INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

Relative contingency learning in Pavlovian conditioning

Robin A. Murphy

Department of Psychology

McGill University, Montreal

January 1999

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfilment of the requirements of the degree of Doctor of Philosophy

Abstract

Five experiments with rats were conducted to determine the extent to which associative processes could be used to explain how rats seem able to learn complex CS-US contingencies during Pavlovian conditioning. Rats were exposed to positive, zero and negative CS-US contingencies and conditioned behaviour was compared with predictions derived from both associative models of conditioning and nonassociative normative theories of causal reasoning. A common measure of contingency, Δp, when used to analyze Pavlovian conditioning requires defining the likelihood of the US in the presence and absence of the CS. Experiments 1 and 2 involved a novel preparation in which, in addition to standard CS presence trials, the absence of the CS was signalled by a second CS, called the trial marker (a lever). Rats were trained to learn relationships in which the CS was either a positive predictor of the US or in which it was unrelated to the US. More conditioned tray entries were observed when the CS signalled an increased likelihood of the US (positive contingency). Consistent with the associative explanations, the trial marker elicited conditioned lever pressing when the CS signalled no change in the likelihood of the US (zero contingency) . Experiments 3, 4 and 5 extended the analysis with multiple CSs. These experiments examined whether learning about one CS was determined by its contingency relative to the contingency of other concurrently trained CSs. In experiments 3 and 4 conditioned responding to a moderately predictive CS was determined by its contingency relative to a perfectly predictive CS. Experiment 5 extended this effect to a case in which conditioning was influenced by the presence of a perfect predictor of the absence of the US.Together these results support the hypothesis that relative contingencies determine the strength of conditioned responding. The results are discussed from the perspective of both associative and nonassociative theory.

Résumé

Les cinq expériences exécutées dans cette thèse ont démontré la mesure dans laquelle des processus associatifs peuvent expliquer le facon dont le rat semble être capable d'apprendre des complexes contingences SC-SI durant le conditionnement pavlovien. Les rats furent exposés à des contingences SC-SI positives, négatives et neutres et leur comportement fut comparé aux prédictions dérivées des théories associatives et nonassociatives (normatives) du raisonnement causal. L'index normalement utilisé pour mesurer le degré de contingence entre un SC et un SI, le ΔP, requiert que la probabilité de la présence du SI soit mesurable autant durant la présence du SC que durant son absence. Les expériences 1 et 2 incorporaient une nouvelle procédure à l'intérior de laquelle l'absence du SC était signalée par un signal-essaie (l'apparition d'un levier). Les sujets furent entraînés sur des contingences où le SC était soit un bon prédicteur du SI (la contingence SC-SI était positive) ou bien un prédicteur sans valeur ($\Delta P = 0$). Les rats ont visité la mangeoire plus souvent lorsque la contingence SC-SI était positive. Le signal-essai supportait plus de touches du levier quand la contingence SC-SI était à zéro, un résultat conforme aux prédictions d'un modèle associatif. Les expériences 3, 4, et 5 ont approfondi cet analyse avec une procédure incluant plus d'un SC. Ces expériences ont examiné si l'apprentissage supporté par un SC était déterminé par la contingence entre ce SC et le SI relative ou aux contingences entre le SC et le même SI. Dans les expériences 3 et 4, la réponse conditionnée supportée par un SC suggérant de façon modérée que l'apparition du SI était elle-même influencée par la contingence parfaite entre un deuxième SC et le même SI. L'expérience 5 a observé cet effet avec un deuxième SC qui suggérait parfaitement l'absence su SI. Ensemble ces résultats supportent l'hypothèse où le degré de contingence relatif entre un SC et un SI détermine le degré de conditionnement supporté par ce SC. Les résultats des théories associatives et nonassociatives-normatives sout púseutés dans cette thèse.

Table of Contents

Acknowledgements vi

Introduction 1

Conditioning and Causal Induction 2

Normative Contingency Theory 6

Effects of Degrading CS-US Contingency 9

Modifying the p(US|CS) 9

Modifying the p(US|-CS) 13

The role of contextual associations 18

Evidence for Context Associations 20

Theories of Context /CS interaction 23

Associative Competition 23

Theories of Associative Comparison 30

Nonassociative Theory 32

Challenges to Contingency Theory 38

Necessity 40

Sufficiency 45

Rationale for the Present Experiments 46

Predictions for Contextual conditioning 50

Single CS Contingency Learning Experiments 52

Experiment 1 52

Method 59

Results 61

Discussion 70
Experiment 2 73
Method 75
Results 77
Discussion 85

Multi-CS Experiments 88
Experiment 3 90
Method 92
Results and Discussion 93
Experiment 4 106
Method 111
Results and Discussion 113
Experiment 5 122
Method 127
Results and Discussion 128

General Discussion 136
Causal Knowledge 143
Implications 145

Statement of Original Contribution 148

Conclusion 149

References 151

Acknowledgements

I would like to thank my supervisor Andy Baker for his help, for sharing his knowledge of associative learning and his careful reading of this manuscript. Others that have listened to and/or encouraged me include; Mark Bouton, Tony Dickinson, Peter Milner, John Pearce and Robert Rescorla. They provided valuable insight into the interpretation of these experiments. They always encouraged my interpretations, if they were not always convinced. I thank Tony Marley and Peter Milner for editorial comments on the final version of my thesis. Important support came from my fellow students and the administrative staff at McGill. Two students contributed to my progress towards completion. They were helpful and I learned from them both. Rob Mcdonald provided the Yin to and Fred Vallée-Tourangeau's Yang. For technical support I wish to thank Adrien Pitrat and Fay Guarraci for help changing the rats.

Personally I would like to thank two strangely related individuals for motivating me to study animal learning. I thank Ron Weisman who gave me a kick when I was an undergraduate at Queen's. I needed it. Andy Baker was an intellectual and personal inspiration to me and I thank him deeply for my wonderful experience in his lab (as well as a kick or two). I thank my parents for encouraging my interest in thinking and arguing. Finally, Vicki for her love and encouragement to whom I dedicate this work.

Introduction

Pavlov's (1927) study of the salivary reflex in hungry dogs showed that repeated pairings of an initially neutral stimulus (e.g., metronome) with food endowed the metronome with the ability to elicit salivation. He called this a conditioned rather than unconditioned reflex because the salivary response to the metronome was conditional upon the metronome-food pairing. In addition to an extensive description of the experimental parameters which influence this type of learning, Pavlov proposed a theory to explain it. He suggested that conditioned responses (CR) were the product of excitatory associations between a neural representation of the conditioned stimulus (CS; e.g., the metronome) and the representation of the unconditioned stimulus (US; e.g., the food). According to this theory, excitatory conditioning is the result of excitatory associations formed due to CS-US pairings.

Subsequent experiments (e.g., Rescorla, 1968; Dweck & Wagner, 1970; Hearst & Jenkins, 1974) suggested an alternative interpretation. Rather than simply acquiring associative memories on the basis of CS-US pairings, animals might actually be learning about the more complex statistical contingency between a CS and US (Baker, Murphy & Vallée-Tourangeau, 1996; Cheng & Holyoak, 1995). The most convincing data for this idea come from experiments which equate the number of CS-US pairings but vary how often the CS or US are presented alone. In spite of consistent CS-US pairings these manipulations often seriously affect the strength of the CR (Gibbon, Farrell, Locurto, Duncan & Terrace, 1980; Rescorla, 1968). This type of training and subsequent changes in behaviour may expose learning about CS-US contingency or predictiveness (Mackintosh, 1974; Miller & Schachtman, 1985; Prokasy, 1965) and therefore a different theory of learning. I will begin by reviewing evidence favouring a contingency learning perspective followed by criticisms described clearly by Papini and Bitterman (1990) who have argued that the notion of the CS-US contingency is not helpful for understanding the processes underlying Pavlovian conditioning. A series of experiments designed to answer this debate

is proposed and described.

Conditioning and Causal Induction

A stimulating approach to the study of conditioning has been to propose an analogy between animal conditioning and optimal or normative causal induction (Baker, Murphy & Vallée-Tourangeau, 1996; Cheng, 1997; Gormenzano & Kehoe, 1981; Mackintosh, 1977; Tolman & Brunswik, 1935). During Pavlovian conditioning animals are exposed to the arranged correlation between CS and US and learn to detect, or are sensitive to, the underlying CS-US relationship. While it is true that causal induction in humans may involve causal inferences that are independent of covariation information it is also true that much causal reasoning in humans involves determining whether two events covary (Baker et al.). To induce whether event A causes another event B, humans are sensitive to what is referred to as the *positive statistical relevance* of event A (Sosa & Tooley, 1993). Event A is positively relevant if, much like an excitatory CS, (1) event A occurs prior to event B and (2) the probability of event B is greater when event A occurs than is the overall probability of event B. Much like this type of causal learning, animal conditioning may be a form of cause-effect learning (Cheng & Holyoak, 1995). Therefore nature may equip organisms with a simple causal learning mechanism that learns to respond to the most highly contingent stimulus for an important event. This idea may help explain why animals exposed to pairings between a CS and US sometimes fail to learn about the relation. A CS may be paired with the US but not be correlated with it.

There is some resistance to this idea because it suggests that animals might be required to perform complex mental calculations (Papini & Bitterman, 1990). It is possible however, that animals may have evolved simple mechanisms that approximate these normative calculations. Comparing behaviour with optimal contingency theory may help us understand the learning mechanism. An analysis of animal behaviour in terms of computational norms has proven quite useful in other research areas (e.g., Church, 1997;

Herrnstein, 1961) without proposing that animals actually use a normative style algorithm. For example, foraging behaviour (Lea, 1981) has been described with reference to optimization models. Many animals forage for food in patterns that mimic an optimal strategy for minimizing energy expenditure and maximizing food intake. Research on foraging has compared actual behaviour with the behaviour of the hypothesized ideal forager without positing that animals perform a full cost-benefits analysis.

One principle in normative theorising about how causes are discovered is that causes of events tend (although not always) to occur just before their effects. They are temporally contiguous. Similarly temporal contiguity has been identified as a primary determiner of associative learning.

While contiguity is one cue to causality a significant problem for normative models of causal induction has been to understand how to select from among the many events that may be contiguous with an outcome the single event that should be attributed true causal status. Some events have multiple causes, but a system that induces causal relationships on the basis of contiguity alone will spend limited resources learning about events that are irrelevant, but because of serendipitous pairing seem to be causally related. Scientists and Philosophers have pondered this problem with respect to human causal reasoning. They have been puzzled by how people seem able to induce causes in the face of ambiguous data (for a review see Sosa & Tooley, 1993).

One solution, captured by the scientific method is to compare the relative likelihoods of the to-be-predicted outcome in the presence and absence of a set of possible causes. A standard experiment with an experimental and control treatment uses this logic to identify whether an experimentally manipulated variable causes a change in a dependent variable. A scientist compares the effect of the crucial putative causal variable on the experimental group with a control condition that holds all other factors (i.e., possible causes) constant. According to this analysis, causal factors can be induced by comparing the relative validity

of all possible causal factors with an outcome. The event which signals the largest relative increase in the likelihood of the outcome will be identified as having caused the outcome. This 'Method of Difference' retains the simple contiguity principle while allowing the reasoner to disambiguate redundant cues in the environment (Mill, 1843/1973).

The learning mechanism at work during conditioning procedures may have evolved to solve this type of problem (Cheng & Holyoak, 1995). In contrast, traditional associative approaches to Pavlovian conditioning that stress the role of contiguity have largely ignored any global relationships. Early theories of conditioning described how the strength of a learned response was primarily a function of the number of CS-US pairings (Hull, 1952). Stimulus-Response theories of conditioning formalised this idea in general laws of learning, for example, Thorndike's Law of effect (Thorndike, 1911/1965). However, if CS-US pairings were all that was important for developing a learned response, then any manipulation of CS-US covariation which holds the joint occurrence of CS and US constant but varies whether they occur separately should produce equivalent learning. Three findings, which I will argue are related, have challenged this assumption and encouraged the idea that conditioning involves a principle that is similar to the 'Method of Difference'. These are that: 1) the strength of a CR is influenced by how often the US is presented in the absence of the CS (e.g., Baker, 1977; Durlach & Shane, 1993; Gamzu & Williams, 1973; Rescorla, 1968), 2) relative to the number of CS-US pairings, a CR is influenced by the relative frequency with which the CS is presented without the US (e.g., Gibbon, Farrell, Locurto, Duncan & Terrace, 1980) and 3) the amount of exposure to the absence of both CS and US affects the strength of the CR (e.g., Kremer, 1974; Baker, 1977; Miller, Barnet & Grahame, 1995).

One interpretation of these findings is that they are unrelated. Each one may describe a separate variable that influences conditioned responding. Alternately, they may represent related findings each demonstrating sensitivity to a different component of the

overall CS-US relationship or contingency. One way to explore this idea further is to see whether the results of these three treatments are consistent with the changes in behaviour expected to occur if animals were sensitive to changes in CS-US contingency. To see how these treatments may involve changes in CS-US contingency requires defining more clearly the notion of a normative measure of contingency.

It may be worth clarifying at the outset in what way I am implicating the notion of a statistical contingency in the conditioning process. A useful categorization of the levels at which a function can be described and analyzed is provided by Marr (1982; see also Schmajuk, 1997; Church, 1997). Marr identified three levels of analysis, originally to describe an analysis of the visual system. These levels are, the computational level, the representation and algorithm level and the level of implementation. Each level describes a method of analysis of the same process. Furthermore, he proposed that it may be possible to ignore two levels while trying to describe the other.

Pavlovian conditioning has, until recently, been explained in terms of the middle representational and algorithmic level. Learning theorists have spent much effort describing possible representations of learning (e.g., habits, associations) and the algorithms that give the precise input/output relationships (Hull, 1952). More recently, with the advent of neural measurement techniques, researchers have been able to explore how the brain implements learning processes. However, descriptions of learning at the computational level have often been excluded from psychological investigation (although see Tolman & Brunswik, 1935). According to Marr, the computational level involves describing the goals of the computations, why they are appropriate and their underlying logic. Furthermore, all three levels are required for a complete understanding of a phenomenon. It will be argued that a normative theory of causal induction provides a computational description of why animals may acquire conditioned responses to stimuli, and what guides this selection. However, to make this point clear, even if one argues that computationally animals are learning about a

specific CS-US contingency, it is still possible that at the level of representation and algorithm they are learning associations in a manner that bears more resemblance to the mechanisms described by Pavlov (1927) or Hull than by those used to describe causal reasoning. It is important throughout this discussion to keep separate the idea of contingency as a computational description and the representational and algorithmic theories upon which much of the focus will rest.

Normative Contingency Theory

One common normative estimate of contingency combines four types of conjunction between two binary events. These are summarised in a two-way contingency table, an example of which is shown in Figure 1. In conditioning terms, cell A represents the frequency of CS-US pairings, cell B the frequency of CS presentations in the absence of the US (i.e., extinction), cell C the frequency of unsignaled US presentations, and cell D the frequency of the absence of both events (i.e., the inter-trial interval).

There are several ways of combining these cells to estimate the overall contingency between a CS and US, each of which is equally valid mathematically (Granger & Schimmler, 1986). The different contingency metrics differ by stressing the relative importance of one or another cell or combination of cells (Hammond & Paynter, 1983). One commonly employed measure is the Chi-Square (Allan, 1980). The Chi-Square though, is a two-way measure of contingency. This may be inappropriate for describing relationships in which there is an underlying causal frame. In the situation described here, the CS comes before or predicts the US. One measure of the one-way contingency which uses an unbiased weighting of the four cells is known as Δp (delta p) and was originally used for the description of human judgements of the strength of event-outcome relationships (Allan, 1980). Δp is the difference between two ratios (equation 1). In its application to conditioning, one ratio can be considered as the proportion of training trials in which the CS is paired with the US [p_1(US|CS)] and the other the proportion of trials

Figure 1. Two-Way contingency matrix describing the possible event conjunctions between a CS and US in Pavlovian conditioning. Below the table are the two conditional probabilities that comprise Δp, a measure of the one-way contingency between two binary variables (Allan, 1980).

	US	No US
CS	a	b
No CS	c	d

$p(US|CS)=a/(a+b)$
$p(US|\text{-}CS)=c/(c+d)$

$\Delta p=[a/(a+b)]-[c/(c+d)]$

without the CS paired with the US [p_0(US|-CS].

$$p_1(US|CS)=A/(A+B)$$

$$p_0(US|\text{-}CS)=C/(C+D)$$

(1) $$\Delta p=p_1 - p_0$$

Δp is not the only contingency metric which has been used to describe Pavlovian conditioning. For example, Gibbon, Berryman and Thompson (1974) have compared various metrics and settled on a version of the correlation coefficient ϕ, shown here for illustrative purposes only (equation 2).

$$p_2(CS|US)=(A/A+C)$$

$$p_3(CS|\text{-}US)=(B/B+D)$$

(2) $$\phi = \sqrt{(P_1-P_0)(P_2-P_3)}$$

I will not debate the relative merits of these or other contingency metrics, in general in this paper Δp will be discussed as the normative measure of contingency for two reasons. The various metrics make similar predictions in the treatments that we will discuss and one prominent associative theory of conditioning (i.e., the Rescorla-Wagner model; Rescorla & Wagner, 1972; Wagner & Rescorla, 1972) seems to share some features with Δp (Baker, Murphy & Vallée-Tourangeau, 1996).

Δp can take on values continuously between 1 and -1, covering the complete range of covarying relationships between two binary variables. Positive Δp values are appropriate for describing CSs that signal an increase in the likelihood of the US while negative Δp values can be used to describe CSs that signal a decrease in the likelihood of the US. A Δp of zero is a special case in which a CS signals no change in the likelihood of US delivery from that observed in its absence. Δp then, can be taken as an estimate of the predictiveness of a CS for the US.

The idea that Pavlovian conditioning involves contingency learning means that

animals perform a given CR because there is a positive or negative contingency between the CS and US and that the animal learns this relationship during training. The evidence that can decide this hypothesis comes from studies which have systematically tested whether changing the overall CS-US Δp systematically changes the strength of conditioned responding.

Effects of Degrading CS-US Contingency

Modifying the p(US|CS)

If the conditioned response generated by a CS is determined by the CS-US contingency then changing the contingency by decreasing either the likelihood of the US during the presence of the CS or the likelihood of the US during the absence of the CS should have a measurable effect on the CR. These two manipulations are similar to changing the two probabilities described by Δp. Decreases in the p(US|CS) or increases in the p(US|-CS) should decrease the strength of the conditioned response.

I will consider the evidence concerning decreases in p(US|CS) first. This change is equivalent to the difference between a continuous and partial reinforcement (PR) schedule (Mackintosh, 1974 pp. 72-75). If animals are sensitive to changes in CS-US contingency then in comparison to continuous reinforcement (100%; all CSs are followed by the US) an animal might be expected to acquire a weaker response following a PR schedule in which only some of the CS exposures are followed by reinforcement (e.g., 50%). However, it is difficult to compare continuous and partial reinforcement while controlling for two possible effects on the conditioned response independent of the change in contingency. Table 1 presents the event frequencies of a continuous reinforcement schedule and two possible control treatments.

To compare a 50% PR schedule with a continuous reinforcement schedule the experimenter can modify the continuous schedule by omitting half of the US presentations [Partial-(a)]. This treatment maintains the total number of CS presentations (10) but reduces

Table 1. Two partially reinforced (50%) experimental groups [Partial(a) and Partial(b)] designed to test for an effect of partial reinforcement on Pavlovian conditioning. These treatments are compared with conditioning in a continuously reinforced group. Partial(a) controls for the total number of CS presentations and Partial(b) for the total number of US presentations.

	p(US|CS)	CS-US pairing	CS alone	Total US	Total CS
Continuous	100%	10	0	10	10
Partial(a)	50%	5	5	5	10
Partial(b)	50%	10	10	10	20

the number of CS-US pairings (from 10 to 5). Any change observed in the CR may occur for two reasons. It may be affected by the reduced probability of reinforcement in the presence of the CS or because of the reduced number of US presentations. One way to control for this confound is to reduce the contingency by increasing the total number of CS presentations to 20 while keeping the number of CS-US pairings constant [Partial-(b)]. While this treatment reduces the p(US|CS) to 50% and maintains the same experience with the US, it increases experience with the CS. Neither control alone can be used to investigate the effects of PR. If both partial reinforcement schedules produce a similar change in responding from that found with continuous training, it could be argued that it was changes to p(US|CS) and not simply to the number of CS-US pairings, the number of CS or the number of US occurrences which was responsible for the strength of the conditioned response. In summary, one test for the position that animals are sensitive to CS-US contingency is to investigate how partial reinforcement influences the conditioned response.

Early investigations of partial reinforcement often maintained experience with the CS and reduced experience with the US. PR schedules produced in this manner often produce weaker conditioned responding than continuous reinforcement (ex., Fitzgerald, 1963; Pavlov, 1927; Wagner, Siegel, Thomas & Ellison, 1964). For instance, Fitzgerald studied the conditioned salivary response in hungry dogs. Each of three treatments received 24 fifteen second tones each day for 10 days. The groups differed with respect to the proportion of these trials that were paired with oral presentations of dilute acid (US). Either 100%, 50% or 25% of the CS trials were paired with the US. The partially reinforced groups showed inferior acquisition of the conditioned salivary response and this response was more resistant to extinction.

The differences in acquisition make intuitive sense from theories of learning which predict that the CR develops as a direct consequence of CS-US pairings (Hull, 1952). These theories would also seem to predict that continuous training results in a stronger CS-

US association and consequently it produces stronger conditioned responding. A stronger association might also be more resistant to extinction or reduction. It might, for example, take longer to reduce the effect of the conditioning treatment if training involved continuous pairings. Fitzgerald (1963) presented the same tone CS in the absence of any acid and found that continuous reinforcement produced *faster* extinction of the conditioned salivary response than either 50% or 25% reinforcement schedules. However, because all three groups received different numbers of US exposures these results do not unambiguously answer whether it was the probability of US following the CS or simply the total number of US presentations which produced this partial reinforcement effect.

Gibbon, Farrell, Locurto, Duncan & Terrace (1980) conducted a more extensive investigation of partial reinforcement. They studied PR with a classically conditioned keypecking response in pigeons. Pigeons were exposed to an illuminated keylight (CS) which signalled the delivery of grain (US). They trained pigeons with five reinforcement probabilities from 100% to 10%. Consistent with the findings of Fitzgerald (1963), reducing the number of US deliveries slowed acquisition of the CR (treatment Pa). They went on to test partial reinforcement schedules that were generated by holding the number of CS-US pairings constant and increasing the number of nonreinforced CS presentations (treatment Pb). When this was done the differences in acquisition speed and asymptotic levels of responding were shown to be determined by the absolute number of CS-US pairings. Gibbon et al., showed how this finding was consistent with a meta-analysis of data involving a variety of species including chicks, rabbits, rats, dogs and goldfish using a variety of appetitive and aversive conditioning procedures. In spite of these effects on acquisition the studies also concluded that both types of PR schedules increased resistance to extinction independent of the absolute number of CS-US pairings.

