, 187.64) = 19.68$, $p < .001$, $\eta_P^2 = .22$.

Differences between the groups at Test 1 and at Test 2

A two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-1 data revealed:

- a significant effect of cue, $F(3.04, 206.66) = 314.93$, $p < .001$, $\eta_P^2 = .82$,
    - ratings between B and X, and between X and Y, did not differ significantly, ts ≤ 2.57, ps ≥ .167, rs ≤ .3, while ratings between all other cues did, ts ≥ 3.98, ps ≤ .002, rs ≥ .43.
- no significant effect of group, $F(1, 68) = .14$, $p = .714$, $\eta_P^2 < .01$,

- no significant interaction, $F(3.04, 206.66) = .88$, $p = .502$, $\eta_p^2 = .02$.

A two-way (cue [A, B, C, D, X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data revealed:

- a significant effect of cue, $F(3.83, 260.47) = 464.98$, $p < .001$, $\eta_p^2 = .87$,
    - while ratings for X and Y did not differ significantly, $t(69) = 2.38$, $p = .282$, $r = .28$, ratings between all other cues did, ts ≥ 3.35, ps ≤ .018, rs ≥ .37,
- a significant effect of group, $F(1, 68) = 13.78$, $p < .001$, $\eta_p^2 = .17$,
- a significant interaction, $F(5, 340) = 23.63$, $p < .001$, $\eta_p^2 = .26$,
    - paired comparisons indicated that the groups did not differ in ratings for A, D, and X, ts ≤ .68, ps ≥ .5, rs ≤ .08, but Group Inhibitory had lower ratings for B, $t(68) = 3.18$, $p = .002$, $r = .36$, lower ratings for C, $t(33.83) = 8.95$, $p < .001$, $r = .74$, and higher ratings for Y, $t(68) = 3.57$, $p = .001$, $r = .4$ than Group Neutral.

Differences between the tests in each group

In Group Neutral, a two-way (cue [A, B, C, D, X, Y] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(3.47, 121.55) = 253.78$, $p < .001$, $\eta_p^2 = .88$,
- a non-significant trend for the effect of test, $F(1, 35) = 3.51$, $p = .069$, $\eta_p^2 = .09$,
- a significant interaction, $F(2.65, 92.74) = 10.24$, $p < .001$, $\eta_p^2 = .23$,
    - paired comparisons indicated that at Test 2 the ratings for B, $t(35) = 3.52$, $p = .001$, $r = .51$, and C, $t(35) = 4.17$, $p < .001$, $r = .58$, were higher, and ratings for Y, $t(35) = 2.92$, $p = .006$, $r = .44$, were lower than at Test 1. Ratings did not significantly differ between the tests for the other cues, including X, ts ≤ 1.47, ps ≥ .152, rs ≤ .24.

In Group Inhibitory a two-way (cue [A, B, C, D, X, Y] by group [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(3.19, 105.35) = 321.11$, $p < .001$, $\eta_P^2 = .91$,
- a significant effect of test, $F(1, 33) = 17.44$, $p < .001$, $\eta_P^2 = .35$,
- a significant interaction, $F(2.64, 87.21) = 10.27$, $p < .001$, $\eta_P^2 = .24$,
    - paired comparisons indicated that at Test 2 the ratings for B were lower, $t(33) = 2.82$, $p = .008$, $r = .44$, and ratings for C were lower, $t(33) = 4.51$, $p < .001$, $r = .62$, and there was a non-significant trend for higher ratings for cue Y, $t(33) = 1.93$, $p = .062$, $r = .32$, relative to Test 1, while ratings between tests did not differ for the other cues, $ts \leq 1.64$, $ps \geq .112$, $rs \leq .27$.

Exploring the redundancy effect with a three-way ANOVA

A three-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(1, 68) = 8.7$, $p = .004$, $\eta_P^2 = .11$,
- a significant interaction between group and test, $F(1, 68) = 4.66$, $p = .034$, $\eta_P^2 = .06$,
- a significant three-way interaction, $F(1, 68) = 14.09$, $p < .001$, $\eta_P^2 = .17$,
- no other significant main effects or interactions, $Fs \leq 1.54$, $ps \geq .22$, $\eta_P^2 \leq .02$.

Differences in X and Y ratings between tests in each group

In Group Neutral, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed:

- a significant effect of cue, $F(1, 35) = 6.37$, $p = .016$, $\eta_P^2 = .15$,
    - ratings for X were higher than for Y,
- a significant effect of test, $F(1, 35) = 4.48$, $p = .041$, $\eta_P^2 = .11$,
- a significant interaction, $F(1, 35) = 8.72$, $p = .006$, $\eta_P^2 = .2$,

- simple main effects analyses indicated that ratings for X did not differ between the tests, $t(35) = .35$, $p = .732$, $r = .06$, while ratings for Y were lower at Test 2, $t(35) = 2.92$, $p = .006$, $r = .44$, than at Test 1.

