Extinction and Renewal of Human Causal Learning

Adrian Camilleri

Submitted in partial fulfilment of the requirements for the Degree of

Bachelor of Psychology (Hons.)

School of Psychology

University of New South Wales

October 2007

Certificate of Originality

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any other degree or diploma of the university or other institute of higher learning, except where due acknowledgement is made in the text.

I also declare that the intellectual content of this thesis is the product of my own work, even though I may have received assistance from others on style, presentation and language expression.

Signature: ________________________________

Student's Name: ____________________________

Date: ____________________________________

Acknowledgements

I offer my sincerest thank-you to my supervisor, Dr. Chris Mitchell, for his continued guidance, patience, and wit throughout the year. I appreciate the advice, discussion, and company of all the members of Lab 726, most particularly Susan Wardle for her talent in writing the test program with me. A special thanks also to Dr. Juan Rosas for his insightful emails from halfway across the planet, to Dr. Melanie Gleitzman for her shared statistical know-how, and to Chris Tynan for donating his time to read the final draft.

To the ‘psych-crew’ who have travelled on this 4-year adventure with me, I hope the experience has been equally rewarding and memorable, and I wish you the best of luck. Most importantly, I appreciate the faith, love, and support shown to me by Joy Francisco, my mum and dad, and the rest of my family.

And to Mr. X, who can still sometimes be found at the Shady Bar.

~ Table of Contents ~

List of Tables and Figures

Introduction

Experiment 1

Experiment 2

Experiment 3

List of Appendices

Abstract

Associative learning in humans is often assessed with a causal learning task, such as the allergist task. In the present study, three experiments used an allergist task in which a number of food cues were paired with three possible allergic reaction outcomes. At test, participants were asked to recall the outcome paired with each cue, and to judge the extent to which the cue caused that outcome. In general, cues paired with outcomes were judged with a high causal rating, and participants could accurately recall which outcomes were paired with those cues. Some cues that were paired with an outcome during initial acquisition training were subsequently paired with no outcome during extinction training. In all three experiments an extinction effect was observed in that these extinguished cues were judged with a low causal rating. This occurred despite participants possessing good memory for the outcome initially paired with the extinguished cues. In addition, the context where extinction training took place was manipulated. In Experiments 1 and 3, when extinction training occurred in a context different from acquisition training and test, then a renewal effect was observed in that the extinguished cues were judged with a high causal rating. No extinction or renewal effects were observed on the memory measure. These findings cast doubt on the traditional assumption that extinction and renewal reflect retrieval interference in memory. It is suggested that the results are more likely the result of a reasoning process, based on very good memory for the events that occurred during training.

What is true in one context at one time may not be true in another context at another time. For example, a dog-owner walking their Chihuahua through the local park might be attacked and stung by a venomous wasp, causing them much pain and distress. Subsequent to the experience they may learn to avoid, and possibly even fear, wasps of all kind. The phobic may seek out the assistance of a therapist who, through repeated and controlled exposure to a variety of harmless wasps, helps the phobic to overcome their fear. Once outside the therapists' office, however, the phobic's confidence might evaporate and they may resolve to avoid wasps, and their local park, at all costs.

In this example the phobic has learned at least two conflicting things about wasps that consequently influence his or her behaviour. At the park the phobic learns that wasps should be feared and avoided; at the therapists office the phobic learns that wasps need not be feared and avoided. How are these pieces of knowledge represented in memory? What psychological mechanisms underpin the phobic's learning and behaviour? Two models have emerged to answer these questions (see Pineno & Miller, 2007, for a review). The associative approach argues that the phobic's behaviour will be determined by their current context, which will automatically retrieve memories of previous experiences that have taken place in that context (Bouton, 1993). In contrast, the inferential approach argues that the phobic will behave in accordance with a conscious reasoning process based on their experiences at the park and at the therapists office (Mitchell, De Houwer, & Lovibond, submitted). The debate between these approaches forms the backdrop for the current study.

Assuming that fear learning is equivalent to other types of learning, it is possible to investigate the mechanisms involved by way of a human contingency learning task. One of the most common of these tasks, and the one to be used in this thesis, is the allergist task (Larkin, Aitken, & Dickinson, 1998). In this task a participant is presented with trial-by-trial information about the relationship between a number of different food cues and the allergic reaction outcomes that sometimes follow from eating these foods in the hypothetical patient Mr. X. Just as two things may be learned about the relative hazards of encountering a wasp, two things may be learned about a particular food cue by manipulating the outcomes the cue is paired with. For example, garlic may be paired with nausea initially, but paired with no reaction near the end of training. At test the participant is asked to judge the probability of an allergic reaction given a food cue, or to rate the causal status of the cue.

The associative and inferential approaches make different predictions about participants' ability to recall the outcome paired with an extinguished cue, and the impact of instructions. The major aim of the current study is to test these opposing predictions.

Extinction

Extinction is one of the simplest and well-studied examples of Pavlovian conditioning. In the first systematic investigation of extinction Pavlov (1927) rang a bell before presenting food to his dog subjects. After repeated pairings of the two stimuli Pavlov observed that merely ringing the bell was sufficient to cause the dogs to salivate, even if food was not present. He argued that the dogs had formed a mental link between the ringing bell and the food, such that presentation of the bell triggered the expectation that food would also be presented. During an extinction phase, the dogs were again presented with the ringing bell, but it was no longer followed by food. After many such presentations the dogs ceased to salivate at the sound of the ringing bell alone – the original learning had been made 'extinct'.

More formally, extinction involves three experimental phases: acquisition, extinction, and test. A neutral stimulus (the conditioned stimulus [CS] e.g., a ringing bell) is first paired with a biologically significant stimulus (the unconditioned stimulus [US] e.g., food) during phase 1. As a consequence of this the CS alone comes to elicit a conditioned response (CR, e.g., salivation) appropriate to the US (Pearce, 1997). In the extinction phase the CS is then presented alone, which results in a decrease in CR at test. Presumably the link between the CS and US has in some way become extinguished, but what exactly is learned during the extinction phase? This question is explored in the following sections.

Extinction does not destroy the original trace

An obvious explanation, and one that has been assumed in some models of learning and memory (e.g., Rescorla & Wagner, 1972), is that the original CS-US memory trace is destroyed as a consequence of presenting the CS alone during the extinction phase. Data from both animal and human experiments, however, suggest that extinction actually involves new learning that is context dependent (Bouton, 2004). These data include observations of spontaneous recovery, reinstatement, and renewal (Bouton, 2002). The first two phenomena will be described briefly before moving on to a more extended examination of renewal, which is the focus of the current thesis.

Spontaneous recovery refers to the finding that when a sufficient passage of time between extinction and test is allowed to pass, the CS can again come to elicit the CR despite no further learning trials. This recovery of learning has been replicated many times in animal learning (for a review, see Rescorla, 2004), and also in human contingency learning (Rosas, Vila, Lugo, & Lopez, 2001). Reinstatement is a related phenomenon in which the CR re-emerges as a consequence of the presentation of the US in the learning context. Again this phenomenon has been demonstrated in animal learning (Bouton & Bolles, 1979b), and in human contingency learning (Garcia-Gutierrez & Rosas, 2003).

Renewal of the CR can also occur if there is a change in context between the extinction and test phases. For example, Pavlov's dogs may have learned in week 1 that during the day (context Y) the bell predicts food, and then learned in week 2 that during the night (context Z), the bell predicts nothing. If the dogs are then presented with the ringing bell during the day, it is likely that they will again show the CR, but this is unlikely to occur if the ringing bell is presented during the night. Renewal has been replicated in almost every animal conditioning preparation in which it has been tested (for a review, see Bouton, 2002), and has also been observed in several human learning studies (e.g., Rosas et al., 2001).

These three examples demonstrate that even after extinction training, the original CS-US link remains intact in memory, but may be poorly retrieved (Rescorla, 1996; 2001). Moreover, it appears that context plays an important role in determining which memories are retrieved, where context refers to the set of stimuli that surround the CS and US. The context may be internal, such as a drug state (Overton, 1964), external, such as a scent, (Bouton & Ricker, 1994), temporal, such as the passage of time (Pavlov, 1927), and associative, such as associations the CS has with other stimuli (Garcia-Gutierrez & Rosas, 2003).

Renewal

Evidence for renewal in animals

One of the first experimental demonstrations of renewal in the animal literature was in a study by Bouton and Bolles (1979a), where rats were presented with tone-shock pairings in a Skinner box (context Y) before experiencing tones in the absence of shock in an activity box (context Z). When the rats were returned to the training context Y, fear of the tone was observed to have returned to levels nearly equal to controls that had never been exposed to tone-alone extinction trials. Context exposure for all rats was equivalent. Bouton and Bolles called this ‘YZY renewal’[1] since acquisition, extinction and testing occurred in contexts Y, Z, and Y respectively. There have also been demonstrations of ‘YZX renewal’(Bouton &

[1] Bouton and Bolles actually called this “ABA” renewal, but this labelling was confusing with the current design and was therefore changed appropriately.

Brooks, 1993) and ‘YYZ renewal’(Bouton & Ricker, 1994). It appears that subjects need to be in the extinction context in order to show responding consistent with extinction training. That fact that renewal does not occur using a YYY or YZZ paradigm (e.g., Corcoran & Maren, 2004) suggests that there must be a switch out of the extinction context for renewal to occur. Renewal, it appears, is quite a robust, general effect (for a review, see Bouton, 2002).

Evidence for renewal in humans

Renewal has also been observed in human memory and learning tasks. Initial demonstrations can be found in the early verbal learning literature investigating memory interference (see Slamecka & Ceraso, 1960 for a review). For example Greenspoon & Ranyard (1957) had participants learn one list of nonsense syllables, list A, in context Y, and a second list, list B, in context Z. The context was manipulated by conducting the experiments in different rooms, the experimenter taking up a different posture and position within the room, and by using a different sized stimulus exposure stand in each room. Memory for list A was significantly better when the test was conducted in the same context as initial learning, and when list B was learned in a different context (YZY), compared to when both lists were learned in the same context (YYY, YYZ). This memory effect, of course, parallels the observations of extinction and renewal in learning studies. The authors suggested that the results were due to memory interference from list B, which shared the same contextual cues as list A. Participants’ that experienced YZY, in contrast, demonstrated good retrieval for list A as this list was uniquely cued by the context. These findings provided initial evidence for the important role played by context in memory retrieval processes. Indeed, such evidence served in the development of the encoding specificity principle of memory, which states that recall is improved when information available at the time of encoding is also available at the time of retrieval (Tulving & Thomson, 1973).

In recent decades the importance of context has been further investigated in a varied number of human extinction and renewal tasks (e.g., Baeyens et al., 2005; Lovibond, Davis, & O'Flaherty, 2000; Mineka, Mystkowski, Hladek, & Rodriguez, 1999; Neumann, 2006; Paredes-Olay & Rosas, 1999; Romero, Vila, & Rosas, 2003; Rosas & Callejas-Aguilera, 2006; Rosas et al., 2001). For example, Paredes-Olay and Rosas (1999) asked participants to predict the likelihood of illness in a hypothetical patient after the patient consumed a particular compound. For nine of the first 12 presentations of each cue, consuming compounds A or B resulted in illness. In the subsequent extinction phase, all 12 trials where compound A was consumed resulted in illness, and compound B was not presented at all. Extinction was observed in that participants predicted that illness would not occur if compound A was consumed, compared to compound B. Renewal was then investigated by manipulating the learning and test contexts. During acquisition training participants' learnt that compound A resulted in illness at hospital Y, and compound B resulted in no illness at hospital Z. During extinction training participants learnt that compound A now resulted in no illness at hospital Z, and compound B continued to result in no illness in hospital Y. At test participants were asked to predict the probability of illness when given compound A at the acquisition context (hospital Y; YZY condition) and the extinction context (hospital Z; YZZ condition). Paredes-Olay and Rosas found evidence for renewal in that compound A was predicted to be significantly more likely to result in illness when tested in the acquisition training context than in the extinction training context.

The literature review presented above suggests that extinction training, in both animal and human paradigms, leads to new learning that supplants original learning. Original learning, however, may be renewed when testing occurs in a context different from the one experienced during extinction training. The psychological mechanisms thought to be responsible for renewal are described in the following section.

The Associative Approach

The majority of theories proposed to account for simple conditioning, such as extinction and renewal, have taken an associative approach. Note that defenders of the associative approach do not deny that high-level cognition, such as reasoning, can and does occur. They argue instead for a dual system whereby learning can be accomplished by a ‘low-level’ associative system that operates independently but in parallel with a sometimes conflicting ‘higher-level’ reasoning system. That is, learning takes place within a ‘levels of learning’ framework (Dickinson, 2001).

The central concept of associative theory is that learning involves the formation of a hypothetical mental link between distinct nodes stored in memory (Pearce & Bouton, 2001). A node is envisioned to represent a certain specific stimulus, and the mental link connecting the nodes can be likened to a copper wire. Just as a copper wire may transmit electricity, the mental link may transmit ‘activation energy’ between linked representations (Wagner, 1981). Thus, the experience of a CS and US paired together will automatically cause a link to form between the appropriate two nodes in memory. For example, Pavlov’s dogs, after hearing the bell and receiving food, would have automatically formed a mental link between their mental representation of food and their mental representation of the bell. Mental links need not only be excitatory; inhibitory associations may also form. For example, a stimulus may become a conditioned inhibitor if its presentation predicts the *absence* of the US, which otherwise might be expected to occur. Within certain constraints imposed by different models, this process is envisioned to be bottom-up and automatic, that is, unintentional, unconscious, goal independent, purely stimulus driven, and fast (Moors & De Houwer, 2006).

According to associative theory, after successful conditioning between the CS and US, future presentations of either stimulus will activate that stimuli’s memory node and also, via established mental links, activate other strongly associated mental nodes (Shanks, 2007).

Thus, after Pavlov's dogs learned the bell-food association, whenever they heard the ringing bell the node corresponding to the bell would activate and, via the CS-US link, the food node would also become activated. This in turn would result in the dog thinking about food, so they begin to salivate. The salivation is taken to indirectly reflect the level of US activation, which is in turn proportional to the strength of the CS-US association (Pearce, 1997).

Most associative models (e.g., Bouton, 1993, 1997; Mackintosh, 1975; Pearce, 1987, 1994; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981) share these assumptions, but impose different restrictions as to how stimulus contiguity generates a CS-US link, or how this link functions to produce behaviour. Of the various low-level associative models, the most famous and intuitive is the Rescorla-Wagner model, and the most comprehensive is Bouton's (1993, 1997) retrieval-failure model.

The Rescorla-Wagner model

The Rescorla-Wagner model (Rescorla & Wagner, 1972) conceptualises extinction training as causing a reduction in the strength of the link between the CS and US until that link is effectively destroyed. That is, extinction leads to unlearning of the original CS-US association. Data from renewal, reinstatement, and spontaneous recovery clearly show that it cannot be the case that extinction destroys the original CS-US link because, under the right conditions, that link can be reactivated and retrieved (Bouton, 2004).

Bouton's retrieval-failure model

If extinction does not cause the CS-US link to become broken, then the next most likely proposal is that extinction results in the formation of new learning between the CS and US. This idea is no better captured than in Bouton's (1993; 1997) retrieval-failure model of learning, which posits that two mechanisms are responsible for forgetting: retroactive interference and context change. Both of these mechanisms are characterised as failures of retrieval. That is, learning experiences are all successfully encoded and stored, and failures of memory actually reflect failures to successfully retrieve the relevant memories from storage.

During acquisition training, CS and US pairings result in the formation of an excitatory link between the CS and US memory nodes. During extinction training a second inhibitory link is thought to form between the CS and US memory nodes, because the CS now predicts the absence of the expected US. For example, Pavlov's dogs may have initially formed an excitatory link between the ringing bell and food, but extinction training causes a subsequent inhibitory link to form between the ringing bell and food (Figure 1). At this point the status of the ringing bell is ambiguous since originally it predicted the presentation of food (excitatory link), but later did not (inhibitory link). In fact, the bell now has the potential to either excite or inhibit activation of the food node. Extinction, according to Bouton, reflects the second-learnt, inhibitory link interfering with retrieval of the first-learnt excitatory link. At test, this is observed as a failure to produce the CR (i.e., salivation). Extinction is thus understood as a failure to retrieve the original CS-US excitatory link.

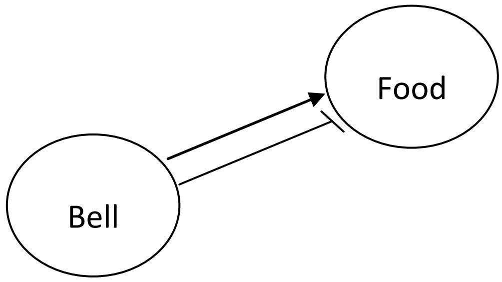

Figure 1. Bouton's (1993) model for extinction. Circle indicates mental representation. Arrow indicates excitatory association. Blocked line indicates inhibitory association.

This explanation satisfactorily accounts for extinction, but how then might it explain renewal? According to Bouton (1993), when extinction training begins, and the subject's expectations are contradicted, they actively search for some cause or factor that might help them to disambiguate the uncertainty produced by the sudden change in outcome. One factor that they may focus on is the context, especially in the YZY renewal design, in which the context changes as well as the reinforcement contingency. As a result, the subject directs attention to the context, and the inhibitory link that forms between the CS and US becomes context dependent. Bouton (1994) likens this situation to the ambiguous meaning of the word "*Fire*", which could mean "pull the trigger" or "flame", depending on the context of its verbalization. In both cases the ambiguity is resolved by context.

For example, if Pavlov's dogs learnt that food followed the ringing bell during the day and then found that during the night that this association did not hold during the night, a second inhibitory link would form between the bell and food nodes that was mediated by the night context. This can be likened to an AND-gate mechanism that only permits activation to pass across the inhibitory link if both the CS and context nodes are activated (Figure 2). Note that context is not encoded during the initial learning trials, but only during the extinction trials when the subject begins to pay attention to the context. Following from this, ringing the bell during the night will pass activation from both CS and context nodes to the AND-gate, where they will sum and inhibit the US such that the dogs do not salivate. Any test conducted outside of the extinction context will fail to pass activation from the context node and thus fail to summate with the CS at the AND-gate. Instead, activation will pass across the CS-US excitatory link, causing the appropriate CR of salivation. Renewal is thus understood as a failure to activate and therefore retrieve the second-learned, context mediated inhibitory link.

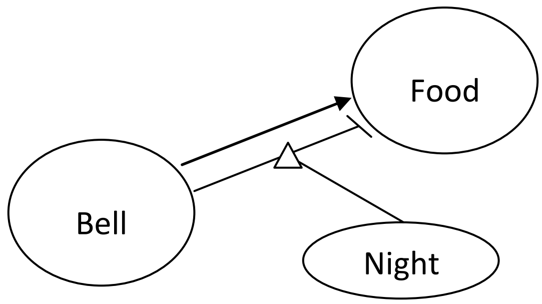

Figure 2. Bouton's (1997) model for renewal. Circle indicates mental representation. Arrow indicates excitatory association. Blocked line indicates inhibitory association. Triangle indicates AND-gate mechanism.

Human causal learning as an analogue of Pavlovian conditioning

Associative models, particularly Bouton's (1993) retrieval-failure model, have also been successfully applied to areas of human learning, particularly causal learning. Causal learning occurs when an individual learns whether one stimulus produces another stimulus, and is important for everyday functioning because it allows an individual to predict and control their environment (De Houwer & Beckers, 2002). Many learning and memory phenomena observed in animal behavioural research have successfully transferred to causal learning research, suggesting that the same underlying mechanisms are responsible for both. These include the importance of event contiguity (Shanks & Dickinson, 1987), event contingency (Shanks & Dickinson, 1988), acquisition functions (Dickinson, Shanks, & Evenden, 1984), and trial order effects (Dickinson & Burke, 1996), as well as demonstrations of blocking (Dickinson et al., 1984), inhibitory learning (Chapman & Robbins, 1990), reinstatement (Vila & Rosas, 2001a), spontaneous recovery (Vila & Rosas, 2001b), and extinction and renewal (Paredes-Olay & Rosas, 1999; Rosas et al., 2001).

Dickinson, Shanks, and Evenden (1984) were the first to suggest that human causal judgements may operate according to the principles of associative theory. For example, in the allergist task described previously, presentation of the food cues are likened to a CS and the presentation of the allergic reaction outcomes are likened to a US. The repeated pairing of the cues and outcomes over trials is hypothesised to create an associative link between the representational nodes in memory. Consequently, presentation of a cue will pass activation, via the associative link, to the allergic reaction node paired with it. This in turn will cause the subject to provide a high causal judgement, which is analogous to a CR.

According to this associative view the causal status of a particular cue is determined by the degree to which the presentation of that cue activates a representation of the outcome in memory. For example, the renewal in predictive learning effect observed by Paredes-Olay

& Rosas (1999), described above, was interpreted in terms of Bouton's (1997) retrieval-failure model. They argued that during acquisition training, an excitatory link formed between compound A and the illness. During extinction training, an inhibitory link also formed between compound A and illness, but this second link was specific to the context in which it was formed, hospital Z. At test, the context served to disambiguate the status of compound A. No illness was predicted in the YZZ condition because compound A and hospital Y activation successfully gated the inhibitory link and therefore the illness outcome was not activated in memory. In contrast, illness was predicted in the YZY condition because compound A activation successfully passed across the context-free excitatory link and therefore activated the illness outcome in memory.

Due to the fact that there were only two cues, two outcomes, and two contexts, which constitutes a small number of levels within these variables, the conclusions of this study must be approached with caution. An obvious alternative explanation is that participants could recall all of the cue-outcome learning trials at test and were actually responding based on a reasoning process. For instance, over the course of the first two phases the participants may have formulated the following propositional hypothesis: "*Compound A results in illness at hospital Y, but results in no illness at hospital Z*". Use of such a hypothesis could easily account for the observations. This sort of reasoning process forms the heart of the inferential approach, the major alternative to associative theory, to which we will now turn.

The Inferential Approach

Failures of the associative approach

A core assumption of the associative approach is that extinction and renewal reflect a phenomenon of memory. This is most readily appreciated in Bouton's (1997) retrieval-failure model, which explicitly argues that extinction is the result of a failure to retrieve first-learnt information due to retroactive interference from second-learnt information. Similarly, renewal is thought to be the result of a failure to retrieve context-dependent second-learnt information, but successful retrieval of first-learnt information. According to the inferential approach, however, extinction and renewal may not reflect memory phenomena at all, but could instead be the result of a reasoning process. Thus, in the Paredes-Olay and Rosas (1999) study for example, subjects were likely completely aware of the fact that a cue was paired with an outcome in one context, and not paired with an outcome in a second context. They may have reasoned their responses based on a complete memory for the learning trials.