The reduced rate of extinction following partial reinforcement is a robust effect and suggests that partial reinforcement produces a change in behaviour which is different from

continuous reinforcement. The likelihood of the US in the presence of the CS is learned and does effect behaviour (Mackintosh, 1974) independent of individual CS or US or compound CS-US exposure. Intuitively one might expect that if a weaker CS-US association was formed as a result of partial reinforcement, then this might result in behaviour which extinguished more rapidly. One explanation offered for this seemingly paradoxical finding involves the similarity between the partial reinforcement training experience of nonreinforced exposure to the CS and the nonreinforced exposure to the CS during extinction. Capaldi (1966) has evidence to support the view that the similarity between PR and extinction serves to maintain the conditioned response during extinction training. Clearly though, the simple associative account which ties learning to the number of CS-US pairings is not tenable.

In terms of contingency, so far we have only considered the relative number of A to B cells (from Figure 1). In terms of Δp, being sensitive to these two cells is only half the story. The second source of information about the contingency arises from sensitivity to the density of US presentations in the absence of the CS [ie., p(US|-CS)]. Manipulating experimentally the relative number of A to B trials was simply matter of varying the number of reinforced to nonreinforced CS trials, one of the difficulties with manipulating precisely the likelihood of the US in the absence of the CS involves identifying the set of cues which constitute the absence of the CS.

Modifying the p(US|–CS)

Rescorla (1968) showed that increasing the frequency of US presentation in the absence of the CS, while keeping the frequency of USs during the CS fixed, could reduce the strength of the CR (see also Prokasy, 1965). These experiments involved the use of the conditional emotional response (CER) procedure in which pairing a tone with shock endowed the tone with the ability to suppress ongoing lever pressing behaviour. Rescorla assessed the effect of unsignaled shock presentations on the strength of a conditioned

response elicited by a tone. Regardless of the absolute number of CS-US pairings, increasing the relative number of shock presentations in the absence of the tone CS decreased fear of the tone (Rescorla, 1966, 1967; see also Rescorla & Lolordo, 1965). In one experiment rats were trained to bar press for food, and then received tone and shock pairings (Experiment 2 from Rescorla, 1968). Ten different groups of rats received different combinations of rates of shock in the presence and absence of the tone. Four treatments received uncorrelated CS-US training p(Shock|Tone) = p(Shock|No Tone)] while six received training in which the tone was a positive predictor of the shock [p(Shock|Tone)>p(Shock|NoTone)]. The actual rates per minute during the tone and during the absence of the tone are represented by the following pairs (0.8-0.4, 0.8-0.2, 0.8-0, 0.4-0.2, 0.4-0, 0.2-0, 0.8-0.8, 0.4-0.4, 0.2- 0.2, and 0-0). Rescorla argued that these differences in the relative likelihood of shock, rather than any absolute likelihood of shock in the presence or absence of the CS determined whether the rats would show conditioned fear to the tone. From the perspective of CS-US contingency one could argue that this was analogous to training rats with either a positive or zero Δp contingency.

Rescorla's (1968) result is surprising when considered from the perspective of the traditional contiguity-based theories of conditioning proposed to account for Pavlovian conditioning (Hull, 1952; Spence, 1940). These accounts assumed that the strength of the tone-shock association should have been the same in each group with the same shock density during the CS [e.g., 0.8-0,0.8-0.2, 0.8-0.4, 0.8-0.8] and different in those groups that shared different shock density during the CS [e.g., 0.4-.04, 0.2-0.2].

Like the problems with interpreting differences in conditioning between continuous and partial reinforcement schedules referred to earlier, differences between correlated and uncorrelated training may not reflect contingency sensitivity but rather sensitivity to either the number of CS-US pairings or differences in CS or US experience. To show that the likelihood of the US in the absence of the CS has an effect on a CR, Rescorla needed to

show that, for example, the conditioned response in treatment .8-0 was greater than responding in the zero contingency .8-.8 that received the same rate of CS-US pairing but received more US experience. Additionally he needed to show that responding was greater in the .8-0 group than in the 0.4-0.4 that received fewer CS-US pairings but the same absolute amount of US experience. Rescorla did not report this crucial analysis although visual comparison of the means from the data from the original paper suggests that the changes in contingency and not CS-US pairing, CS exposure or US exposure alone explains the pattern of results (although see Jenkins & Shattuck, 1981). One of the goals of this thesis is to conduct these tests.

Subsequent experiments using different conditioning procedures have shown that sensitivity to the presentation of the US in the absence of a CS occurs in different species and conditioning preparations. Using an active avoidance response procedure, Weisman and Litner (1969) showed that rats learn to shuttle from one corner of a conditioning chamber to another when presented with a signal that is paired with shock. If the shock occurs with an equivalent likelihood in the absence of the signal then the conditioned response is reduced. In appetitive conditioning, sensitivity to the overall contingency has been shown both using key pecking for food with pigeons (Gamzu & Williams, 1973; Lindblom & Jenkins, 1981; Durlach, 1983; Durlach, 1989a) as well as using a tray flap push response for food with rats (Durlach & Shane, 1993). In an experiment using the key peck procedure Gamzu and Williams (1973) showed that pigeons developed an autoshaped keypecking response to a keylight which signalled food delivery. When the probability of grain was the same regardless of whether the keylight was lit or dark (the absence of the CS) then the pigeons failed to acquire the autoshaped response (uncorrelated condition) (see also Gibbon, Locurto & Terrace, 1975). In a related finding from instrumental conditioning, superimposing unsignaled US presentations on an instrumentally conditioned bar pressing response reduces bar pressing (Dickinson & Charnock, 1985; Hammond,

1980; Hammond & Weinberg, 1984; Rescorla & Skucy, 1969).

The wealth of data on this manipulation, supports the notion that conditioned responses are influenced by both the likelihood of the US in the presence of a CS as well as the likelihood in the absence of the CS. The previous discussion has been confined to excitatory conditioning, in which a positive CS-US relation seems to be learned; however contingencies can be constructed so that a CS signals the absence of the US. One might propose from the previous argument that animals might be able to learn that a CS was contingently related to US absence. The notion of inhibitory conditioning has been known since Pavlov (1927). Rescorla showed that it was negative contingencies that allowed a CS to acquire inhibitory control over responding.

Using a CER conditioning procedure Rescorla showed that a tone could inhibit a conditioned fear response for shock if the presence of the tone signalled the absence of shock (Rescorla, 1969a, 1969b, 1969c). Under these conditions, subsequent attempts to condition a Tone-US relation were retarded. In addition, pairing the tone with a known excitor resulted in a reduction in the strength of the excitatory response from that observed during the excitor alone. These retardation and summation tests together support the idea that the negative contingency produced inhibitory conditioning. The CS was never itself paired with the US and therefore these experiments represent a limited range of negative contingencies. Subsequent experiments demonstrated inhibitory conditioning even when the CS itself was paired with the US, as long as the CS was paired with a lower proportion of USs than was the absence of the CS (Witcher & Ayres, 1980). In addition to the CER, inhibitory conditioning following negative contingency training has also been demonstrated using an aversive conditioned lick suppression procedure with rats (Miller & Schachtman, 1985). Finally, Cotton, Goodall and Mackintosh (1982) have demonstrated inhibitory conditioning in spite of consistent CS-US pairings. In these experiments, a CS was a perfect predictor of a reduction in the magnitude of the US. This inhibitory CS, which

always signalled the delivery of a single food pellet (A+), acquired inhibitory properties when training included compound training with a second CS (B) signalling two food pellets (AB++). Similarly, discriminative stimuli for instrumental lever pressing for access to sucrose in rats have also been shown to develop inhibitory properties if they signal a reduction in the probability of reinforcement signalled by a second cue (also Nelson, 1987).

Together these experiments show that even with consistent CS-US pairing, a CS can come to inhibit behaviour if it signals a relative decrease in the likelihood or even the amount of the US. Although, in general, inhibitory conditioning arising from negative contingencies is more difficult to demonstrate than excitatory conditioning from positive contingencies, most findings are consistent with the generalisation that it is the overall predictiveness captured by the global CS-US contingency rather than simply the absolute number of CS trials, US experiences, number of CS-US pairings or CS-US probability that determines both the strength of the conditioned response and its valence.

Finally, if animals can learn about a positive or negative contingency between two events then what do they learn when there is a zero contingency between CS and US ? There are at least three possibilities. Early theories of conditioning assumed that zero contingency training would result in excitatory conditioning because the CS and US would still be paired together by chance occasionally. These chance pairings would cause the formation of an excitatory association. Alternatively, Rescorla's findings about the relation between conditioning and contingency sensitivity suggested that uncorrelated training might leave a CS neutral since it would be a poor predictor of both US presence and absence. Rescorla (1967), thus argued that uncorrelated training was the correct control treatment for conditioning. Others have suggested that animals might learn specifically that both CS and US are unrelated or that the CS is irrelevant (Baker & Mercier, 1989; Mackintosh, 1974; Seligman, 1992).

Quinsey (1971; Kremer, 1971) showed that uncorrelated CS-US training can

endow a CS with excitatory properties. Others claim to have found evidence that animals can learn the specific uncorrelated relationship. For example, Baker and Mackintosh (1977) demonstrated using appetitive conditioned licking procedure with rats that a tone trained with a zero contingency with water was difficult to condition in a subsequent conditioning phase. Because the zero contingency training interfered with both excitatory and inhibitory conditioning they took this to mean that the animals had learned explicitly that the tone and shock were unrelated (see also, Bennet, Malado & Mackintosh, 1995; Matzel, Schachtman & Miller, 1988). While this research confirms that animals are sensitive to uncorrelated training, it would seem that experience with a zero contingency may not leave the CS associatively neutral as suggested by Rescorla. While this finding and its interpretation is not free from controversy (see Bonardi & Hall, 1996; Killcross & Dickinson, 1996; Mackintosh, 1983) animals at least seem to learn something during zero contingency training.

In summary, the reason why any of the preceding manipulations involving CS and US presentation were effective in producing a change in conditioned behaviour could be related to their influence on the CS-US contingency. Furthermore, the normative statistical measure Δp captures the computational requirements that contribute to the formation of a conditioned response.

The role of contextual associations

While these data support the idea that conditioning is a product of CS-US contingency sensitivity they do not suggest an algorithm by which this comes about. One possibility proposed by Cheng (1997) is that animals actually use the Δp algorithm. They compute Δp from a memory of events. An alternative is to explain these effects in terms of a representational account involving associations. There are several associative theories of Pavlovian conditioning (e.g., Gibbon & Balsam, 1981; Jenkins, Barnes & Barrera, 1981; Mackintosh, 1975; Miller & Shactman, 1985; Pearce, 1987, 1994; Pearce & Hall, 1980;

Rescorla & Wagner, 1972; Wagner, 1981). Although these approaches differ with respect to the precise learning algorithm, they share an important assumption. They assume that the contingency manipulations affect CS conditioning by influencing the signal value (Wagner, 1969) or predictiveness of all the cues other than the CS. This occurs as a result of changes to the associative status of the experimental context. It is context-US associations that mediate the strength of the CR to the CS. How this change to the conditioned response emerges separates the theories into two general classes.

Theories of associative *competition* propose that conditioned stimuli actually interfere with or compete with each other for association with the US during training thereby affecting each others' ability to elicit a CR (e.g., Mackintosh, 1975; McLaren, Kaye & Mackintosh, 1989; Pearce, 1987; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Schmajuk, 1997; Wagner, 1981). In contrast, *comparator* theories describe how the associations are formed independently and how associations compete for control of responding (e.g., Gibbon & Balsam, 1981; Jenkins, Barnes & Barrera, 1981; Miller & Schachtman, 1985). Therefore, contingency effects can be conceived as reflecting differences in either acquisition or performance (Baker & Mercier, 1989; Durlach, 1989b).

Finally, in contrast to associative accounts that only posit associations there are nonassociative accounts involving other, possibly cognitive, variables (e.g., Baker & Mackintosh, 1977; Baker & Mercier, 1989; Baker, Singh & Bindra, 1985) that may mediate contingency learning. Explicit reference has been made to production systems (Holland, Holyoak, Nisbett & Thagard, 1986; Holyoak, Koh & Nisbett, 1989) and normative statistical reasoning modules (Cheng & Holyoak, 1995; Cheng, 1997). In spite of their differences the associative and nonassociative approaches acknowledge a crucial role for the signalling value of the experimental context in learning about CS-US contingencies.

I will briefly describe empirical evidence for the role of contextual cues in Pavlovian

conditioning. I will then describe how the different theories account for basic sensitivity to contingency. Throughout this review I will argue that the crucial experiments that will test the extent to which the CS-US contingency is an appropriate theoretical tool for understanding conditioning require measuring and controlling the predictive status of the contextual cues in conditioning. Following this review I will propose a series of experiments designed to test contingency sensitivity by using a procedure that attempts to explicitly define the CS-US contingency.

Evidence for Context Associations

Whenever a CS is paired with a US, there are inevitably several sources of sensory stimulation, in addition to that provided by the CS, which could become associated with the US. However, standard descriptions of the programmed events presented to animals during conditioning often reduce to the two discrete events programmed by the experimenter, the CS and US (Dickinson, 1980, pp. 30-36). This view neglects an important component of training, namely the conditioning chamber itself. The space in which the experiment takes place contains tactile, auditory, visual, olfactory and possibly gustatory cues which could, according to the principle of contiguity, become associated with the US.

Pavlov (1927, p. 115) recorded anecdotal evidence of contextual conditioning. The dogs in his experiments became visibly excited by the sight of the experimental apparatus and room. He argued that these visual contextual cues were signals for the imminent presentation of food. A number of procedures for measuring conditioning to contextual cues have been reported. Balsam (1985) describes direct and indirect methods with which to infer contextual stimulus control over responding. The context's associative strength can be directly measured using a preference test in which animals are allowed to choose one of several previously conditioned compartments (Bouton, 1984). Behavioural measures of activity can also be used to assess contextual conditioning. For example, Blanchard and Blanchard (1969) presented unsignaled shocks to rats and found they exhibited increased

crouching. Crucial for any demonstration of contextual conditioning is that the conditioned behaviour must be specific to the training context in which the USs are received. Blanchard and Blanchard showed just this by shifting the rats to a second conditioning chamber and showing a reduction in crouching. Crouching was not the result of a nonassociative process arising from experience with the shock but reflected a conditioned response arising from sensitivity to the context-shock pairings (see also Bouton, 1984; Dweck & Wagner, 1970).

Some claim that contextual conditioning can also be inferred indirectly by measuring the manner in which procedures, designed to produce contextual conditioning, are able to influence conditioning during standard discrete CS conditioning (Balsam, 1985; Grau & Rescorla, 1984; Hall, 1991; Hall & Honey, 1989; Odling-Smee, 1975a; Siegel, 1977; Tomie, 1976). The context has been assigned a role in the development of habituation (Hall & Channell, 1985) and conditioning phenomena including the retardation of conditioning following CS pre-exposure (latent inhibition; Baker & Mercier, 1982; Hall & Honey, 1989; Lovibond, Preston & Mackintosh, 1984) the reduction of the conditioned response following repeated nonreinforced exposure (extinction; Bouton & Bolles, 1979), the reduction of the conditioned response following US preexposure (Baker, Mercier, Gabel & Baker, 1981[1]), the return of an extinguished response following extinction (renewal; Bouton & King, 1986; Bouton & Peck, 1989), inhibitory conditioning (Bouton & Nelson, 1994) and the return of a conditioned response following presentation of the US after extinction (US reinstatement; Bouton & King). If contextual cues play an important role then why are they often omitted from standard descriptions of conditioning preparations?

Pavlov (1927) suggested that contextual cues probably play a relatively minor role during many conditioning procedures. He proposed that during CS conditioning any contextual conditioning is normally counteracted by the gradual development of internal

[1]Although they argue that the effects of US pre-exposure on conditioning to a CS are almost certainly not simply the results of contextual conditioning.

inhibition because the context is extinguished (i.e., it receives repeated nonreinforced exposure) during the inter-trial-interval. Since the context is often experienced without the US it gradually becomes a signal for the absence of the US. However, this theory of contextual involvement in conditioning is too simple. In many situations the context retains the ability to influence the CR.

Following sufficient CS-US pairings the discrete CS acquires control over a CR and this may occur because the context-US association becomes behaviourally silent as a result of extinction. There is however, strong evidence that animals are sensitive to the relative predictiveness of the cues (e.g., Wagner, Logan, Haberlandt & Price, 1968). Therefore it might be the relative signalling value of the CS+Context versus the Context which determines whether a stimulus acquires a conditioned response. Manipulations which influence the relative predictiveness of the context or the CS through pre-training or extinction might be expected to influence the course of CS conditioning. This analysis implies that rather than describing the standard conditioning training as involving CS trials ($\mathbf{A}^+$) a truer description requires including the context (**X**). Thus a simple single CS conditioning procedure is really a discrimination between reinforced experience with the CS and context ($\mathbf{AX}^+$) and nonreinforced context experience ($\mathbf{X}^-$; Rescorla, 1969c). When such a discrimination is conducted with two discrete cues in the role of **A** and **X** it is called a feature positive discrimination.

It is clear that contextual cues can have a significant role in conditioning and should be included in a description of the events experienced during training. This seems particularly important for arguing that contingency is a measure of the difference between the likelihood of the US in the presence and absence of the CS. I will now discuss the three theoretical approaches with explicit reference to the emphasis that they place on processing of contextual cues. The theories are; 1) theories of associative competition 2) theories of associative comparison and 3) finally a more general class of nonassociative theories.

Theories of Context /CS interaction

Associative Competition: Rescorla-Wagner model

Prior to Rescorla's (1968) experiments, theories of learning assumed that the response control acquired by a CS was the direct result of CS-US experience (this experience was assumed to influence the accumulation of habit strength or response probability; Hull, 1952; Spence 1940). In the spirit of quantification, Bush and Mosteller (1951) derived a mathematical instantiation of Hull's version of this principle (equation 3) which could be used for calculating changes in the probability of a given response on a trial by trial basis. These changes were assumed to be a weighted difference between the maximum probability and the accumulated probability (p) from previous trials.

(3) Δ in response probability=$k(\lambda-p)$

The parameter for the US ($\lambda<=1$) sets the upper limit for responding. The formula is weighted by a learning rate parameter (k) for learning about specific stimuli. The model predicts a negatively accelerating learning curve characteristic of many learned responses. With enough training a given stimulus will elicit a response with probability λ and learning stops.

The problem with this theory is that it assumes that the conditioned response elicited by one CS should be independent of any responding elicited by another simultaneously trained CS. The conditioned response is determined solely by the number of times that a stimulus is reinforced and extinguished. As a consequence a number of basic conditioning effects including, overshadowing (Pavlov, 1927), contingency sensitivity (Rescorla, 1968) or stimulus selection effects are not anticipated (e.g., blocking; Kamin, 1969; Wagner, Logan, Haberlandt & Price, 1968). In a major conceptual shift Rescorla and Wagner (1972; Wagner & Rescorla, 1972) altered Bush and Mostellor's model (equation 4) changing the focus from response probability to associative strength (V). This shifted the focus of investigation from the behaviour to internal cognitive mechanisms of learning. Behaviour is

assumed to reflect learned associations in some monotonic fashion but the exact mapping has still to be addressed. The second major difference from previous formulations was that the associative strength acquired by a single CS was determined not only by its own previous association with the US, but by the associations acquired by all other cues present on a given trial (ΣV).

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

α and β represent learning rate parameters for the CS and US respectively. The formula can be used to calculate predicted changes in associative strength for a given cue presented on a given trial. This change allowed their new theory to accurately predict some previously unexplained findings. One example is the Rescorla-Wagner model's (RWM) account of Kamin's blocking procedure.[2]

In Kamin's blocking procedure the experimental treatment receives initial training during phase 1 with a single cue paired with the US (**A+**) . During the second phase the same cue is compounded with a second cue (**B**) and the compound is paired with the US (**AB+**). The interesting result is that learning about **B** in the experimental group is attenuated compared with a control group that only receives the phase 2 compound experience. Previous theories assumed that the amount of **B**-US training would determine the level of responding to **B** and since it is the same in the two groups these models predict the same level of learning. The RWM assumes that during phase 1, **A** acquires associative strength and this associative strength interferes with acquisition of associative strength by **B**. **B** is unable to acquire as much associative strength as it would otherwise in the control

[2]In spite of its many successes the model also has a number of failures (see Miller, Barnet & Cole, 1995). For example, the model predicts that neutral cues which are presented without reinforcement in an extinction procedure paired with a previously trained inhibitor will actually gain associative strength (Rescorla, 1971). This prediction has not been supported empirically (Baker, 1974; Zimmer-Hart & Rescorla, 1974). Regardless of any shortcomings of the model as a general theory of conditioning it is able to predict contingency sensitivity without requiring that the animals actually calculate conditional probabilities.

group that received no pre-training of **A**. Previously trained CSs or ones that are more contiguous or those that are more contingent with the US are thus able to block or interfere acquisition of associative strength by other less valid CSs. In support of this interpretation Wagner (1971) showed that this effect was a direct result of prior conditioning and not simply a function of the increased exposure to the US in the experimental group.

Supporting the idea that the amount of associative strength previously acquired by **A** influences how much associative strength can accrue to **B**, Wagner (1968) performed several experiments demonstrating that the conditioned response acquired by a CS (**B**) trained in compound with **A** was monotonically related to the amount of pre-training with (**A**). This finding supported the idea that the strength of **A**'s association with the US interfered with **B**'s association. Subsequent experiments showed that **A** could interfere with **B** even when the two trial types (**A+**, **AB+**) were presented intermixed during the same training phase (Rescorla, 1969c).

From the perspective of the Rescorla-Wagner model, Rescorla's contingency experiments can be seen as an example of associative competition between the experimental context and the discrete CS. The RWM can account for the reduced conditioning to a CS following zero contingency training by assuming that the increase in US presentations in the absence of the CS increases the context-US association which in turn interferes with the acquisition of an association by the discrete CS. The CS can only win this competition when p(US|CS)>p(US|-CS) or when **AX** is a better predictor of the US than **X** alone.

Interestingly, the RWM's predictions concerning which cue should acquire more associative strength and a normative calculation of contingency Δp make very similar predictions under some conditions. Assuming that the learning rate parameters for the presence and absence of the US (β_1 and β_0) are equal, the ordinal asymptotic predictions from RWM and Δp are identical (see Chapman & Robbins, 1990). However, this assumption, that $\beta_1 = \beta_0$ is not commonly accepted in the animal learning literature. The

more standard assumption, is that reinforcement has a greater impact on behaviour than nonreinforcement, and therefore the standard usage of the model requires that the learning rate parameter for reinforcement is set greater than that for nonreinforcement (i.e., $\beta_1 > \beta_0$). Even with this difference the model still assumes that CSs that are correlated with a US should acquire more associative strength than uncorrelated CSs (Wasserman, Elek, Chatlosh & Baker, 1992). With unequal betas Wasserman et al. have shown that the model predicts that CS-US contingencies with low US density, that have fewer pairings with the US, should acquire slightly more associative strength than the same overall CS-US contingency with more US experience. In spite of these departures from the normative Δp, this similarity between the computational description and an algorithm designed to account for empirical findings in Pavlovian conditioning provides an interesting example of the convergence between two independent theoretical traditions.

It is possible then that at the computational level animals have evolved a mechanism for assessing contingencies, but that the mechanism or algorithm that performs this task involves accumulation of associative strength using a competitive associative accumulator such as the RWM. The experiments described in this thesis will compare how well an associative model like RWM predicts the results of experiments designed to vary the CS-US Δp contingency.