- As a result, the redundancy effect was larger at Test 2, $t(35) = 3.8$, $p = .001$, $r = .54$ (X ratings – Y ratings: $M = 2.79$, $SEM = .76$), than at Test 1, $t(35) = .89$, $p = .378$, $r = .15$, $BF_{01} = 3.86$ (X ratings – Y ratings: $M = .74$, $SEM = .85$), in Group Neutral.

In Group Inhibitory, a two-way (cue [X, Y] by test [1 vs 2]) ANOVA revealed:

- no significant effect of cue, $F(1, 33) = 2.57$, $p = .119$, $\eta_p^2 = .07$,
- no significant effect of test, $F(1, 33) = .63$, $p = .433$, $\eta_p^2 = .02$,
- a significant interaction, $F(1, 33) = 5.89$, $p = .021$, $\eta_p^2 = .15$,
    - simple main effects analyses indicated that ratings for X did not differ between the tests, $t(33) = 1.64$, $p = .112$, $r = .27$, while there was a non-significant trend for lower ratings for Y at Test 2, $t(33) = 1.93$, $p = .062$, $r = .32$, than at Test 1.
- As a result, the redundancy effect was smaller at Test 2, $t(33) = -.22$, $p = .827$, $r = .04$, $BF_{01} = 5.32$ (X ratings – Y ratings: $M = -.18$, $SEM = .8$), than at Test 1, $t(33) = 3.18$, $p = .003$, $r = .48$ (X ratings – Y ratings: $M = 1.99$, $SEM = .62$), in Group Inhibitory.

The redundancy effect at Test 2 was larger in Group Neutral than Group Inhibitory, $t(68) = 2.73$, $p = .008$, $r = .31$.

Differences in X and Y ratings between the groups at Test 1 and Test 2

A two-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-1 data revealed:

- a significant effect of cue, $F(1, 68) = 6.81$, $p = .011$, $\eta_p^2 = .09$,

- ratings for X were higher than for Y,

- no significant effect of group, $F(1, 68) = .002$, $p = .965$, $\eta_P^2 < .001$,
- no significant interaction, $F(1, 68) = 1.44$, $p = .235$, $\eta_P^2 = .021$.

A two-way (cue [X, Y] by group [Group Inhibitory vs Group Neutral]) ANOVA on Test-2 data revealed:

- a significant effect of cue, $F(1, 68) = 5.8$, $p = .019$, $\eta_P^2 = .08$,
  - ratings for X were higher than for Y,
- a significant effect of group, $F(1, 68) = 5.38$, $p = .023$, $\eta_P^2 = .07$,
- a significant interaction, $F(1, 68) = 7.47$, $p = .008$, $\eta_P^2 = .1$,
  - simple main effect analyses indicated that, similarly to the previous experiment, the redundancy effect was observed in Group Neutral, $t(35) = 3.8$, $p = .001$, $r = .54$, but it was not significant in Group Inhibitory, $t(33) = .22$, $p = .827$, $r = .04$, $BF_{01} = 5.32$.

References

Aitken, M. R. F., Larkin, M. J., & Dickinson, A. (2000). Super-learning of causal judgements. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *53*(1), 59–81.

Allan, L. M., Williams, J. H., Wellman, N. A., Tonin, J., Taylor, E., Feldon, J., & Rawlins, J. N. P. (1995). Effects of tobacco smoking, schizotypy and number of pre-exposures on latent inhibition in healthy subjects. *Personality and Individual Differences*, *19*(6), 893–902.

Arcediano, F., Matute, H., & Miller, R. R. (1997). Blocking of Pavlovian Conditioning in Humans. *Learning and Motivation*, (28), 188–199.

Beckers, T., De Houwer, J., Pineño, O., & Miller, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(2), 238–249.

Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: forward blocking in Pavlovian Animal Conditioning is sensitive to constraints of causal inference, *Journal of Experimental Psychology: General, 135*(1), 92–102.

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*(1), 114–120.

Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *The Quarterly Journal of Experimental Psychology*, *68*(11), 2175–99.

Bower, G. H., & Trabasso, T. R. (1964). Concept identification. In R. C. Atkinson (Ed.), *Studies in mathematical psychology* (pp. 32–96). Stanford: Stanford University Press.

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. London: Chapman and Hall Ltd.

Byrom, N. C. (2013). Accounting for individual differences in human associative learning. *Frontiers in Psychology*, *4*, 1–8.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, *18*(5), 537–545.

Cheng, P. W. (1997). From Covariation to Causation: A Causal Power Theory. *Psychological Review*, *104*(2), 367–405.