The inferential approach is supported by a large body of data unexplained by the associative approach. One such piece of evidence is the effect of additivity in blocking. In the blocking paradigm a specific cue, A, is paired with an outcome during phase 1 (A+; where the ‘+’ denotes reinforcement). In a second phase, cue A is presented in compound with a second cue, B, and this compound is paired with the outcome (AB+). When the cues are subsequently independently tested, the subject typically shows strong conditioning to A and little or no conditioning to B, compared to appropriate controls (Kamin, 1969). According to the automatic link-formation mechanism endorsed by the associative approach, this occurs because cue A already predicts + in the AB trials and so nothing is learned about cue B.

Lovibond, Been, Mitchell, Bouton, and Frohardt (2003) suggested that blocking might be better explained by an inferential reasoning process, based on a natural expectation that two simultaneous causes should produce an effect of greater magnitude than one cause alone Thus, if cue A is a cause (A+) and cue B is a cause (B+), then simultaneously presenting both cues should result in a larger overall effect (AB++). In the blocking paradigm participants are only presented with A+ and AB+, from which it is logical to conclude B-. Indeed, Lovibond et al. only observed blocking when it was logical to make this inference. This, they argued, demonstrated that participants had attempted to figure out an underlying causal structure.

The propositional nature of human associative learning

The inferential approach argues that learning that two or more events in the world are associated with one another can only be accomplished by a higher-order reasoning process (Mitchell et al., submitted). This process is believed to be deliberate, requiring effort, and results in conscious, declarative, propositional knowledge about those events (De Houwer, Beckers, & Vandorpe, 2005). Although a variety of automatic, bottom-up processes inevitably play a role in learning, they do so merely as a consequence of limiting what information is perceived, remembered, and made available to top-down reasoning processes.

For example, subsequent to experiencing a ringing bell followed by a shock, future presentations of the ringing bell will automatically retrieve that painful experience from memory. It might also retrieve any past hypotheses the subject has made with regard to the likelihood of shock following the ringing bell. Although the subject is able to remember the event, they will not yet show a CR to future presentations of the ringing bell. First the subject must hypothesise, all things being equal, that the next time they hear the ringing bell they will receive another shock. If their expectations are indeed fulfilled then they will encode this confirmed hypothesis into memory in propositional form, such as, "*I believe that when I hear a bell, I will receive an electric shock*". This encoded proposition, which is based on the now confirmed hypothesis, constitutes learning.

Note from this example that bottom up, automatic processes such as the recollection of previous experiences and recollection of hypotheses serve as necessary inputs into the inferential reasoning processes. These inputs, however, in and of themselves do not amount to learning, for example, by forming a mental link between the ringing bell and shock. What is required to learn that a cue and outcome are associated is to entertain a hypothesis that describes how these stimuli are related, and storing this hypothesis as a propositional statement. Hence, learning requires inferential reasoning.

This framework has led to a number of novel predictions that distinguish it from the associative approach (De Houwer, Beckers et al., 2005). For example, according to the inferential approach learning is the result of a conscious reasoning process, which suggests that learning should always be accompanied by contingency awareness. In contrast, the associative model describes learning as automatic and may occur outside of awareness. In support of the inferential approach, to date there is no convincing evidence that Pavlovian conditioning can occur outside of awareness (Lovibond & Shanks, 2002).

The inferential approach argues that reasoning is an effortful process, whereas the associative approach suggests that learning is automatic and does not depend on the amount of cognitive resources available. In support of the inferential model, it has been observed that difficult secondary tasks have a greater negative impact on learning than easy secondary tasks (De Houwer & Beckers, 2003). Finally, the inferential approach states that learning involves the formation of propositional knowledge based on the use of abstract rules and deductive reasoning processes. In contrast, the associative approach holds that learning is non-propositional. In support of the inferential model, it has been observed that informing participants verbally of an association between stimuli is sufficient to produce learning (Lovibond, 2003), and that participants can learn and apply complex rules to successfully perform a task (Shanks & Darby, 1998).

The inferential approach can thus provide an alternative explanation to the observations of extinction and renewal described earlier. Participants may simply be testing hypotheses about the association between various cues, outcomes, and contexts and storing these hypotheses in propositional form, such as, "*I believe that when Mr. X eats mushrooms, then he will suffer from an allergic reaction, but only if he was eating at restaurant Y*". One straightforward way of testing between Bouton's (1997) retrieval-failure model, and the reasoning explanation posed by the inferential approach, is to test memory directly.

Cued recall as a direct measure of associative strength

According to the associative approach, the strength of the link between the CS and US nodes can be indirectly measured by the strength of the CR: a strong CR reflects a strong CS-US link. In the case of the allergist task, a high causal judgement is assumed to reflect strong activation of the outcome node, which in turn is an indication of a strong link between the cue and outcome. Thus, for example, after establishing a cue-outcome association, presentation of the cue will excite the outcome memory node, and result in high causal judgement. Note that this description makes the explicit prediction that patterns of causal judgement reflect the degree to which the outcome node is activated when a cue is presented. If the causal rating is low, then this is assumed to reflect a weak link between the cue and outcome memory nodes. This conceptualisation suggests an alternative, more direct, way of measuring associative strength – to directly test memory for, or activation of, the outcome when the cue is presented. That is, to use a cued recall task (Mitchell, Lovibond, & Gan, 2005). According to the associative approach, because recall and judgement both reflect the associative strength of a cue-outcome link, they should show a strong correlation. In contrast, the inferential approach argues that causal judgements are made by a reasoning process based on good memory for cues-outcome learning trials; it may therefore be possible to dissociate the two.

For example, Mitchell et al. (2005) used a cued-recall task to measure the associative strength between a food cue and an allergic reaction outcome in the allergist task. They used a cue-competition design in which causal compounds AB+ and CD+ were intermixed with A+ and C- training. In total eight foods were paired with four fictitious allergic reactions. At test participants were presented with each cue individually and were required to identify which outcome that cue had been paired with (i.e. a cued recall test). In addition, they were asked to make a causal judgement reflecting the extent to which they believed the cue had caused the outcome. Mitchell et al. found that recall for the outcome paired with cue B was better than recall for the outcome paired with cue D. In spite of the better memory for the cue B-outcome association, cue B was judged to be less causal than cue D. The authors interpreted this double dissociation as evidence for an inferential account of causal judgment. They argue that, at test, participants could recall that cue B had been paired with an allergic reaction when presented with cue A, but also recall that cue A was paired with an allergic reaction when it was presented alone. These two memories allowed the subjects to perform well on the cued-recall measure. However, these same memories also allowed the subjects to draw the inference that cue A, not cue B, was responsible for the allergic reaction in the AB+ trials, thus causing the subject to give a low causal rating to cue B.

It thus appears that, at least in certain circumstances, learning is due to inferential processes and not to the ability of the cue to activate the outcome. Could other psychological phenomena be similarly explained? Extinction is particularly interesting in this line of investigation since it has been explicitly described as a failure to retrieve original learning (Bouton, 1997). Could it, too, be the result of an inferential process based on good memory? Scully and Mitchell (in press) investigated this question using the allergist task with an extinction design, and measured both cued recall and causal judgement. Participants were asked to take on the role of an allergist and were presented with 15 common foods, some of which caused one of three fictitious illnesses or no reaction. During an extinction phase some of the previously hazardous foods were presented alone. On test, participants were required to rate both their memory of cue-outcome as well as their causal rating of that cue.

As is typical with causal judgement in extinction, participants rated the extinguished cue as less causal than a control cue that had not been presented during the extinction phase. Interestingly, extinction was also observed in cued-recall: compared to the controls, memory for the first-learned cue-outcome relationship was poor. These results appear to support Bouton's (1993; 1997) retrieval-failure hypothesis that extinction in causal judgement, at least partly, reflects a failure to strongly activate the outcome representation when presented with the cue at test. The present study extends this finding to examine renewal.

The Present Study

Given the importance of context effects to (a) the interpretation of the mechanism responsible for extinction, and (b) the methods used in clinical settings, it is vital to find out whether renewal is indeed a release from retrieval interference in memory (Bouton, 1997), or simply the encoding of a propositional statement describing the different outcomes of certain cues, produced in different contexts (Mitchell et al., submitted).

In clinical settings, for example, fear disorders – such as the phobia suffered by the dog-owner in the first example – are believed to be acquired through a process of Pavlovian conditioning (Wolpe, 1958). Treatments were developed based on the extinction design and the assumption that extinction worked via an unconscious automatic learning mechanism. One example of this was systematic desensitization, which comprised of repeated and controlled exposure to the feared stimulus, while the patient remained completely relaxed, until a reduction of the conditioned fear response was observed (e.g., Craske & Rowe, 1997). Although this extinction-based treatment was successful in some attempts (Chambless & Ollendick, 2001), there are several examples of fear returning after therapy (i.e., renewal, Rachman, 1989). If Pavlovian conditioning is in fact due to a conscious propositional reasoning system, then a better understanding of how, and what, learning experiences impact on propositional knowledge may help to develop more effective clinical interventions.

The broad aim of the current study is to further investigate the role of higher-level processes in causal learning. More particularly, the current study seeks to test predictions derived from Bouton’s (1997) retrieval-failure hypothesis against predictions derived from the inferential reasoning approach (e.g., Mitchell et al., 2005), using a renewal paradigm.

Experiment 1

Following from Bouton's (1997) associative model, causal judgements are made based on the strength of a cue-outcome association, and extinction and renewal phenomena reflect failures of retrieval due to memory interference. As argued above, causal judgement is believed to reflect the level of activation of the outcome node, which in turn indicates the strength of the cue-outcome association. It was also suggested that outcome activation may be measured more directly with a cued recall task. Thus, the associative approach predicts a strong correlation between cued recall and causal judgement. For instance, extinction in causal judgement is expected to be accompanied by extinction in memory. In contrast, the inferential approach argues that causal judgements occur via a reasoning process based on good memory for the training trials. As a result, this approach predicts that it should be possible to dissociate these two measures. For instance, extinction in judgement is expected to occur despite good cued recall, that is, even when outcome activation is high.

The methodology used by Scully and Mitchell (in press), which measured both cued recall and causal judgement, provides one way of evaluating these conflicting predictions. The broad aim of Experiment 1 then was to introduce a context manipulation into the allergist-task paradigm used by Scully and Mitchell. A much simpler version of the task was used in order to maintain a design comparable to previous causal judgement renewal studies (e.g., Paredes-Olay & Rosas, 1999) that interpret their findings in terms of retrieval failure.

The associative and inferential approaches make the same predictions with regards to the causal judgement measure. Participants in the Extinction group are expected to judge extinction cue A as non-causal, or give it a low causal rating, compared to control cue B. In contrast, participants in the Renewal group are expected to judge cue A equivalently to cue B.

The associative and inferential approaches make different predictions with regards to memory measure. According to the associative approach participants in the Extinction group should show poor recall for the outcome paired with extinction cue A, compared to the outcome paired with control cue B. In addition, participants in the Renewal group are expected to show equivalently good recall for the outcome paired with cue A as to the outcome paired with cue B. That is, the associative approach predicts a strong similarity in the pattern of results observed in the judgement and memory measures because causal judgement is based on memory. In contrast, the inferential approach allows that participants might have good recall for the outcomes paired with both cues A and B in both groups, even though causal judgements for cue A may be lower than for cue B.

Method

Participants

The sample consisted of 42 undergraduate first year University of New South Wales psychology students, 13 males and 29 females, participating in exchange for course credit. The average age was 19.6 years, with a range of 18 to 43 years.

Design

The experiment was a 2 x (2) mixed design. The between-subjects factor was the context during the extinction phase. Participants in the Extinction group received acquisition, extinction and test phases all in the same context, whereas participants in the Renewal group received extinction training in a different context. The within-subjects factor was cue presentation during the extinction phase. Specifically, the reinforcement contingency of extinction cue A was changed, while control cue B was not presented at all.

The experiment contained three phases: acquisition, extinction, and test (Table 1). All participants received training in context Y in the acquisition phase, where cue A was paired with outcome O_1, cue B paired with outcome O_2, and cue C with no outcome. In the subsequent extinction phase, the Extinction group received extinction training in context Y, whereas the Renewal group received extinction training in context Z. During this phase the previously causal cue A was not followed by an outcome. The control cue B was not presented. Filler cues C and D ensured that each phase contained at least two different outcomes. Each trial type was presented six separate times per phase, resulting in 36 trials in total. Trials were presented in an intermixed, random order within phases, and there was no break or discernable change between the acquisition and extinction phases.

Table 1

Outline of the Design of Experiment 1.

	Acquisition Phase		Extinction Phase		Test Phase	
	Context Y	Context Z	Context Y	Context Z	Context Y	Context Z
Extinction Group	A-O_1	*	A-	*	A?	
	B-O_2	*			B?	
	C-	*	C-	*		
			D-O_2	*	D?	
Renewal Group	A-O_1	*	*	A-	A?	
	B-O_2	*			B?	
	C-	*	*	C-		
			*	D-O_2	D?	

Note: Letters A – D refer to food cues. Letters O_1 – O_2 refer to the allergic reaction outcomes. Contexts Y and Z refer to the restaurant environments. * refers to simple exposure to the alternative context. Cues, outcomes, and contexts were randomised at the beginning of the experiment.

Presentation of cue-outcome pairings in one context were randomly intermixed with simple exposure to the alternative context. Simple exposure to a context was identical to normal learning trial screens, with the exception that the place where cues and outcomes were usually presented was replaced with the sentence *"Mr. X ate nothing"*. In total, there was the same number of simple exposure trials as there were training trials in the alternative context. This ensured that at the end of the learning phases each context was equally familiar.

Two measures were assessed in the test phase, which was conducted in context Y for both groups. First, a causal judgement rating measured participant's belief as to the causal status of cues A, B, and D. Second, a cued recall test measured participants' memory for the outcome paired with the cues A, B, and D. "No reaction" was removed as an option because participants may have selected this option for the extinction cue A, which was technically correct, even if they could recall the outcome paired with cue A during the acquisition phase (the specific aim of the measure). As a result, cue C was omitted from test because it was paired with "no reaction" in both phases. The memory test always preceded the judgement test and each cue was assessed on both measures before moving on to the next cue. The order of the cues tested was randomised.

Apparatus and stimuli

The experiment took place in a 2m x 1m room furnished with a desk, chair, lamp, and IBM-compatible computer. All instructions, stimuli and questions were presented on the 15" computer screen, and programmed with Revolution Studio Max. Participants were able to move between screen displays and respond to stimuli using the computer mouse.

The four food cues were cheese, mushrooms, ham, and banana. The three outcomes were headache, nausea, and no reaction; the two contexts were the fictitious restaurants, the Shady Bar and Ocean's Pearl. The cues, outcomes, and contexts were all randomly allocated at the beginning of each experiment. All stimuli were presented in lower-case, Arial, size 36 font on the computer screen. The foods were presented in blue text, the allergic reactions in red text, and the "*no reaction*" outcome in green text. All instructions were written in black, Arial, size 24 font. The background was coloured white unless in a particular restaurant context. The Ocean's Pearl background was a picture of a restaurant on the water's edge; the Shady Bar's background was a picture of some chair stools at a bar (see Appendix B).

Procedure

The experiment was conducted in a single, half-hour session. One participant was tested at a time. After reading a brief summary and agreeing to take part in the experiment, the participant sat at the computer terminal and entered their age and gender. Participants were told to read through the instructions and to consult the experimenter, who remained in the room while the participant read the instructions, with any problems or questions. Navigation through this section was achieved by clicking on a "*Next*" button at the bottom of the screen; there was also a "*Back"* button if needed. After reading through all the instruction screens, the participant clicked on the "*Start*" button and the experimenter left the room.

The instructions informed the participant that they would be accompanying the suspected mafia boss Mr. X to two restaurants where he would be meeting some business associates. The instructions asked the participant to assume the role of an allergist and to attempt to determine which foods caused which outcomes in Mr. X (see Appendix A for instruction screen). Two practice screens were presented during the instruction phase. At the top of the practice screen the participant was presented with the sentence "*When Mr. X ate*" and the specific food cue was presented beneath. Below this, after a two second delay, the participant was presented with the sentence "*Did Mr. X experience*" along with the three possible outcomes – headache, nausea, and no reaction, in that order. In one example bread caused nausea, and in the second hamburger caused no reaction. Participants made their selection by using the mouse to click on one of the three outcomes, which highlighted after selection. If the correct outcome was selected then the words "*Correct*" appeared at the bottom of the screen with a happy face, as well as arrows indicating the correct response. If the incorrect outcome was selected, then the words "*Incorrect*" appeared at the bottom of the screen with a sad face, as well as arrows indicating the correct response.

The procedure in the actual learning phases was exactly the same as that in the earlier practice screens, except that cues and outcomes were presented in a restaurant context. The context was established in two ways. First, the words "*While Mr. X was conducting business at the Shady Bar* [or *Ocean's Pearl*] *…"* were added at the top of each screen. Second, the background picture continually changed between trials. After feedback was presented there was a one-second delay before a "*Continue*" button appeared allowing progression to the next trial. In the case of a simple exposure screen, the background picture changed and the place where cue and outcomes were usually presented was replaced with the sentence *"Mr. X ate nothing"*. After a two-second delay the "*Continue*" button automatically appeared.

Following the final training trial in the extinction phase, participants were presented with the test instructions (see Appendix C). On each test trial the participant was presented with the sentence "*When Mr. X ate"*, beneath which the particular cue appeared. After a 2 second delay, the participant was presented with the sentence "*What did Mr. X experience?*", beneath which the two outcomes appeared. Once the participant had made their response, they were presented with the question "*To what extent did* [the cue] *cause the allergic reaction?*", beneath which appeared an 11-point Likert scale ranging from 0 ("*this food was definitely NOT the cause*") to 10 ("*this food was definitely the cause*"). A "*Continue*" button appeared after a response had been made (see Appendix D for test screens). Once participants had made their response, they were not able to change their answers. No feedback was given.

At the completion of the test phase a screen appeared informing the participant that the experiment was complete, thanked them, and asked them to press the "*End*" button. The experimenter then gave the participant a verbal debrief as to the purpose of the experiment.

Results

Training

Figure 3 displays the percentage of participants who predicted an allergic reaction outcome, correct or not, for cues A, B, C, and D across the two learning phases for the Extinction and Renewal groups. It can be seen that the participants in both groups rapidly learned and remembered the experimental contingencies. Participants correctly predicted the outcome paired with cues A, B, and C with 89% accuracy by the second trial, and this improved to 99% by the end of the acquisition phase. Participants in both groups quickly learned to predict “no reaction” when the cue A contingency changed during the extinction phase. No participants in the Extinction group, and only one participant in the Renewal group, predicted that an outcome would follow the last cue A presentation in phase 2.

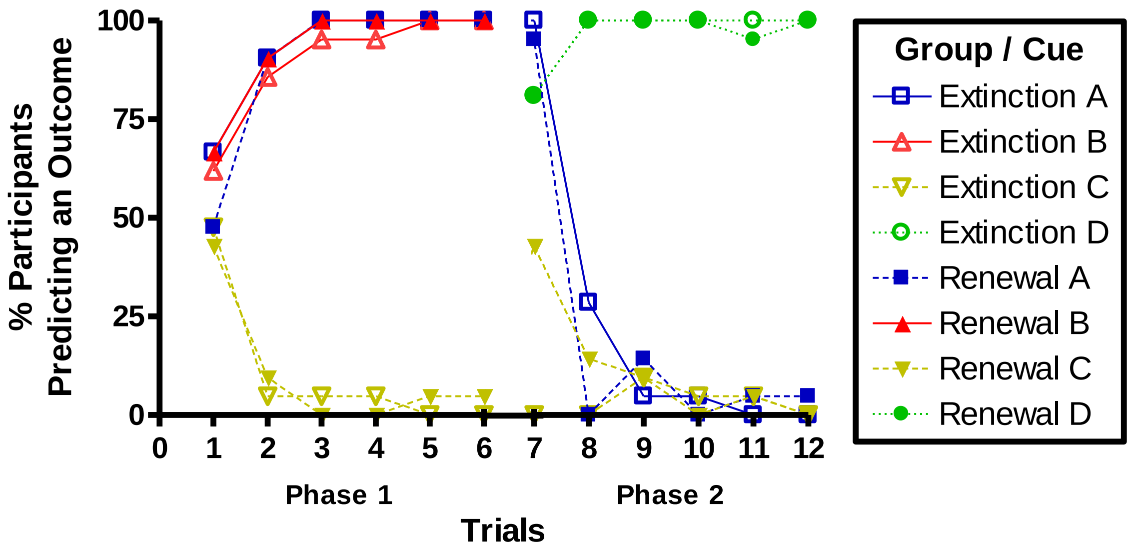

Figure 3. Percentage of participants in the Extinction (open symbols) and Renewal (filled symbols) groups who predicted an allergic reaction on each trial in phases 1 and 2 of Experiment 1.

Unexpectedly, it appears that more participants in the Extinction than in Renewal group predicted an outcome for cue A on trial 8. Post-hoc analyses confirmed a difference on trial 8 ($\chi2_{(1)} = 6.83$, $p < .01$), but no difference on trial 7 ($\chi2_{(1)} = 1.00$) or on trial 9 ($\chi2_{(1)} = 1.08$). This result will be discussed further in the General Discussion.

Judgement

Participants’ scores on the judgement measure for cues A, B, and D, which ranged between 0 and 10 for each cue, were averaged across participants. These data are presented in Figure 4. It appears that there is a difference in causal judgment between cues A and B in the Extinction group, but little difference between cues A and B in the Renewal Group. To investigate these differences, four planned contrasts in a two-factor mixed analysis of variance (ANOVA), including simple effects, were tested. A Bonferroni adjustment controlled the family wise error rate (FWER) at $\alpha = 0.05$ (see Appendix F). Note that the most convincing evidence for renewal would be an interaction, coupled with the simple effect of extinction cue A rated as more causal in the Renewal group than in the Extinction group.

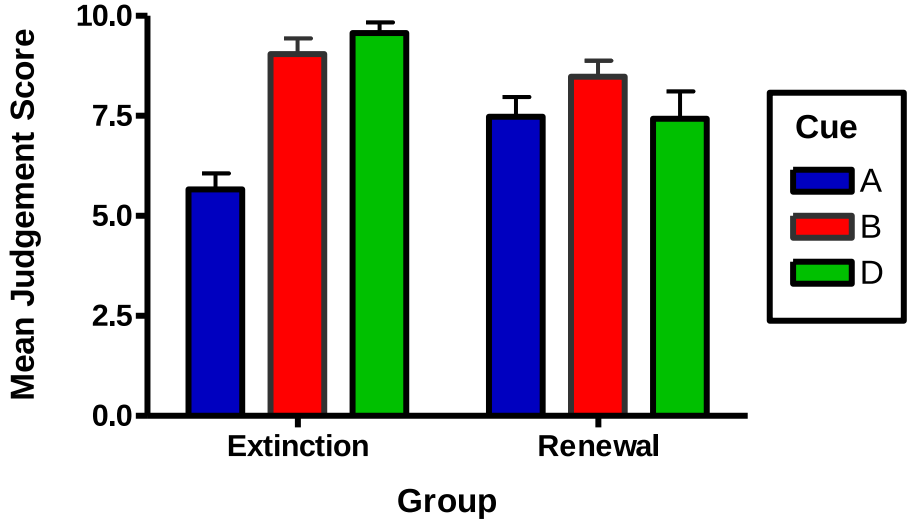

Figure 4. Mean causal judgement scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 1. Error bars denote standard errors of the mean.