Wagner's SOP theory

An alternate associative theory has been proposed by Wagner (1981). His standard operating procedures (SOP) model, like the RWM, also assumes that pairing a CS and US allows the CS to reduce the potential of the US to enter into association with other CSs. This theory is described by a set of rules which determine when memories for events can and can not enter into association with one another. According to the model, experiences with conditioning events activate any internal representations or memories of these events. Events become associated if they occur at the same time. This part of the theory is

consistent with the RWM. SOP also describes how associations can interfere with one another, and in this manner allows the model to account for contingency sensitivity. Rather than competing for associative strength, the model assumes that memory representations compete for activation in a limited capacity memory state.

Representations of events can be either inactive (I) or in one of two active memory states (A1, A2). The presentation of a CS itself brings its representation into A1. The activation of a representation solely via an association brings the representation into A2. These two states have different properties. Following repeated CS-US training, presentation of the CS will invoke both its representation into A1 and the representation of any other stimuli with which the CS has been paired only into A2. Excitatory associations between two representations can only occur when both elements of the association are in the A1 state[3]. When some events are in A1 and others are simultaneously in A2, inhibitory associations are formed. If pairing of a CS and US forms a strong excitatory association then subsequent presentation of the CS will invoke the representation of the US into A2 which will effectively limit its ability to enter into any new excitatory associations and may even result in inhibitory associations with new stimuli.

This set of rules predicts contingency sensitivity. Pairing a CS and US will lead to the development of an association between the two, and like the RWM any unsignaled US experience will form an association between the context and the US. During zero contingency training the CS acquires a weaker association with the US because the context can more readily activate the US representation into the A2 state on later CS-US trials. This result limits the CS-US association and its subsequent ability to elicit a response.

Pearce's (1987) model of Stimulus Generalization

---

[3] On the basis of data on retrospective processing Dickinson & Burke (1996) have suggested a modification of these rules which permits an associative account of retrospective learning effects by assuming that excitatory associations can also form when both representations are in A2.

A third example of how the notion of competing associations might be used to account for contingency sensitivity is Pearce's (1987; 1994) model of stimulus generalisation. The RWM and SOP both involve the assumption that animals represent stimuli as elements. In contrast, Pearce has argued that learning involves representations of entire stimulus configurations. If training involves an **AX+**, **X-** discrimination animals do not form a separate **A** and **X** association with the US, rather they form one association with a representation of the entire **AX** configuration and a second with the **X** representation. Furthermore the response elicited by any stimulus is a function of the excitation and/or inhibition that may have been acquired by that stimulus as well as any generalised excitation/inhibition from all representations of stimulus configurations which are similar to the stimulus.

Changes in levels of excitation can be described by equation (5). On any trial the amount of excitatory associative strength acquired by a stimulus configuration **AX** is a weighted function of the difference between the total amount of possible associative strength sustainable by the US (λ) and the sum of excitation and inhibition that the configuration has already acquired [$E_{AX} - I_{AX}$] either directly or through generalisation from other similar stimulus configurations. Therefore generalised associative strength can reduce the level of associative strength that a CS can acquire. Changes in inhibitory strength are calculated using a similar formula (see Pearce, 1987 for more details).

(5) $\Delta E_{AX}=B(\lambda-[E_{AX}-I_{AX}])$

Pearce proposes that the degree to which configuration **A** is similar to or excites the same units as configuration **AX** ($_{A}S_{AX}$) is represented by the degree to which two configurations share common elements. This approach may seem to bring us back to the elemental framework proposed in RWM, however, Pearce describes two levels of

description of a stimulus representation. A stimulus has a configural representation which mirrors the stimulus programmed by the experimenter. But also representations can be thought of as comprising multiple elements. These elements are either common to other configurations or unique. The degree to which **A** is similar to **AX** ($_{A}S_{AX}$) for example, can be captured by equation (6). The formula multiplies the proportion of elements in each stimulus that are common. The proportion of common elements in **A** and **AX** is a ratio of the common element (P_{com}) with the total elements for each stimulus ($P\Sigma A$ and $P\Sigma AX$ respectively).

$$(6) \qquad {}_{A}S_{AX} = \frac{P_{com}}{P\Sigma A} \cdot \frac{P_{com}}{P\Sigma AX}$$

The model has been applied with some success to account for a number of conditioning effects that seem beyond the scope of the simple nonconfigural version of RWM, including phenomenon such as negative patterning (Pearce & Redhead, 1993; Redhead & Pearce, 1995; Wilson & Pearce, 1989) which requires a configural solution. The model has recently been successfully applied to human contingency learning (Vallée-Tourangeau, Murphy, Drew & Baker, 1998; López, Shanks, Almáraz & Fernández, 1998). Like the previous two theories, Pearce's model accurately predicts that during an **AX**, **X** discrimination, learning about **A** will be directly related to its contingency with the US relative to **X**.

The RWM, SOP, and Pearce's Stimulus Generalization model involve different terminology and different associative mechanisms, but they each assume that sensitivity to **A**'s contingency following **AX**, **X** discrimination is directly determined by learning either an **X**-US (Rescorla-Wagner, 1972, Wagner, 1981) or **AX**-US (Pearce, 1987) association. With regard to the normative statistical measures of contingency such as Δp, associations are particulary useful psychological entities because they economically generate the total

frequencies of the event conjunctions required for the calculation of the one-way Δp contingency described in Figure 1. Because the predictions of levels of associative strength map quite easily onto Δp, associative models offer a parsimonious and plausible solution to a complex normative calculation.

Theories of Associative Comparison

Others have proposed that the decrease in conditioned responding observed following a zero CS-US correlation may reflect a performance failure rather than an effect on acquisition (Baker & Mercier, 1989; Durlach, 1989b; Miller & Matzel, 1988). These performance based theories assume that context-US associations are learned independently of CS-US associations and that these associations compete for control of behaviour (Bouton, 1993; Gibbon & Balsam, 1981; Jenkins, Barnes & Barrera, 1981; Miller & Schachtman, 1985). Jenkins et al. proposed the Waiting Time Hypothesis, Gibbon and Balsam, Scalar Expectancy Theory (SET) and Miller and Schachtman, the Comparator Hypothesis. To account for contingency effects comparator theories suggest that responses are the result of a process that compares the relative strengths of two associations, one between the CS and the US and one between the Context and the US.[4] This hypothesis is similar to the computational notion described by Δp.

The comparator models, like the competition models, are able to account for Rescorla's contingency learning experiments. Tone-shock experience establishes tone-shock and context-shock associations, while unsignaled shock experience contributes to the context-shock association. The relative strengths of the tone-US and context-US associations determines whether the CS or the context will control conditioned responding. One distinguishing feature of this type of theory is that animals keep complete memories of

---

[4] Gallistel (1990) has proposed a similar nonassociative comparator account of Pavlovian conditioning. His theory involves representations of the relative prevalence of the US. These representations are compared much like the associative comparator approach.

CS-US contiguity but sometimes fail to display knowledge if training included a relatively more contiguous and therefore more highly associated context. Evidence in support of this theory is derived primarily from experiments that involve post training manipulations of the associative strength of comparison stimuli. Sometimes, latent responding to the target stimulus is revealed. For example, if two stimuli are paired with a US in compound (**AB+**), theories of competition assume that this overshadowing treatment will result in weaker associations and therefore less conditioned responding to either element than if they had received training alone. In contrast, the comparator approach suggests that the association formed between **A** and the US during an overshadowing training procedure would be the same as that formed when training only involves **A**. In support of Miller's comparator model a test of **A** following overshadowing training will show relatively less responding to **A** because **B** will represent a stronger comparison stimulus for **A** than does the context which is the only comparison stimulus in **A+** training. Furthermore, unlike the competition models, any post training extinction of **B**, is predicted to influence the response controlled by **A** because of **B**'s reduced association. Miller & Schachtman (1985) found evidence for this type of post training inflation effect.

Post-training extinction effects have been demonstrated to reveal latent conditioned responding with a range of other conditioning phenomenon including: blocking (Balaz, Gustin, Cacheiro & Miller, 1982), relative validity (Cole, Denniston & Miller, 1996; Cole, Gunther & Miller, 1997), and conditioned inhibition (Schachtman, Brown, Gordon, Catterson & Miller, 1987). In contrast, theories of associative competition are unable to account for this type of post training inflation effect since **A**'s association should be untouched by new training of **B** (although see Van Hamme & Wasserman, 1994 and Dickinson & Burke, 1996).

In spite of their differences, both competition and comparator approaches treat conditioning as the learning of associations between events and argue that CS-US

contingency sensitivity is a direct consequence of interference by context associations on either the formation or the expression of CS-US associations. Both approaches predict that conditioned behaviour will often track the CS-US contingency making Δp a good computational description of learning behaviour even if Δp fails to capture the algorithm underlying this effect.

Nonassociative Theory

A third approach has its roots in normative or statistical models of human cognition (Cheng & Holyoak, 1995). From this perspective, learning involves acquiring global information about the causal structure of the environment. Organisms are argued to calculate the covariation between all possible causes and an event. Cheng and her colleagues (Cheng, 1997; Cheng & Novick, 1992; Waldmann & Holyoak, 1992) have developed a general nonassociative, normative account of causal reasoning that could apply to both human causal reasoning and to animal Pavlovian conditioning. One implication of this idea is that during conditioning conditioned responses directly reflect the result of multiple CS-US contingency calculations. One of the goals of these experiments is to assess the extent to which conditioned responses reflect this type of learning. If it was found that animals did display fine sensitivity to contingencies then it might provide an interesting framework for understanding a variety of conditioning phenomenon. Conditioning might reflect an adaptive computational system for inferring causal structure based on statistical regularity.

There is an important problem with any normative account that simply relies on calculations of Δp. In many cases the same Δp contingency between a CS and US results in very different conditioned responses because of the presence of alternate cues for the US (Mackintosh, 1983). There are many examples of these effects in conditioning which are called selective association phenomena. Selective associations are recruited as evidence that the Δp contingency alone is not a good theory of the mechanism underlying a conditioned

response (Mackintosh, 1977; Papini & Bitterman, 1990). However, it is possible that selective associations reflect a more complex computation than the simple unconditional Δp, rather the conditioning system may be sensitive the relative contingencies or conditional Δp (Shanks, 1995).

Rather than simply calculating Δp as described in Figure 1, the normative Probabilistic Contrast Model (Cheng, 1997; PCM) proposes that individual cause-effect relationships must be calculated conditional upon the presence and absence of other conditionalizing cues. In conditioning this might mean that individual CSs are compared with other CSs on the basis of their contingencies. Kamin's blocking effect (1969) is one example of a selective association effect in which a cue's unconditional contingency fails to predict the strength of the response that it acquires. As described earlier the effect involves repeated **A+** training followed by **AB+** training. The surprising result is that **B** acquires relatively weak control over a conditioned response compared to an **AB+** only control group. Recall that the Rescorla-Wagner model predicts a reduced association between **B** and the US in the blocking treatment due to reduced effectiveness of the US as a reinforcer following the **A+** training in phase 1. Several writers (e.g., Mackintosh, 1983) have suggested that both **B**'s contiguity and contingency are the same following these two treatments, however, this ignores the phase 1 training. If contingency is calculated conditional upon the presence of **A**, then **B** is a poor or at best ambiguous predictor of the US since the US is as likely in the presence of **B** as it is in the absence of **B** [$p(US|\mathbf{B}._{\mathbf{A}})=p(US|no\mathbf{B}._{\mathbf{A}})$] (Waldmann & Holyoak, 1992). Therefore, by choosing a specific comparison stimulus (or focal set) it is possible to calculate normatively whether or not a CS is a good causal candidate for the US. Blocking may be thus seen as normative and hence not a challenge to contingency theory.

Unlike the concept of the experimental context as described in the conditioning literature, a focal set for a CS, like a comparison stimulus (Miller & Schachtman, 1985) or

a temporal frame (Baker & Mercier ,1989), refers more generally to a specific stimulus or collection of stimuli that can act as a conditionalizing cue for the calculation of contingency. There is a great deal of evidence on context dependent learning in Pavlovian conditioning. During conditioning animals are able to keep track of separate experiences with the same CS in different physical contexts (Bouton, 1993). Cheng (1997) argues that keeping track of and responding on the basis of separated experiences is not a result of multiple associations but rather the result of multiple calculations of the conditional contingencies of each cause and its effects in different focal sets. Therefore, with respect to conditioning, within one focal set a CS may have a positive contingent relationship with a US, but within a second focal set a CS may have no relation or even a negative relationship. In logical terms, this involves calculating a contingency conditional upon the presence of a second cue or set of cues (Cheng & Novick, 1992).

As an example of how this idea works in human reasoning, consider how the importance of the conditionalizing context can greatly influence the attribution of the causal status of smoking as a cause of lung cancer. In general the overall, unconditional contingency in the general population between smoking and lung cancer has some positive value since the likelihood of smokers developing lung cancer is greater than that for non-smokers. In probabilistic terms the probability of lung cancer is greater in the sample of smokers (p(lung cancer|smokers)>(p(lung cancer|non smokers). However, one could consider the causal status of smoking conditional upon the presence or absence of other possible causal factors. For example, in a sample of chemical factory workers exposed to a chemical which when inhaled induces lung cancer, smoking may have little or minimal causal influence (p(lung cancer|smoking.$_{chemical}$)=(p(lung cancer|No Smoking.$_{chemical}$), because the chemical is a stronger causal factor. In spite of a moderate unconditional relationship with cancer in the general population, smoking may seem irrelevant as a cause for the factory workers. Depending upon the set of cues upon which you conditionalize

your calculations, estimates of the causal contingency between two variables can vary greatly (Shanks, 1995).

It is possible that during conditioning a similar causal induction computation occurs. In fact Cheng (1997)has referred to this possibility although there have been no tests of this theory with animals other than humans. Organisms may 'decide' which CS is the best predictor within different possible focal sets. By calculating contingencies in this way the animal "reasoner" may be able to make more precise estimates of causal effectiveness by eliminating redundant or irrelevant causes. It provides a framework for considering contextually mediated learning within causal learning.

In summary, an account of selective associations which uses a comparison of conditional contingencies, like that found in the causal induction literature, predicts that CSs should be ascribed different causal status on the basis of their relative relationships to the US. This idea predicts some of the selective association effects such as blocking (Cheng & Holyoak, 1995; Waldmann & Holyoak, 1992; Shanks, 1995).

Although no experiments have been conducted to test this theory directly there is some evidence that might be consistent with this hypothesis. Some experiments have been conducted which indicate that animals may be sensitive to relative relationships. Like Rescorla's work on contingency learning with a single CS, Wagner, Logan, Haberlandt & Price (1968) describe a simple but elegant design for contrasting the effects of CS-US contiguity and CS-US predictiveness with multiple CSs. They showed that when exposed to multiple stimuli signalling the occurrence of a US, animals learned to respond to a stimulus on the basis of its relative validity as a signal for the US. Table 2 presents the design of Wagner et al.'s experiments. Animals were trained with two compound cues (***AX***, ***BX***) each comprising a unique element (***A*** or ***B***) and a common element (***X***). In the True-Discrimination (TD) group one compound (***AX***) was always paired with the US and another (***BX***) was never paired with the US (***AX***$^{+}$, ***BX***$^{-}$). In the Pseudo-Discrimination

Table 2. Trial types and frequencies for the two compounds in the design of Wagner, Logan, Haberlandt and Price's (1968) experiment on relative validity. The conditional probabilities for the occurrence of the reinforcement (US) are those needed to calculate the Δp contingency.

	True-Discrimation	Pseudo-Discrimination
AX+	10	5
AX-	0	5
BX+	0	5
BX-	10	5
p(US|**A**)	1.0	.5
p(US|**A**)	0	.5
$\Delta p_{\mathbf{A}}$	1.0	0
p(US|**B**)	0	.5
p(US|-**B**)	1.0	.5
$\Delta p_{\mathbf{B}}$	-1.0	0
p(US|**X**)	.5	.5
p(US|-**X**)	undefined	undefined
$\Delta p_{\mathbf{X}}$	undefined	undefined

(PD) group half of the trials involving each compound were paired with the US ($AX^{+/-}$, $BX^{+/-}$). The common element ***X*** was paired with the US on 50% of the trials in both groups and thus, in isolation, was equally informative. In Wagner's terminology it was equally valid as a predictor of the US. Wagner et al. (1968) use validity to mean the likelihood that the CS will be paired with the US [p(US|CS)=0.5]. During an extinction test in which ***X*** was presented alone, the response elicited by ***X*** was much weaker in group TD than in the PD group. They reported this effect using an appetitive instrumental conditioning procedure with rats (Experiment 1) and using aversive Pavlovian conditioning procedures with rats and rabbits (Experiment 2 & 3 respectively). Subsequent research has replicated the effect using an appetitive Pavlovian conditioning procedure with rats (Murphy & Baker, 1996) and pigeons (Udell & Rescorla, 1979, Wasserman, 1974), using a conditioned taste aversion procedure (Luongo, 1976) and using a conditioned lick suppression paradigm with rats (Cole, Barnet and Miller, 1995; Cole, Denniston & Miller, 1996; Cole, Gunther and Miller, 1997). Recently a similar result has been reported in the human casual reasoning (Van Hamme & Wasserman, 1994).

One possible interpretation of these findings, that would be consistent with Cheng's probabilistic contrast approach is that the cue with the strongest positive contingency elicits a conditioned response in these treatments because it has the strongest relative conditional contingency. In the True-Discrimination treatment, **A** signals all of the US presentations. It is a perfect predictor while **X** is only a moderately accurate signal. In contrast, in the Pseudo-Discrimination group, **A** and **B** are uncorrelated with the US while **X** is still moderately correlated. While this design demonstrates that the absolute validity of a stimulus does not in of itself determine the strength of a conditioned response, it does suggest that animals might be able to learn about the relative validities of cues. Unfortunately, there were never any trials in the absence of **X** and therefore the exact contingency of **X** was undefined in these experiments. One of the goals of this thesis is to

develop a preparation that would allow presentation of multiple CSs each having a defined CS contingency. Regardless of the deficiency in Wagner et al.'s design, the experiment does raise the possibility that conditioning involves assessment of individual CS-US contingencies. Unfortunately, little experimental work has tested this position properly and therefore the whole question of contingency learning remains unanswered.

In summary, conditioning can be described at multiple levels of analysis. The CS-US contingency is a convenient description of the operational relationship guiding conditioned responding in Rescorla's experiments. It may also be an emergent phenomenon of a simple associative process reflecting a specialised adaptive neural function which deals with relative frequencies of discrete variables. The three theoretical approaches described provide a framework from which to consider contingency sensitivity. All three stress the notion that unsignaled US presentations should increase either the actual association between context and US (Rescorla & Wagner, 1972; Miller & Schachtman, 1985) or the relative causal role of all cues including the conditionalizing contexts (Cheng & Holyoak, 1995). In the single CS case it is interesting to note that they all agree that Rescorla's rats exposed to a CS-US contingency will behave in a manner which is virtually indistinguishable from that predicted on the basis of a normative calculation of Δp. They differ to the extent that they involve either a greater computational load (Cheng & Holyoak) or involve idiosyncratic rules of association formation (Wagner, 1981). Contingency learning is considered an important ability for the survival of all organisms and it is interesting to understand the psychological mechanisms behind it. Recently however, Papini & Bitterman (1990) have written a powerful critique of the role of contingency in Pavlovian conditioning in a flagship journal in the field of psychology, *Psychological Review*.

Challenges to Contingency Theory

Papini & Bitterman (1990) discourage the use of contingency as an explanation of

conditioning. They suggest that many of the findings which purport to demonstrate contingency learning are misrepresentations of a larger conditioning literature inconsistent with a contingency analysis. If conditioning is determined by establishing a CS-US contingency they argue, it follows that a CS-US contingency should be both necessary and sufficient for establishing a conditioned response. Papini and Bitterman cite a number of experiments which seem to violate these logical conditions. As a consequence they call for the abandonment of Δp and other contingency metrics for both computational or algorithmic descriptions of conditioning.

Before considering their evidence there are a number of reasons why their argument may not be valid *a priori*. The first is that they suggest that only one factor can be responsible for conditioned behaviour. This, in of itself, is an empirical question and not a logical necessity. One can easily provide support for the idea that multiple factors may facilitate or hinder the development of a conditioned response that emerges following CS-US contingency training. For example, there is evidence with both Pavlovian and Instrumental conditioning that relations between some stimuli and USs or response reinforcers are easier to learn than others (Breland & Breland, 1965; Garcia & Koelling, 1966; Sevenster, 1973). Sevenster describes how the male stickleback biting response is associated with fighting. Biting can be reinforced with access to a male opponent, but not by access to a sexually active female. Sevenster argues that the stimulus and reward must be motivationally compatible for an association to form. Similar findings have been reported in simple Pavlovian conditioning preparations (e.g., Rizley & Rescorla, 1972). Using a Pavlovian taste aversion preparation, Garcia & Koelling demonstrated that illness and flavour are more readily associated than are illness with visual or auditory stimuli. This has been interpreted as consistent with the idea of an initial bias for certain associations, a biological preparedness for learning. Therefore, this factor may mediate the strength of the conditioned response independent of changes to CS-US contingency. These findings do

not exclude the possibility that Pavlovian conditioning involves sensitivity to CS-US contingency.

A second objection to their argument is whether the experimental evidence they report actually unambiguously falls on either side of the contingency debate. The bulk of their evidence is derived from studies in which experimenters claim to have programmed a zero CS-US contingency but still find evidence for excitatory conditioning to the CS (e.g., Benedict & Ayres, 1972; Kremer, 1971; Quinsey, 1971). The second part of their challenge describes experiments that test whether CS-US contingency is sufficient for a conditioned response. In these studies researchers have programmed positive CS-US contingencies and shown that the individual unconditional CS-US contingency does not seem to determine the strength of the conditioned responding (Jenkins & Shattuck, 1981; Kamin, 1969; Wagner, Logan, Haberlandt & Price, 1968). I will deal with these two challenges in turn and describe the development of a novel preparation for studying contingency learning.

Necessity

Papini & Bitterman (1990) argue that if a CS-US contingency is necessary for conditioning then random or uncorrelated training, in which the CS does not signal a change in the likelihood of US delivery from that experienced in the absence of the CS, should be unable to support a conditioned response. Rescorla's (1968) experiments showed that equating the likelihood of US delivery in the presence and absence of the CS interfered with conditioned responding. Subsequent studies have also found this result (Ayres, Benedict & Witcher, 1975; Benedict & Ayres, 1972; Keller, Ayres & Mahoney, 1977; Kremer, 1971). The question then is whether this reduced responding can be interpreted as a response generated by learning the zero contingency. One might predict that the zero contingency should produce zero behaviour, however this is again an empirical question. Zero contingency training appears to be more likely to result in excitatory

conditioning with relatively few training trials particularly if this training starts off with relatively more 'random' pairings of the CS and US (Benedict & Ayres). This suggests that at least part of the response generated by the zero contingency is the result of a possibly misperceived positive contingency by the animal subjects. It also points to a significant problem with this research. The design of these experiments requires accepting the hypothesis that the actual contingency perceived by the animals is the same as the zero contingency operationalized by the experimenter.