Cheng, P. W., & Holyoak, K. J. (1995). Complex Adaptive Systems as Intuitive Statisticians: Causality, Contingency, and Prediction. In J. A. Meyer & H. Roitblat (Eds.), *Comparative approaches to cognition* (pp. 271–302). Cambridge: MIT Press.

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*(1), 42–45.

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

De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, *56*(4), 345–357.

De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *The Quarterly Journal of Experimental Psychology*, *55A*(3), 965–985.

Denniston, J. C., Savastano, H. I., & Miller, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65–117). London, UK: Lawrence Erlbaum Associates.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of Act-Outcome Contingency: The role of Selective Attribution. *Quarterly Journal of Experimental Psychology*, *36A*(January), 29–50.

Dwyer, D. M., Haselgrove, M., & Jones, P. M. (2011). "Cue interactions in flavor preference learning: A configural analysis": Correction to Dwyer et al. (2011). *Journal of Experimental Psychology. Animal Behavior Processes*, *37*(2), 188.

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*(1718), 2553–61.

George, D. N., & Pearce, J. M. (1999). Acquired distinctiveness is controlled by stimulus relevance not correlation with reward. *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 363–373.

Gray, N. S., Fernandez, M., Williams, J., Ruddle, R. A., & Snowden, R. J. (2002). Which schizotypal dimensions abolish latent inhibition? *The British Journal of Clinical Psychology*, *41*, 271–284.

Hall, G., & Pearce, J. M. (1979). Latent inhibition of a CS during CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*(1), 31–42.

Haselgrove, M., & Evans, L. H. (2010). Variations in selective and nonselective prediction error with the negative dimension of schizotypy. *The Quarterly Journal of Experimental Psychology*, *63*(6), 1127–1149.

Hinchy, J., Lovibond, P. F., & Ter-Horst, K. M. (1995). Blocking in human electrodermal conditioning. *The Quarterly Journal of Experimental Psychology Section B*, *48*(1), 2–12.

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: the role of uncertainty. *The Quarterly Journal of Experimental Psychology*, *61*(11), 1658–68.

Jeffreys, H. (1961). *The theory of probability* (3rd ed.). Oxford: Oxford University Press.

Jones, P. M., & Pearce, J. M. (2015). The fate of redundant cues: Further analysis of the redundancy effect. *Learning & Behavior*, *43*, 72–82.

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

Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review*, *7*(4), 636–645.

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*(5), 830–845.

Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *57*, 193–243.

Le Pelley, M. E., Beesley, T., & Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*(2), 220–229.

Le Pelley, M. E., & Mclaren, I. P. L. (2001). The mechanics of associative change. *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*.

Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *56*(1), 68–79.

Le Pelley, M. E., & McLaren, I. P. L. (2004). Associative history affects the associative change undergone by both presented and absent cues in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *30*(1), 67–73.

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), 1111–1140.

Le Pelley, M. E., Vadillo, M., & Luque, D. (2013). Learned predictiveness influences rapid attentional capture: evidence from the dot probe task. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *39*(6), 1888–1900.

Lewandowsky, S. (2011). Working memory capacity and categorization: individual differences and modeling. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 37*(3), 720-738.

Livesey, E. J., & Boakes, R. A. (2004). Outcome additivity, elemental processing and blocking in human causality judgements. *The Quarterly Journal Of Experimental Psychology*, *57B*(4), 361–379.

Livesey, E. J., Lee, J. C., & Shone, L. T. (2013). The relationship between blocking and inference in causal learning. *Proceedings of the 35th Annual Meeting of the Cognitive Science Society* (pp. 2920–2925).

Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In Schmalhofer, F., Young, R. M., & Katz, G. (Eds.) *Proceedings of EuroCogSci 03: The European Cognitive Science Conference* (pp. 217–222). New Jersey: Lawrence Erlbaum Associates.

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 & Cognition*, *31*(1), 133–142.

Lovibond, P. F. (2003). Causal beliefs and conditioned responses: Retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(1), 97–106.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276–298.

Maes, E., Boddez, Y., Alfei, J. M., Krypotos, A., D 'Hooge, R., De Houwer, J., & Beckers, T. (2016). The elusive nature of the blocking effect: 15 failures to replicate. *Journal of Experimental Psychology: General*, *145*(9), e49–e71.

Melchers, K. G., Wolff, S., & Lachnit, H. (2006). Extinction of conditioned inhibition through nonreinforced presentation of the inhibitor. *Psychonomic Bulletin & Review*, *13*(4), 662–7.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 51–92). San Diego, CA: Academic Press.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *The Behavioral and Brain Sciences*, *32*(2), 183-98-246.

Mitchell, C. J., & Lovibond, P. F. (2002). Backward and forward blocking in human electrodermal conditioning: Blocking requires an assumption of outcome additivity. *The Quarterly Journal of Experimental Psychology: Section B*, *55*(4), 311–329.