The Group x Cue interaction was significant ($F_{(1, 40)} = 12.21$, $p < .01$), suggesting a renewal effect. Simple effect contrasts revealed that participants in the Extinction group exhibited extinction in that they judged extinction cue A to be less causal than control cue B ($F_{(1, 40)} = 49.25$, $p < .001$), but no such difference was observed in the Renewal group ($F_{(1, 40)} = 4.31$). Participants in the Renewal group exhibited renewal in that they judged cue A to be significantly more causal than did participants in the Extinction group ($F_{(1, 40)} = 8.09$, $p < .01$). Unexpectedly, it appears that participants in the Extinction group judged cue D to be more causal than participants in the Renewal group. This difference was confirmed in a post-hoc analysis ($F_{(1, 40)} = 8.52$, $p < .01$) and will be considered further in the General Discussion.

Memory

Participants' cued recall scores for the allergic reactions paired with cues A, B and D, recorded as 0 or 1 for each, were averaged across participants. These data are presented in Figure 5. It appears that there is little difference in recall for the allergic reactions paired with cues A and B for either group. Because the data were not normally distributed the non-parametric Kruskal-Wallis and Friedman tests were used to test the between and within subjects contrasts respectively. Contrasts revealed that there was no difference in the accuracy of recall for the outcomes paired with cues A and B in the Extinction group ($\chi 2 < 1$), nor in the Renewal group ($\chi 2_{(1)} = 1.00$). Additionally, there was no difference between the Extinction group and the Renewal group in outcome recall accuracy when presented with cue A ($H_{(1)} = 1.00$).

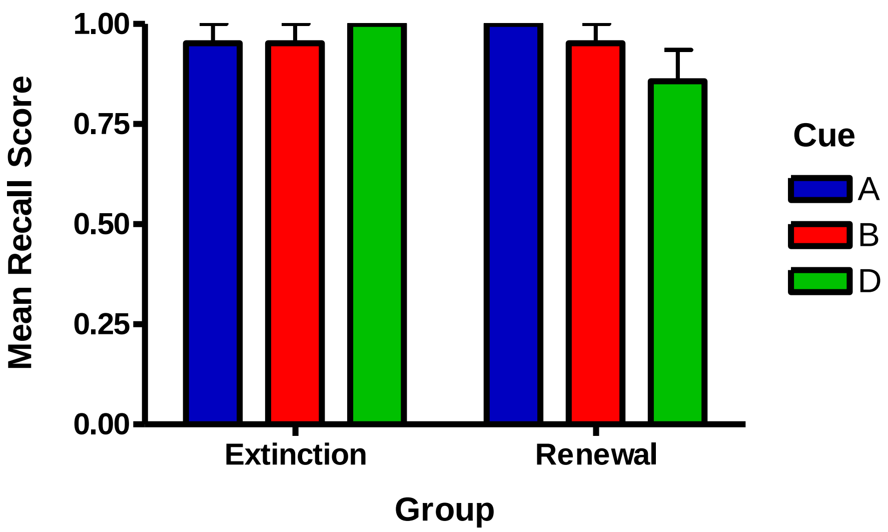

Figure 5. Mean cued recall scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 1. Error bars denote standard errors of the mean.

Discussion

The data from the causal judgement measure replicated previous observations (e.g., Paredes-Olay & Rosas, 1999; Scully & Mitchell, in press) that a cue-outcome association can be extinguished if the cue is later paired with no outcome (the Extinction group). Moreover, this extinguished cue-outcome association may be recovered if there is a context switch between the extinction phase and the test phase, or if extinction occurs in a different context from acquisition training and test (the Renewal group), again replicating previous findings (e.g., Paredes-Olay & Rosas, 1999; Rosas et al., 2001). One explanation for these results is provided by Bouton's (1997) retrieval-failure model of learning. This model asserts that, during acquisition, an excitatory association between cue A and O_1 is formed. Then, during extinction, a context-mediated inhibitory association is also formed between cue A and O_1. For the Extinction group, this inhibitory link is thought to interfere with retrieval of the excitatory A-O_1 link. This is reflected in the low causal ratings of cue A at test (i.e., extinction). For the Renewal group there is a failure to retrieve the inhibitory link when tested in context Y because the inhibitory link was formed in context Z and is specific to that context. As a result, the original A-O_1 link is retrieved, as it suffers no interference, and this is reflected in the high causal ratings of cue A on test (i.e. renewal).

The data from the memory measure, however, cast doubt on this explanation. It follows from Bouton's (1997) retrieval-failure model of learning that failure to retrieve a cue-outcome relationship will also result in poor memory for that outcome, given the cue. Thus, it should be extremely difficult to dissociate cued recall from causal judgements. The data from the memory measure show, however, that recall was equally good for both groups and for both cues A and B. This finding suggests that the results of the causal judgement measure are due to the work of a high-level reasoning process, not the degree to which the cue activates the outcome. Indeed, post-experimental discussion with the participants revealed that most found the task extremely easy and made their responses based on a specific reasoning strategy that did or did not incorporate the role of context (see General Discussion for an evaluation). Thus, the results of Experiment 1 are best accounted for by the inferential approach to causal learning (Mitchell et al., submitted).

A problem with this conclusion, however, is that the memory measure was at ceiling, and may therefore not have been sensitive enough to detect differences in recall between the groups. A second experiment was designed to rule out this explanation.

Experiment 2

The first experiment found that extinction and renewal effects in causal judgement can be obtained with the current allergist task procedure. A dissociation between causal judgement and memory was observed, contrary to Bouton's (1997) memory interference interpretation of extinction and renewal in human causal judgement. One problem with the first experiment was that it appeared to be too easy and the cued recall measure was at ceiling. Indeed, Shanks' (2007) has argued that causal judgement tasks with only a few cues are tests of reasoned as opposed to intuitive, that is, associatively-based, judgment. Indeed, defenders of the associative approach do not debate that causal learning can be accomplished by high-level reasoning; they adhere to a 'levels of learning' framework. The debate centres on whether causal learning can also be exclusively accomplished by a low-level associative, or 'intuitive', process. In order to test this, Shanks recommends the use of a complex task.

In response, the major difference between Experiments 1 and 2 was the inclusion of several more food cues in Experiment 2. This was to ensure an increase memory load so that cued recall performance would not be at ceiling, as it appeared to be in Experiment 1. Moreover, the simple exposure technique was replaced by presenting cues and outcomes in both contexts throughout the experiment. As a result, not only were contexts equated for familiarity, but also in terms of the events that took place in them.

The current experiment also used a more salient context manipulation. It combined elements from previous renewal designs including the manipulation of the screen background (Havermans, Keuker, Lataster, & Jansen, 2005), presentation of the context name on screen (Garcia-Gutierrez & Rosas, 2003), and manipulation of the room lighting and music (Neumann, 2006). A salient context change was not used in Experiment 1 in order to maintain a methodology comparable to previous studies of renewal in causal learning (e.g., Paredes-Olay & Rosas, 1999) but was included in Experiment 2 to ensure that the experiment had greater real-world validity and was more similar to the contextual manipulations used in animal behavioural research, the literature from which associative theory was derived.

Method

Participants

The sample consisted of 48 undergraduate first year University of New South Wales psychology students, 19 males and 29 females, participating in exchange for course credit. The average age was 19.0 years, with a range of 17 to 26 years.

Design

Design details not mentioned here were the same as those in Experiment 1. The experiment was divided into three phases, and each phase was divided into two blocks marked by a context change (see Table 2). The two blocks in any one phase contained equivalent cue-outcome contingencies, but different cues and context. For example, the experience of a participant in the Extinction group for cue A may have been as follows: garlic is paired with headache (A_1-O_1) at the Shady Bar (context Y) in the first block of acquisition training, and then paired with no reaction (A-) in the first block of extinction training in the Shady Bar (context Y). Likewise, cheese is paired with nausea (A_2-O_2) at the Ocean's Pearl (context Z) in the second block of acquisition training, and then paired with no reaction (A_2-)

in the second block of extinction training at the Ocean's Pearl (context Z). A context change occurred mid-way through each of the three phases, which also represented the beginning of a new block. The order in which the blocks were presented is demonstrated on Table 2, that is, YZZY during the learning phases, and YZ during the test phase. This ensured that each participant experienced three context changes during the experiment. Each complete block was presented before moving on to the next block.

Table 2

Outline of the Design of Experiment 2.

	Acquisition Phase		Extinction Phase		Test Phase	
	Context Y	Context Z	Context Z	Context Y	Context Y	Context Z
Extinction Group	A_1-O_1	A_2-O_2	A_2-	A_1-	A_1?	A_2?
	B_1-O_2	B_2-O_1			B_1?	B_2?
	C_1-	C_2-	C_2-	C_1-		
			D_2-O_1	D_1-O_2	D_1?	D_2?
	E_1-	E_2-	E_2-	E_1-		
	F_1-O_1	F_2-O_2	F_2-O_2	F_1-O_1	F_1?	F_2?
			G_2-	G_1-		
Renewal Group	A_1-O_1	A_2-O_2	A_1-	A_2-	A_1?	A_2?
	B_1-O_2	B_2-O_1			B_1?	B_2?
	C_1-	C_2-	C_2-	C_1-		
			D_2-O_1	D_1-O_2	D_1?	D_2?
	E_1-	E_2-	E_2-	E_1-		
	F_1-O_1	F_2-O_2	F_2-O_2	F_1-O_1	F_1?	F_2?
			G_2-	G_1-		

Note: Letters A_1 – G_2 refer to food cues. Letters O_1 – O_2 refer to allergic reaction outcomes. Contexts Y and Z refer to restaurant environments. All were randomised at the beginning of the experiment.

Simple exposure to the context was replaced with presenting additional cues in the alternative context. Thus, for each group two cues were trained and extinguished, (A_1 and A_2), two cues served as controls (B_1 and B_2), and so forth for every cue. For the Extinction group, the cues were trained and extinguished in the same context; for the Renewal group, they were extinguished in a different context from that in which they were trained and tested.

For ease of explanation, the two individual exemplars of each cue class shall be referred to as a single cue. For example, cues A_1 and A_2 will be referred to as cue A.

In addition to the extinction cue A, control cue B, and filler cues C and D, Experiment 2 used additional filler cues E, F, and G. Cue E was paired with no outcome in both phases, cue F was paired with an allergic reaction outcome in both phases, and cue G was paired with no outcome and only presented in the second phase. It is important to note that apart from the experimental cues A and B (whose contingencies were identical to those of cues A and B in Experiment 1) the remaining filler cue-outcome contingencies remained the same for the entire experiment and were always presented in the same context.

Each cue-outcome pair was presented five times in each block that it appeared. As a result there were 50 cue-outcome pairings per phase, and 100 cue-outcome pairings in total. Trials in each block were presented in a random order. Throughout the two learning phases there was an equivalent number of outcome-causing and non-outcome-causing cues in each context. Thus experience with cues, outcomes, and contexts were equated throughout the acquisition and extinction phases.

Apparatus and stimuli

Apparatus and stimuli details not mentioned here were the same as those in Experiment 1. A CD player in the corner of the room played background music in each context. The Ocean’s Pearl context was created by turning the lights on, playing classical piano music (Beethoven: Moonlight Sonata and Chopin: Nocturne in G minor), and changing the screen background. In contrast, the Shady Bar context was created by turning off the lights, playing pub rock music (Cold Chisel: Khe Sanh and Cold Chisel: Cheap Wine), and changing the screen background. The lamp remained on at all times.

An additional twelve food cues (lemon, yoghurt, garlic, steak, chicken, tomato, fish, lettuce, beef, cucumber, beans, and turkey) were used in Experiment 2.

Procedure

Procedural details not mentioned here were the same as those in Experiment 1. After the participant had read through the instructions and worked through the two examples (just prior to pressing the "*Start*" button), the experimenter explained that in order to make the two restaurants seem "more real" the music and lighting would be manipulated. At this point the music was started, the overhead light switched off (if the Shady Bar was the first context experienced), and the experimenter left the room.

All instructions were identical to Experiment 1, except those requiring the participant to fetch the experimenter in the middle of each phase to change the context (see Appendix A). At this point a screen was presented asking participants to contact Mr. X's "*right hand man*" to transport them to the next restaurant. The experimenter silently came into the room, changed the lights, music, and then pressed a key on the keyboard allowing the participant to continue. A similar procedure occurred mid-way through the test phase.

Results

In order to demonstrate an extinction and renewal effect at test, it was important that all participants actually learned the contingencies of cues A and B during the acquisition phase. Thirteen participants who consecutively failed to correctly predict the outcome paired with either cue A or cue B on trials 4 and 5 were removed from the study and replaced.

Training

Figure 6 displays the percentage of participants who predicted an allergic reaction outcome, correct or not, for cues A, B, C, and D across the two learning phases for both groups. It can be seen that the participants rapidly learned and remembered the experimental contingencies. This is not a surprise given that thirteen participants were replaced in order to achieve this performance. Participants correctly predicted the outcome paired with cues A, B, and C with 82% accuracy by the second trial, and this improved to 96% by the end of the acquisition phase. Participants in both groups very quickly learned to predict "no reaction" when the contingency of cue A changed during the extinction phase, although it appears that this occurred more slowly for those in the Extinction group. Indeed, a post-hoc analysis revealed that on trials 6 and 7 significantly more participants in the Extinction group predicted an allergic reaction outcome when presented with cue A ($F_{(1, 46)} = 5.31$, $p < .05$ for trial 6 and $F_{(1, 46)} = 10.62$, $p < .01$ for trial 7). By trial 8 this difference had disappeared ($F < 1$). This unexpected finding will be discussed further in the General Discussion.

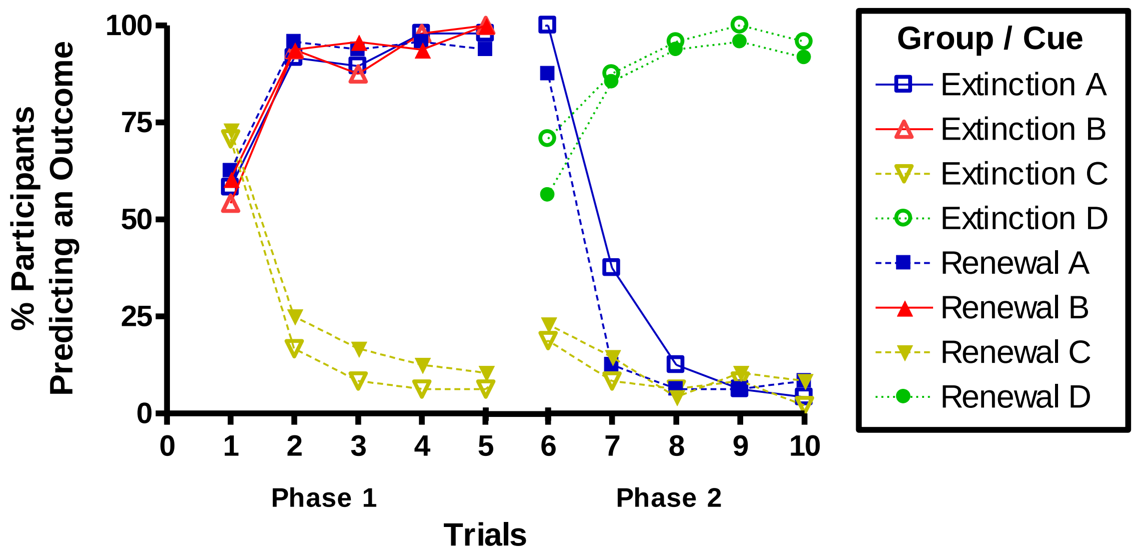

Figure 6. Percentage of participants in the Extinction (empty symbols) and Renewal (filled symbols) groups who predicted an allergic reaction on each trial in phases 1 and 2 of Experiment 2.

Judgement

Participants' scores on the judgement measure for cues A, B, and D, which ranged between 0 and 10 for each cue, were averaged across participants. These data are presented in

Figure 7. It appears that there is a large difference in causal judgment between cues A cue B in both groups. The same set of planned contrasts used in Experiment 1 was used to test the data of Experiment 2. The Group x Cue interaction was not significant ($F < 1$), suggesting no renewal effect. Simple effect contrasts revealed that participants in both groups exhibited an extinction effect. Specifically, cue A was judged to be less causal than cue B in the Extinction group ($F_{(1, 46)} = 10.28$, $p < .01$), and in the Renewal group ($F_{(1, 46)} = 9.49$, $p < .01$). There was no evidence for a renewal effect as participants in the Renewal group judged cue A to be just as causal as did those in the Extinction group ($F < 1$).

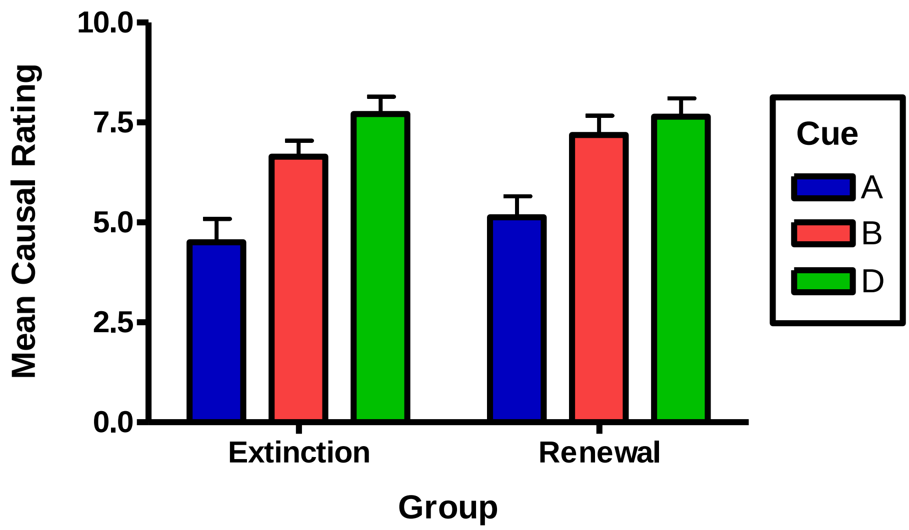

Figure 7. Mean causal judgement scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 2. Error bars denote standard errors of the mean.

Memory

Participants' cued recall scores for the allergic reactions paired with cues A, B, and D, recorded as 0, 0.5, or 1 for each, were averaged across participants. These data are presented in Figure 8. It appears that recall for the allergic reaction paired with cue B was slightly better than recall for the allergic reaction paired with cue A in both groups. As there were three possible scores for each cue, the assumption of normally distributed data was sufficiently satisfied and therefore the most appropriate analysis was a two-factor mixed ANOVA. Four planned contrasts were tested, and a Bonferroni adjustment controlled the FWER at $\alpha = 0.05$ (see Appendix F). The Group x Cue interaction was not significant ($F < 1$). Simple effect contrasts revealed that there was in fact no difference in the accuracy of recall for the reactions paired with cues A and B in the Extinction group ($F_{(1, 46)} = 1.71$), nor in the Renewal group ($F_{(1, 46)} = 1.71$). Between the Extinction group and the Renewal group, there was no difference in the accuracy of recall for the reactions paired with cue A ($F < 1$).

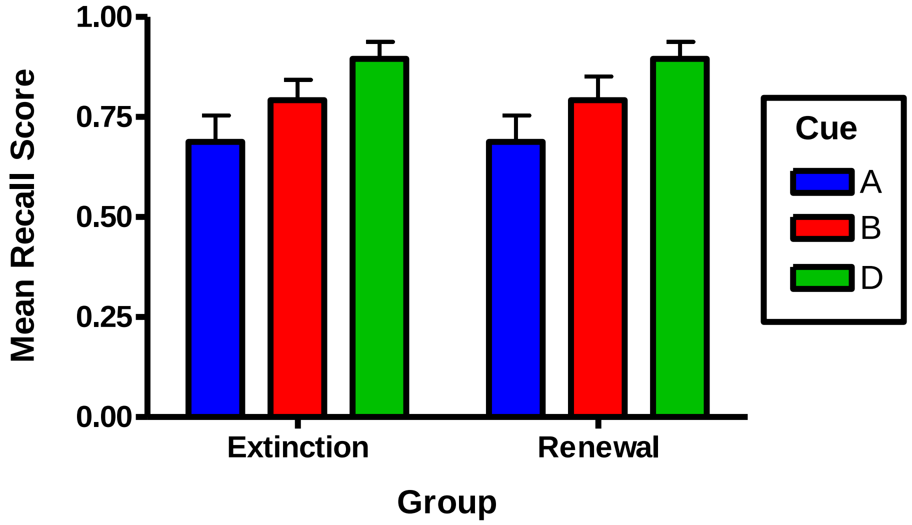

Figure 8. Mean cued recall scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 2. Error bars denote standard errors of the mean.

Discussion

The goal of the second experiment was to test for renewal under conditions that were less likely to produce a ceiling effect in recall. Informal comparisons of the memory data in Experiments 1 and 2 (Figures 4 and 8) suggest that this goal was achieved. For example, 98% of participants in Experiment 1, averaged across groups, were able to recall the outcome paired with cue A, compared to just 69% in Experiment 2. Moreover, this 69% did not reflect a ceiling effect since the outcome paired with cue D in Experiment 2 was recalled with 90%

accuracy, nor did it reflect chance responding (50%). Consistent with the inferential approach, there was no difference in recall for the outcomes paired with cues A and B.

In spite of the fact that recall for the outcomes paired with cues A and B were equivalent, both groups demonstrated an extinction effect in causal judgement. These results replicate those observed in the Extinction group participants of Experiment 1. Note that this dissociation between judgement and memory is consistent with the inferential approach, but not with the associative approach. In contrast to the findings of Experiment 1, however, there was no evidence for renewal following a context switch back into the original training context at test for participants in the Renewal group. The absence of a renewal effect is perplexing because it is contrary to findings of previous studies (e.g., Paredes-Olay & Rosas, 1999), and to the predictions of both the associative and inferential approaches. The most straight-forward explanation is that the context was not salient enough and therefore was not encoded during the learning trials. This explanation was tested in the third experiment.