That there is room for ambiguity in many conventional preparations is highlighted by an example. If a rat receives 10 shocks during 20 one minute Tone presentations and 10 shocks during 20 one-minute inter-trial-intervals (ITI) it might be argued that the likelihood of receiving a shock during the tone is equal to the likelihood of shock during the ITI. This type of assumption is behind many of the experimental procedures which have sought to test animals' sensitivity to CS-US contingency (ex., Rescorla, 1968; Hallam, Grahame & Miller, 1992). However, if we try to map our normative Δp metric on to these experimental preparations it is not entirely clear what CS-US contingency the rats are being exposed to. The frequencies of occurrence of the CS, and ITI are ambiguous and involve specific assumptions about how to partition experiences which are extended in time. If time is partitioned into units which are similar to the one minute time intervals of the Tone then the relationship between the Tone and shock in the chamber might be zero since the $p(US|Tone)=10/20=0.5$, $p(US|\text{-}Tone)=10/20=0.5$ and thus $\Delta p=0.5-0.5=0$. However, even if one accepts that animals treat one minute of CS time and one minute of ITI time equivalently, and this assumption is required for this calculation to be an accurate reflection of the relationship perceived by the subjects, animals are not permitted to include any other 'no CS' experience without changing the perceived contingency. Animals might include other experience in their estimation of this frequency. For example, they may include time in their home cage. If so, then the animals' 'estimate' of the likelihood of the shock in the

absence of the tone would necessarily decrease, and the perceived contingency between CS and US would increase. For example, if 10 minutes of a rat's life outside the chamber is included then the resulting contingency between CS and US will necessarily increase [p(US|noCS)=10/30=.33 and Δp=0.5-.33=.17]. Therefore, simply changing the events which might be recruited converts a putative zero contingency into a positive contingency. This positive contingency might be sufficient to generate the excitatory conditioned response found following supposed zero CS-US contingency exposure. Similar accounts of excitatory conditioning following zero contingency training have been proposed by Quinsey (1971) and others (Baker, 1977; Baker, Murphy & Vallée-Tourangeau, 1996).

For this account to be plausible it must be shown that the conditioned response is influenced by events occurring outside the conditioning chamber or at least outside of the conditioning sessions. As partial support for this idea it is common practice for experimenters to expose animal subjects to the conditioning chamber in the absence of the CS or the US or both before any actual experimental training begins. Habituation sessions prior to conditioning training reduce unconditioned responses to the context, a similar exposure is often used with CSs. Animal subjects are routinely trained to retrieve food pellets from food hoppers prior to CS training (magazine training) or habituate to novel contexts before training with shock (e.g., Miller, Hallam, Hong & Dufore, 1991). Any of these treatments might be expected to change the CS-US contingency, but more importantly demonstrate the extent to which transfer of experience has been accepted as being effective for modifying conditioned responding.

Baker (1977) showed how context experience outside of the CS training exposure might influence the strength of CR to that CS. He showed using the conditioned emotional response procedure that the level of suppression of lever pressing to a CS was determined not only by experiences with that CS in the conditioning chamber but also by other experience with shock received during other sessions in the absence of the CS. Baker

trained rats with either a negative or zero contingency between a noise CS and shock. The negative contingency group never received the CS and US in the same session. Only by combining experience from the two sessions could the CS be interpreted as an inhibitor. The results were consistent with the interpretation that the CS had become an inhibitor for US occurrence. Miller and his colleagues have also conducted experiments designed to show that the strength of a conditioned response to a CS can be increased simply by post-training extinction of the contextual cues (Miller & Shactman, 1985; Miller, Barnet & Grahame, 1995).Thus it is quite possible and, in fact, likely that some programmed zero CS-US contingencies are perceived as positive because of the influence of experience outside of the CS-US training. Any experiment designed to demonstrate learning about CS-US contingencies requires control of both CS presence and importantly CS absence. One of the primary goals of the experiments described in this thesis was to develop a preparation that permitted a less ambiguous interpretation of the CS-US contingency, by defining CS absence

Papini and Bitterman's (1990) also assume that zero contingency training should produce no change in behaviour. A zero contingency may or may not result in a change in responding and whether it does is an empirical question. From the perspective of contingency theory there is no reason why a zero contingency should be behaviourally silent. One can easily argue that knowing that two events are uncorrelated is not the same as knowing nothing about their relationship. It is sufficient to show that negative, zero and positive contingencies produce ordinally different conditioned response strengths to claim that animals are sensitive to contingencies. Therefore even though an experimenter may find excitatory conditioning following zero contingency training, as long as the response is stronger than that produced by a negative contingency and weaker than that produced by an excitatory contingency then the animals could be argued to be sensitive to changes in CS-US contingency.The question of zero responding may have important implications for

interpreting conditioning findings. For example, excitatory conditioning following zero contingency training may rule it out as a proper control procedure for tests of associative learning (Rescorla, 1967) however, it does not rule out the idea that animals are learning CS-US contingencies.

A number of other possible factors may be responsible for changes in behaviour other than CS-US contingency following CS-US training. Simple experience with any of the CS, US or context regardless of any relationships between the events could be responsible for eliciting unconditioned changes in behaviour to the CS. Zero contingency CS-US training may also increase the associative status of the context but the CS may act as an appropriate substrate for responding controlled by context-US associations (e.g., Gewirtz, Brandon & Wagner, 1998). Alternatively, animals may acquire context-US associations and second order CS-context associations. Either of these types of association may facilitate CS elicited responses via the context. Any of these possibilities might be expected to produce some level of behaviour, but none of them rule out animal sensitivity to Δp contingencies. The argument that zero CS-US contingency training should leave an animal in a state as if it had never been exposed to either event (i.e., as if it has zero knowledge about the two events) is of relevance to some associative theories (ie., Rescorla & Wagner, 1972) but this is not a litmus test for contingency theory.

Testing whether it is the contingency which controls the conditioned response might be facilitated by using a more explicitly programmed CS-US relationship. This test might have the desired effect of reducing the possibility that the conditioned response emitted by the animal will be controlled by various experiences beyond direct experimental control. One way of doing this might be to programme the CS-US contingency relative to a second CS rather than defining it relative to the temporally extended static contextual cues. The contextual cues are normally considered to be everything in the chamber except the CS and US and, in the previous examples, may include other life experience. Choosing a

comparable trial context CS may bring the contingency under better experimental control. The experiments described in this thesis all define CS-US contingencies by using a discrete CS in the role of the trial context.

Sufficiency

In addition to the criticism outlined above, Papini and Bitterman (1990) pointed out that there are many examples in which a programmed positive CS-US contingency is clearly an insufficient criterion for the formation of a conditioned response (see also Mackintosh, 1983 pp.182). One category of experiment consistent with this claim comes from the selective association experiments. In some cases more valid predictors are able to reduce the strength of the response generated by relatively less valid CSs (Kamin, 1969; Rescorla, 1972; Wagner, 1969; Wagner, Logan, Haberlandt & Price, 1968) or contextual cues (Murphy, McDonald & Baker, 1998). In these experiments CS-US contingency alone does not predict whether a CS acquires a conditioned response, however many selective association experiments might be consistent with a contingency analysis if it is assumed that animals are sensitive to the relative contingencies and therefore sometimes fail to learn CS-US contingencies when training includes more contingent predictors.

If one assumes that selective associations reflect sensitivity to relative CS-US contingencies as described by Cheng (1997), then these effects are not failures of a contingency analysis but are rather consistent with this analysis. The previous discussion of Cheng & Holyoak's normative model assumes that assessing relative validity is a normative solution for discovering the true cause of an event. If it could be demonstrated that during selective association experiments animals were responding to stimuli as if they were learning to respond to the most contingent CS then selective association data would provide support for the general contingency sensitivity hypothesis.

A related set of findings which also call into question the sufficiency of contingency to account for conditioned behaviour, involves only a single discrete cue. Jenkins and

Shattuck (1981) demonstrated that conditioning to a CS can be attenuated by *increasing* the amount of CS-US experience (see also Jenkins & Lambos, 1983; Jenkins, Barnes & Barrera, 1981; Rescorla, 1968; Quinsey, 1971). Adding extra shock (US) presentations during the inter-trial-interval degraded the conditioned response controlled by a CS which was paired with shock. However, this reduction occurred both with unsignaled USs and with USs which were signalled by the same CS. In this treatment the increased CS-US pairings should not have resulted in a decrease in CS-US contingency. However, like Kamin's blocking procedure, increasing the US density or base rate does increase the relative contiguity and contingency of the contextual cues. It is possible that the relative validity of the context can mediate the strength of the conditioned response to the CS (Durlach, 1989a; Goddard & Jenkins, 1987). There is good empirical evidence that this effect might be a function of increased associative competition from the contextual cues (Baker, Singh & Bindra, 1985; Tomie, 1976). If it could be demonstrated that the increased validity of the contextual cues could reduce responding to a CS then these findings would be supportive of the relative contingency hypothesis.

Both responses to Papini & Bitterman (1990) that I have described involve a common underlying issue, namely that the contextual cues play a crucial role in mediating the conditioned response to a CS. Any test of whether CS-US contingency is important for determining the strength of a CR may require simultaneous assessment of the conditioned responses elicited by the CS and the context. The method proposed here to accomplish this involves using a discrete cue in the role of the trial context.

Rationale for the Present Experiments

A standard solution for designing an experiment that presents different CS-US contingencies and tests an animal's sensitivity to them has been to vary the rates of reinforcement per unit of time in the presence and absence of a CS and interpret this ratio as the ratio of frequencies required for Δp calculations (ex., Miller, Hallam & Grahame, 1990;

Rescorla, 1968). For example, in Figure 2 the standard procedure for defining a CS-US contingency is shown. In the first time line, four ten second CS presentations are randomly programmed to occur, two are paired with a US and two are not. The likelihood of the US during the CS is .5. There are also 2 US presentations in the absence of the CS. For these purposes the standard assumption is that animal subjects have both a timing and counting device (e.g., Gallistel, 1990) that accumulate the number of US experiences per unit of time during both presence and absence of the US. The actual time unit the animal might employ is unknown but, as long as the animal uses the same unit in different positive and zero contingency treatments during both CS and no CS experience, the different programmed contingencies should still be discriminable. Assume that the animal divides the ITI into bins that are equivalent in length to the CS (10 seconds). Using this method two of the context bins are paired with the US (labelled C) and 10 of the bins are not paired with the US (labelled D). Calculating the contingency this way the CS is predicted to be a moderately positive predictor of the US.

Another method of programming a CS-US contingency that might avoid any potential problems might be to reevaluate how the CS and the context cues are presented. If the trial context was a discrete cue of the same duration as the CS, then the likelihood that the absence of the CS would include other events might be reduced. Consider the second time line shown in Figure 2. In addition to the standard CS_1 and US presentation a second CS_2 represents the trial context and is always presented in conjunction with CS_1 and in addition, CS_2 sometimes occurs in the absence of the CS_1. While CS_1's contingency relative to the background contextual cues is the same as it was in the first example, its contingency with the US relative to cue CS_2 is much reduced because there are fewer no CS_1 -noUS events. Furthermore because CS_2 serves to mark the no CS_1 experience, it might be more salient and therefore easier to learn about and less likely to be combined with other experience.

Figure 2. The standard and proposed method of calculating CS-US contingency in Pavlovian conditioning. The standard (Top) assumes that noCS experience comes from subdividing time in the context into the four types of events that describe the contingency table (A,B,C,D). The proposed method (Bottom) involves using a discrete trial context (**X**) stimulus to define the no CS experience.

	US	-US
CS	A	B
-CS	C	D

Standard Procedure for defining CS-US contingency

16 bins of equal context experience

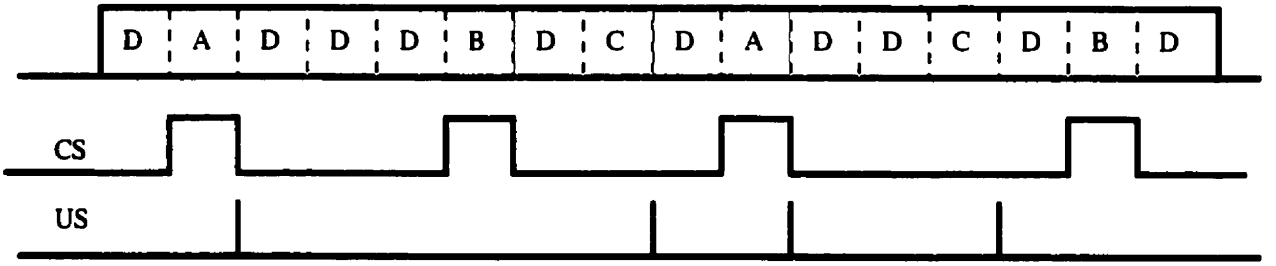

$$\Delta p = (p(US|CS) - p(US|-CS)$$
$$= [(2/2+2) - (2/2+10)]$$
$$= .5 - .17$$
$$= .33$$

Proposed Method of Defining CS-US contingency

Background contextual cues

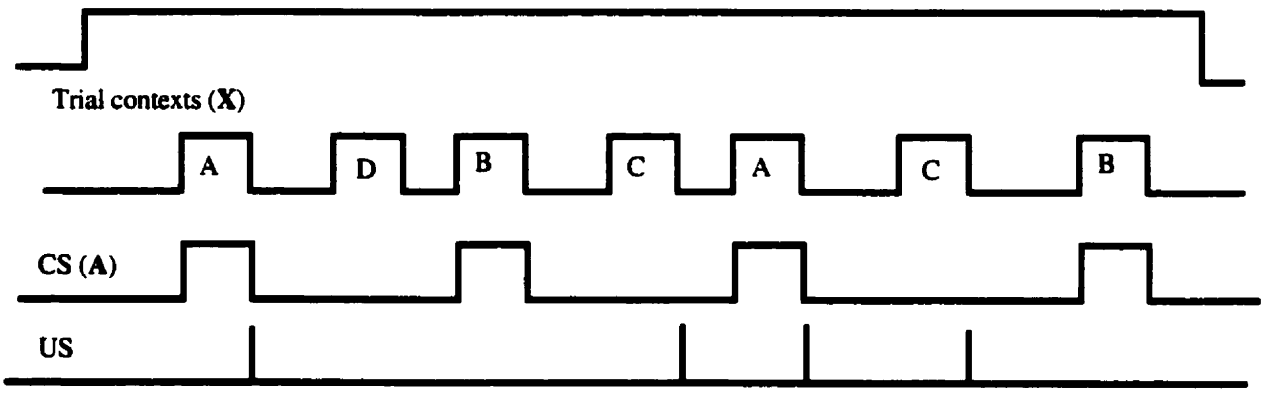

$$\Delta p = (p(US|CS) - p(US|-CS)$$
$$= [(2/2+2) - (2/2+1)]$$
$$= .5 - .67$$
$$= -.17$$

The experiments in this thesis use this method of evaluating the contingency learning hypothesis. CS-US contingency sensitivity was evaluated by training rats with either positive, negative or zero contingency CS-US relationships. The CS was always 10 seconds in length and it was paired with a second trial marker CS. This Trial Context was always presented for the same duration as the CS. Therefore, this design involves CS and US contingencies defined relative to a second CS.

In addition to a possibly less ambiguous contingency, an important feature of this design is that it allows direct and sensitive measure of conditioned responses acquired by the trial context. The interpretation of Rescorla's (1968) contingency experiments hinges on the assumption that the contextual cues did not acquire strong control over conditioned freezing (Papini & Bitterman, 1990). Decreasing CS-US contingency may not have reduced the CR to the CS as Rescorla proposed. Rather, the increases in shock frequency in the absence of the CS may have simply served to increase the associative strength and conditioned behaviour controlled by the context (i.e., freezing; see also section 1.4.2 on comparator theories, Papini & Bitterman). If this was the case then the suppression ratio used as the measure of conditioning may simply reflect changes in conditioning during the pre-CS interval, not changes in conditioning to the CS (see also Hurwitz & Davis, 1983 for a discussion of the problem with relative measures of conditioning). According to this analysis it is not clear whether reducing CS-US contingency by increasing the p(US|noCS) simply increases the response suppression of the context or actually increases lever pressing to the CS.

In his defense, the tests of conditioning in Rescorla's experiments were done after two sessions designed to extinguish fear of the contextual cues. Any subsequent measures of responding to the CS assume that the contextual cues were sufficiently extinguished. However, if the context was not extinguished and differences in contextual conditioning produced by the training were maintained, then the suppression ratios would be difficult to

interpret. There is data to suggest that Rescorla's extinction procedure may not have been sufficient to extinguish the context because extinction itself is context specific. Extinction of a CS in one context usually does not transfer to a second context (Bouton & Bolles, 1979). In Rescorla's contingency experiments the context may not simply be the set of static spatial cues in the conditioning chamber. There is research to show that other cues including the presence of the CS may serve as part of the general training context (Hall, 1991; Bonardi & Hall, 1996). In Rescorla's experiment it is possible that the response generated by the cues of the conditioning chamber may have become conditional upon the presence of the CS. In other words, extinction itself may have been context specific and the presence of the CS(even the US may serve as a conditioning context) defined the learning context and this context was different during the extinction sessions than during training in Rescorla's experiments. One way to avoid these problems is to measure any contextual conditioning directly during training. The present experiments were designed to measure both conditioned responses to the CS and the trial context against which the CS-US contingency was defined.

Predictions for Contextual conditioning

In addition to testing whether Pavlovian conditioning involves learning about the CS-US contingency as defined by Δp, these experiments test predictions of associative theories. Most associative theories of conditioning predict that as the contingency between CS and US is diminished and the strength of the conditioned response generated by the CS decreases there should be a complementary increase in the conditioned response controlled by the context. Odling-Smee (1975a, 1975b) found evidence for increasing contextual conditioning following deceases in CS-US contingency. He showed using an aversive conditioning procedure with rats that an avoidance response to shock was controlled by the cues of the experimental context. Furthermore, this control increased as the contingency between a discrete tone CS and the shock decreased.

There is other evidence consistent with this reciprocal conditioning prediction. One experiment using the logic of Kamin's blocking procedure demonstrated that unsignaled US preexposure can block acquisition of a conditioned response by discrete CSs (Tomie, 1976). Tomie argued that the unsignaled US presentations increased the associative value of the context which interfered with acquisition of associative strength by the CS. Consistent with the blocking interpretation of the interference, signalling the US preexposure with a second CS which should reduce contextual conditioning, reduces the interference effect. Baker, Mercier, Gabel and Baker (1981) demonstrated this finding in series of conditioned emotional response experiments with rats (although they also found evidence that contradicted the context conditioning hypothesis). Other experiments have also used this logic and demonstrated that reducing contextual conditioning can enhance CS conditioning. Extinguishing the contextual cues (Dweck & Wagner, 1970), preexposing the context (Wagner), signalling the ITI USs with a second CS (Durlach, 1983, 1989a; Goddard & Jenkins, 1987) all reduce the apparent attenuation in CS conditioning.

However, rather than indirectly demonstrating the potentially disruptive influence of contextual cues on conditioning by reducing the influence of the context with various training manipulations, the present design can directly assess contextual conditioning as a function of changes in CS contingency. There have been other attempts to directly measure contextual conditioning during changes to CS contingency. Like the assumptions used in the standard definition of CS-US contingency they have used contextual cues which were of longer duration than the CS (Baker, Singh & Bindra, 1985; Tanner, Rawlins and Mellanby, 1987) with mixed results. While Baker et al. found some evidence for reciprocal conditioning between CS and context, Tanner et al found evidence of reciprocal conditioning but only limited sensitivity to different partial reinforcement schedules ($p(US|CS)$ and no evidence of sensitivity to CS-US contingency as defined by Δp. Neither experiment systematically examined the role of reciprocal conditioning during contingency

training.

Single CS Contingency Learning Experiments

Experiment 1

The first experiment described here was designed to compare the conditioned response to a CS (**A**) with either a positive or zero contingency with the US using a procedure which defined the CS's contingency relative to a trial context (**X**).The likelihood of US delivery in the presence of **A** on **AX** trials, and in the absence of **A**, during **X** trials was varied. The discrimination between **AX** and **X** has much in common with a number of experiments which have studied feature positive and feature negative discriminations (e.g.Holland, 1992; Wilson & Pearce, 1989). In these experiments a feature stimulus (**A**) is sometimes paired with a second stimulus (**X**). **AX** signals something different from **X** by itself. For example, in a feature positive discrimination ($\mathbf{AX}^{+}$,$\mathbf{X}^{-}$) the feature, **A**, signals the delivery of the US and **X**, the absence of the US. Animals will readily learn to respond in the presence of **A** but withhold responding in the presence of **X**. In contrast, during a feature negative discrimination ($\mathbf{AX}^{-}$, $\mathbf{X}^{+}$), the feature signals the absence of US delivery. Animals learn to respond on **X** trials but not on **AX** trials, and **A** usually becomes an inhibitor.

From the perspective of contingency theory, feature positive and feature negative discriminations can be thought of as two boundary conditions in which **A** is either perfectly positively correlated with the US or perfectly negatively correlated. In the feature positive discrimination the **A**-US contingency is perfectly positive since **A** perfectly signals the presence of the US [p(US|A=1 and p(US|-A)=0 Δp_A=1]. In the feature negative case the **A**-US contingency relative to **X** is -1 as the probability of reinforcement in the presence of **A** is 0 and in the absence of **A** is 1 (0-1=-1). The finding of excitatory conditioning in the first instance and inhibitory conditioning of **A** in the latter is consistent with the contingency hypothesis. There are a number of other possible interpretations. Rather than assuming that

animals calculate Δp, these results are also consistent with the hypothesis that learning is dependent upon sensitivity to CS-US contiguity. It is impossible with this design to differentiate these two alternatives. The experiments in this dissertation attempt to identify whether contingency or contiguity alone can account for the strength of the CR in this type of discrimination procedure. Experiments 1 and 2 compare the strength of the conditioned response to a light CS (**A**) when it was trained with either an imperfect moderate positive contingency (Δp_A=.50) or a zero contingency (Δp_A=0) with the US. In order to control for differences in contiguity unrelated to differences in contingency three different US density treatments were included.

The two contingency conditions were crossed with three levels of US density. The three density levels were chosen to control for differences in CS-US pairings and US density (Gibbon, Locurto & Terrace, 1975). One possibility is that conditioned responding is controlled by the number of CS-US pairings and so two of the positive contingency groups have the same number of CS-US pairings as two zero contingency groups. Alternatively, Jenkins & Shattuck (1981) have suggested that increasing US experience independent of changes in CS-US contingency can reduce the strength of the conditioned response. In order to control for US exposure in the present experiment three groups were given positive contingency (ΔP_A=.50) training with different levels of US density (High, Medium and Low US density; **H.5**, **M.5**, **L.5**).

Table 3 shows how the contingencies were defined for the 6 treatments. For example, in the **H.5** group all the **AX** trials were paired with the US and 50% of the **X** trials were paired with the US. The overall Δp_A contingency was 0.50 because p(US|**AX**)-p(US|**X**)=1-0.50=0.50 and the average probability of the US or US density [p(US)], was .75. The moderate density treatment **M.5** received the same contingency, but received fewer **A**-US pairings and fewer US overall. Although group **M.5** received the same number of **AX** and **X** trials only 75% and 25% of the **AX** and **X** trials respectively were

Table 3. Trial types, frequencies, conditional probabilities and Δp contingency calculation for cue **A** in Experiment 1. The six treatments cross three levels of US density (High, Medium and Low) with two levels of Δp_A contingency (.50 and 0).