Mitchell, C. J., Lovibond, P. F., & Condoleon, M. (2005). Evidence for deductive reasoning in blocking of causal judgments. *Learning and Motivation*, *36*(1), 77–87.

Moran, P. M., Al-Uzri, M. M., Watson, J., & Reveley, M. A. (2003). Reduced Kamin blocking in non paranoid schizophrenia: Associations with schizotypy. *Journal of Psychiatric Research*, *37*(2), 155–163.

Pavlov, I. P. (1927). *Conditioned reflexes*. (G. V. (Trans. . Anrep, Ed.). Oxford: Oxford University Press.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, *94*(1), 61–73.

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, *101*(4), 587–607.

Pearce, J. M. (2002). Evaluation and development of a connectionist theory of configural learning. *Animal Learning & Behavior*, *30*(2), 73–95.

Pearce, J. M., Dopson, J. C., Haselgrove, M., & Esber, G. O. R. (2012). The fate of redundant cues during blocking and a simple discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*(2), 167–179.

Pearce, J. M., Esber, G. R., George, D. N., & Haselgrove, M. (2008). The nature of discrimination learning in pigeons. *Learning & Behavior, 36*(3), 188-199.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*(6), 532–552.

Pearce, J. M., & Mackintosh, N. J. (2010). Two theories of attention: A review and possible integration. In M. E. Le Pelley & C. J. Mitchell (Eds.), *Learning and Attention* (pp. 11–39). Oxford, England: Oxford University Press.

Posner, M. I. (1980). Orienting Attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.

Premack, D. (2007). Human and animal cognition: Continuity and discontinuity. *Proceedings of the National Academy of Sciences*, *104*(35), 13861–13867.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*(3), 372–422.

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*(1), 1–41.

Rescorla, R. A. (1969). Pavlovian Condition Inhibition. *Psychological Bulletin*, *72*, 77–94.

Rescorla, R. A. (2001). Unequal associative changes when excitors and neutral stimuli are conditioned in compound. *The Quarterly Journal of Experimental Psychology*, *54B*(1), 53–68.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Projasy (Eds.), *Classical Conditioning II: Current Research and Theory* (pp. 64–99). New York, US: Appleton Century Crofts.

Sauce, B., & Matzel, L. D. (2013). The causes of variation in learning and behavior: Why individual differences matter. *Frontiers in Psychology*, *4*(395), 7–16.

Savastano, H. I., Cole, R. P., Barnet, R. C., & Miller, R. R. (1999). Reconsidering conditioned inhibition. *Learning and Motivation*, *30*, 101–127.

Schultz, D. H., & Helmstetter, F. J. (2010). Classical conditioning of autonomic fear responses is independent of contingency awareness. *Journal of Experimental Psychology. Animal Behavior Processes*, *36*(4), 495–500.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, *4995*(February 2015), 37–41.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*(3), 759–783.

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, *88*(2), 135–170.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monograph Supplements*, *2*(4), 1–8.

Trabasso, T. R., & Bower, G. H. (1986). *Attention in learning: theory and research*. New York, US: Wiley.

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*(4), 323–33.

Urcelay, G. P., & Miller, R. R. (2006). A comparator view of Pavlovian and differential inhibition. *Journal of Experimental Psychology: Animal Behavior Process*, *32*(3), 271–283.

Vandorpe, S., & De Houwer, J. (2006). People want to see information that will help them make valid inferences in human causal learning. *Memory & Cognition*, *34*(5), 1133–1139.

Vogel, E. H., & Wagner, A. R. (2017). A theoretical note in interpretation of the " Redundancy effect " in associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition*, *43*(1), 119–125.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, *76*(2), 171–180.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*(2), 222–236.

Waldmann, M. R., & Walker, J. M. (2005). Competence and performance in causal learning. *Learning & Behavior : A Psychonomic Society Publication*, *33*(2), 211–229.

Wasserman, E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science*, *1*(5), 298–302.

Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention, Part 4*, (4), 96–104.

Wills, A. J., Barrasin, T. J., & McLaren, I. P. L. (2011). Working memory capacity and generalization in predictive learning. *Proceedings of the 33rd Annual Meeting of the*

Cognitive Science Society (pp. 3205-3210).

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*(5), 843–854.

Wuthrich, V., & Bates, T. C. (2001). Schizotypy and latent inhibition: non-linear linkage between psychometric and cognitive markers. *Personality and Individual Differences, 30*, 783–798.

Yarlas, A. S., Cheng, P. W., & Holyoak, K. J. (1995). Alternative approaches to causal induction: The probabilistic contrast versys the Rescorla-Wagner model. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Meeting of the Cognitive Science Society* (pp. 431–436). NJ: Erlbaum.

Zaksaite, T., & Jones, P. M. (2017). The redundancy effect in human causal learning: Evidence against a Comparator Theory explanation. *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*.