Experiment 3

Experiment 2 succeeded in pulling the memory performance away from the ceiling effect observed in the Experiment 1, and revealed that there was no significant difference in cued recall accuracy between the groups or cues. Moreover, although Experiment 2 replicated the extinction effect observed in causal judgment in Experiment 1, there was no evidence for renewal. An explanation consistent with both the associative and inferential approaches is that the context was not salient enough.

For example, according to Bouton’s (1997) retrieval-failure model of learning, context is automatically encoded during the extinction trials when the contingency of cue A changes. At this point the participant begins to attend to the context as they search for something to help them resolve the ambiguity caused by the contingency change. This process of attending to the context, however, depends on the relative salience of the context with respect to the cues and outcomes. It is possible that significantly increasing in the number of cues in Experiment 2, and thus the number of cue-outcome relationships, dramatically decreased the relative salience of the context.

An alternative explanation emerges from the inferential account and centres upon the assumption that learning is cognitively demanding, effortful, and therefore limited. Based on this idea it is apparent that increasing the number of cues, and thus the number of cue-outcome relationships, makes the task difficult for the participant, who is likely to attempt to simplify the problem. Since context has no obvious relation to cue-outcome associations during the acquisition trials, it is possible that the participant notices the context change but simply discontinues paying attention because they do not believe the context is relevant.

One way to increase the salience of the context is to explicitly direct participants' attention to context by instruction. A similar manipulation has been used to achieve 'instructed extinction', in which participants are forewarned before extinction training that that contingency of the CS will change (Grings, Schell, & Carey, 1973). Compared to uninformed controls, forewarned participants show rapid and enhanced extinction. It is possible that an instructional manipulation may produce a renewal effect with a complex design, such as that used in Experiment 2. It is difficult to see how this manipulation, if successful, would be explained by an automatic, link-formation account, as proposed by the associative approach. In contrast, this finding would fit quite naturally within an effortful, propositional reasoning account of learning, as proposed by the inferential approach.

The goal of Experiment 3 was to replicate the Experiment 2 exactly, with the addition of explicit instructions directing the participants' attention to the context. Interestingly, this is the first human causal learning experiment to actually directly manipulate attention to the context by way of instruction, although this is quite a common phenomenon in the memory literature (e.g., directed forgetting, David & Brown, 2003). An observation of renewal in causal judgement in the present experiment would be strong evidence that higher-level, effortful reasoning processes are responsible for performance in this task.

Method

Participants

The sample consisted of 44 undergraduate first year University of New South Wales psychology students, 12 males and 32 females, participating in exchange for course credit. The average age was 19.7 years, with a range of 17 to 40 years.

Design, instrumentation, stimuli, and procedure

The design, instrumentation, stimuli, and procedure used in Experiment 3 were the same as Experiment 2. The instructions were also the same, except for directing participants to pay attention to the restaurant contexts. During the initial instructions screens participants were presented with the instructions, "*it is VERY important that you pay attention to what restaurant he* [Mr. X] *is eating at when he does suffer from an allergy*" (see Appendix A). During the test instruction screens participants were presented with the instructions, "*He* [Mr. X] *hopes that you paid attention to the restaurants where he was getting sick, and suggests that you use this knowledge when making your ratings*" (see Appendix C).

Results

Based on the same exclusion criteria used in Experiment 2, six participants were removed from the study and replaced.

Training

Figure 9 displays the percentage of participants who predicted an allergic reaction outcome, correct or not for cues A, B, C, and D across the two learning phases for both groups. Again, participants in both groups rapidly learned and remembered the experimental contingencies. This is not a surprise given that six participants were replaced in order to achieve this performance. Participants correctly predicted the outcome paired with cues A, B, and C with 81% accuracy by the second trial, and this improved to 98% by the end of the acquisition phase. Participants in both groups very quickly learned to predict "no reaction" when the cue A contingency changed during the extinction phase, although it appears, as in Experiment 2, that this occurred more slowly for those in the Extinction group. Post-hoc analyses revealed that on trials 6 and 7 significantly more participants in the Extinction group predicted an allergic reaction outcome when presented with cue A ($F_{(1, 42)} = 15.02$, $p < .001$ for trial 6 and $F_{(1, 42)} = 10.66$, $p < .01$ for trial 7). By trial 8 this difference had disappeared ($F < 1$). A cross-experimental interaction contrast with Experiment 2 comparing cue A predictions on trial 6 was also significant ($F_{(1, 88)} = 3.98$, $p < .05$). This suggests that participants in the Renewal group of Experiment 3 were faster than participants in the Renewal group of Experiment 2 at learning the new reinforcement schedule of cue A.

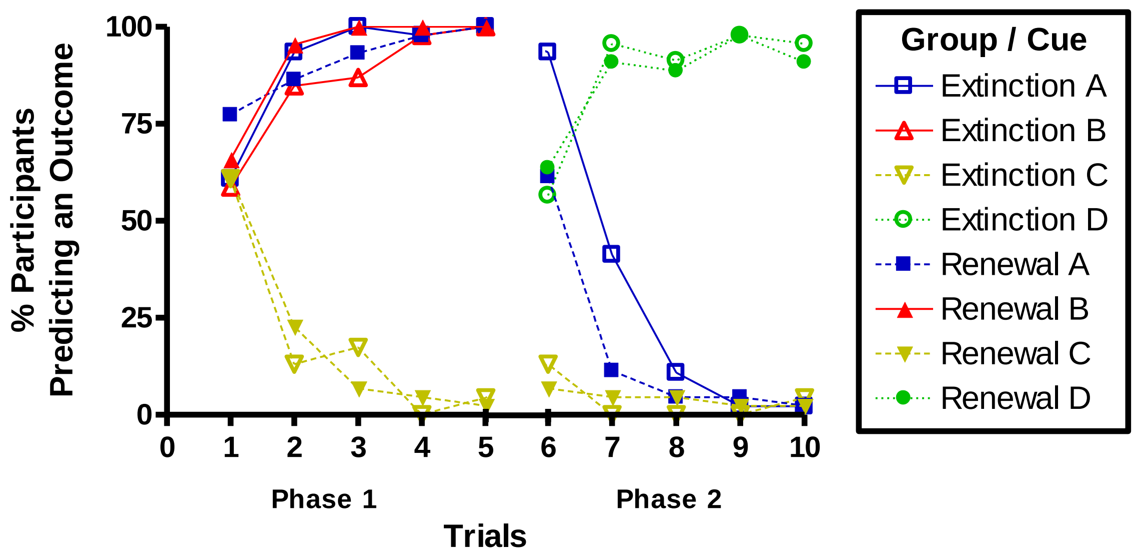

Figure 9. Percentage of participants in the Extinction (empty symbols) and Renewal (filled symbols) groups who predicted an allergic reaction on each trial in phases 1 and 2 of Experiment 3.

Judgement

Participants' scores on the judgement measure for cues A, B, and D, which ranged between 0 and 10 for each cue, were averaged across participants. These data are presented in Figure 10. It appears that there is a large difference in causal judgment between cues A and B in the Extinction group, but a much smaller difference between cues A and B in the Renewal Group. The same set of planned contrasts used in Experiment 1 was used to test the data of Experiment 3. The Group x Cue interaction was significant ($F_{(1, 42)} = 10.15$, $p < .01$), suggesting a renewal effect. Simple effect contrasts revealed that participants in the Extinction group exhibited extinction in that they judged cue B to be significantly more causal than cue A ($F_{(1, 42)} = 37.97$, $p < .001$), but no such difference was observed in the Renewal group ($F_{(1, 42)} = 1.60$). Participants in the Renewal group exhibited renewal in that they judged cue A to be significantly more causal than did participants in the Extinction group ($F_{(1, 42)} = 8.73$, $p < .01$).

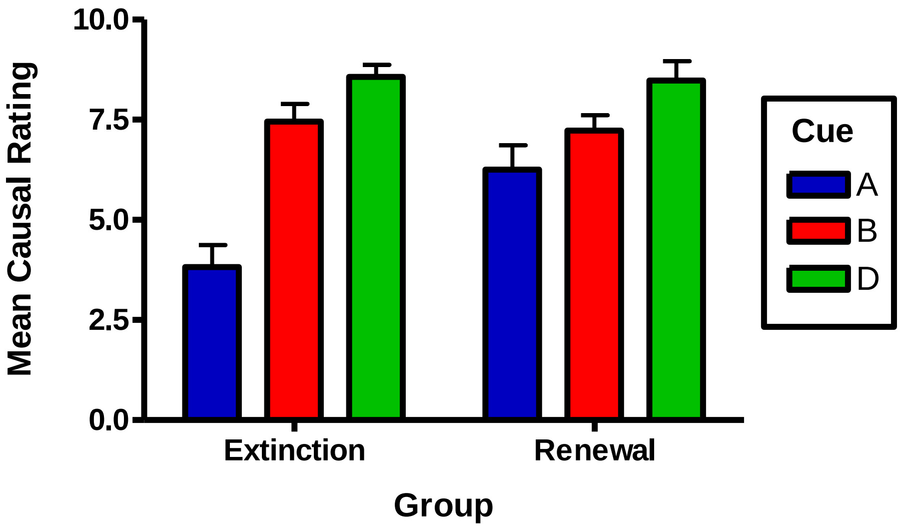

Figure 10. Mean causal judgement scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 3. Error bars denote standard errors of the mean.

Memory

Participants' cued recall scores for the allergic reactions paired with cues A, B, and D, recorded as 0, 0.5, or 1 for each, were averaged across participants. This data is presented in Figure 8. It appears that there is a small difference in recall for the allergic reaction paired with cues A and B in the Extinction group, but no such difference in the Renewal group. The same set of planned contrasts used in Experiment 2 was used to test the data of Experiment 3. The Group x Cue interaction contrast was not significant ($F < 1$). Simple effect contrasts revealed that that there was no difference in the accuracy of recall for the reactions paired with cues A and B in the Extinction group ($F_{(1, 42)} = 1.60$), nor in the Renewal group ($F < 1$). Between the Extinction group and the Renewal group, there was no difference in the recall for the reactions paired with cue A ($F < 1$).

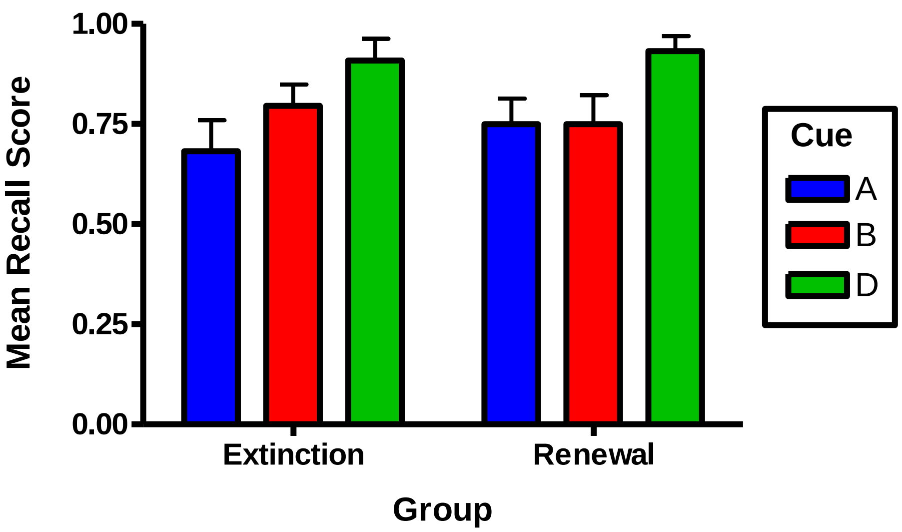

Figure 11. Mean cued recall scores for cues A, B, and D in the Extinction and Renewal groups, in Experiment 3. Error bars denote standard errors of the mean.

Cross-experimental analysis

A cross-experimental analysis investigating the three-way interaction between cues A and B, Extinction and Renewal groups, and Experiments 2 and 3 were conducted on the causal judgement measure. As the only difference between the two experiments was the instructional manipulation, any observed differences can be tentatively interpreted as due to the additional instructions. Specifically, the aim of the analysis was to test whether the renewal effect observed in the third experiment was the result of directing participants' attention to the context. The Experiment x Group x Cue three-way interaction was significant for the causal judgment measure ($F_{(1, 88)} = 4.11$, $p < .05$). That is, the Group x Cue interaction was stronger in Experiment 3. This is direct evidence to support the conclusion that the addition of explicit instructions in Experiment 3 led to the renewal effect observed in causal judgement in Experiment 3.

Discussion

The goal of Experiment 3 was to increase the salience of the context by instruction in order to obtain a renewal effect in a complex causal learning task. A renewal effect was observed in participants that experienced extinction in a different learning context to that of acquisition training and test (Renewal group). As with Experiments 1 and 2, an extinction effect was observed in causal judgement for participants who experienced acquisition, extinction, and test all in the same context. These results suggest that participants can indeed processes and encode multiple cue-outcome associations and context when they pay attention to the context during training, and can use this knowledge at test. There were no significant differences in the accuracy of recall for the outcomes paired with cues A and B between or within groups, replicating Experiments 1 and 2. This finding does not appear to be the result of a ceiling effect, as in Experiment 1, because accuracy of recall for the outcome paired with filler cue D (91.8%) was much higher than the average of cues A and B (74.3%).

The best explanation for the results derives from the inferential approach and argues that the cognitive processes required for learning are limited. Therefore, participants faced with many cue-outcome associations are likely to ignore all factors that they believe are irrelevant to solving the task. In Experiment 2, it is likely many participants believed the context to be unimportant with respect to their goal of predicting what cues caused what outcomes. The instructions provided in Experiment 3, however, indicated that the context could actually aid them in their goal and as a result participants chose to pay attention to, and learn about, the context as well as the to the cue-outcome contingencies.

A second line of support for this conclusion comes from the unexpected observation that participants in the Renewal group learnt that cue A was paired with no outcome much more quickly than participants in the in the Extinction group did. On the final trial of the acquisition phase all participants in both groups predicted an outcome. However, on the very next trial (the first trial of the extinction phase) 39% of participants in the Renewal group spontaneously changed their response to no outcome, compared to only 7% of participants in the Extinction group (see Figure 9). A similar observation was made in Experiment 2. This is surprising since feedback had not yet been given, and so the contingency change should have come as a complete surprise for both groups. Note, however, that there were actually two separate cues that comprised of the cue A condition (A_1 and A_2). Inspection of the learning data suggests that after being surprised by the change in reinforcement contingency of cue A_1 several participants hypothesised that the reinforcement contingency of cue A_2 might also change, and so predicted no outcome on the first trial of the extinction phase. This effect was stronger in the Experiment 3 than in Experiment 2. This may be due to participants actively looking to the context to mediate cue-outcome contingencies, as a result of the instructional manipulation. These data thus support the conclusion that participants were actively hypothesising about the nature of the task during the actual learning phases, a conclusion that is entirely inconsistent with the associative approach.

General Discussion

Summary of the major outcomes

The goal of the present study was to test predictions made by the associative and inferential approaches to learning, especially with regard to the correspondence between judgement and memory, and the impact of instructions. Both approaches made analogous predictions with respect to the causal judgement measure. Specifically, it was hypothesised that participants who experienced a cue paired with an outcome during acquisition training, which was subsequently paired with no outcome during extinction training, would rate that cue as less causal compared to a control cue that was not presented during the extinction phase. An extinction effect in causal judgement was indeed observed using both a simple design (Experiment 1) and a complex design (Experiments 2 and 3). This effect, however, was limited in each case to the Extinction group, who experienced acquisition training, extinction training, and test all in the same context. The Renewal group, in contrast, experienced acquisition training and test in one context, and extinction training in a different context. It was hypothesised that these participants would rate the extinguished cue as more causal than those who had not experienced a shift out of the extinction context. A renewal effect in causal judgement was indeed observed using a simple design (Experiment 1), but not in a complex design (Experiment 2), except when the instructions explicitly directed participants attention to the importance of context (Experiment 3).

It was on the second measure of outcome activation, cued recall, that the two approaches made alternative hypotheses. The associative approach predicted a strong correlation in the pattern of results between causal judgement and cued recall, since both are thought to reflect associative strength. The specific expectation was poor recall for the outcome paired with the extinction cue but good recall for outcome paired with the control cue in the Extinction group, and good recall for outcomes paired with both cues in the

Renewal group. In contrast, the inferential approach predicted a dissociation between causal judgement and cued recall, since judgements are thought to be made by reasoning processes privileged to good recall for all experienced associations. The specific expectation was equivalently good recall for the outcome paired with both cues in both groups. In support of the inferential approach, there were no differences in cued recall between the extinction and control cues for either group in Experiment 1. In fact, memory was at ceiling for all cues. In response, Experiments 2 and 3 deliberately multiplied memory load by increasing the number of cues. In spite of this manipulation there remained no difference in recall between the cues, or across the groups, in Experiments 2 and 3.

Explanations of the Data

Causal judgment

The results of Experiment 1 replicate those from animal behavioural research (see Bouton, 1993, for a review), human learning tasks in general (e.g., Baeyens et al., 2005; Greenspoon & Ranyard, 1957; Mineka et al., 1999; Neumann, 2006) and human contingency learning tasks in particular (Paredes-Olay & Rosas, 1999; Rosas & Callejas-Aguilera, 2006; Scully & Mitchell, in press). Contingency learning tasks, such as the one conducted by Paredes-Olay & Rosas, have been largely interpreted with Bouton’s (1997) retrieval-failure model. Shanks (2007), however, has argued that associative models are only applicable to human contingency learning when the task is highly complex, because simple tasks can be solved by higher-order reasoning processes (Dickinson, 2001). One confound with the majority of human contingency learning tasks, and certainly all of the studies conducted by Rosas and his colleagues (e.g., Paredes-Olay & Rosas, 1999; Rosas et al., 2001), is that they have used simple designs. For example, Paredes-Olay and Rosas used just two cues, two outcomes, and two contexts. It is likely that these experiments, as well as Experiment 1, are better interpreted as the result of participants using different reasoning strategies.

A strategy is an effortful, goal-directed process that participants use to enhance their memory performance and to solve tasks (Bjorklund & Douglas, 1997). Strategy use may be taken as an indicator of higher-level reasoning. Scrutiny of individual participant's cue-by-cue test predictions, supplemented by an informative, semi-structured debrief that involved asking the participants the reason behind their causal judgements, suggest that two different averaging strategies were used. An averaging strategy is one that that integrates cue-outcome information from across trials (Collins & Shanks, 2002). Context was not a relevant factor for those in the Extinction group, and as a result 17 out of the 21 of these participants gave a causal rating to the extinction cue of 5/10. In contrast, context was a relevant factor for those in the Renewal group, and at least half spontaneously used an averaging strategy that took context into account and correspondingly gave a causal rating to the extinction cue of 10/10. The context-sensitive averaging strategy may have implemented itself by augmenting the premises used in the reasoning process the participants used to solve the task. For example, instead of hypothesising that "*garlic goes with headache*" during acquisition training, the participant may have hypothesised that "*in the Shady Bar, garlic goes with headache*".

This averaging strategy is consistent with statistical models of reasoning, such as Cheng's Power PC model (Cheng, 1997). The PC model is a normative theory based on the ΔP model, which states that during learning the participant forms a mental matrix that records the frequency of cue and outcome presentations (Allan, 1980). These frequencies are then used to compute ΔP, which is equal to the probability of the outcome given the cue ($P[O|C]$) minus the probability of the outcome in absence of the cue ($P[O|\sim C]$).

In order to minimise participant's ability to reason, and to further investigate any underlying associative learning mechanism, Experiment 2 used a much more complex design.

Curiously, there are no examples of complex renewal learning tasks to which Experiment 2 might be compared. There are, however, a number of cue competition studies that have varied the complexity of their designs. For instance, Mitchell et al. (2005) used a relatively simple design in which compounds AB+ and CD+ were intermixed with A+ and C- training. In total there were just eight cues and four outcomes. They observed a double dissociation in that recall was better for cue B than for cue D, even though cue B was judged to be less causal than cue D. In a subsequent study Mitchell et al. (2006) attempted to reduce the likelihood of participants using reasoning by increasing the complexity of the task. In total participants saw 25 cues and five outcomes. In contrast to the simple experiment, there was a strong correlation between recall and judgement: in each case cue B was observed to be less than cue D. From these two studies it is clear that increasing task complexity may be one useful approach to investigating low-level, associative mechanisms in learning.

In Experiment 2, an extinction effect was observed for participants in the Extinction group, replicating the results obtained by Scully and Mitchell (in press). Contrary to expectations, and inconsistent with previous findings in more simple contingency learning tasks (e.g., Paredes-Olay & Rosas, 1999), renewal was not observed when extinction training took place in a different context to acquisition training and test. It was suggested that this absence of renewal was due to the context not being salient enough and therefore was not encoded. This may have been the result of decreased relative salience of the context as the number of cue-outcome associations increased, or simply the result of the design complexity and apparent incidental presentation of the context causing participants to ignore the context.

In order to increase the context salience in a way that would differentiate between the associative and inferential approaches, a third experiment was run. Experiment 3 replicated Experiment 2, with the addition of instructions directing participants to pay attention to the context and to use it to help them make their judgements at test. The additional instructions were sufficient to restore the renewal effect observed in Experiment 1, demonstrating that, when prompted, participants can indeed use the context to decide the causal status of a cue, even within a relatively complex design. The question remains, by what mechanism did the instructions cause renewal to occur?

One explanation is that the attentiveness to the context, due to the instructions, increased the relative salience of the context, such that it was able to be encoded by a low-level link-formation account of learning. It is difficult to understand however, from a strict associative standpoint, how mere instructions influenced this automatic process. Less strict, more hybrid accounts, do, however, provide a better account, as will be discussed below.

Alternatively, a more likely explanation is that the instructions caused participants to adopt a reasoning strategy similar or identical to the one used in Experiment 1, that is, an averaging strategy that incorporated the role of context. Specifically, the instructions provided the participants with the expectancy that the context could help them to predict what outcome was paired with each cue. Following from this premise, participants may have paid attention to the context and formed hypotheses about the nature of the cue-outcome associations that were context-sensitive. Participants in Experiment 2, in contrast, did not have this expectation to begin with and may have discarded any strategy that was context-sensitive in preference for a cognitively less-demanding strategy, such as one that ignored the context. Based on this interpretation of the data, the pattern of results obtained in Experiments 2 and 3 are most consistent with the proposal that the judgement task was solved by an effortful, goal-directed reasoning process.