	Positive Contingency (Δp=.5)			Zero Contingency(Δp=0)		
US Density	High (**H.5**)	Medium (**M.5**)	Low (**L.5**)	High (**H0**)	Medium (**M0**)	Low (**L0**)
AX-->US	32	24	16	24	16	8
AX-->no US	0	8	16	8	16	24
X-->US	16	8	0	24	16	8
X--no US	16	24	32	8	16	24
Total US	48	32	16	48	32	16
Total no US	16	32	48	16	32	48
p(US|A)	1	.75	.5	.75	.5	.25
p(US|-A)	.5	.25	0	.75	.5	.25
Δp	.5	.5	.5	0	0	0

paired the US. Again Δp_A is still 0.50 (0.75-0.25=0.50) but the outcome density was .5. For the low density group only 50% **AX** trials were paired with the US (0.50-0.0=0.50) and they received the lowest US probability [p(US)=.25]. In contrast three groups were given zero contingency training with the same US density levels as in the positive contingency treatments (**H0**, **M0**, **L0**). In the high density condition **H0**, 75% of **AX** and **X** trials were paired with the US. In the moderate density condition **M0**, this was reduced to 50% and in the low density **L0** it was only 25%. In spite of these US density differences each group had a zero contingency between **A** and the US.

In terms of control groups the level of responding in group **H.5** can be compared with **H0** to compare the effects of US presentations since they received the same number of US presentations (i.e., US density). **M.5** has both a CS-US pairing control (i.e., **H0**) and an overall density control (i.e., **M0**) that received the same overall level of US exposure. Finally, group **L.5** also has both a CS-US pairing control (**M0**) and an overall US density control (**L0**).

In summary, if the *A*-US contingency determines the strength of a conditioned response, then all three positive contingency groups should demonstrate strong conditioned responding to cue **A** while zero contingency training should favour conditioning of **X**. However, if the decrease in the conditioned response to **A** is a result of increased experience with the US, or differences in **A**-US pairings, conditioning to **A** will be a function of US density or **A**-US pairing. If the Δp_A contingency controls conditioned responding then **A** should control a stronger response in all three positive contingency groups with little difference between them. In contrast, the context **X** was equally valid in each pair of High, Medium and Low density conditions with its contingency increasing with increases in US density. These predictions were tested by assessing the conditioned response elicited by the nominal CS (**A**) and the trial context (**X**). The actual stimuli used in the roles of **A** and **X** were chosen for their ability to elicit different conditioned responses.

In this experiment, **A** was always a light that when conditioned, elicited tray entry behaviour (Kaye & Pearce, 1984), while **X** was the entrance of a retractable lever that elicited lever pressing as a CR (Peterson, Ackil, Frommer & Hearst, 1972). Changes in either of these two behaviours were predicted to match changes in the CS-US contingency.

Theories of classical conditioning generally predict that positive contingency **A**-US training should produce a stronger conditioned response to **A** than zero contingency training. Associative theories differ with respect to how the effect of US density is predicted to influence this sensitivity. They also differ in their predictions in contextual conditioning. The normative contingency account described by Cheng (1997) predicts that as US density increases, the response generated by **X** in general will increase because responding to **X** will be determined by the p(US|**X**)[5]. An animal that was sensitive to the conditional contingencies between **A**, **X** and the US might be expected to show more conditioned responding to **X** simply as the likelihood of the US being delivered on an **X** trial increased.

An associative model also predicts this pattern (Rescorla-Wagner, 1972) Because of its mathematical formalization, the Rescorla-Wagner model (RWM) can be used to generate very specific predictions about the relative associative strengths of the six groups in Experiment 1. Simulations of the six treatments in Experiment 1 are shown in Figure 3. The formula for updating associative strength (Equation 2) was used to calculate the predicted associative values of the two cues **A** and **X** following the relative frequency of trial types used in this experiment (see Table 3). The choice of values for the free parameters effecting the rate and level of learning were those recommended by Rescorla

---

[5] Cheng's(1997) Power PC theory also predicts that changes in US density will influence the perceived causal strength of **A** independent of any change in Δp. She proposes that contingency is a function of both an assessment of Δp and an estimate of causal power (see Vallée-Tourangeau, Murphy, Drew & Baker, 1998 for a failure to confirm the theory in human contingency judgements). For the present discussion, I do not refer to causal power.

Figure 3. Predictions of the cumulative final associative strength for cues **A** and the trial context **X** in Experiment 1 using the Rescorla-Wagner (1972) model. Cue **A** was presented with either a positive or zero CS-US contingency at each of three US density levels (High, Medium and Low). Density refers to the overall probability that the US would be delivered on one of the 64 trials [p(US)=.75, 50 and .25 respectively].

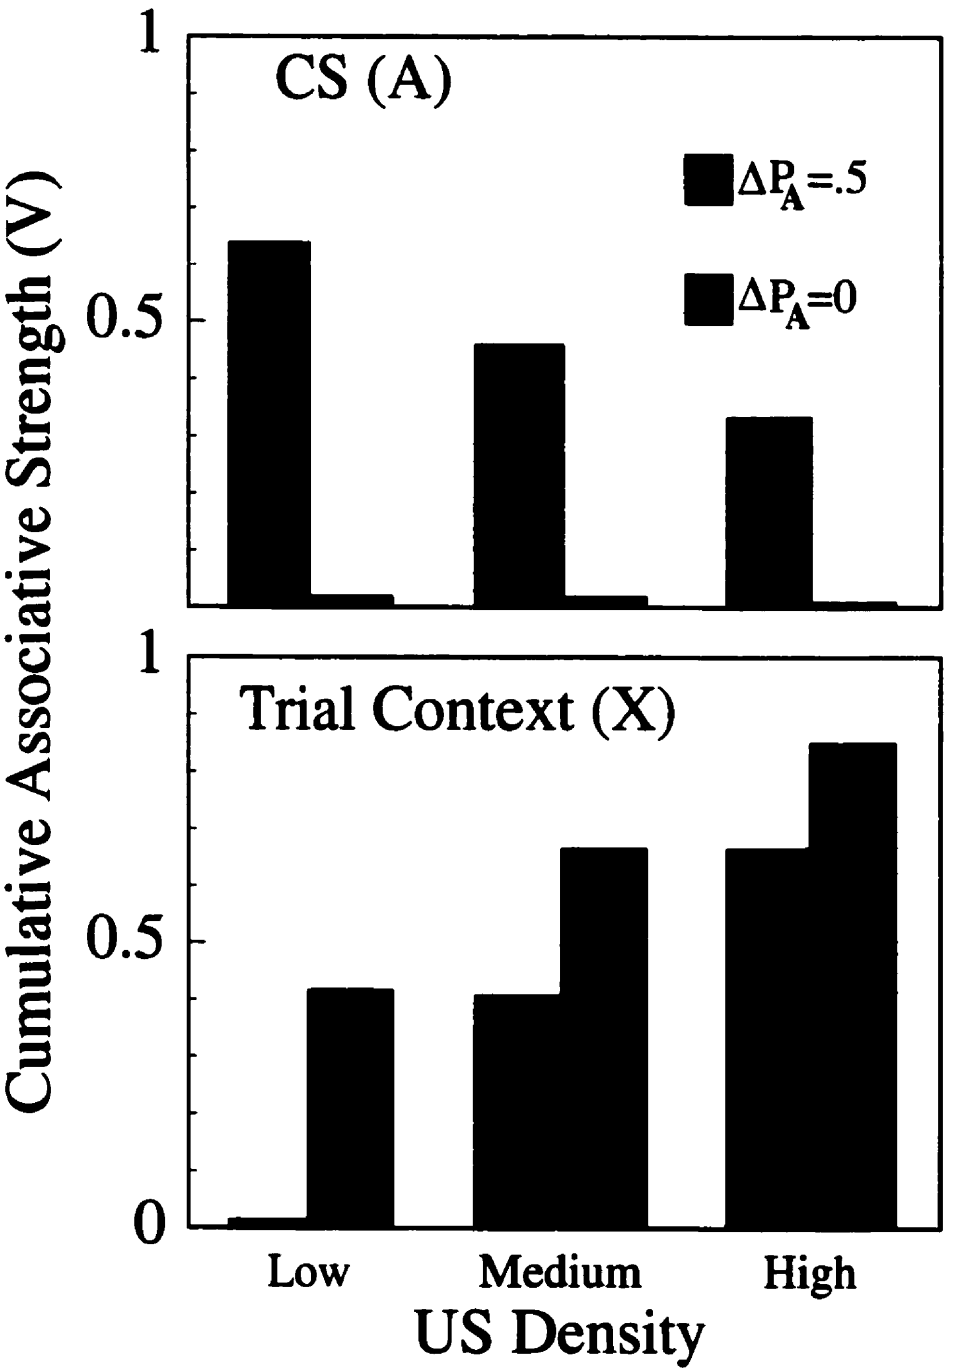
1
CS (A)
ΔP_A=.5
ΔP_A=0
0.5
1
Trial Context (X)
0.5
0
Low
Medium
High
US Density
Cumulative Associative Strength (V)

-Wagner (1972, p. 88). The parameter coding for the salience or learning rate for the two stimuli was set equal ($\alpha_A=\alpha_X=.5$). The asymptotic value (λ) for US trials and nonUS trials was 1 and 0 respectively. Following convention described by Rescorla and Wagner, the parameter associated with reinforcement was set at twice that associated with nonreinforcement ($\beta_1=.1$, $\beta_0=.05$). All simulations involved 5 epochs of 500 trials. The average final predicted associative strengths for *A* and *X* in the six treatments are shown in the two panels of Figure 3.

The general pattern of predictions for cue **A** matches those derived from the normative measure of contingency Δp. The associative strength of *A* in each of the positive contingency groups is higher than in the zero contingency groups. The model also predicts that the context cue *X* should acquire more associative strength when *A* has a zero contingency with the US. This increase occurs since **A** is less able to retain its associative strength following zero contingency training because there are relatively more trials in which **A** is extinguished allowing **X** to gain associative strength. In addition, unlike the normative Δp predictions, RWM predicts that the strength acquired by *X* and **A** will also be sensitive to the US density. However, the effect of US density is the opposite on the two cues, **X** should acquire more associative strength with increases in US density, **A** should acquire less association with increases in US density. For comparison, Pearce's (1987) model predicts that US density should have the same effect on both **A** and **X**. His model predicts that increases in US density should increase the associative strength to both **A** and **X**.

In summary, an associative model of conditioning (i.e., Rescorla & Wagner, 1972) predicts that cues will acquire associative strength on the basis of their relative validity as predictors of the US. This prediction is also consistent with general predictions derived from a normative estimate of the one-way contingency between two binary events, Δp (Allan, 1980).

Method

Subjects: Thirty six naive male Wistar rats were obtained from a local breeder (Charles River, Lacolle QC). They weighed between 275-300 g and were reduced to 80% of their free feeding body weight and maintained at this weight for the duration of the experiment. They were housed individually in wire cages in a room lit on a 12 hour light-12 dark cycle.

Apparatus: The six experiments described in this thesis used the same six standard sound and light attenuating conditioning chambers, 24 (H) x 22 (L) x 19 (W) cm (Campden Instruments). The opening to a food tray was located in the middle of one wall of the chamber, flush with the wall and covered by a 5.5 x 6.0 cm perspex flap. The animals could retrieve 45 mg food pellets (Noyes, Formula A) by pushing the flap approximately 20-40 degrees from the vertical exposing the food tray. A standard Campden Instruments retractable lever, which acted as the trial context (*X*), was located on the wall to the left of the food tray flap. The houselight (G.E. no. 1820, 28 V. bulb) which acted as the variable cue (*A*) was located in the centre of the ceiling. Background noise level was produced by a ventilating fan (75-80 dB). A 286 PC microcomputer (Hewlett Packard) running version 2.0 of Med-PC (Med associates Ltd.) controlled the events in the chamber and recorded the data.

Procedure: During a single 30 minute session, rats were trained to retrieve food pellets from the food tray. The lever was retracted for this session. Single pellets were delivered on a variable time 60 second schedule (range 5 s - 115 s).

The following 7 conditioning sessions consisted of 64 ten second trials presented on a variable interval 30 second schedule. The start of each trial was signalled by the entry of the retractable Lever (*X*) into the chamber. On half these trials the Light (*A*) was paired with the Lever for the full 10 second interval. Therefore, there were 32 trials each of the Light +Lever (*AX*) and Lever alone (*X*). All trials ended with the termination of the stimuli, coinciding on some trials with the delivery of a single food pellet into the tray. The

distribution of the US during the two types of trial was different in each of the six treatments. Six rats were assigned to each of six treatments which crossed two levels of Light-Food contingency (Δp_A= .50 or 0) with three levels of food density [High (**H**), Medium (**M**) or Low (**L**)]. The letter designation refers to the density or absolute number of food deliveries during the entire session (48, 32 or 16 food pellets respectively; Density:.75, .50 and .25). In the three positive contingency groups **AX** trials signalled an increased likelihood of food compared with that signalled by **X** alone. For the three zero contingency groups both ***AX*** trials and ***X*** trials signalled the same likelihood of food. Therefore, there was a positive and zero **A** contingency treatment at each US density. The conditional probabilities for the occurrence of food on the ***AX*** and ***X*** trials for the three positive contingency groups were [**H.5** : p(US|AX)=1.0 and p(US|X)=.50], [**M.5**: p(US|AX)=.75 and p(US|X)=.25], [**L.5**: p(US|AX)=.50 and p(US|X)=0]. The conditional probabilities for the occurrence of food on the ***AX*** and ***X*** trials for the three zero contingency groups were [**H0** : p(US|AX)=.75 and p(US|X)=.75], [**M0**: p(US|AX)=.50 and p(US|X)=.50], [**L0**: p(US|AX)=.25 and p(US|X)=.25]. The exact trial frequencies and proportion of trials paired with the US are reported in Table 3.

Following training with the US there was a single extinction test session during which **AX** and **X** were presented according to the same schedule as during training but there were no US deliveries.

Response Measures and Statistical Analysis:One of the dependent variables in this experiment was the number of pushes of the flap covering the food tray. In order to compare the response generated by the discrete CSs with an appropriate baseline, tray entries were recorded both during the CSs and during the 10 second interval immediately preceding each trial (Pre-CS). In addition to tray entries, presses of the retractable lever (***X***) were recorded during the last training session and during the test. Lever pressing was observed in spite of the absence of any programmed contingency between lever pressing

and food delivery. Because lever pressing was correlated with the predictiveness of the lever for food, lever presses were assumed to be a Pavlovian conditioned response measuring the strength of the Lever-US relation (Peterson, Ackil, Frommer & Hearst, 1972). To control the effects of inter-subject variability a natural log transform was performed on the data. Statistical analyses was carried out on the data collapsed in two session blocks but because there was an uneven number of training sessions, session 1 was omitted from the analysis. Unfortunately, lever pressing data was only recorded on the final session of training and during the test. All statistical tests in this and in subsequent experiments used a .05 rejection level except where stated. Where interactions were reliable simple effects were analyzed by single df comparisons using Scheffe's (1953) method to control for Type I errors except where indicated.

Results

The statistical analyses are consistent with the interpretation that animals were sensitive to CS-US contingency. There was an increase in tray entries to **A** when rats received a positive **A**-US contingency, regardless of the number of **A**-US pairings or overall US density. Consistent with the prediction that the context mediates contingency learning, there was more lever pressing to the trial context when rats received a zero **A**-US contingency. In addition to the effect of contingency there was evidence that increases in US density increased rate of lever pressing during the trial context. There was additional support for the role of the context in mediating sensitivity to **A**'s contingency. The general contextual cues of the chamber, like the discrete trial context **X**, elicited more tray behaviour in the zero contingency groups. In the following analysis tray entries are reported first followed by lever pressing.

Tray Entries

Pre-CS intervals. The mean natural log tray entry rates during the Pre-CS interval are shown in Figure 4. There appears to have been increased tray behaviour during the pre

Figure 4. The Mean natural log tray entries per minute during the 10 second Pre-CS interval in two session blocks during training and the extinction test in Experiment 1. Data are plotted separately for the three positive A-US contingency treatments (filled symbols) and three zero A-US contingency treatments (open symbols).

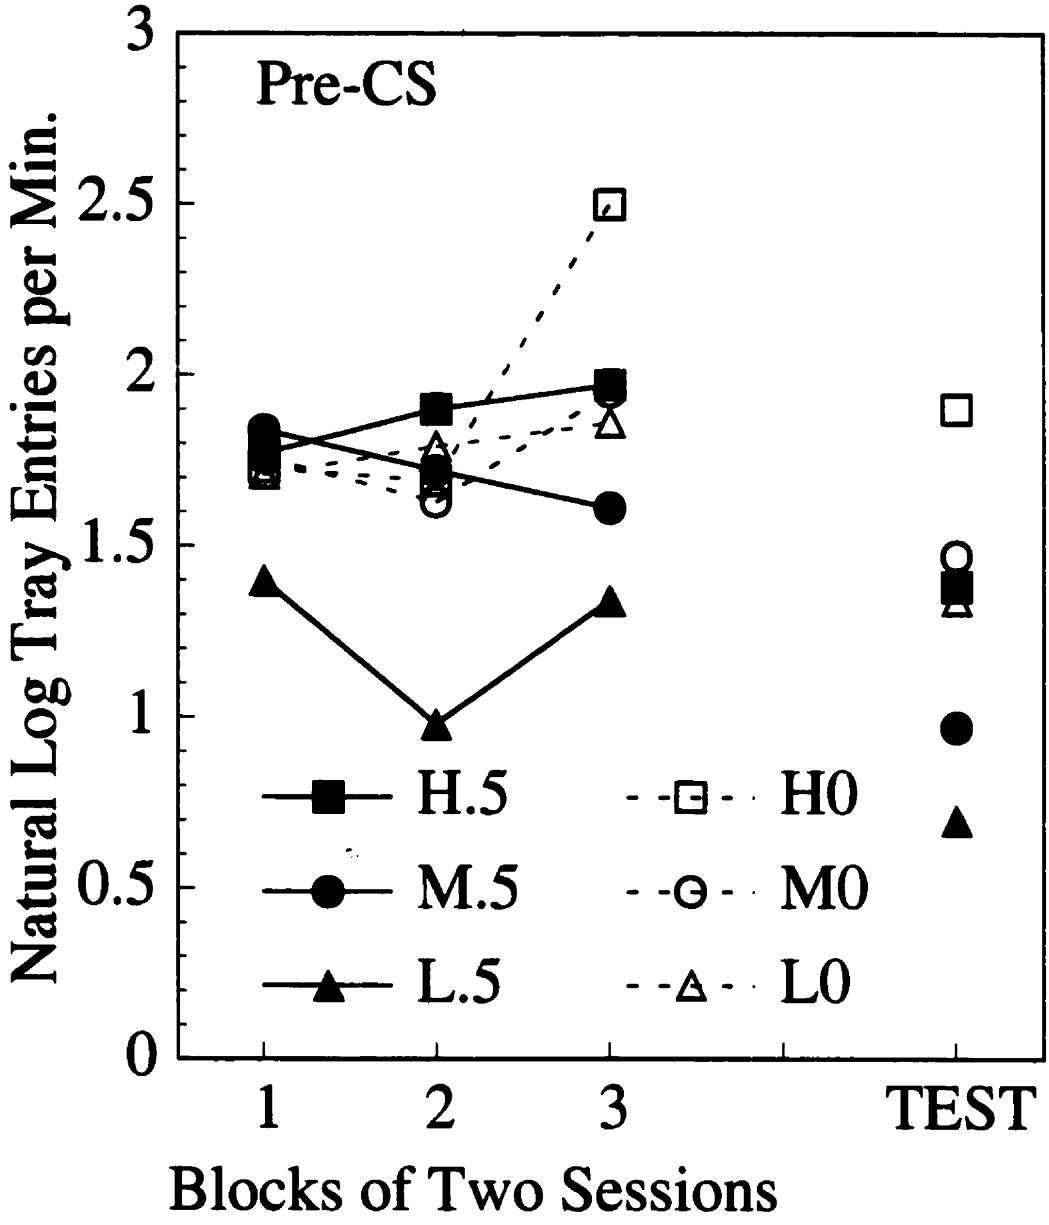
Pre-CS
Natural Log Tray Entries per Min.
3
2.5
2
1.5
1
0.5
0
H.5
M.5
L.5
H0
M0
L0
1
2
3
TEST
Blocks of Two Sessions

-CS period in the groups receiving zero contingency A-US training, but the statistical evidence for this is weak. An overall analysis of variance (ANOVA) of the three factors, *A* Contingency (positive or zero), US density (High, Medium, Low) and training Blocks (1-3) found a reliable main effect for blocks [$F(2,60)=5.01$] and interaction of Blocks and *A* Contingency [$F(2,60)=3.16$]. The effect of A contingency was not particularly robust. Post hoc tests conducted on the data from each block failed to reveal reliable differences.

AX and X trials. Figure 5 shows the development of tray entries during **AX** trials and **X** trials as well as a difference between these two periods (**AX-X**). The difference scores are included to show how much **A** elevated entries above the levelsupported by **X**. The top panel indicates that the mean natural log tray entries during **AX** trials increased over blocks for all three positive contingency groups (**H.5**, **M.5** and **L.5**) but decreased in the three zero contingency groups (**H0**, **M0**, **L0**). An overall 2 x 2 x 2 x 3 ANOVA of the two trial types (***AX*** and **X** trials), contingency ($\Delta p=.5$ or 0), US density (H, M, L) and Session Blocks (1-3) confirmed the observation that rates of tray entries were different during the two trial types and that the A-US contingency contributed to this difference. The main effect for trial type and the trial type by contingency interactions were both reliable, [$F(1,30)=26.3$] and [$F(1,30)=37.5$]. The three-way trial type by contingency by Blocks interaction was also reliable [$F(2,60)=7.91$].

A separate analysis of rates of entries during **AX** and **X** trials supports the claim that tray entries on **AX** trials was sensitive to **A**'s programmed contingency while on **X** trials this was not the case. The analysis of entries on **AX** trials found reliable main effects for *A* Contingency and the interaction of *A* Contingency and Blocks [$F(1,30)=14.11$ and $F(2,60)=5.93$ respectively]. None of the other effects was reliable [maximum $F=1.49$]. Individual comparisons at each block found that the three positive contingency groups showed reliably more tray entries than the zero contingency groups only on the third and final block of training sessions [$F(2,60)=7.64$].

Figure 5. Mean natural log tray entries per minute in two session blocks during training and the extinction test in Experiment 1 shown separately for **AX** trials (top panel) and **X** trials (middle panel) as well as the difference scores (**AX-X**; bottom panel). The three positive **A**-US contingency treatments (filled symbols) and the three zero contingency (open symbols) are shown separately.

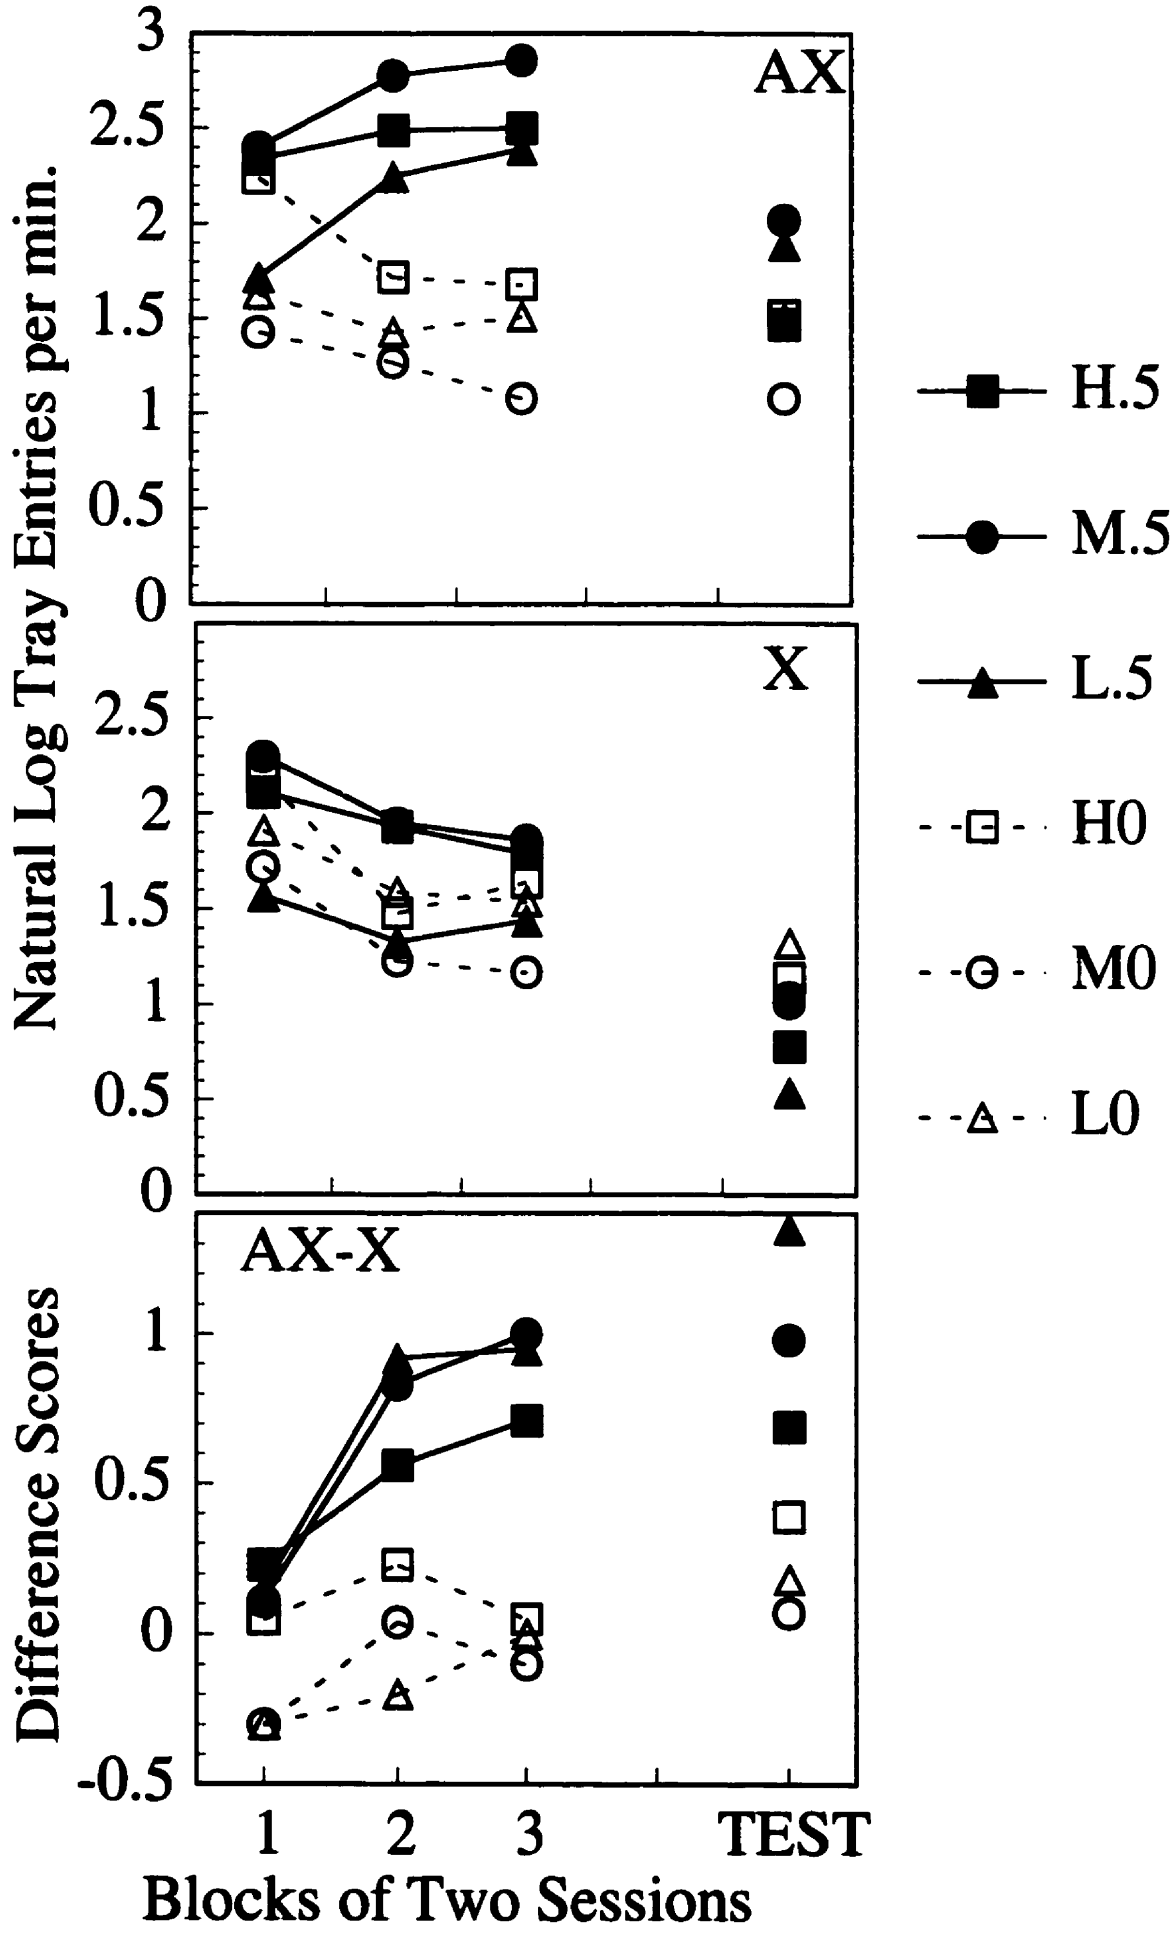
AX
X
AX-X
Natural Log Tray Entries per min.
Difference Scores
3
2.5
2
1.5
1
0.5
0
-0.5
1
2
3
TEST
Blocks of Two Sessions
H.5
M.5
L.5
H0
M0
L0

In contrast, the middle panel of Figure 5 shows that all six groups showed similar levels of tray behaviour on **X** trials and a general decline over the three blocks of sessions. The analysis supports this assertion. Only the main effect for blocks was reliable [$F(2,60)=7.65$].

Difference Scores. It is possible that any differences in conditioning on **AX** trials may reflect differences in conditioning to the general contextual cues, to the trial context **X** rather than differences in conditioning to **A**. One way to compensate for any differences in baseline rates of tray entries is to remove the contribution that all the cues except ***A*** make to the response on **AX** trials by subtracting responding on **X** trials from responding on **AX** trials (Bouton, Rosengard, Achenbach, Peck & Brooks, 1993). The bottom panel of Figure 5 presents the difference scores (**AX-X**). In agreement with the previous analysis, **A** elicited an increase in tray behaviour over and above that elicited by **X** alone but only in the positive contingency groups. An ANOVA found reliable main effects for **A** Contingency [$F(1,30)=37.54$], Blocks [$F(2,60)=22.20$] and the interaction between these two effects [$F(2,60)=8.21$]. Individual contrasts at each block only found a difference between the two contingencies on block 3 [$F(2,60)=22.38$].

Extinction. Following training, a test in the absence of the US was performed. Entries during the Pre-CS in extinction are shown in Figure 4. There were more tray entries during the Pre-CS in the Zero Contingency groups. There was also evidence that this contextual conditioning was positively correlated with US density. The 2 x 3 analysis, Contingency (Pos, Zero) and US density (H, M, L) found both main effects were reliable [$F(1,30)=9.21$ and $F(2,30)=4.00$ respectively]. Individual comparisons were carried out to determine which density conditions differed from one another. Only the Low and High density condition differed reliably [$F(2,30)=3.09$]. Therefore, there was evidence for more conditioning to the general cues of the context in the high US density condition.

Figure 5 presents natural log tray entries from the test on **AX** and **X** trials as well as

the difference scores. As with the training data the contingency of **A** influenced entries on **AX** trials but not during **X** trials. An overall Trial type, Contingency and Density ANOVA found a reliable trial type effect , [$F(1,30)=37.5$] and trial type by **A**-US contingency interaction, [$F(1,30)=15.9$]. The individual tests of **A**-US contingency and US density for **AX** and **X** trials separately only suggested reliable effects for **A**-US contingency but in both analyses the main effects for contingency failed to pass the .05 criterion [maximum $F(1,30)=3.45$, $p<.07$].

Figure 5 also displays the difference scores (**AX-X**) from the test session. This measure controls for differences in baseline rates of tray behaviour. The three Positive contingency groups showed a greater increase in tray behaviour during *AX* trials. The overall 2 x 3 analysis of variance for the difference scores during the test found a reliable main effect for *A* Contingency [$F(1,30)=16.19$] but both the main effect for density and the contingency by density interaction were not reliable. Therefore, in comparison to responding on the *X* alone trials responding on *AX* trials was higher only if *A* was trained with a positive contingency with food.

The three Zero contingency control groups received the same total number of US presentations and the same total number of non-US trials as the Positive contingency groups. This suggests that these results can not be due to differences in exposure to the US. However, the positive contingency groups did differ with respect to the total number of **AX**-US and **AX**-noUS pairings. One possible explanation of the main effects attributed to **A**'s contingency found in the previous analyses is that the positive contingency groups as a whole received more **AX**-US pairings and fewer **AX**-noUS pairings than the zero contingency groups that shared the same US density. One planned test of this explanation is to compare directly the level of responding in the **M.5 (75:25)** and the **L.5 (50:0)** with the **H0 (75:75)** and the **M0 (50:50)** groups. These four groups act as a direct test of the **A**-US pairing hypothesis. If the effect for contingency reported in the previous analysis is a

function of the overall higher number of A-US pairings in the positive contingency treatments then a contingency difference should not be found in this sub-analysis. Table 3 shows that both **M.5** and **H0** groups each received 24 **AX**-US trials per session and 8 **AX**-noUS trials (a 75% probability of food in the presence of **A**) while **L0** and **M0** each received 16 **AX**-US trials per session and 16 **AX**-noUS trials (a 50% probability of food in the presence of **A**). A 2 x 2 analysis of variance of **A** contingency (Pos or Zero) and **A**-US pairings (24 or 16) found only the main effect for **A** contingency reliable [$F(1,20)=14.0$]. This result supports the hypothesis that it is neither the number of CS-US pairings nor the amount of experience with the US alone which determined the number of tray entries, but rather **A**'s overall CS-US contingency. This finding generalizes Rescorla's (1968) finding with direct tests of alternative explanations for the effect.

Lever Pressing

AX and X trials. Because the lever was only present during the discrete trials there was no Pre-CS lever pressing data. The results of the lever pressing data from the **AX** and **X** trials were consistent with the prediction of associative theories of conditioning (e.g., RWM and Pearce), that zero contingency **A**-US training should result in a reduction in associative strength to **A** with a corresponding increase in the associative strength of the trial context. When **A** had a zero contingency with the US there was less tray entering and more lever pressing in the presence of **X**. Indeed the three zero contingency groups showed more lever pressing than the positive contingency groups.

Extinction. Figure 6 presents the natural log number of lever presses per minute on **AX** and **X** trials during the test. The analysis from the final training session which is not reported here showed similar results. During the test, the zero contingency groups pressed the lever more than the positive contingency groups with more lever pressing on **X** trials than on **AX** trials. This result is consistent with idea that the decreased tray behaviour to *A* accompanying zero contingency training was accompanied by increased conditioning

Figure 6. Mean Log lever presses per minute during the extinction test in Experiment 1 on **AX** trials (top panel) and **X** trials (bottom panel).

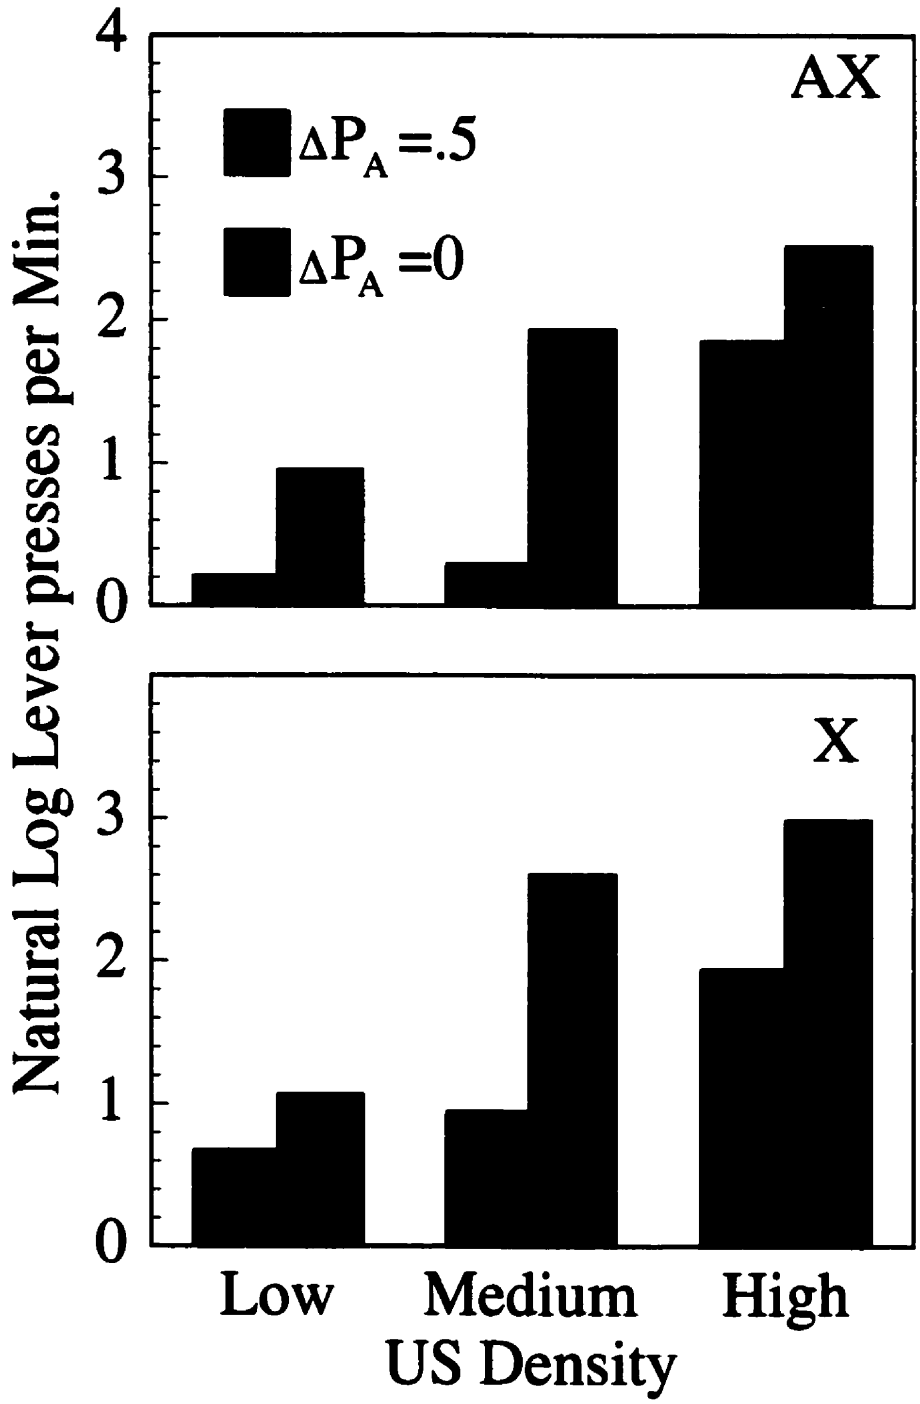
AX
ΔP_A =.5
ΔP_A =0
X
Natural Log Lever presses per Min.
0
1
2
3
4
Low
Medium
High
US Density

to the trial marker. Lever pressing rates were also a function of US density. Higher US density resulted in more lever pressing. The analysis of variance for Trial type (**AX**, **X**) **A**-Contingency (.5,0) and US density (H, M, L) found all three main effects were reliable. The reliable main effects for trial type confirmed that lever pressing was greater on **X** trials [$F(1,30)=22.9$]. The reliable main effect for **A** Contingency [$F(1,30)=17.5$] suggests that lever pressing rates were lower when rats were trained with zero **A** contingency while the reliable main effect for US density [$F(2,30)=14.2$] supports the observation that lever pressing increased with higher US density. None of the interactions were reliable [max $F(2,30)=2.2$].

As with the analysis of the tray entry data, comparing the 6 treatments for **A**'s contingency and US density does not control for the absolute number of pairings between **X**, the trial marker, and the US. To investigate this issue a second analysis analogous to that reported for tray entries was conducted. First, for lever pressing on **AX** trials in both **M.5** and **H0** treatments, there were 24 pairings of the lever with the US while in **L.5** and **M0** there were 16 lever-US pairings. The analysis of lever pressing during the test in these four groups for **A**-Contingency (Pos, Zero) and Pairing (24 or 16) found a reliable effect for contingency, [$F(1,20)=57.4$] but not for the number of **AX**-US pairings, [$F(1,20)=1.59$]. A similar analysis was performed on a different subset of treatments for lever pressing on **X** only trials. In both **H.5** and **M0** there were 16 Lever US pairings, while in **M.5** and **L0** there were 8 Lever US pairings (See Table 3). The analysis of these four groups for **A**-Contingency (Pos, Zero) and Pairing (16 or 8 US pairings) during the test of ***X*** alone found a reliable main effect for number of Pairings but there were no reliable effects involving the ***A***-contingency factor [$F(1,20)=14.5$ and maximum $F(1,20)=1.5$]. These results suggest that the differences in lever pressing on **AX** trials were a function of **A**'s contingency, while on **X** only trials the number of pairings between **X** and the US in the absence of **A** determined whether rats would press the lever.

Discussion

This experiment is the first attempt to use a discrete trial marker to investigate contingency sensitivity. Previous experiments on the role played by CS-US contingency in conditioning have done so primarily by manipulating contingency by varying the ratio of signalled to unsignaled US deliveries in aversive conditioning procedures (e.g., Dweck & Wagner, 1970; Quinsey, 1971; Rescorla, 1968; Miller, Grahame, & Hallam, 1990). The procedure described here, at the very least complements these previous experiments and may have the advantage of allowing better control over the presentation of a CS-US contingency. In addition it allows a direct test of the role of context in contingency learning. The differences in tray entry responses on **AX** and **X** trials showed clearly sensitivity to **A**'s moderate .5 Δp contingency. In addition to contingency sensitivity based on differential responding to the variable cue **A**, the discrete trial marker **X** elicited a separate conditioned response, which showed the characteristic reciprocal conditioning predicted by many models of conditioning.

Previous experimental evidence has been equivocal with regards to the role played by the overall CS-US contingency in Pavlovian conditioning. This may partially be due to the difficulty in comparing the level of responding in a group trained with a positive CS-US relationship with one trained with a zero CS-US contingency. There are always at least two confounding variables (Gibbon, Locurto & Terrace, 1975) that make interpretation difficult. Any differences in conditioned responding could be a function of the difference in contingency or equally be due to differences in total experience with the US and total CS-US pairings. The control treatments used in this experiment allowed a direct comparison between positive contingency treatments with zero contingency controls which received either the same total amount of experience with the US or the same number of pairings between the CS and US.

One simple explanation of the reduction in the strength of conditioned responding to

a CS following zero contingency CS-US training involves experience with the US. Zero contingency treatments are often programmed simply by increasing the number of unsignaled US. It has been argued that reductions in conditioned responding to the target **A** is a direct consequence of these increased presentations (Jenkins & Shattuck, 1981). There were zero contingency control groups in this experiment that received the same number of US presentations as comparable positive contingency groups yet there were still differences in responding to **A**. This rules out this density account as a unitary explanation of the contingency learning effect.

However, even if increased US experience alone were to result in a reduction in the strength of a conditioned response this does not necessarily mean that animals are unable to learn a CS-US contingency as well. It is still possible that zero contingency training can influence the relative value of the contextual cues (Jenkins & Shattuck, 1981; Jenkins & Lambos, 1983; Kremer, 1974). There is strong evidence in the present experiment that whether the trial marker acquired control over a conditioned response was a function both of the discrete cue's (**A**) contingency and the overall US density.

Jenkins proposed the waiting time hypothesis as an alternative to contingency learning (Jenkins, Barnes & Barrera, 1981; see also, Gibbon & Balsam, 1981; for the similar comparator hypothesis see Cole, Gunther & Miller, 1995). According to this hypothesis, the crucial variable controlling the Pavlovian response is the ratio of the density of the US during the CS to density of the US during the experimental context. Animals learn and compare these two densities. The level of tray entries to **A** in Experiment 1 is consistent with this hypothesis because in each of the positive contingency groups the waiting time for the US in the presence of the **A** was less than the waiting time in the presence of **X**. However, it is unclear how differences in waiting time should effect responding to **X**. The overall waiting time for the US in the presence of **X** is the same in each pair of density conditions (**H.5** an **H0**, **M.5** and **M0** and **L.5** and **L0**). In spite of

this, lever pressing to **X** was quite different in these groups. An analysis demonstrated that responding to **X** on **AX** trials was determined by **A**'s contingency but on **X** trials by the likelihood of the US during **X** [p(US|**X**)]. This result is not easily accounted for by the waiting time hypothesis or SET theory because waiting time is defined during **X** as a function of p(US|**X**) and p(US|**AX**).

In contrast to the comparator models, the RWM provides precise and accurate predictions for the six treatments described in experiment 1. A comparison of the model's predictions and the experimental results can be made by comparing the simulation data from Figure 3 with the data from Figures 5 and 6. The model's predictions match the data closely. The ordinal position of the six positive and zero contingency treatments matches the model's predictions capturing both the contingency sensitivity and the sensitivity to US density. The model does seem to predict that conditioning to the trial context will be equivalent in two pairs of conditions (**H.5** and **M0** as well as **M.5** and **L0**) whereas the data seems to indicate otherwise. However, the reported statistical analysis of these groups failed to find any main effect for contingency between these four treatments only the main effect for number of pairings. This difference is reflected in the simulations. In general, the model provides a fairly accurate mapping of both tray and lever data. One difference is that the RWM using the standard assumptions predicts that increases in US density should result in decreases in conditioning to A. There was no evidence for this difference in the data. Pearce's (1987) stimulus generalization model also predicts the basic contingency difference but predicts that increased US density should result in increased conditioning to A. While both of these mathematical models predict US density effects on conditioning to A, the normative Δp model does not. So only the normative contingency model predicts the present set of data. Although this conclusion assumes accepting the null hypothesis that there were no differences in responding to **A** within the two contingency treatments. It is interesting to point out the success of the normative theory of causal learning.

One possible explanation of the persistent excitatory conditioning which sometimes occurs following zero contingency training may be related to the type of stimuli normally used in contingency experiments. If the context elicits a response which is compatible with the response generated by the CS, then the response generated by the two cues may summate (Rescorla, 1997; Weiss, 1972). Even with an absolute decrease in CS conditioning, zero contingency CS training may produce increased contextual conditioning. If contextual conditioning and CS conditioning are manifested and measured using the same response (e.g., conditioned freezing during the CER) then it might be difficult to separate the contributions of the CS and context, without context switches. In the present experiment the context cue and the CS elicited different responses, allowing independent analyses of the reciprocal conditioning pattern between CS and Context.