Cued recall

Cued recall was relatively good and equivalent across both extinction and control cues, regardless of design complexity. These results are in contrast to those obtained by

Scully and Mitchell (in press), since they observed recall for the outcome paired with the extinction cue to be lower than recall for the outcome paired with the control cue. One explanation for this difference is that the Scully and Mitchell experiment had more statistical power to detect these differences, since they used a completely within subjects design and tested 60 participants. The current experiment used a mixed subjects design and the largest number of participants in any one cell was 24. In support of this explanation Scully and Mitchell found ~70% correct recall for the outcome paired with the extinction cue and ~81% correct recall for the outcome paired with the control cue (Experiment 1). These percentages are comparable to the present Experiment 3: ~68% correct recall for the outcome paired with the extinction cue, and ~80% correct recall for the outcome paired with the control cue.

An alternative explanation can be made in light of the sensitivity of the memory measure. Scully and Mitchell (in press) used four outcomes, whereas the current study used just three. It is possible that there were too few outcomes used in the present study, and, as a consequence, the memory test was too easy. Indeed, there is some evidence to suggest that three, and sometimes even four, outcomes may not be sufficiently sensitive to observe differences in cued recall (Mitchell & Scully, in preparation).

A third explanation for the good memory observed in the current experiment is that the context served as an additional retrieval cue that aided memory. This would be expected in light of the encoding specificity principle, which states that recall is better when information available at the time of encoding is also available at the time of retrieval (Tulving & Thomson, 1973). Indeed, the most obvious difference between the current experiments and those of Mitchell and Scully (in preparation; in press) are the context changes. In the Mitchell and Scully papers, the only cue that participants were given to help them retrieve the outcome was the food cue. In contrast, all participants in the current experiment received cue-outcome trials and the cued recall tests in the same context. For those in the Extinction group, however, this same context was also where the cue was paired with no reaction, but since "no reaction" was not an option at test, the context may still have helped with the retrieval of the outcome. Context may therefore have provided an additional retrieval cue to help participants recall the outcome that was paired with the cue. This may partially account for the across-the-board good memory performance in the current study. Note that this explanation suggests that memory should have been superior in the Renewal group; the lack of overall sensitivity in the cued recall measure may account for the absence of this difference.

Dissociation between cued recall and causal judgment

In the current study a single dissociation was observed in that extinction cue A was rated low on causal judgment compared to the control cue B, but the outcomes paired with each were recalled with equivalently high levels of accuracy. This pattern of results is in contrast to the parallel pattern of results observed by Scully and Mitchell (in press). As explained above, this dissociation may have been an artefact of the memory measure being too easy due to additional retrieval cues, or too insensitive due to too few outcomes. In fact, in Experiments 2 and 3 the accuracy of cued recall was to some extent paralleled in the results of the causal judgement measure, as predicted by the associative approach. Thus, when recall for the outcome paired with the cue A was lower than recall for the outcome paired with the cue B, cue A was judged to be significantly less causal than cue B. Although the parallel was not significant in any case, it was always in the right direction.

It might also be informative to compare the present single dissociation with the double dissociation observed by Mitchell et al. (2005). The striking difference between this study, those conducted by Mitchell and Scully (in preparation; in press), and the current study, is that Mitchell et al. used cue compounds during training. In fact, these compounds were chosen explicitly so that the food cues were already highly associated (e.g., fish and chips) to produce an ‘augmentation effect’. Augmentation occurs when A+ trials increase, rather than reduce, the associative strength accrued to cue B on AB+ trails, by way of the strong within-compound link between cues A and B (Batsell, Paschall, Gleason, & Batson, 2001). This augmentation, on test, allowed cue B to be more easily recalled due to its link with the outcome and to cue A, the latter of which further activated the outcome node as a result of the link between cue A and the outcome (established on the A+ trials).

The current experiments did not include a manipulation designed to increase cued recall for the extinction cue, such as within-cue augmentation. As a result, it is extremely unlikely that a double dissociation could have emerged, and so the “best-case” scenario was a single dissociation like the one observed. Assuming the two dependent measures were in fact sufficiently and similarly sensitive, it may be the case that there was indeed a genuine dissociation between cued recall and causal judgement. It must be noted, however, that compared to a double dissociation or a parallel result in cued recall and causal judgement, a single dissociation is the weakest form of evidence. In light of this, little can be firmly concluded with regard to the single dissociation observed in the current study.

Differential learning curves during extinction training

In all three experiments, participants in the Renewal group learnt the changed contingency of cue A faster than participants in the Extinction group. In Experiment 1 this difference was observed on the second trial of the extinction phase. In Experiments 2 and 3 this was observed on the first and second trials of the extinction phase, and was stronger in Experiment 3. A plausible explanation for these findings is that participants were actively hypothesising about the contingencies between the cues and outcomes during the learning trials. For example, participants in the Renewal group may have noticed that the contingency of cue A changed when it was presented in a different context. This would have allowed these participants to form the hypothesis that some cue-outcome contingencies change when the context changes and, accordingly, adjust their future predictions to presentations of cue A. This hypothesis, however, could not be formulated based on the experience of a participant in the Extinction group. This explanation presents a plausible mechanism as to why participants in the Extinction group appeared to be slower at learning the new contingency of cue A.

In Experiments 2 and 3 the difference between groups was observed on the very first trial of the extinction phase. Presumably, this was possible because in these two experiments each cue was represented by two food cues (A_1 and A_2), which allowed participants to adjust their hypothesis and responses on the first trial of the extinction phase. For example, participants in the Renewal group may have noticed that the contingency of cue A_1 changed when it was presented in a different context. This may have allowed them to update their hypotheses, and therefore adjust future responses to presentations of both A_1 and A_2. In fact, the instructions in Experiment 3 actually provided participants with the hypothesis that context was important and, consequently, learning of the new contingency of cue A for those in the Renewal group proceeded even faster than had occurred in Experiment 2. These data support the conclusion that participants were actively hypothesising about the nature of the task during the actual learning phases, a finding that is only consistent with an intentional, controlled, goal-directed reasoning process.

Theoretical Considerations

The Associative Approach

As predicted by Bouton’s (1997) retrieval-failure model, in all three experiments the extinction cue was judged with a lower causal rating than the control cue. Contrary to expectations, renewal did not occur in Experiment 2 when the number of cue-outcome associations was increased. This is a problem for the associative approach since context is thought to be automatically encoded when the meaning of a cue become ambiguous, as at the beginning of extinction training (Bouton, 1993). An associative-based explanation for this absence of renewal is that by increasing the complexity of the design, the relative salience of the context decreased, and was therefore not encoded. The results of Experiment 3, however, cast doubt over this explanation. In the third experiment a renewal effect was observed, using the exact same context manipulation as in Experiment 2, merely as a result of changing the instructions. From a strict associative point of view, instructions should have no effect on the associative strength of a cue-outcome link, since all that matters in determining the associative strength are the contingency, contiguity, trial number and trial order of the cue-outcome pairings. Therefore, according to this approach there should have been no difference between Experiments 2 and 3, since the additional instructions had no impact upon these factors. Similarly, there should have been no difference in the results between Experiment 1 and 2, since manipulating the total number of cues also had no impact upon these factors.

Contrary to the associative prediction, there were no differences in the ability to recall the outcomes paired with the extinction and control cues in any of the experiments. This failure to find a difference poses a problem for the associative approach, since it predicts that after extinction training the associative link between the extinguished cue and the outcome paired with it will be weakened. Since cued recall may be taken to be a measure of associative strength, it was expected that recall for the outcome paired with the extinguished cue would be significantly lower than recall for the outcome paired with a cue never extinguished. One associative explanation is that both the extinction and control cues reached a ‘threshold level’ of activation that permitted accuracy on the cued recall measure, but that the cues passed this threshold by different amounts. It is possible that a more sensitive measure of memory may have revealed this critical difference.

Finally, instead of the parallel predicted between the two dependent measures, a single dissociation was observed. As noted by previous investigators (e.g., Mitchell et al., 2005) a single dissociation is relatively weak evidence compared to a double dissociation, since it may be the case of differential sensitivity in the measures. Differential sensitivity cannot be ruled out in the current study, given that the cued recall measure appears to have been insensitive. Moreover, inspection of the descriptive statistics in Experiments 2 and 3 show a parallel between the two dependent measures that did not reach significance. It should be noted, however, that cued recall in every case was well above chance, which suggests that most participants possessed adequate memory for which outcomes were paired with which the cues when they made their judgements.

In summary, few of the critical predictions made from the associative approach were supported in this study. Two pieces of evidence are particularly damaging: first, the apparent causal role played by the instructions to obtain a renewal effect in Experiment 3; second, the observation that participants were actively hypothesising about the causal structure of the task during the learning trials in all three experiments. It seems safe to conclude that participants in the current study did not base their causal judgements solely on the level of activation of the outcome memory node. In the case of both the simple and complex designs there appeared to be an underlying and effortful reasoning strategy that guided learning and was impacted upon by task complexity and instruction.

The Inferential Approach

Like the associative approach, the inferential approach correctly predicted that in all three experiments the extinction cue would be given a low causal rating compared to the control cue. Contrary to expectations, renewal did not occur in Experiment 2 when the design complexity was increased. The absence of renewal was likely due to the increased cognitive load associated with the additional cue-outcome associations that caused participants to simplify the task by ignoring the context. This is consistent with the inferential account that learning is effortful, takes up cognitive resources, and that causal judgements are the end product of a higher-level reasoning process based on good recall for all cue-outcome associations. Support for this conclusion was obtained by observing a renewal effect when instructions were given directing participants attention to the context in Experiment 3. It is likely that the instructions altered the reasoning strategy that the participants used to solve the task, by means of the hypotheses they made about how the cues and outcomes were being paired (De Houwer, Beckers et al., 2005). This result also provides evidence against the claim that high-order reasoning only occurs in relatively simple tasks (Dickinson, 2001).

As predicted by the inferential approach, the outcomes paired with both the extinction and control cues were recalled with good accuracy, and without any difference between cues or groups. The single dissociation between judgement and memory is also consistent with the inferential approach. According to this approach, contiguous pairings of cue and outcome allow the subject to recall that the two stimuli ‘go together’, but the judgement that the cue causes the outcome involves an extra inferential step (Mitchell, Livesey, & Lovibond, 2007).

The strongest piece of evidence in support of the inferential approach was the observation that participants were actively hypothesising during the learning phase. As a result of this hypothesising, participants in the Renewal group were faster at learning the new contingency of the extinction cue during extinction training, because it was logical to form the hypothesis that the contingency changed with the context. Strong evidence for this interpretation was obtained in Experiment 3, which actually provided participants with this hypothesis, and consequently resulted in even faster learning for those in the Renewal group.

In summary, all of the observations made in the current study are consistent with the inferential approach to learning. Causal judgments need not reflect the strength of the association between cue and outcome, but are best understood as the cooperative effort of a low-level memory system and a higher-level reasoning system.

Hybrid model

It is becoming increasingly evident that traditional associative models of learning cannot, by themselves, account for the data obtained here and elsewhere (see De Houwer, Vandorpe, & Beckers, 2005, for a review). Similarly, there are a number of examples that do not permit an inferential interpretation (e.g., Karazinov & Boakes, 2007; Le Pelley, Oakeshott, & McLaren, 2005; Mitchell et al., 2006). Moreover, the inferential approach suffers from an explanatory gap as to how low-level processes encode associations that are then made available to the higher-level inferential processes in a single, cooperative system. Several authors have suggested that a hybrid model incorporating two coordinated, synergistic mechanisms may provide a more thorough account for the data (De Houwer & Beckers, 2002; De Houwer, Vandorpe et al., 2005; Lovibond, 2003; Pineno & Miller, 2007).

Indirect evidence for this hybrid model is provided by the pattern of studies that lend support to either associative or reasoning processes. Typically, studies that provide the most evidence against the associative approach have used simple experimental design (e.g., Larkin et al., 1998), such as in Experiment 1. In contrast, studies that provide the most evidence against the inferential approach have used complex designs (e.g., De Houwer, 2002), such as in Experiment 3. It may perhaps be most parsimonious to conclude that subjects prefer to use a higher-level, reasoning system based on memory stored in an associative network, but that this higher-level system is compromised when taxed, causing the bottom-up, automatic, associative system to assume ascendancy. Note that such a cooperative hybrid model is different from the ‘levels of learning’ approach endorsed by most associativists, which argues that associative and inferential systems are distinct, separate, and competitive.

In spite of the face validity of this hypothetical cooperative hybrid model, the task of properly characterizing and testing it will likely be difficult. A hybrid model would have to specify and account for the numerous interactions between the associative and inferential systems, and clearly define how and when each system operated within a single framework. A fully-specified hybrid model will likely incorporate low-level associative processes, higher-level cognitive processes, and have important roles for reasoning and attention.

Attentional theory of context processing

Recently Rosas and his colleges have tentatively postulated the outline of a possible hybrid model, which attempts to integrate low-level associative processes and attention. It is based on their observation that a context switch experienced in a contingency learning task causes everything learned subsequently to become context-specific (Rosas & Callejas-Aguilera, 2006). In fact, a similar observation was made in the judgement data for cue D in Experiment 1. Although cue D maintained the same contingency throughout the experiment, participants trained in context Z and then tested in context Y judged cue D to be less causal than those who were trained and tested in context Y. Because cue D was only presented during the extinction phase, it appears that the 'cue D-outcome' association became context specific at the point that participants began to pay attention to the context, that is, when the causal status of cue A became ambiguous. The implication is that cue-outcome associations may become context dependent due to other factors than causal status ambiguity. This finding is inconsistent with strict associative models, such as Bouton's (1997) retrieval-failure model, which argues that that learning becomes context dependent only when second-learnt information conflicts with first-learned information. Such observations have led Rosas and his colleges to propose an extension of Bouton's model, which they call the 'Attentional Theory of Context Processing' (Rosas, Callejas-Aguilera, Álvarez, & Abad, 2006).

According to attentional theory, cue-outcome associations become context dependent whenever the context is attended to during learning, regardless of the nature or order of that learning. Attention to the context is thought to be determined by five factors: a) experience with the contexts, b) instructions in human participants, c) the informative value of the context, d) the presence of ambiguous information, and e) the relative salience of the context with respect to the cues (Rosas et al., 2006). The second and fifth factors allow attention to the context to be modified by way of instruction, as, for example, in Experiment 3, when the instructions manipulated participants' perception of the salience of the context. For example, attentional theory would argue that the instructions used in Experiment 3 caused the participants to attend to the context throughout the experiment, which resulted in the context being automatically encoded, via associative links, in all of the cue-outcome associations that formed during the learning trials. Subsequently, the context then mediated retrieval of all cue-outcome associations at test, resulting in the observations of extinction and renewal.

Rosas et al. (2006) conceptualise their theory as an extension of retrieval theory, and therefore assume that learning is inherently associative in nature. However, the theory is silent with respect the nature of the mechanisms that regulate attention to the contexts, the mechanisms that regulate the role of context retrieval, nor how propositional instructions causally interact with attention and learning. The current study suggests that these explanatory gaps may be filled by a cooperative, higher-order, propositionally-based reasoning system. In any event, the still-developing attentional theory provides a good account for the causal judgement data, and demonstrates the potential of hybrid models.

Limitations

The main limitation of the current study was the questionable validity of the cued recall measure, which may have been insensitive to differences in memory. A second problem with the cued recall measure is that it might actually reflect the ability of an episodic memory system. If this were the case, then the dissociation observed in the current study may be the result of two dependent variables measuring two different things, rather than two dependent variables measuring associative strength.

The test order of the two dependent measures used in the current study was not counter-balanced. This may be a source of confound if the impact of conducting the cued recall test prior to the causal judgement test was larger or smaller for one particular condition or group. For example, participants faced with cued recall for the extinction cue may have been surprised to see “no reaction” absent as a test response, causing them to pause and think back to the beginning of the experiment when the extinction cue was paired with an outcome. In contrast, no such surprise would have occurred when faced with cued recall for the control cue. As a result causal judgement scores for the extinction cue may be lower than they would have been had causal judgement been measured first.

An inherent confound within the YZY renewal paradigm is that the context change is correlated with the reinforcement schedule change. As a result, it is unclear as to whether renewal was due to a switch out of the extinction context (Bouton, 1993, 1994), or as a result of the acquisition and test contexts being the same (Tulving & Thomson, 1973).

Lastly, despite claiming that the instructions in Experiment 3 changed participants level of attention towards the context, no direct evaluation of attention was conducted.

Future Studies

Future experiments should ensure that the cued recall test is more difficult, and therefore more sensitive, by increasing the number of outcomes to at least four. Some participants should be given the option of “no reaction” at test and compared to those not given the option, thereby investigating the possible confound associated with its surprising absence on test. The provision of the “no reaction” option will also allow Bouton’s (1997) retrieval-failure model to be tested in a second way. Just as the participants demonstrating extinction in causal judgement are expected to fail to retrieve the allergic reaction outcome paired with the extinction cue, participants demonstrating renewal in causal judgement would be expected to fail to retrieve the “no reaction” outcome paired with the extinction cue.

Of course, it has also been argued that the cued recall task itself is inherently flawed because it may not in fact reflect associative strength, but merely episodic memory ability. Future studies should investigate this claim by adding a third, non-episodic-dependent measure of associative strength. One likely candidate is the categorization task used by Mitchell et al. (2007). The logic of this measure is that the formation of an associative bond by pairing together a cue and outcome should allow these stimuli to be categorised together more easily at test, and without the necessity of remembering the specific episode in which they were paired. If this is the case, then a parallel between this measure, cued recall, and causal judgment would be strong evidence for the associative approach.

Future human contingency learning studies investigating renewal should attempt to discover if the effect is due to a switch out of the extinction context, the lack of interference from learning information in different contexts, or the ability of the context to serve as a retrieval cue on test. Important for these investigations will be comparing groups that experience YZY, YZX, and YYZ renewal experiences. A growing body of evidence suggests that renewal is more difficult to achieve in the case of YYZ but no associative account can satisfactorily explain why (Rosas et al., 2006). It may be possible to explain these findings in terms of a higher-order inferential account. For example, in the case of YYZ, the participant may hypothesise that in a new context it is statistically more probable that the contingency they learned first is more prevalent, and therefore likely to be encountered in new contexts. If this propositionally stored hypothesis does in fact exist, then it may be possible to verbally provide additional premises that will change the outcome, such as, "*Context Z is the Bizarro World, where everything is opposite*". An observation of weak, or no renewal, would lend support to the inferential account of extinction and renewal.

Manipulations designed to encourage a double dissociation will also be important to the current debate. Although augmentation is exclusively a cue compound phenomenon, it does suggest avenues for potentially increasing the accuracy of cued recall. One option for the renewal paradigm is to pair the extinction cue with an outcome that it is already linked with. Thus, for example, the food cue "cheese" might be paired with outcomes associated with lactose intolerance, such as "diarrhoea". In contrast, the control cue (e.g., "broccoli") would be paired with something unrelated to it (e.g., "heartburn"). The pre-existing association between cheese and diarrhoea may make cued recall easier than in the case of broccoli and heartburn; however it is possible that participants will, after extinction training, still rate cheese to be less causal than heartburn. Such a double dissociation would suggest that judgements are not made solely on the level of activation of the outcome by the cue.

More generally, perhaps the most exciting theoretical direction for future research is in the specification of a complete hybrid model. Because it is assumed that inferential processing only occurs when the subject possesses the motivation and opportunity to engage in rational thinking, one prospect for separating out the two systems will be to manipulate these factors. Opportunity may be manipulated by introducing a time limit at test, such that judgments must be made within a very small time frame. Another manipulation might control the number of cue-outcome relations, like in the current experiment, but in a more systematic fashion. Motivation might be manipulated by potential rewards the participant might obtain with accurate performance. These manipulations will presumably cancel out the opportunity for judgements to be made by the inferential system and prove enlightening with respect to the factors required to fail the higher-level system, and the ability of the lower-level system.

The current experiment also suggests applied future endeavours in the clinical setting. Perhaps, the most important finding in the current study in this respect was the absence of the renewal observed in Experiment 2. This is important because in many cases, such as with the treatment of phobias, the goal of clinical treatment is to prevent renewal (Bouton, 2002). The absence of renewal was due to the participant ignoring the context. This effect was reversed when the participants were directed to attend to the context. This suggests that treatment designed to reduce the likelihood of renewal may be more effective by eliminating or reducing the possibility of context being encoded with the therapeutic learning. A range of practices may be used to discourage attention to the context. For example, cognitive-behavioural therapy may be conducted in the client's home, where, presumably, the context is habituated to and, therefore, muted in salience.

Implications

The theoretical implication of the present study is that purely associative theories cannot account for human causal learning. From the observations made it appears that causal judgements do not directly reflect the strength of the cue-outcome association, but are influenced by the attention paid to the context, the reasoning strategy used, instructions provided, and task complexity. These factors suggest that associative learning is not an automatic process, but is crucially dependent on a conscious, effortful, hypothesis-testing, higher-order reasoning system. The repercussions of this conclusion are that pure associative models that have traditionally been used to explain human contingency learning (Bouton, 1997; Rescorla & Wagner, 1972) must be forgone or remodelled to incorporate higher-order cognition, a process that has already begun (Mitchell et al., submitted; Rosas et al., 2006). A second implication is that findings of classical conditioning in animal behavioural research may not always be compatible with human causal learning (De Houwer & Beckers, 2002), despite beliefs to the contrary (Dickinson, 2001; Dickinson et al., 1984).

The practical implication of the present study is that clinical therapeutic techniques based on the assumption that reinforcement and extinction work on an unconscious, nonverbal, automatic system are misguided. The findings of the current study support the view that many disorders are likely to be due to false conscious beliefs and impaired reasoning. This is supported by observations that anxious patients overestimate the probability and cost of future harm, and that patients with anorexia falsely perceive their bodies to be overweight (Clark, 2004). The present study supports the use of cognitive-behavioural therapy as an effective form psychotherapy, and suggests that this methodology focus on integrating experience and language, since learning and reasoning are based on propositionally stored beliefs and hypotheses (Lovibond, 1993; Mitchell et al., submitted).

Conclusions

Associative learning is inseparably related to contingency learning, habit formation, motivation, addiction, the formation of phobias and a range of other psychopathologies (Pearce & Bouton, 2001). This study has provided evidence to suggest that human associative learning can, and in many cases is, accomplished by a higher-order inferential reasoning process that is sensitive to context and is influenced by task complexity and instruction. It is suggested that the current findings guide future theoretical research to form a hybrid model integrating low-level associative networks with higher-level reasoning and attentional processes. The potential application of these findings relate primarily to treatments in the clinic for a range of psychopathologies including anxiety disorders and addiction.

References

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgement tasks. *Bulletin of the Psychonomic Society, 15,* 147-149.

Baeyens, F., Vansteenwegen, D., Beckers, T., Hermans, D., Kerkhof, I., & De Ceulaer, A. (2005). Extinction and renewal of Pavlovian modulation in human sequential feature positive discrimination learning. *Learning & Memory, 12*, 178-192.