Although the use of two CSs that elicit different responses has some advantages, the stimuli chosen here might expose this experiment to a possible criticism. Experiment 2 will be a partial replication of Experiment 1 with different stimuli to test the issue of generalisability.

Experiment 2

It was argued that the findings in Experiment 1 reflected a phenomenon based on the relative validity of the CS and context. There are at least two possible criticisms of this conclusion. First, there was no counterbalancing of the stimuli so it is possible that the effects were stimulus specific. Secondly, these results might be influenced by the spatial location of the cues. Experiment 2 sought to partially replicate and extend the findings of Experiment 1. Three pairs of groups received training with different stimuli in the roles of the CS and Context. Each pair received treatments similar to **M.5** and **M0** of Experiment 1. One pair received the same stimuli in the same roles as Experiment 1. The Houselight was the variable cue (**A**) and the Lever the trial context (**X**). The second pair of groups received the same stimuli but their roles defining the contingency were reversed. The Lever was the

variable cue (**A**) and the Houselight was the trial context (**X**). Finally, a third pair of treatments involved a different light source located inside the food tray as the variable cue (**A**) with the Lever as the context cue. The choice of the tray light was motivated by a second goal, to strengthen the difference between positive and zero contingency training. In addition, the tray light was located adjacent to the lever rather than in the ceiling making it less likely that animals would fail to see the light as they pressed the lever, and make it less likely that the light itself would draw them away from the food tray.

One reason for thinking that the difference found in Experiment 1 might have been stronger with another pair of stimuli relates to the relationship between CSs and the form of the conditioned response. The conditioned response to stimuli is more complex than simply tray entries or lever pressing data would suggest. Visual stimuli quite often elicit at least two identifiable components. With relatively long CSs (greater than 5 seconds) the initial portion of the CR involves orientation and rearing to the light source (Holland, 1977). However, as the time to the US delivery approaches, or with increased training rats will show increased orientation to the US source (Kaye & Pearce, 1984). To the extent that the CS and US sources are differentiated in space the two components of the response, CS directed responding and US directed responding, might be incompatible. In Experiment 1 this incompatibility may have interfered with the attempt to measure conditioned behaviour to the light since any rearing or orienting to the light source would not have been recorded. A visual cue emanating from the tray area might elicit an orienting response which was compatible with the US directed behaviour (Kaye & Pearce, 1984) and, therefore, might increase the magnitude of conditioning to the light in the positive contingency groups.

An alternative interpretation of the contingency effect found during the test session in Experiment 1 involves response competition at the time of the CS test. The test of the strength of **A**'s control over tray entries was always conducted in the presence of **X** and **X** sometimes controlled lever pressing which is physically incompatible with entering the tray.

Thus, there may have been response competition at the test. One way to avoid response competition is to perform the extinction test of **A** in the absence of the trial context (**X**). This was done in Experiment 2.

Method

Subjects and Apparatus: Naive rats were used as subjects, all details are the same as those described in Experiment 1. The training chambers and stimuli in Experiment 1 were also used here except that two groups received training with a tray light instead of the overhead houselight. The tray light (Spectro no. 1448, 28V bulb) was located behind and immediately above the tray flap in the food tray. When illuminated the light could be seen through the clear perspex tray flap.

Procedure: The animals were trained to retrieve food pellets from the food tray during a single session as in Experiment 1. Conditioning followed and involved either positive contingency or zero CS-US contingency training with a CS (**A**) whose contingency was defined relative to a trial marker (**X**). Three different pairs of stimuli were used in the role of **A** and **X** and this assignment is shown in Table 4. Twelve animals received training with the Houselight as **A** and Lever as **X** as in Experiment 1, twelve animals were trained with the Tray light as **A** and the Lever as **X** and the final twelve animals had the roles of the Houselight and Lever reversed. Conditioning involved daily sessions of 64 ten second trials randomly presented on a variable 30 (2-58) second schedule. Half of the trials (32) were signalled by the **AX** compound while the other half were signalled only by **X**. Each of the positive contingency groups received the same distribution of USs as treatment **M.5** from Experiment 1, 24 of the 32 **AX** trials (75%) but only 8 of 32 of the **X** trials (25%) ended with delivery of food. These groups had a moderate positive contingency between **A** and the US [$p(US|\mathbf{AX})=.75$ and $p(US|\mathbf{X})=.25$; $\Delta p_{\mathbf{A}}=.75-.25=0.5$. The zero contingency groups received the same training as group **M0** from Experiment 1, 50% of both **AX** and **X** trials (16 out of 32) ended with the delivery of a single food pellet [$p(US|\mathbf{AX})=0.5$ and

Table 4. Stimulus counterbalancing for the target CS (**A**) and the trial context (**X**) in Experiment 2.

	A-US Δp contingency (.5 or 0)					
	.5Hlight	0HLight	.5Lvr	0Lvr	.5Tlight	0Tlight
CS (**A**)	Houselight	Houselight	Lever	Lever	Traylight	Traylight
Trial Context (**X**)	Lvr	Lvr	Houselight	Houselight	Lever	Lever

$p(US|X)= 0.5$; $\Delta p_A=0.5-0.5=0$](See Table 3).

Following 12 conditioning sessions a single nonreinforced extinction test was conducted; 32 10-second trials of **A** were presented in the absence of **X**.

Results

The overall pattern of results was consistent with Experiment 1. The strength of the conditioned response elicited by stimulus **A** was a function of the **A**-US contingency, furthermore, the form of the conditioned response controlled by **A** was determined by its stimulus modality. When **A** was a visual stimulus (either houselight or traylight) and was trained with a positive contingency, **A** controlled tray entries, but when A was the lever it controlled lever presses. There were no differences in conditioning to **A** or **X** resulting from the counterbalancing of the houselight and the traylight, so the analysis is presented collapsed across these two groups. To simplify the discussion, results are described first in terms of tray entries and separately for the groups receiving training with either a light stimulus as *A* (n=24) or the lever as *A* (N=12).

Tray Entries

Pre-CS intervals. Unlike in Experiment 1, there was little difference in Pre-CS tray entries as a result of either Positive or Zero contingency training. There was a gradual decrease in Pre-CS entries across the training blocks. The top two panels in Figure 7 show the natural log Pre-CS entry rates across the training blocks for the groups split by whether **A** was a light (Top left panel) or lever (top right panel). An overall analysis of variance of the Pre-CS data for **A** Contingency (Pos, Zero), **A** stimulus (Light, Lever) and the four blocks found only a reliable main effect for blocks [$F(3,96)=5.74$]. The differences in tray entries found during the Pre-CS in Experiment 1 may have been specific to the actual stimuli used in that experiment. It is possible that when either the traylight or the lever were trained in the role of **A**, they may have been more effective at blocking the formation of contextual associations with the static cues of the conditioning chamber that presumably

Figure 7. Mean natural log tray entries per minute in three session blocks during training in Experiment 2. The data are separated by the Pre-CS interval (top panel), the two types of trial (**AX** and **X**) and the difference scores (**AX-X**; bottom panel). Training involved either positive ($\Delta p=.5$) or zero ($\Delta p=0$) CS-US contingency training. The left hand panels show tray entries for the groups receiving a light as the discrete stimulus (**A**) and lever as the trial marker **X**, while the right hand panels show responding for groups receiving training with a lever as the discrete CS (**A**) and light as trial marker **X**.

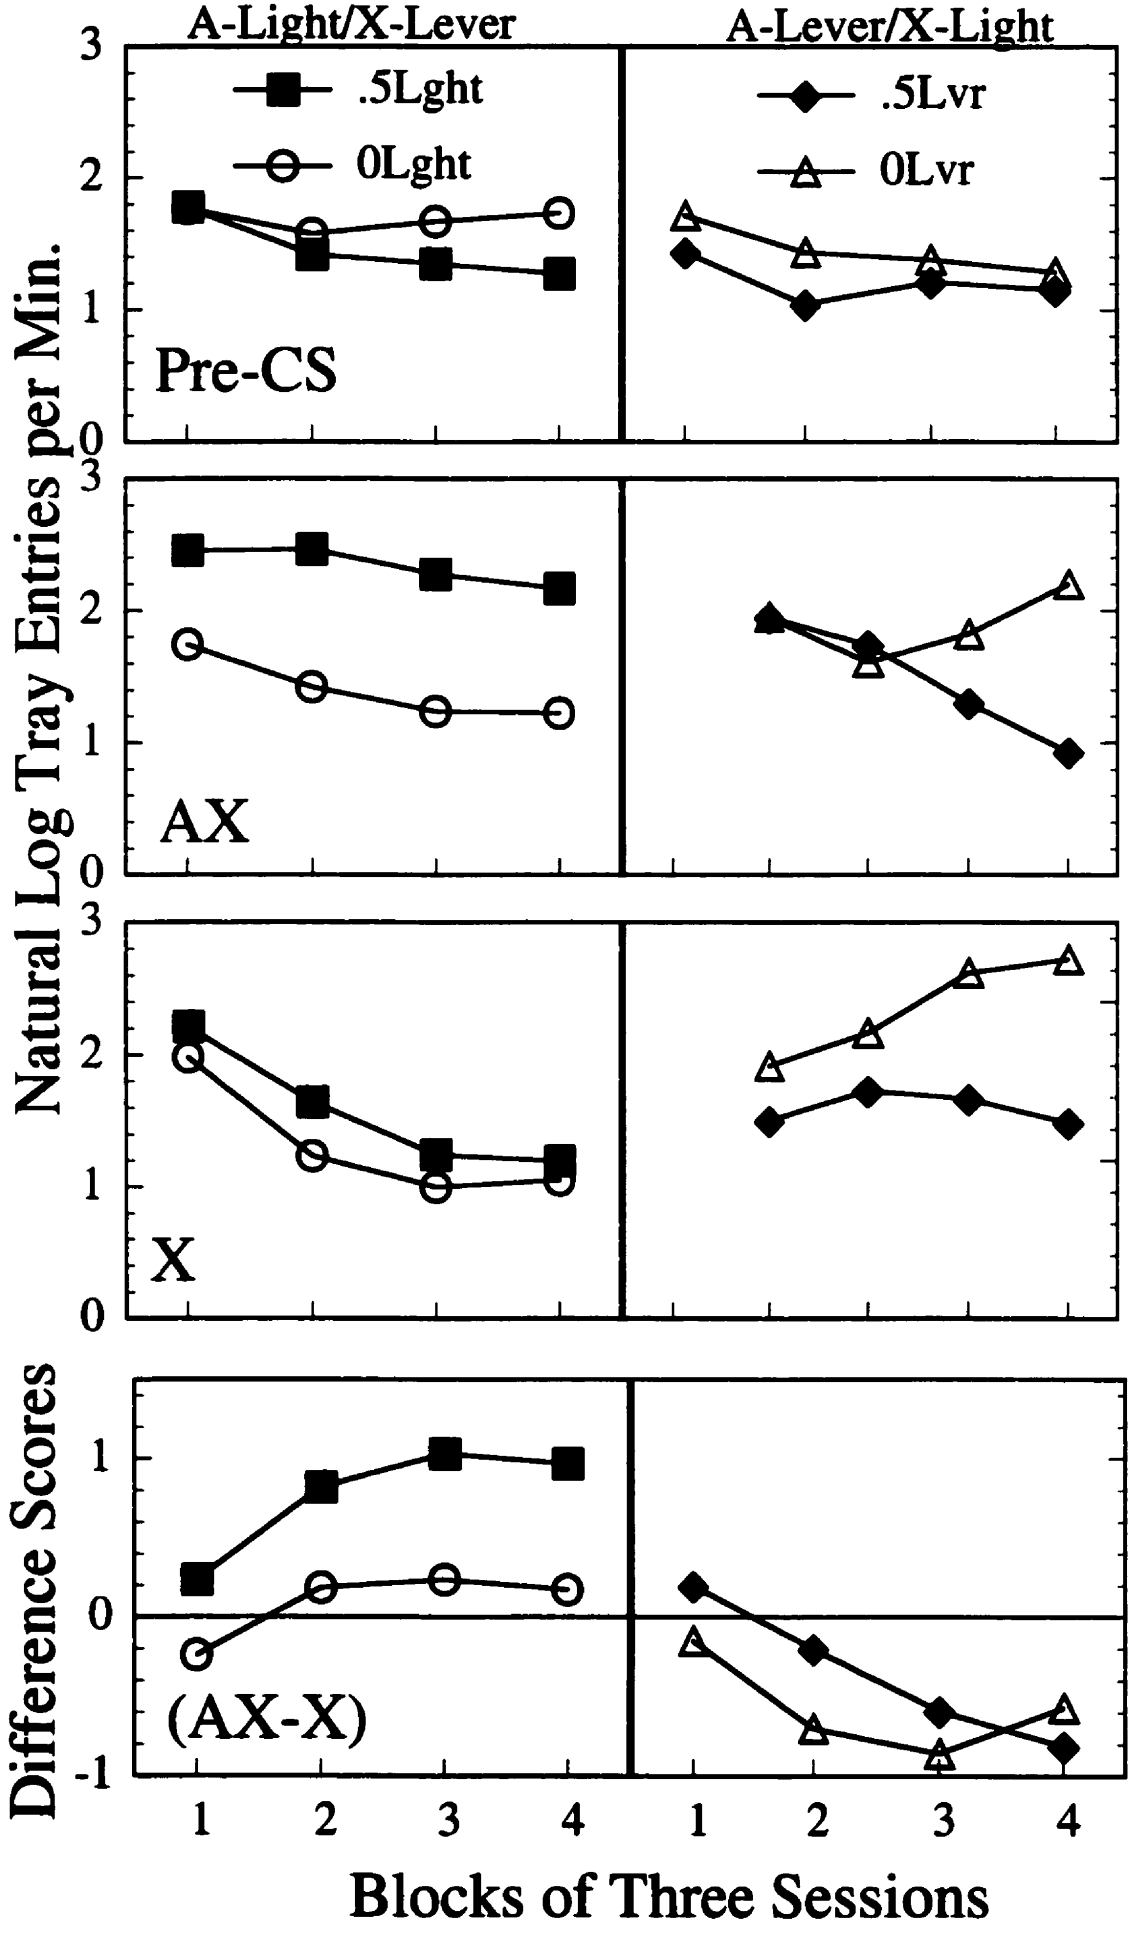
A-Light/X-Lever
A-Lever/X-Light
.5Lght
0Lght
.5Lvr
0Lvr
Pre-CS
AX
X
(AX-X)
Natural Log Tray Entries per Min.
Difference Scores
Blocks of Three Sessions
0
1
2
3
-1
1
2
3
4

control Pre-CS responding. To test this hypothesis a separate analysis on the two groups which used the same stimuli as those in Experiment 1 (the houselight and lever groups as *A* and *X* respectively) was conducted. The interaction between Blocks and *A* contingency just failed to reach the .05 level of significance [$F(3,30)=2.80$, $p<.057$]. Neither main effect for blocks nor contingency were reliable. This analysis generally confirms the small difference in Pre-CS responding found in Experiment 1 with the light as stimulus **A**.

AX and X trials. The middle four panels of Figure 7 display the natural log tray entries during **AX** and **X** trials split by the stimulus counterbalancing roles, **A**-Light, **X**-Lever (the left hand panels) or **A**-Lever, **X**-Light (the right hand panels). When **A** was a light stimulus and contingently trained, **A** controlled tray entries.

The overall three-way analysis of tray entries included the following factors: **A**-Contingency [Pos ($\Delta p=.5$), Zero ($\Delta p=0$)], **A**-Stimulus (Light, Lever), Trial Type (**AX** or **X** trials), and training Blocks, and supports these observations. The three way interaction between **A**-Stimulus, Trial Type and Blocks was reliable confirming that the difference between **AX** and **X** over blocks was influenced by whether **A** was a light or lever [$F(3,96)=19.5$].

Separate analyses of the data were conducted based on whether the **A**-Stimulus was a light or lever. For the **A**-Light groups, positive **A**-US contingency training increased tray entries on **AX** trials but had little effect on the rate of tray entering on **X** trials. The ANOVA conducted on these data found reliable main effects for **A**-contingency, Trial Type and Blocks [$F(1,22)=5.76$, $F(1,22)=33.6$ and $F(3,66)=8.84$ respectively] and found reliable **A**-contingency by Trial Type [$F(1,22)=21.0$] and Trial Type by Blocks interactions [$F(3,66)=15.1$]. An analysis of the tray entries separated by trial type showed that during **AX** trials, more tray entries were observed in the positive **A**-US contingency group, [$F(1,22)=11.2$] but no contingency effects were found on **X** trials [$F(1,22)=1.18$]. The general level of tray behaviour did decline across the four blocks [$F(3,66)=18.4$]. The

difference scores illustrate the interaction between A-contingency and Trial Type. The lower left panel of Figure 7 shows the increased tray entries on **AX** trials for the positive **A** contingency group.

The other groups of rats received training with the lever as the **A** stimulus, for these groups tray entries increased when **A** was uncorrelated with the US. Responding was higher on both type **AX** and **X** trials for this group. The results of the statistical analyses supports this observation. The main effect for A-Contingency was reliable [$F(1,10)=2.95$], as well as the interactions for A-Contingency with Blocks, and the Trial Type x Blocks interaction [$F(3,30)=3.11$ and $F(3,30)=6.17$ respectively]. However, unlike the analysis for the A-Light groups the Trial Type by A-Contingency interaction was not reliable [$F<1$]. Therefore, tray entries were controlled by the light, when the light was the trial marker and **A** was uncorrelated with the US. These results confirm the reciprocal conditioning effect demonstrated in Experiment 1.

Extinction. The test data are consistent with the findings from training and show that the response elicited by **A** (either tray entries or lever pressing) was influenced both by **A**'s-contingency and by its stimulus modality. In the positive A-US contingency groups, **A** elicited an increase in responding in comparison with the zero contingency groups. In the zero contingency groups, the trial marker **X** controlled behaviour. The particular response controlled by **A** or **X** was dependent on whether **A** or **X** was a light or lever. The test of **A** was carried out in the absence of the trial marker in the same conditioning chamber as training. Any tray entries observed during the presentation of **A** may be a function of either the response controlled by **A** or the response controlled by **A** and the static cues of the conditioning chamber. To separate these two possibilities the tray entries were analyzed in two ways.

The test for sensitivity to **A**'s Δp contingency was analyzed by directly comparing the levels of tray entering in the positive and zero contingency groups.

Evidence for the hypothesis that it was the discrete stimulus and not the contextual cues of the chamber that actually controlled tray entering came from comparing tray behaviour during the Pre-CS and the CS. Firstly, the top left hand panel of Figure 8 shows how, as in Experiment 1, when **A** was a light, **A** elicited more tray entries in the positive contingency group than in the zero contingency group. This was true both in terms of the absolute rate of tray behaviour and the relative increase above that elicited by the contextual cues alone (the Pre-CS). In contrast, the top right hand panels of Figure 8 show tray entries for the groups in which **A** was a lever. It is clear that **A** did not elicit an increase in tray entries, in fact **A** may have slightly suppressed tray entries. The analyses support these assertions. The ANOVA for **A**-Contingency (positive or zero), **A**-Stimulus (**A**-Light, **A**-Lever) and Response Period (Pre-CS, ***A***) found a reliable main effect for **A**-Stimulus and the interaction of ***A*** -Stimulus with Response Period [$F(1,32)=9.04$ and $F(1,32)=34.1$ respectively]. The three-way interaction was not reliable, but separate analysis of the tray rates for the **A**-Light and **A**-Lever groups were conducted to test whether the effect found in Experiment 1 was replicated. For the **A**-Light groups the increase in tray entries elicited by **A**, above that recorded during the Pre-CS, was greater in the positive contingency group than in the zero contingency group. The analysis found a reliable main effect for Response Period and an interaction between **A**-Contingency and Response Period [$F(1,22)=30.0$ and $F(1,22)= 6.34$ respectively]. The interaction is reflected in the difference scores between responding to **A** and the Pre-CS which are displayed in the bottom left hand panel of Figure 8. Although the difference between Pre-CS and CS responding was greater in the positive contingency group there was still a reliable increase in the Zero contingency group [minimum $F(1,10)=9.64$], and may suggest why the three way interaction in the overall analysis was not significant. The analysis of the two factors, 1) response period and 2) **A** contingency for the tray entry data, recorded for the two groups receiving training with **A** as a lever, found that there was a reliable decrease in tray

Figure 8.Mean natural log tray entries per minute during the 10 second Pre-CS interval and during the presentation of **A** during the extinction test from Experiment 2. Stimulus **A** either had a positive ($\Delta p=.5$) or zero ($\Delta p=0$) CS-US contingency with the US. The left hand panels show tray entries for the two groups that received a light as the discrete stimulus (**A**), while the right hand panels shows tray entries for the two groups receiving training with a lever as the discrete CS (**A**). The bottom panels show the difference scores (**A**-Pre-CS) which provides an estimate of the tray entries elicited by **A** alone.