Batsell, W. R., Paschall, G. Y., Gleason, D. I., & Batson, J. D. (2001). Taste predonditioning augments odor-aversion learning. *Journal of Experimental Psychology: Animal Behavior Processes, 27*, 30-47.

Bjorklund, D. F., & Douglas, R. N. (1997). The development of memory strategies. In N. Cowan (Ed.), *The Development of Memory in Childhood* (pp. 201–246). Hove East Sussex, UK: Psychology Press.

Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin, 114*, 80-99.

Bouton, M. E. (1994). Context, ambiguity, and classical conditioning. *Current Directions in Psychological Science, 3*, 49-53.

Bouton, M. E. (1997). Signals for whether versus when an event will occur. In M. E. Bouton & M. S. Fanselow (Eds.), *Learning, Motivation and Cognition: The Functional Behaviourism of Robert C. Bolles* (pp. 385-409). Washington, DC: American Psychological Association.

Bouton, M. E. (2002). Context, ambiguity, and unlearning: Sources of relapse after behavioral extinction. *Biological Psychiatry, 52*, 976-986.

Bouton, M. E. (2004). Context and behavior processes in extinction. *Learning & Memory, 11,* 485-494.

Bouton, M. E., & Bolles, R. C. (1979a). Contextual control of the extinction of conditioned fear. *Learning and Motivation, 10*, 445-466.

Bouton, M. E., & Bolles, R. C. (1979b). Role of conditioned contextual stimuli in reinstatement of extinguished fear. *Journal of Experimental Psychology: Animal Behavior Processes, 5*, 368-378.

Bouton, M. E., & Brooks, D. C. (1993). Time and context effects on performance in a Pavlovian discrimination reversal. *Journal of Experimental Psychology: Animal Behavior Processes, 19*, 165-179.

Bouton, M. E., & Ricker, S. T. (1994). Renewal of extinguished responding in a second context. *Animal Learning & Behavior, 22*, 317-324.

Chambless, D. L., & Ollendick, T. H. (2001). Empirically supported psychological interventions: Controversies and evidence. *Annual Review of Psychology, 52*, 685-716.

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

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Clark, D. M. (2004). Developing new treatments: On the interplay between theories, experimental science and clinical innovation. *Behaviour Research and Therapy, 42*, 1089-1104.

Collins, D., J., & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgement. *Memory & Cognition, 30*, 1138-1147.

Corcoran, K. A., & Maren, S. (2004). Factors regulating the effects of hippocampal inactivation on renewal of conditional fear after extinction. *Learning & Memory, 11*, 598-603.

Craske, M. G., & Rowe, M. (1997). A comparison of behavioural and cognitive treatments of phobias. In G. Davey (Ed.), *Phobias: A Handbook of Theory, Research, and Treatment* (pp. 247-280). Chichester: Wiley.

David, D., & Brown, R. J. (2003). The impact of different directed forgetting instructions on implicit and explicit memory: New evidence from a modified process dissociation procedure. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *56A*, 211-231.

De Houwer, J. (2002). Forward blocking depends on retrospective inferences about the presence of the blocked cue during the elemental phase. *Memory & Cognition, 30*, 24-33.

De Houwer, J., & Beckers, T. (2002). A review of recent developments in research and theories on human contingency learning. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, *55B*, 289-310.

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*, *56B*, 345-357.

De Houwer, J., Beckers, T., & Vandorpe, S. (2005). Evidence for the role of higher order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior, 33*, 239-249.

De Houwer, J., Vandorpe, S., & Beckers, T. (2005). On the role of controlled cognitive processes in human associative learning. In A. Wills (Ed.), *New Directions in Human Associative Learning* (pp. 41-63). Mahwah, NJ: Lawrence Erlbaum.

Dickinson, A. (2001). The 28th Bartlett memorial lecture. Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology, 54B*, 3-25.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation. *Quarterly Journal of Experimental Psychology, 49B*, 60-80.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *36*, 29-50.

Garcia-Gutierrez, A., & Rosas, J. M. (2003). Context change as the mechanism of reinstatement in causal learning. *Journal of Experimental Psychology: Animal Behavior Processes, 29*, 292-310.

Greenspoon, J., & Ranyard, R. (1957). Stimulus conditions and retroactive inhibition. *Journal of Experimental Psychology*, *53*, 55-59.

Grings, W. W., Schell, A. M., & Carey, C. A. (1973). Verbal control of an autonomic response in a cue reversal situation. *Journal of Experimental Psychology, 99*, 215-221.

Havermans, R. C., Keuker, J., Lataster, T., & Jansen, A. (2005). Contextual control of extinguished conditioned performance in humans. *Learning and Motivation*, *36*, 1-19.

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

Karazinov, D. M., & Boakes, R. A. (2007). Secondorder conditioning in human predictive judgements when there is little time to think. *Quarterly Journal of Experimental Psychology*, *60*, 448-460.

Larkin, M. J. W., Aitken, M. R. F., & Dickinson, A. (1998). Retrospective revaluation of causal judgements under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 1331-1352.

Le Pelley, M. E., Oakeshott, S. M., & McLaren, I. P. L. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 56-70.

Lovibond, P. F. (1993). Conditioning and cognitive-behaviour therapy. *Behaviour Change, 10*, 119-130.

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,* 97-106.

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

Lovibond, P. F., Davis, N. R., & O'Flaherty, A. S. (2000). Protection from extinction in human fear conditioning. *Behaviour Research and Therapy, 38,* 967-983.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, *28,* 3-26.

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

Mineka, S., Mystkowski, J. L., Hladek, D., & Rodriguez, B. I. (1999). The effects of changing contexts on return of fear following exposure therapy for spider fear. *Journal of Consulting and Clinical Psychology, 67,* 599-604.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (submitted). The propositional nature of human associative learning. *Behavioral and Brain Sciences*.

Mitchell, C. J., Livesey, E., & Lovibond, P. F. (2007). A dissociation between causal judgement and the ease with which a cause is categorized with its effect. *Quarterly Journal of Experimental Psychology, 60*, 400-417.

Mitchell, C. J., Lovibond, P. F., & Gan, C. Y. (2005). A dissociation between causal judgment and outcome recall. *Psychonomic Bulletin & Review, 12*, 950-954.

Mitchell, C. J., Lovibond, P. F., Minard, E., & Lavis, Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue-outcome relationship. *Quarterly Journal of Experimental Psychology, 59,* 830-844.

Mitchell, C. J., & Scully, A. L. (in preparation). Latent inhibition in memory and causal judgement.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological bulletin, 132*, 297-326.

Neumann, D. (2006). The effects of physical context changes and multiple extinction contexts on two forms of renewal in a conditioned suppression task with humans. *Learning and Motivation, 37*, 149-175.

Overton, D. A. (1964). State-dependent or "dissociated" learning produced with pentobarbital. *Journal of Comparative and Physiological Psychology, 57*, 3-12.

Paredes-Olay, M. C., & Rosas, J. M. (1999). Within-subjects extinction and renewal in predictive. *Psicológica, 20*, 195-210.

Pavlov, I. P. (1927). *Conditioned Reflexes*. Oxford, UK: Oxford University Press.

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

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

Pearce, J. M. (1997). *Animal Learning and Cognition: An Introduction* (2 ed.). East Sussex, U.K.: Psychology Press.

Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology, 52*, 111-139.

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*, 532-552.

Pineno, O., & Miller, R. R. (2007). Comparing associative, statistical, and inferential reasoning accounts of human contingency learning. *Quarterly Journal of Experimental Psychology, 60*, 310-329.

Rachman, S. J. (1989). The return of fear: Review and prospect. *Clinical Psychology Review, 9*, 147-168.

Rescorla, R. A. (1996). Preservation of Pavlovian associations through extinction. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 49B*, 245-258.

Rescorla, R. A. (2001). Experimental extinction. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of Contemporary Learning Theories* (pp. 119-154). Mahwah, N.J.: Erlbaum.

Rescorla, R. A. (2004). Spontaneous recovery. *Learning & Memory, 11*, 501-509.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning, Vol.2. Current Theory and Research* (pp. 64-99). New York: Appleton-Century-Crofts.

Romero, M. A., Vila, N. J., & Rosas, J. M. (2003). Time and context effects after discrimination reversal in human beings. *Psicologica, 24*, 169-184.

Rosas, J. M., & Callejas-Aguilera, J. E. (2006). Context switch effects on acquisition and extinction in human predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 461-474.

Rosas, J. M., Callejas-Aguilera, J. E., Álvarez, M. M. R., & Abad, M. J. F. (2006). Revision of retrieval theory of forgetting: What does make information context-specific? *International Journal of Psychology and Psychological Therapy, 6*, 147-166.

Rosas, J. M., Vila, N. J., Lugo, M., & Lopez, L. (2001). Combined effect of context change and retention interval on interference in causality judgments. *Journal of Experimental Psychology: Animal Behavior Processes, 27*, 153-164.

Scully, A. L., & Mitchell, C. J. (in press). Extinction in human learning and memory. *Quarterly Journal of Experimental Psychology*.

Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology, 60*, 291-309.

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

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. *The Psychology of Learning and Motivation, 21*, 229-261.

Shanks, D. R., & Dickinson, A. (1988). The role of selective attribution in causality judgment. In D. J. Hilton (Ed.), *Contemporary Science and Natural Explanation: Commonsense Conceptions of Causality* (pp. 94-126). New York, N.Y.: University Press.

Slamecka, N. J., & Ceraso, J. (1960). Retroactive and proactive inhibition of verbal learning. *Psychological Bulletin, 57*, 449-475.

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352-373.

Vila, N. J., & Rosas, J. M. (2001a). Reinstatement of acquisition performance by presentation of the outcome after extinction in causality judgments. *Behavioural Processes, 56*, 147-154.

Vila, N. J., & Rosas, J. M. (2001b). Renewal and spontaneous recovery after extinction in a causal learning task. *Mexican Journal of Behavior Analysis, 27*, 79-96.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information Processing in Animals: Memory Mechanisms* (pp. 5-47). Hillsdale, N.J.: Erlbaum.

Wolpe, J. (1958). *Psychotherapy by Reciprocal Inhibition*. Stanford, CA: Stanford University Press.

Appendix A
Learning Instruction Screens

Experiments 1, 2 and 3 Learning Instruction Screens:

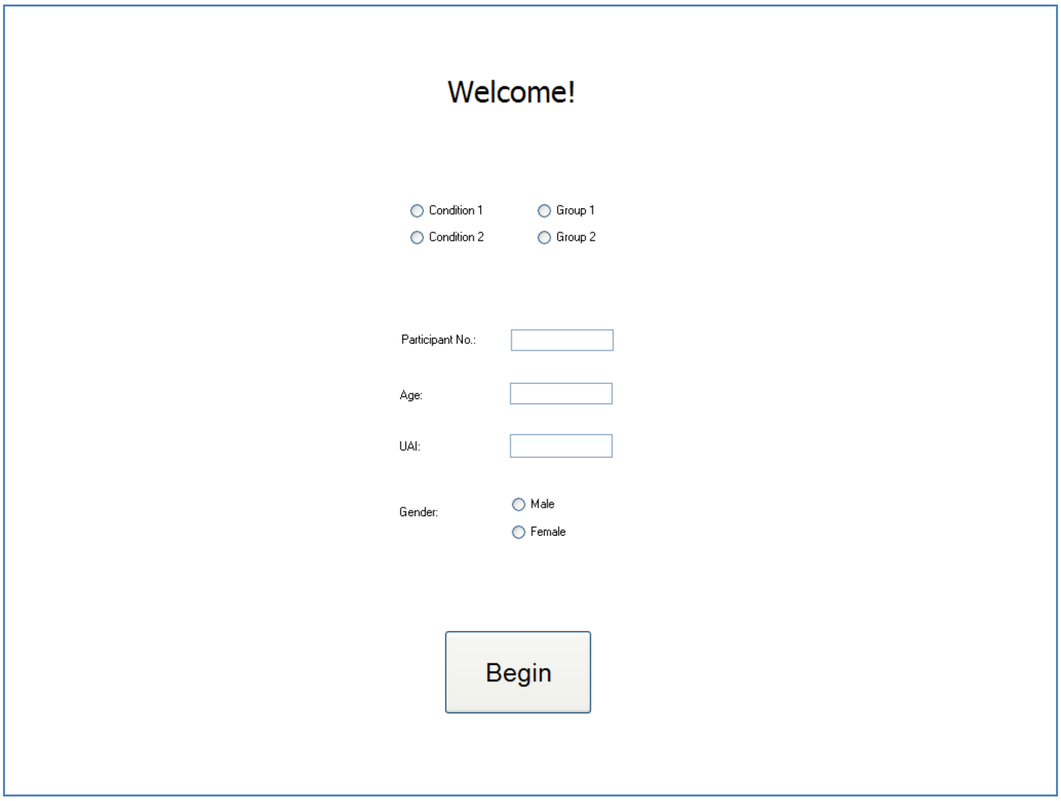

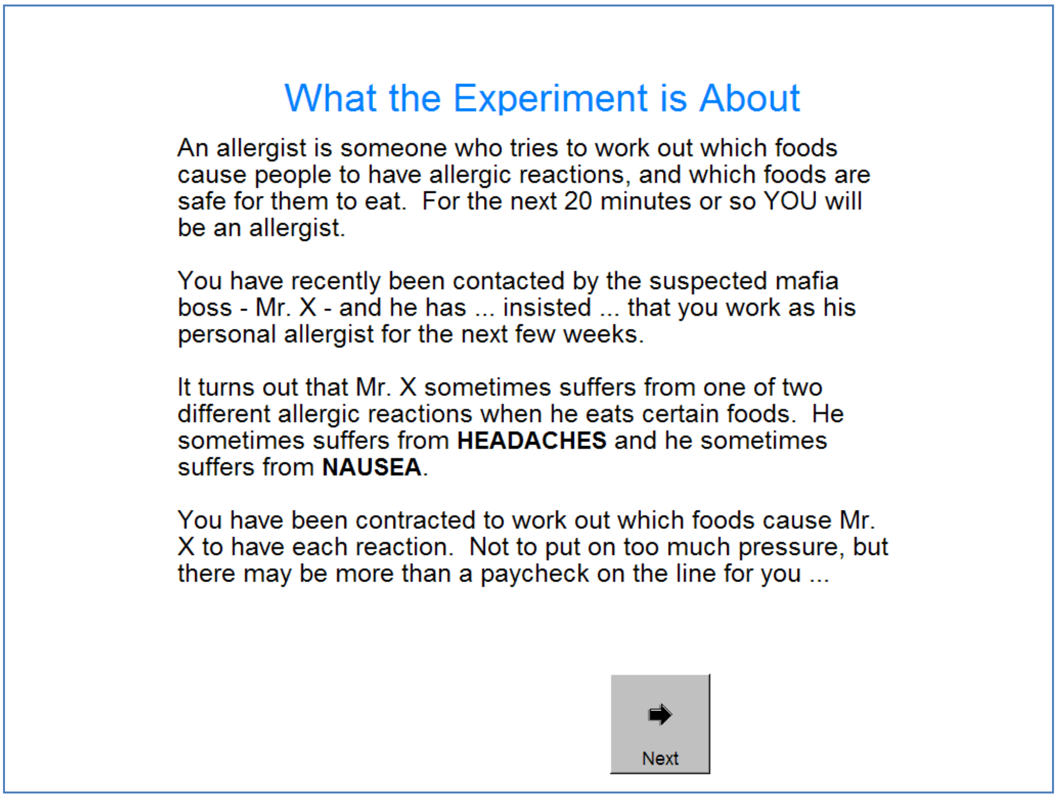

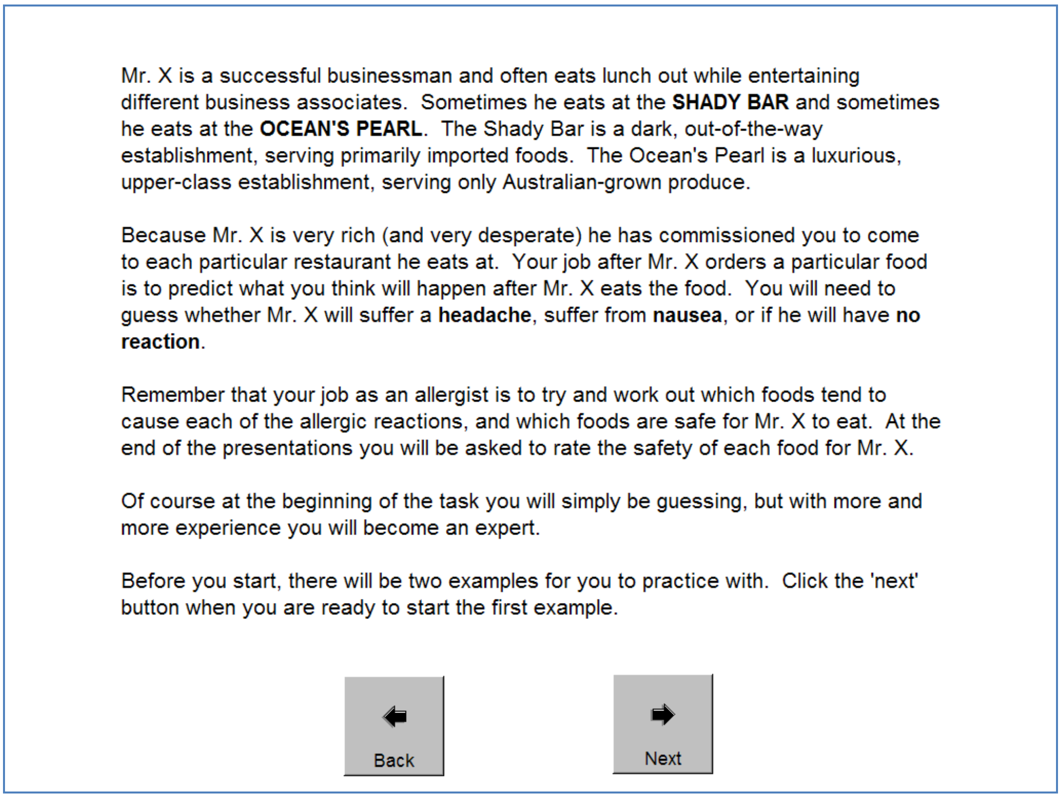
Mr. X is a successful businessman and often eats lunch out while entertaining different business associates. Sometimes he eats at the **SHADY BAR** and sometimes he eats at the **OCEAN'S PEARL**. The Shady Bar is a dark, out-of-the-way establishment, serving primarily imported foods. The Ocean's Pearl is a luxurious, upper-class establishment, serving only Australian-grown produce.
Because Mr. X is very rich (and very desperate) he has commissioned you to come to each particular restaurant he eats at. Your job after Mr. X orders a particular food is to predict what you think will happen after Mr. X eats the food. You will need to guess whether Mr. X will suffer a **headache**, suffer from **nausea**, or if he will have **no reaction**.
Remember that your job as an allergist is to try and work out which foods tend to cause each of the allergic reactions, and which foods are safe for Mr. X to eat. At the end of the presentations you will be asked to rate the safety of each food for Mr. X.
Of course at the beginning of the task you will simply be guessing, but with more and more experience you will become an expert.
Before you start, there will be two examples for you to practice with. Click the 'next' button when you are ready to start the first example.
Back
Next

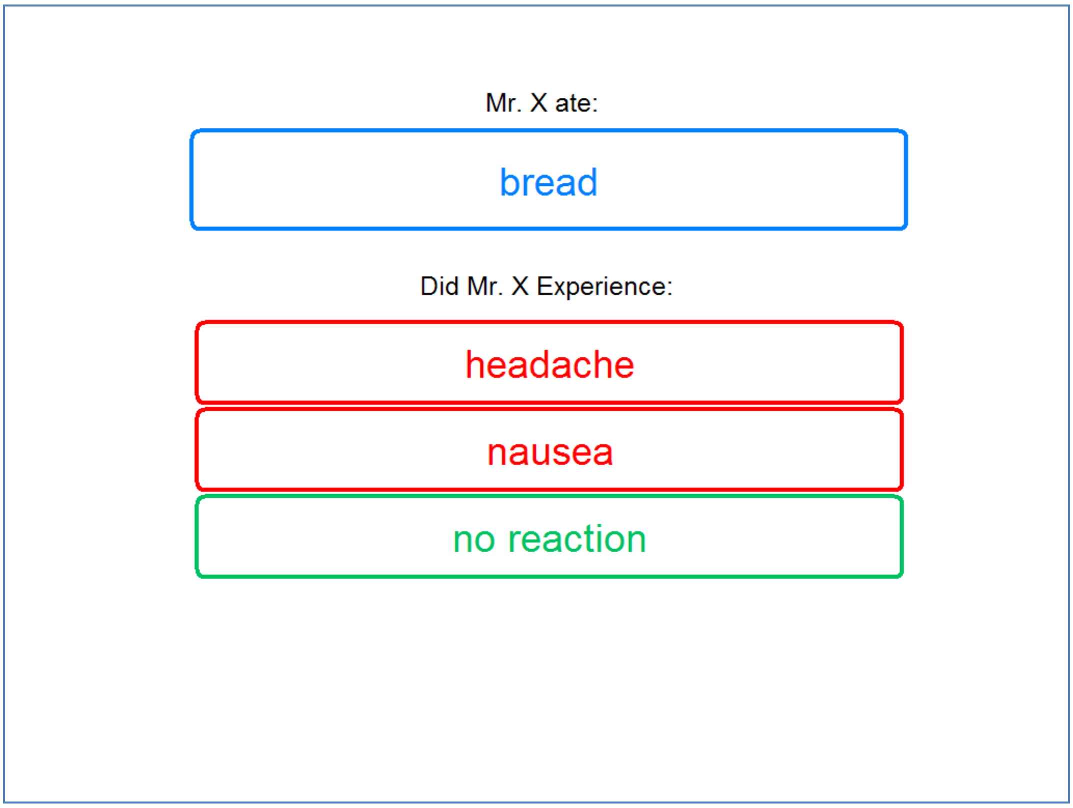
Mr. X ate:
bread
Did Mr. X Experience:
headache
nausea
no reaction

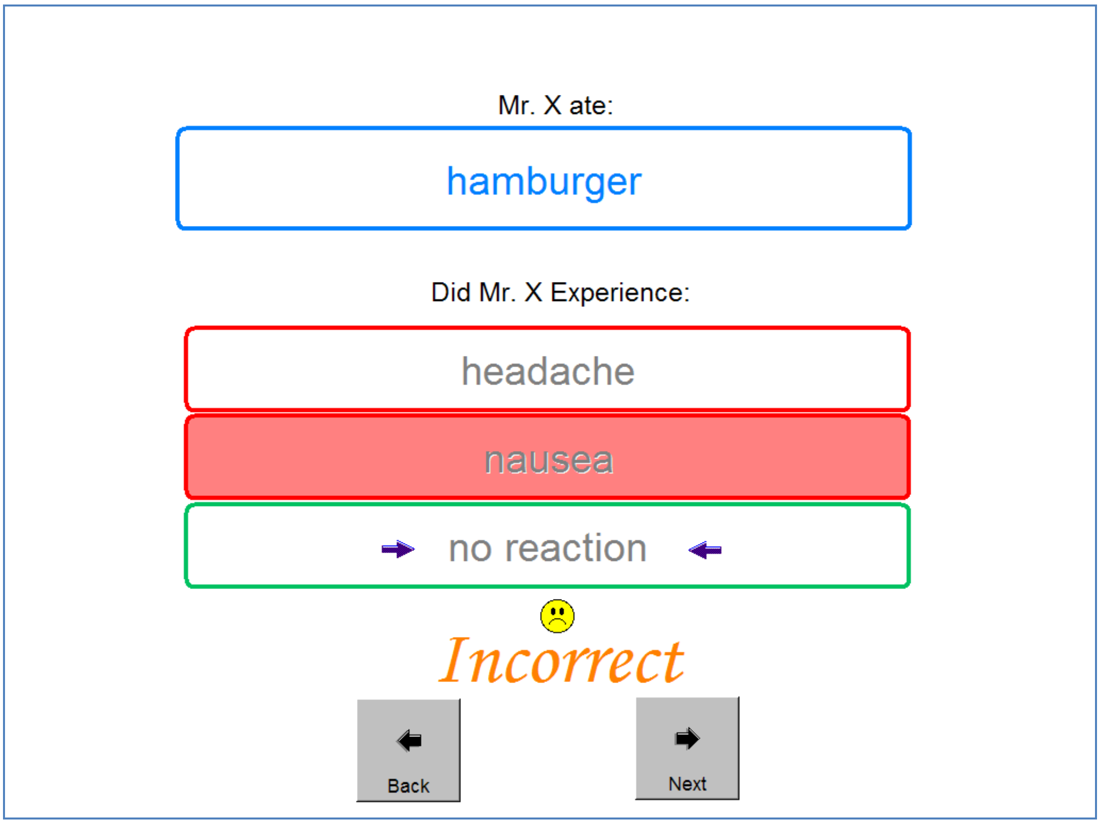
Mr. X ate:
hamburger
Did Mr. X Experience:
headache
nausea
no reaction
Incorrect
Back
Next

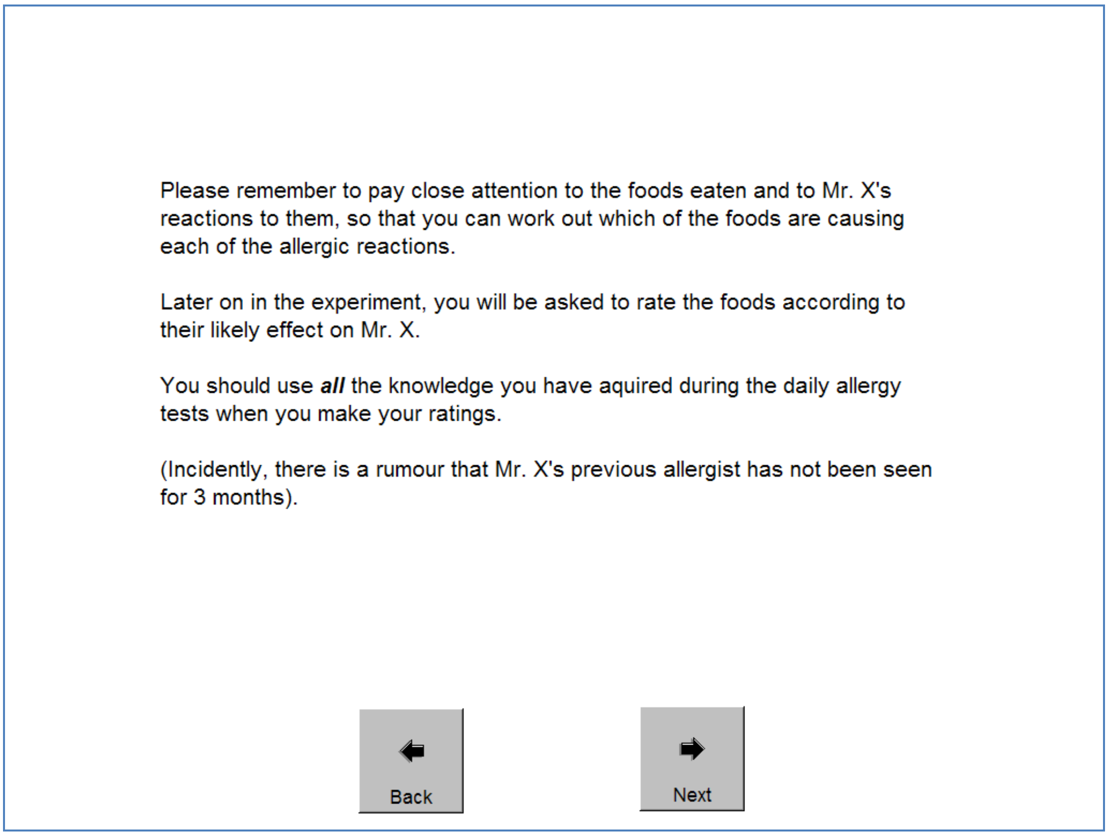
Please remember to pay close attention to the foods eaten and to Mr. X's reactions to them, so that you can work out which of the foods are causing each of the allergic reactions.
Later on in the experiment, you will be asked to rate the foods according to their likely effect on Mr. X.
You should use *all* the knowledge you have aquired during the daily allergy tests when you make your ratings.
(Incidently, there is a rumour that Mr. X's previous allergist has not been seen for 3 months).
Back
Next

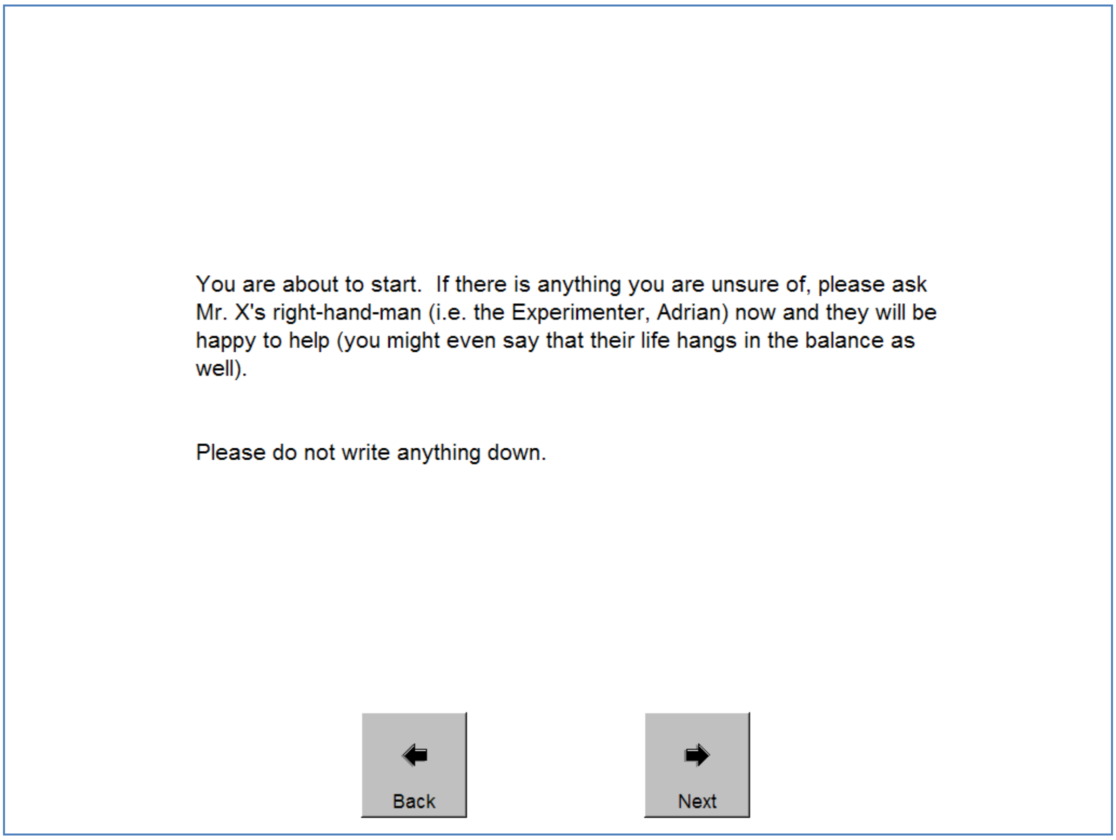
You are about to start. If there is anything you are unsure of, please ask Mr. X's right-hand-man (i.e. the Experimenter, Adrian) now and they will be happy to help (you might even say that their life hangs in the balance as well).
Please do not write anything down.
Back
Next

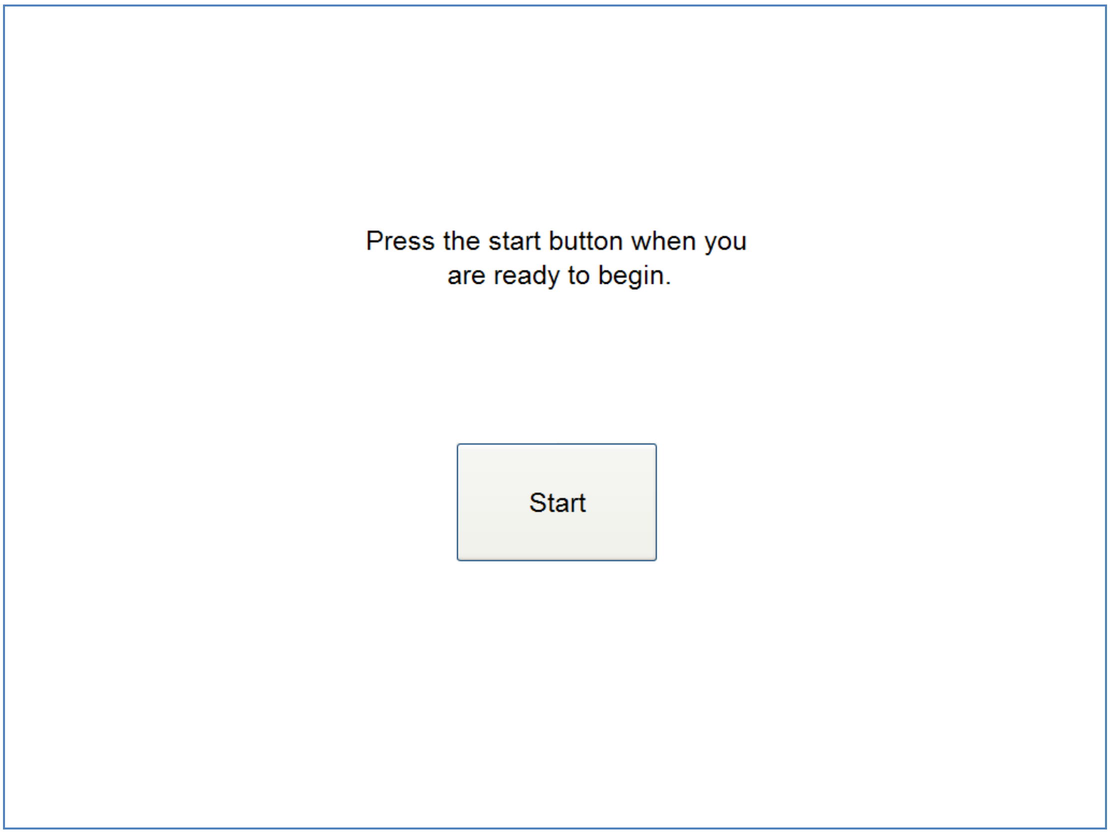
Press the start button when you are ready to begin.
Start

Experiment 3 Additional Instruction Screen:

Finally, Mr. X wants you to know that it is VERY important that you **pay attention to what restaurant he is eating at** when he does suffer from an allergy. Due to the different spices and techniques used by the different chefs at each restaurant, the same foods can cause an allergy or no reaction depending on what restaurant Mr. X is eating at. Thus, it is extremely important that you pay attention to the name of the restaurant at the top of the screen, as well as to the music and lighting of each restaurant, if you intend to help Mr. X (and survive this ordeal).

When you travel to another restaurant a message will come up on screen telling you to call over Mr. X's right-hand-man (i.e. Adrian, the experimenter) to transport you to the next restaurant. If you are unsure about anything, then please ask Adrian about it right now and they will be happy to help (you might even say that their life hangs in the balance as well).

Please do not write anything down.

Appendix B
Learning Screens

Experiment 1, 2, and 3 Learning Screens (with restaurant backgrounds):

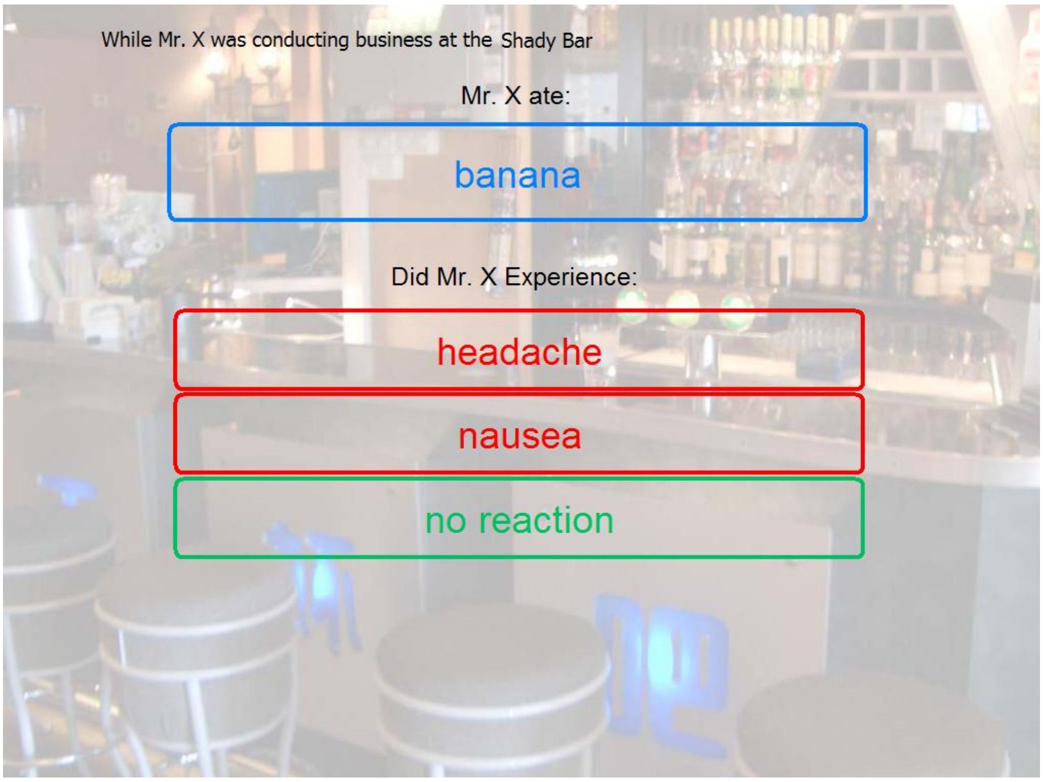

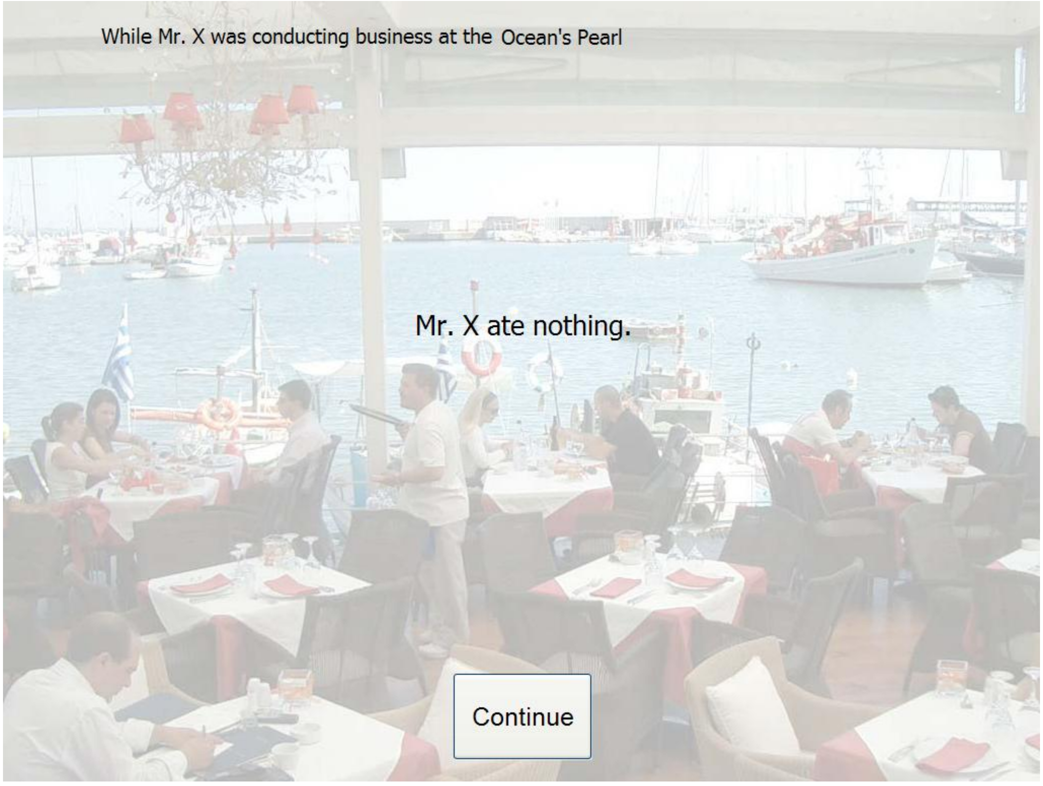

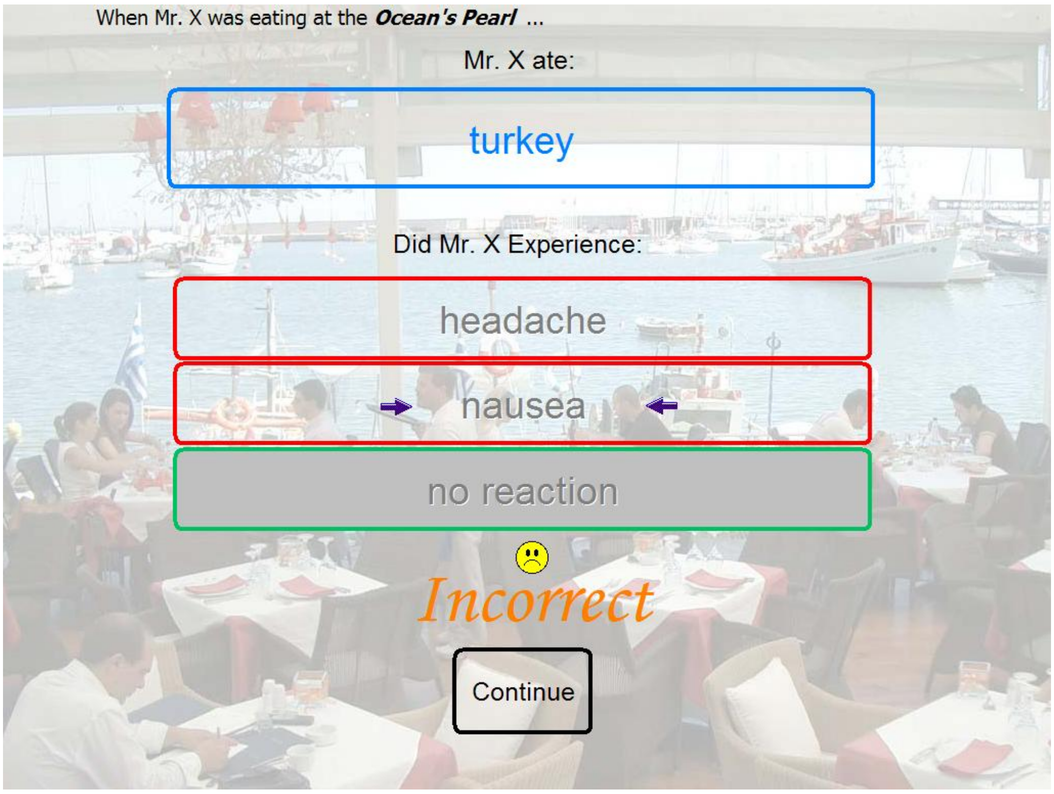
When Mr. X was eating at the ***Ocean's Pearl*** ...
Mr. X ate:
turkey
Did Mr. X Experience:
headache
nausea
no reaction
Incorrect
Continue

Experiments 2 and 3 Travel Screens:

Mr. X has to entertain some... well, less than reputable business associates today at the Shady Bar.

Mr. X has warned you to be alert to any police presence because they are, shall we say, not entirely convinced of the legality of the transaction.

Please call over Mr. X's right-hand-man to arrange transport.

Travel

Mr. X has to entertain some high roller business associates today at the Ocean's Pearl.

Please call over Mr. X's right-hand-man to arrange transport.

Travel

Appendix C
Test Instructions

Experiments 1 and 2 Test Instructions:

At this point Mr. X thinks you should put your expertise into action.

You will be asked to recall which allergic reaction Mr. X experienced after eating certain foods that you have previously seen. You will also be asked to rate to what extent the food caused the allergy.

So, for each food, please indicate which allergic reaction followed that food during the allergy tests, and to what extent that food caused an allergic reaction in Mr. X.

To help boost your memory Mr. X insists that you make your decisions at the restaurants where he ate. He believes that this will increase accuracy (and since he is the one with the money, and the gun, you agree with him).

Please ask Mr. X's right-hand-man now if you have any questions, or press "Begin Test" to start this last phase.

Begin Test

Experiment 3 Test Instructions:

At this point Mr. X thinks you should put your expertise into action.

You will be asked to recall which allergic reaction Mr. X experienced after eating certain foods that you have previously seen. You will also be asked to rate to what extent the food caused the allergy.

So, for each food, please indicate which allergic reaction followed that food during the allergy tests, and to what extent that food caused an allergic reaction in Mr. X.

To help boost your memory Mr. X insists that you make your decisions at the restaurants where he ate. He hopes that you **paid attention to the restaurants where he was getting sick, and suggests that you use this knowledge when making your ratings** (and since he is the one with the money, and the gun, you agree with him).

Please ask Mr. X's right-hand-man now if you have any questions, or press "Begin Test" to start this last phase.

Begin Test

Experiments 2 and 3 Travel Screens:

You have now finished the experiment.

Mr. X thanks you for your help and, right before he dissapears in his limo, reassures you that your "payment" will match your worth ...

Before you relinquish your allergist role, please push the "End" button at the bottom of this screen. THIS IS VERY IMPORTANT.

Thank you once again.

End

Appendix D
Test Screens

Experiments 1, 2, and 3 Test Screens (with context background)

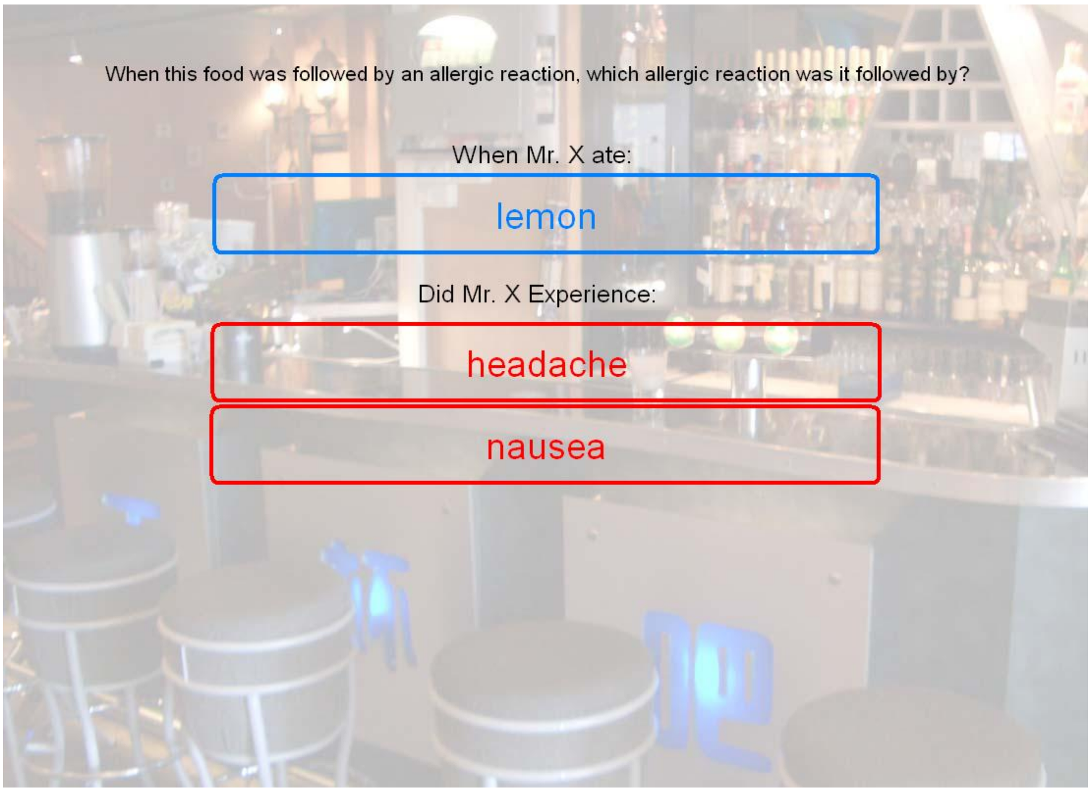

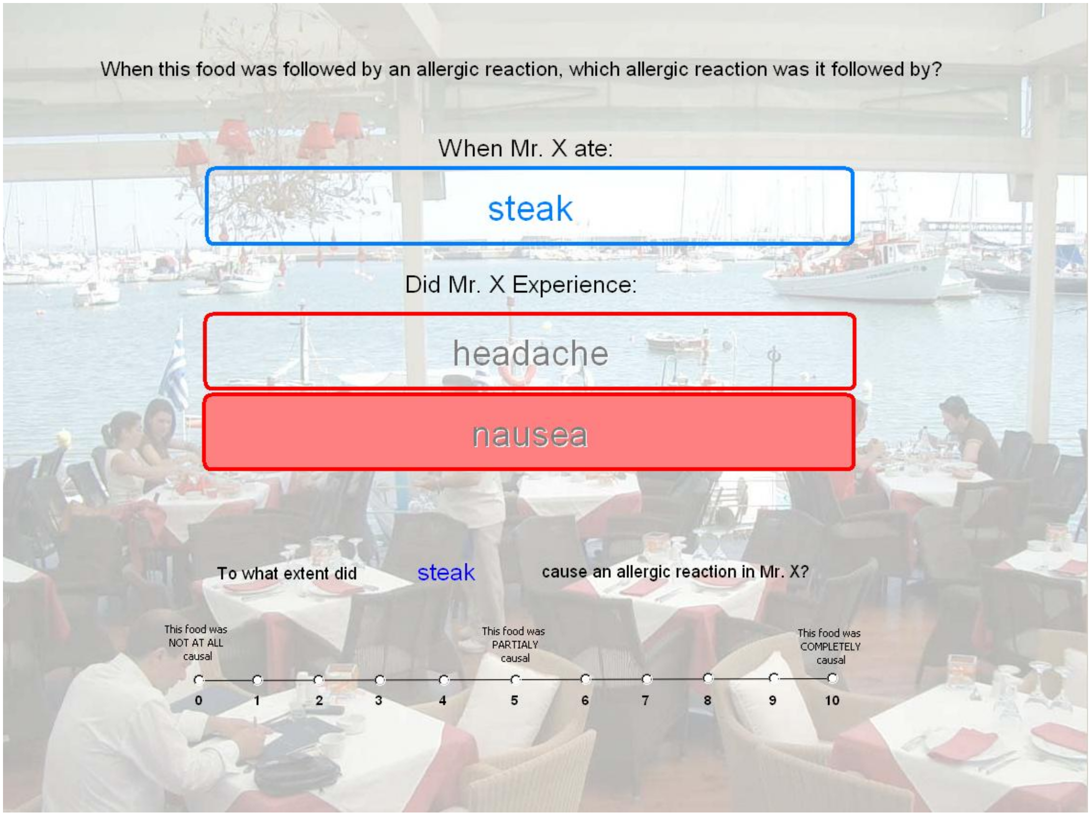

Appendix E
Experiment 1 Raw Data

	Causal Judgment Scores (0 = not causal)			Cued Recall Scores (0 = incorrect prediction)			Cue A Learning Trials (0 = predict no outcome)		
	Cue A	Cue B	Cue D	Cue A	Cue B	Cue D	Trial 7	Trial 8	Trial 9
Extinction	5	10	10	1	1	1	1	0	0
Group	8	8	8	1	1	1	1	0	0
	7	10	10	1	1	1	1	0	0
	5	5	10	0	0	1	1	0	0
	5	10	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	1	0
	10	10	10	1	1	1	1	0	0
	5	9	10	1	1	1	1	1	0
	5	10	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	9	10	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	7	10	10	1	1	1	1	0	0
	7	10	10	1	1	1	1	1	0
	6	8	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	1	0
	5	10	10	1	1	1	1	1	0
	2	5	8	1	1	1	1	0	0
	3	5	5	1	1	1	1	1	1
Renewal	7	9	8	1	1	0	1	0	0
Group	10	10	0	1	1	1	1	0	1
	10	5	5	1	1	0	1	0	0
	8	8	7	1	0	1	1	0	0
	10	10	5	1	1	1	1	0	1
	8	9	0	1	1	1	1	0	0
	10	10	10	1	1	1	1	0	0
	5	6	10	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	10	10	10	1	1	1	1	0	0
	5	9	7	1	1	1	1	0	0
	5	10	10	1	1	1	1	0	0
	7	5	7	1	1	0	1	0	0
	10	10	10	1	1	1	1	0	0
	10	8	8	1	1	1	0	0	1
	5	7	5	1	1	1	1	0	0
	10	10	10	1	1	1	1	0	0
	7	9	9	1	1	1	1	0	0
	4	8	10	1	1	1	1	0	0
	5	5	5	1	1	1	1	0	0
	6	10	10	1	1	1	1	0	0

Experiment 2 Raw Data

	Causal Judgment Scores (0 = not causal)			Cued Recall Scores (0 = incorrect prediction)			Cue A Learning Trials (0 = predict no outcome)		
	Cue A	Cue B	Cue D	Cue A	Cue B	Cue D	Trial 6	Trial 7	Trial 8
Extinction	5	8.5	10	1	1	1	1	0	0
Group	0	10	10	1	0.5	0.5	1	0.5	0.5
	5.5	9.5	10	0.5	0.5	1	1	0.5	0.5
	2.5	6.5	7	0.5	1	1	1	0.5	0
	0	7.5	7	1	1	0.5	1	0	0
	6	5	5	0.5	1	1	1	0.5	0
	7.5	5.5	7.5	0	1	1	1	0	0
	4.5	7.5	8.5	1	1	1	1	1	0
	8	8	9	1	0.5	1	1	0.5	0
	2.5	7.5	10	0.5	0.5	1	1	0.5	0
	10	3.5	2.5	1	1	1	1	0	0
	5	6.5	9	0.5	1	0.5	1	0.5	1
	7	8	6.5	1	0.5	1	1	0.5	0
	4	7	4.5	0.5	1	1	1	0.5	0
	4.5	7	8	1	0.5	1	1	0.5	0.5
	2.5	4	6.5	1	1	1	1	0	0
	4.5	4.5	4.5	1	0.5	0.5	1	0	0
	0	5	9	0.5	1	1	1	0	0
	2.5	6	10	0.5	1	1	1	0.5	0.5
	6	3	5.5	0.5	1	0.5	1	1	0
	9.5	8	10	0.5	0.5	1	1	0.5	0
	0	5	7.5	0	0.5	1	1	0	0
	4.5	6.5	7.5	0.5	0.5	1	1	0.5	0
	6.5	10	10	1	1	1	1	0.5	0
	Cue A	Cue B	Cue D	Cue A	Cue B	Cue D	Trial 6	Trial 7	Trial 8
Renewal	6	5	5	0	0.5	1	1	0	0
Group	5	10	10	0.5	1	1	0	0	0
	4.5	3.5	4.5	0.5	0.5	0.5	1	0	0.5
	6.5	10	9	1	1	1	1	0	0
	5	3.5	4	1	1	1	1	0	0
	3.5	8.5	8	1	1	1	0.5	0	0
	4.5	4.5	7	0.5	0.5	1	0.5	0.5	0
	7.5	10	9.5	0.5	1	1	1	0.5	0
	2	4.5	5	1	1	1	1	0	0
	3.5	5	7	1	1	1	0.5	0	0
	5	10	10	0.5	1	1	1	0	0
	8	6.5	7.5	1	1	1	1	0	0
	10	10	10	1	1	1	1	0	0
	0.5	9	7	0.5	0.5	0.5	1	0.5	0.5
	1	6.5	6.5	0.5	1	1	1	0	0
	2.5	10	10	1	0.5	1	1	0	0
	10	6.5	9	0.5	0.5	1	1	0.5	0
	2.5	5	10	0	1	1	1	0	0.5
	3	5	3.5	0.5	0	0.5	1	0.5	0
	5	6.5	5	1	1	1	1	0	0
	8.5	8	9.5	0.5	1	1	1	0	0
	5	9	10	1	0.5	0.5	1	0.5	0
	7.5	10	10	1	1	0.5	1	0	0
	6.5	6	6.5	0.5	0.5	1	0.5	0	0

Experiment 3 Raw Data

	Causal Judgment Scores (0 = not causal)			Cued Recall Scores (0 = incorrect prediction)			Cue A Learning Trials (0 = predict no outcome)		
	Cue A	Cue B	Cue D	Cue A	Cue B	Cue D	Trial 6	Trial 7	Trial 8
Extinction	6.5	7	10	0.5	1	1	1	0.5	0
Group	0	8.5	8.5	0.5	1	1	1	0.5	0
	4.5	7	8.5	0	0.5	1	1	1	0
	3	6	10	0.5	0.5	1	1	0.5	0.5
	0	7	10	1	0.5	1	0.5	0.5	0.5
	1	6	9	0	1	1	1	1	0.5
	5	10	7.5	1	1	1	1	1	0
	3	9	7.5	0.5	1	0.5	1	0.5	0
	2	7.5	10	0.5	1	1	1	0	0
	0	3.5	8.5	0.5	1	1	0.5	0.5	0.5
	4	10	10	1	1	1	1	0	0
	6	10	10	0.5	0.5	1	1	0	0
	5	7	9	1	0.5	1	1	0.5	0
	7	5.5	7.5	1	0.5	1	1	0.5	0
	5	8	8	1	0.5	1	1	0.5	0
	10	10	10	1	1	1	1	0.5	0
	2.5	6	5.5	0.5	1	0.5	0.5	0	0
	5	10	7.5	1	1	1	1	0	0
	2	3.5	6	0	1	1	1	0	0
	6	7	6.5	1	0.5	0	1	1	0
	5	10	9	1	1	1	1	0	0
	1.5	5.5	10	1	0.5	1	1	0	0
	Cue A	Cue B	Cue D	Cue A	Cue B	Cue D	Trial 6	Trial 7	Trial 8
Renewal	5	7.5	10	0	1	1	0.5	0	0
Group	9	9	9	1	1	1	0	0	0
	10	10	10	1	1	1	1	0	0
	2.5	8	7.5	0.5	0.5	1	0.5	0	0
	8	8.5	9.5	1	1	1	0.5	0	0
	9.5	3.5	8.5	0.5	0.5	0.5	0.5	0.5	0
	3.5	5	3.5	1	1	1	1	0	0
	7.5	5	8.5	0.5	0	1	0	0	0
	2	6	10	0.5	1	1	0	0	0
	9.5	5.5	9.5	1	1	1	1	0	0
	3	7.5	10	0.5	1	1	0.5	0.5	0.5
	1	5.5	10	1	1	1	0.5	0	0
	7.5	10	10	0.5	0.5	1	1	0	0
	6.5	9	10	0.5	0.5	1	0.5	0	0
	3	6	1.5	0.5	0	0.5	1	0.5	0
	4	9.5	9	1	1	1	0.5	0.5	0
	5	6	5.5	1	0.5	1	1	0	0.5
	7.5	8.5	7.5	1	1	1	1	0	0
	6	7	10	0.5	0.5	1	0.5	0.5	0
	8.5	6	7.5	1	1	0.5	0.5	0	0
	9	7	9.5	1	0.5	1	0.5	0	0
	10	9	10	1	1	1	1	0	0

Appendix F
Experiment 1 Analyses

Planned Causal Judgement Data Analysis (ANOVA critical value = $F_{(0.15/4;\ 1,\ 40)} = 4.630$):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	8.048	1	8.048	1.573
Extinction	B2	2273.357	1	2273.357	444.408
Renewal	B3	2672.024	1	2672.024	522.341
	Error	204.619	40	5.115	
	Within				
Cue A vs. B	W1	100.762	1	100.762	41.348
Ψ1	B1W1	29.762	1	29.762	12.213*
Ψ2	B2W1	120.024	1	120.024	49.253*
Ψ3	B3W1	10.500	1	10.500	4.309
	Error	97.476	40	2.437	
Cue A	W2	1813.714	1	1813.714	426.996
Ψ4	B1W2	34.381	1	34.381	8.094*
	B2W2	674.333	1	674.333	158.756
	B3W2	1173.762	1	1173.762	276.334
	Error	169.905	40	4.248	

Planned Cued Recall Data Analysis (Friedman Test critical value = $\chi2_{(0.05;\ 1)} = 3.84$):

Extinction Group cue A vs. cue B:

Ranks

	Mean Rank
CueA	1.50
CueB	1.50

Test Statistics(a)

N	21
Chi-Square	0.000
df	1
Asymp. Sig.	1.000

a Friedman Test

Renewal Group cue A vs. cue B:

Ranks

	Mean Rank
CueA	1.52
CueB	1.48

Test Statistics(a)

N	21
Chi-Square	1.000
df	1
Asymp. Sig.	.317

a Friedman Test

Planned Cued Recall Data Analysis (Kruskal-Wallis Test critical value = $H_{(0.05;\ 1)} = 3.84$):

Extinction Group cue A vs. Renewal Group cue A:

Ranks

	Group	N	Mean Rank
CueA	Extinction	21	21.00
	Renewal	21	22.00
	Total	42	
CueB	Extinction	21	21.50
	Renewal	21	21.50
	Total	42	

Test Statistics(a,b)

	CueA	CueB
Chi-Square	1.000	.000
df	1	1
Asymp. Sig.	.317	1.000

a Kruskal Wallis Test
b Grouping Variable: Group

Post-hoc Causal Judgement Data Analysis (Critical value = $F(_{0.05;\ 1,\ 40)} = 4.085$):

Extinction Group cue D vs. Renewal Group cue D:

```
                    Source         SS       df        MS             F
                    ------------------------------------------------
                    Between
                    ------------------------------------------------
Ext vs. Ren         B1            48.214     1       48.214        8.523*
                    Error        226.286    40        5.657
                    ------------------------------------------------
```

Post-hoc Learning Data Analysis (Kruskal-Wallis Test critical value = $H_{(0.05;\ 1)} = 3.84$):

Prediction for cue A on Learning Trials 7, 8, and 9:

Test Statistics(a,b)

	Trial Seven	Trial Eight	Trial Nine
Chi-Square	1.000	6.833	1.079
df	1	1	1
Asymp. Sig.	.317	.009	.299

a Kruskal Wallis Test
b Grouping Variable: Group

Experiment 2 Analyses

Planned Causal Judgement Data Analysis (ANOVA critical value = $F_{(0.15/4;\ 1,\ 46)}$ = 4.589):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	8.167	1	8.167	1.198
Extinction	B2	1490.755	1	1490.755	218.688
Renewal	B3	1819.172	1	1819.172	266.866
	Error	313.573	46	6.817	
	Within				
Cue A vs. B	W1	106.260	1	106.260	19.774
Ψ1	B1W1	0.042	1	0.042	0.008
Ψ2	B2W1	55.255	1	55.255	10.282*
Ψ3	B3W1	51.047	1	51.047	9.499*
	Error	247.198	46	5.374	
Cue A	W2	1111.688	1	1111.688	147.957
Ψ4	B1W2	4.688	1	4.688	0.624
	B2W2	486.000	1	486.000	64.683
	B3W2	630.375	1	630.375	83.898
	Error	345.625	46	7.514	

Planned Cued Recall Data Analysis (ANOVA critical value = $F_{(0.15/4;\ 1,\ 46)}$ = 4.589):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.000	1	0.000	0.000
Extinction	B2	26.255	1	26.255	254.820
Renewal	B3	26.255	1	26.255	254.820
	Error	4.740	46	0.103	
	Within				
Cue A vs. B	W1	0.260	1	0.260	3.433
Ψ1	B1W1	0.000	1	0.000	0.000
Ψ2	B2W1	0.130	1	0.130	1.716
Ψ3	B3W1	0.130	1	0.130	1.716
	Error	3.490	46	0.076	
Cue A	W2	22.688	1	22.688	216.857
Ψ4	B1W2	0.000	1	0.000	0.000
	B2W2	11.344	1	11.344	108.429
	B3W2	11.344	1	11.344	108.429
	Error	4.813	46	0.105	

Post-hoc Learning Data Analyses (Critical value for all contrasts = $F_{(0.05; 1, 46)} = 4.052$):

Prediction for cue A on Learning Trials 6 Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.188	1	0.188	5.308*
	Error	1.625	46	0.035	

Prediction for cue A on Learning Trials 7 Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.750	1	0.750	10.615*
	Error	3.250	46	0.071	

Prediction for cue A on Learning Trials 8 Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.047	1	0.047	0.945
	Error	2.281	46	0.050	

Experiment 3 Analyses

Planned Causal Judgement Data Analysis (ANOVA critical value = $F_{(0.15/4;\,1,\,42)} = 4.615$):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	26.730	1	26.730	3.609
Extinction	B2	1397.818	1	1397.818	188.742
Renewal	B3	1998.006	1	1998.006	269.783
	Error	311.051	42	7.406	
	Within				
Cue A vs. B	W1	117.071	1	117.071	30.556
Ψ1	B1W1	38.889	1	38.889	10.150*
Ψ2	B2W1	145.455	1	145.455	37.965*
Ψ3	B3W1	10.506	1	10.506	2.742
	Error	160.915	42	3.831	
Cue A	W2	1115.051	1	1115.051	149.553
Ψ4	B1W2	65.051	1	65.051	8.725*
	B2W2	320.727	1	320.727	43.017
	B3W2	859.375	1	859.375	115.261
	Error	313.148	42	7.456	

Planned Cued Recall Data Analysis (ANOVA critical value = $F_{(0.15/4;\,1,\,42)} = 4.615$):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.003	1	0.003	0.026
Extinction	B2	24.006	1	24.006	218.266
Renewal	B3	24.750	1	24.750	225.033
	Error	4.619	42	0.110	
	Within				
Cue A vs. B	W1	0.071	1	0.071	0.799
Ψ1	B1W1	0.071	1	0.071	0.799
Ψ2	B2W1	0.142	1	0.142	1.598
Ψ3	B3W1	0.000	1	0.000	0.000
	Error	3.733	42	0.089	
Cue A	W2	22.551	1	22.551	203.787
Ψ4	B1W2	0.051	1	0.051	0.462
	B2W2	10.227	1	10.227	92.421
	B3W2	12.375	1	12.375	111.829
	Error	4.648	42	0.111	

Post-hoc Learning Data Analyses (Critical value for all contrasts = $F_{(0.05;\ 1,\ 42)}$ = 4.073):

Trial 6 Data Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	1.114	1	1.114	15.022*
	Error	3.114	42	0.074	

Trial 7 Data Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.960	1	0.960	10.658*
	Error	3.784	42	0.090	

Trial 8 Data Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	0.023	1	0.023	0.750
	Error	1.273	42	0.030	

Post-hoc Cross-Experimental Learning Data Analysis (Critical value = $F_{(0.05;\ 1,\ 88)}$ = 3.949):

Trial 6 Data Analysis:

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	1.127	1	1.127	20.933
Exp.1 vs. 2	B2	0.623	1	0.623	11.575
Interaction	B3	0.214	1	0.214	3.977*
	Error	4.739	88	0.054	

Post-hoc Cross-Experimental Causal Judgement Data Analysis (Critical value = $F_{(0.05; 1, 88)}$ = 3.949):

	Source	SS	df	MS	F
	Between				
Ext vs. Ren	B1	26.730	1	26.730	3.609
Exp.1 vs. 2	B2	3369.094	1	3369.094	454.915
Interaction	B3	26.730	1	26.730	3.609
	Error	311.051	42	7.406	
	Within				
Cue A vs. B	W1	117.071	1	117.071	30.556
	B1W1	38.889	1	38.889	10.150
	B2W1	117.071	1	117.071	30.556
Ψ	B3W1	38.889	1	38.889	10.150*
	Error	160.915	42	3.831	