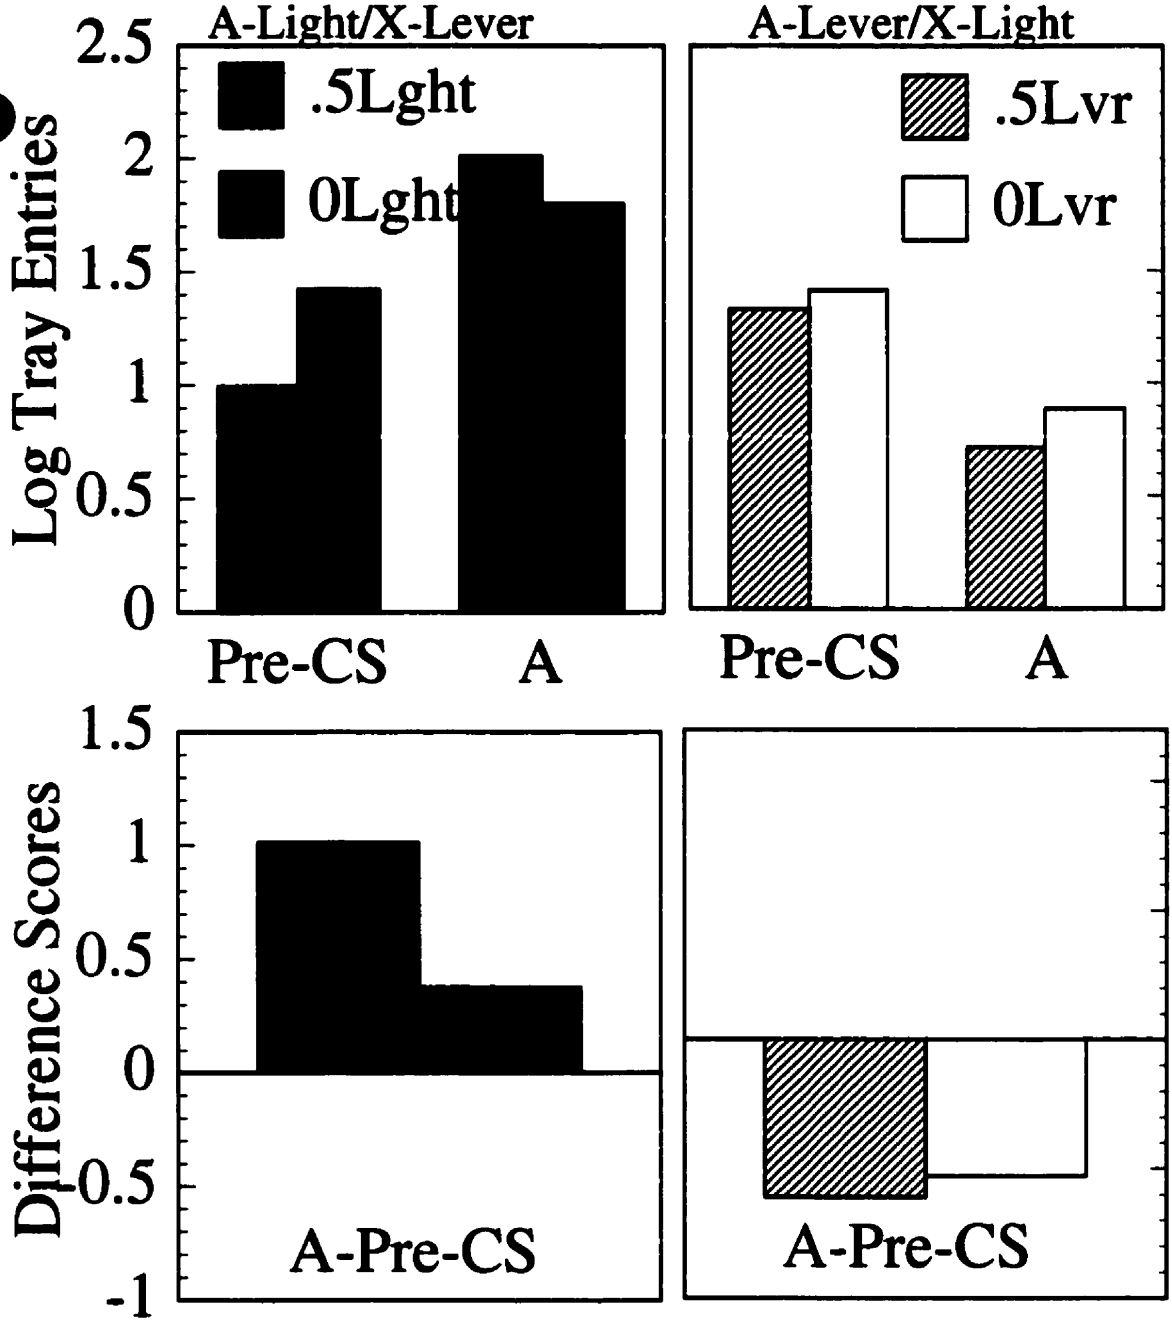
A-Light/X-Lever
A-Lever/X-Light
.5Lght
0Lght
.5Lvr
0Lvr
Log Tray Entries
2.5
2
1.5
1
0.5
0
Pre-CS
A
Pre-CS
A
Difference Scores
1.5
1
0.5
0
-0.5
-1
A-Pre-CS
A-Pre-CS

entering below the Pre-CS level when the lever was inserted [$F(1,10)=11.13$], the interaction was not reliable. Tray entering decreased from Pre-CS levels presumably because the animals were pressing the lever.

Lever Pressing

AX and X trials. Tray entries were recorded throughout the training sessions because the food tray was always present, however, lever pressing could only be recorded on trials in which the lever was extended. Consequently, there are no Pre-CS lever pressing scores and none during presentations of the light alone (i.e., during **X** trials for the **A**-Lever **X**-Light groups). Figure 9 displays the natural log rates of lever pressing during training for the two groups, **A**-Light and **A**-Lever. The analysis of lever pressing suggests the mirror image of the tray data. In the **A**-Lever group lever pressing was greater in the positive **A**-US contingency condition while in the **A**-Light group in which **X** was the lever, lever pressing was greater on both **AX** and **X** trials in the zero contingency group. The overall analysis of variance for lever pressing on **AX** trials conducted with factors **A**-Contingency, **A**-Stimulus and Blocks found a reliable main effect for Blocks [$F(3,96)=15.7$] the **A**-Contingency by **A**-Stimulus interaction [$F(1,32)=14.6$] and three way **A**-Contingency by **A**-Stimulus by Blocks interaction [$F(3,96)=3.27$], none of the other effects were reliable. These results support the observation that lever pressing on **AX** trials was greater in the zero contingency group in the **A**-Light group but higher in the positive contingency **A**-Lever group.

The separate analysis of lever pressing in the **A**-Light groups showed that lever pressing was higher in the zero contingency group during both **AX** and **X** trials. The analysis of **A**-Contingency, Trial Type and Blocks found reliable main effects for **A**-Contingency, Trial Type and Blocks [$F(1,22)=18.6$, $F(1,22)=52.1$ and $F(3,66)=15.4$] as well as the **A**-Contingency by Trial Type interaction [$F(1,22)=10.7$]. Individual analysis of each contingency group found that both positive and zero contingency groups showed

Figure 9.Mean natural log lever presses per minute from training in three session blocks in Experiment 2. Data are separated by the two types of trial (**AX** and **X**; top panels) and the difference scores (**AX**-**X**; bottom left panel). Training involved either positive ($\Delta p=.5$) or zero ($\Delta p=0$) CS-US contingency training. The left hand panels show lever pressing rates for the groups receiving a light as the discrete stimulus (**A**) and lever as the trial marker **X**, while the right hand panels shows lever pressing rates for groups receiving training with the lever as the discrete CS (**A**) and light as trial marker **X**. Because there was no lever present on **X** trials for this latter group there are no difference scores available.

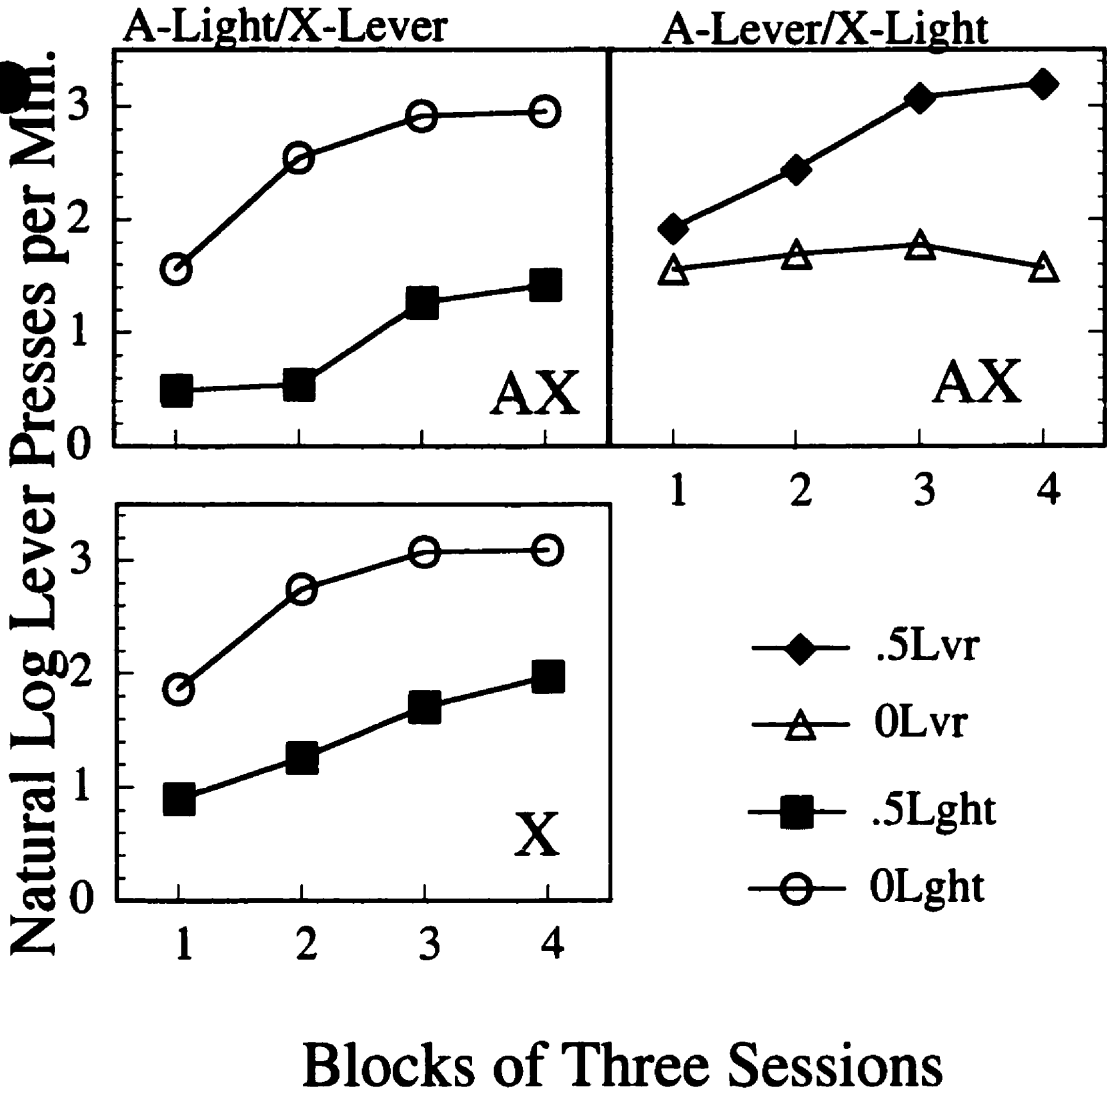
A-Light/X-Lever
A-Lever/X-Light
Natural Log Lever Presses per Min.
0
1
2
3
AX
AX
X
1
2
3
4
.5Lvr
0Lvr
.5Lght
0Lght
Blocks of Three Sessions

greater levels of lever pressing on X trials [$F(1,11)=30.6$, $F(1,11)=38.4$]. Theanalysis of lever pressing in the A-Lever groups on AX trials found lever pressing to be higher in the positive A-US contingency group. The main effect for A-Contingency was not reliable but both the main effect for Blocks and the A-Contingency by Blocks interaction were reliable [$F(3,30)=8.01$ and $F(3,30)=6.22$ respectively]. Individual comparisons on each block showed that while there was no difference on the first block, there was by the fourth [$F<1$ and $F(1,10)=6.26$ respectively].

Extinction. Figure 10 presents Lever pressing during presentation of A for the two groups in which A was a lever. Lever pressing was higher following positive A-US contingency training. While the decrease in tray entries from Pre-CS to A was the same for both positive and zero contingency groups this did not translate into the same rate of lever pressing. Lever pressing was higher in the positive contingency group. The analysis of variance found a reliable interaction between blocks and A-contingency [$F(3,30)=3.53$]. However, none of the individual block comparisons was reliable [maximum $F(1,10)=3.77$, $p<.10$].

Discussion

Training consisted of either positive or zero A-US contingency with the contingency calculated relative to a discrete trial context X. The strength of responding elicited by A was greater in the positive contingency groups. As in Experiment 1 the strength of the response elicited by the trial context X was greater following zero contingency training. This relationship was found regardless of the form of the response elicited by A or X. Furthermore, these results rule out two explanations of the reciprocal conditioning between CS and trial context found in Experiment 1. First, the differences are not stimulus specific. This experiment found a similar pattern of results regardless of whether the entry of the lever or either of two differently located light cues served as cue A. Second, sensitivity to A's contingency in Experiment 1 was probably not an artifact of

Figure 10. Mean natural log lever presses per minute in 5 trial blocks during the extinction test from Experiment 2. For this group training involved the lever in the role of the discrete cue (**A**). The lever either positive ($\Delta P=.5$) or zero ($\Delta P=0$) CS-US contingency training.

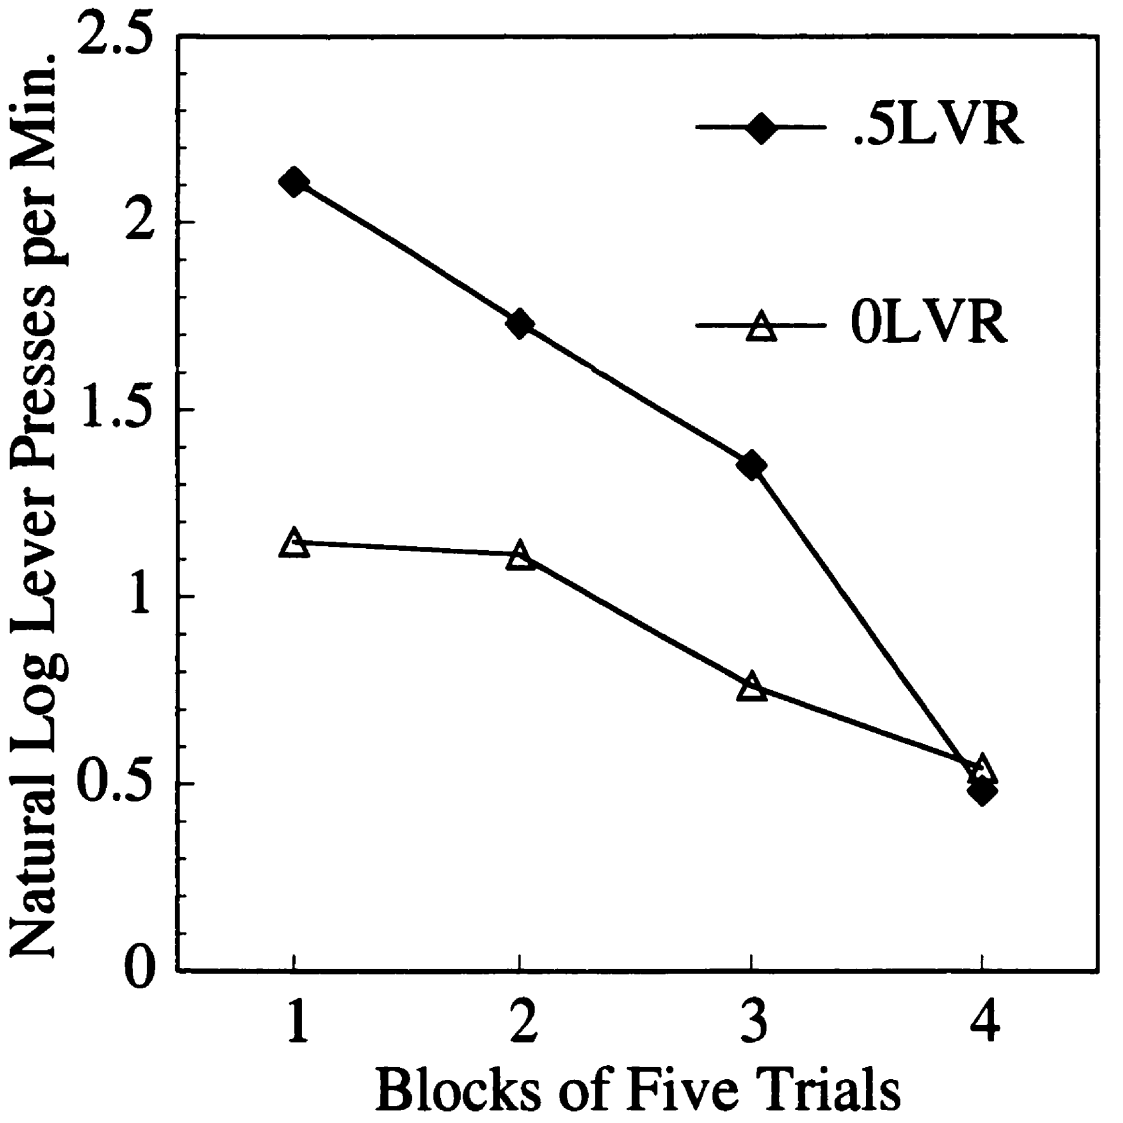
Natural Log Lever Presses per Min.
2.5
2
1.5
1
0.5
0
.5LVR
0LVR
1
2
3
4
Blocks of Five Trials

response competition during the compound test. Tests in the absence of the trial context **X** found a difference in conditioning following positive and zero contingency training. The contingency effects found with cue **A** were not conditional upon the presence during the test of the discrete trial context. Elemental theories of associative learning such as the Rescorla-Wagner model assume that learning involves acquiring associations based on combining experiences with discrete CSs regardless of the accompanying cues. In Experiment 1 and 2, animals received training with both a compound **AX** and element **X**. RWM predicts that despite any influence that either cue may have on the acquisition of associative strength by the other during training, at the time of the test, the response elicited by presentation of **A** or **X** will be independent of the other cue. Pearce's theory assumes that conditioning involves learning about entire stimulus complexes (e.g., Pearce, 1987; 1994). However even though the animals have no training experience with the stimulus **A** in the absence of the trial marker Pearce's model still predicts that there will be more generalized excitation in the positive contingency treatments. In spite of consistent training of **A** with **X**, the test of **A** in the absence of **X** still found a reliable contingency effect on responding to A.

Previous work on the controlling relations in Pavlovian conditioning often referred to the notions of contingency and/or correlation between CS and US. However, there has been little discussion about whether the computation of Δp should be treated as a theoretical construct (Baker, 1974; Cheng, 1997) reflecting internal calculation, or as an empirical generalisation (Rescorla, 1968; Hallam, Grahame & Miller, 1992; although see Baker, Murphy & Vallée-Tourangeau, 1996). The results presented here are consistent with both positions. The CS-US contingency is a controlling variable in Pavlovian conditioning, however an associative algorithm is also a more parsimonious interpretation of the actual calculations performed by the animal because it involves remembering fewer events and performing fewer calculations.

The overall relationship between **A** and the US determined the level of the conditioned response elicited by **A**. The levels of conditioning to the trial marker were also consistent with the idea that rats were sensitive to the contingency between the trial context and the US. Given these data one hypothesis that emerges is that rats may be sensitive to the relative contingencies of **A** and **X** as signals of the US. Miller and Schachtman (1985) have proposed that the relative probability of reinforcements may determine which stimulus elicits a conditioned response, the hypothesis proposed here is that it is the relative contingency that determines whether a stimulus elicits a CR. The present results are consistent with this hypothesis, but no attempt was made to define the contingency of the comparison trial marker. Testing the relative contingency hypothesis requires presenting multiple CSs and varying the relative CS-US contingencies. In Experiments 1 and 2 while **A** was always defined relative to **X**, **X**'s contingency was not defined because there were was no way to define the no**X** experience. Competition between multiple contingent predictors will be the topic of the next three experiments.

Multi-CS Experiments

In their review of the role of contingency in Pavlovian conditioning, Papini and Bitterman (1990) are unequivocal in their position. They argued that if the CS-US contingency is the underlying cause of a conditioned response then it should be both a necessary and sufficient condition to elicit responding. However, this is surely too restrictive. It is similar to arguing that, if smoking causes cancer, it should be necessary and sufficient for the disease to occur and if it fails this test then there is no point studying its causal role. Smoking may cause cancer, but so to may other factors and smoking itself may not always result in cancer. But this argument would only convince a tobacco manufacturer that smoking is not a cause of the disease.

A real test for a causal role of contingency in conditioning requires demonstrating that rats are sensitive to variations in contingency, which Experiments 1 and 2 have

demonstrated. Changes in the CS-US contingency produced changes in the strength of a conditioned response. There is one notable case in which this does not occur. Selective associations are a class of experimental results in which the contingency of a CS can not account for the conditioned response it elicits. However it is worth considering whether this perceived failure of contingency theory can be shown to be consistent with a revised Δp theory.

Selective association phenomenon, are demonstrations in which positive and normally effective CS-US contingency training fails to elicit strong conditioned responding (e.g., Kamin, 1969). From the perspective of associative theories animals seem to be selective in the CS-US associations they form. As a result, the unconditional Δp alone does not predict whether animals develop a conditioned response. With both blocking (Kamin, 1969), and the relative validity effect (Wagner, Logan, Haberlandt & Price, 1968) a cue with a positive contingency with the US acquires a weaker conditioned response when training includes more valid CSs. In blocking procedure the blocked cue is reinforced on every trial [p(US|CS)=1].

These findings have been explained using a number of cognitive mechanisms, some more elaborated than others. For example, with the relative validity effect, recall that the response controlled by the common element **X** following training with two compounds **AX**, **BX** is greater following a pseudo-discrimination ($\mathbf{AX}^{+/-}$, $\mathbf{BX}^{+/-}$) than following a true discrimination between the two compounds ($\mathbf{AX}^{+}$, $\mathbf{BX}^{-}$). This effect is difficult to explain with traditional contiguity based theories of learning because stimulus **X** is reinforced and nonreinforced equally following the two treatments. Wagner et al. (1968) originally suggested that conditioned responding involved sensitivity to the relative validity of the cues, where the validity of each CS was defined by the ratio of reinforced to nonreinforced trials [p(US|CS)]. Others have suggested that the effect may emerge because of different levels of generalization from the training compounds (**AX** and **BX**) to the test

cue X (Pearce, 1987; Thomas, Burr and Eck, 1970; although see Turner & Mackintosh, 1972). Kamin suggested that these effects involve changes in the surprisingness of a CS (see also Pearce & Hall, 1980) while Mackintosh (1974) suggested that they involve changes in attention. The Rescorla-Wagner model, in contrast, anticipates these effects using a competitive mechanism for associative strength as described earlier. Finally, Cheng's (1997) model of causal induction was also designed to account for these type of effects. According to her theory, the animals learn the conditional contingencies between each individual CS and the US, and then selectively respond to the CS with the strongest contingency. Given the debate that surrounds a description of these findings and the possibility that they may involve sensitivity to contingencies as I have described them it is worth exploring these effects further.

Papini and Bitterman (1990) stated that "phenomena such as overshadowing and blocking suggested clearly, at about the time Rescorla's theory was introduced, that a positive contingency is insufficient for excitatory conditioning" (p. 401). The belief that selective association phenomena are completely at odds with a contingency analysis is common. Another example of this belief is found in Mackintosh's (1983) authoritative review of conditioning and associative learning where he states "Contingency theory cannot readily explain ... the blocking experiment" (p. 184). This belief ignores the possibility that it is the relative rather than absolute contingency that controls behaviour. If animals are sensitive to relative contingencies it is possible that selective associations might be interpreted as supporting rather than contradicting the normative contingency analysis of conditioning (Shanks, 1995, p. 43; Waldmann & Holyoak, 1992). The following experiments describe an attempt to use the contingency learning preparation developed in Experiments 1 and 2 and test whether selective associations are examples of relative contingency sensitivity.

Experiment 3

The following experiment investigated a version of Wagner, Logan Haberlandt and Price's (1968) relative validity effect. The basic finding is that conditioned responding to a Pavlovian CS or a discriminative stimulus for instrumental responding is determined by the relative predictive value of the cues. The question of interest is what determines predictive value? In previous demonstrations of the effect the target cue X was always partially reinforced. There were never any US presentations in the absence of X. In the following experiment there were US presentations in the presence and absence of the test cue. As with the previous experiments contingencies were instantiated using a single trial marker. The design involved presenting two CSs and testing whether a strongly contingent CS would reduce the strength of the CR elicited by a CS with a weaker Δp contingency.

The discriminations between the two CSs **A** and **B** were always conducted in the presence of a discrete trial marker (**X**; the lever). As before, the trial marker was included so that the subjective CS-US contingencies might be less ambiguous, and provide a measure of the contextual conditioning. The target cue **A** was involved in either a moderate positive contingency with the US (Δp_A=.5) or a zero contingency (Δp_A=0). Comparisons of responding to **A** in these treatments are tests for sensitivity to **A**'s absolute contingency. In addition, to test whether the strength of the conditioned response to **A** was determined by its relative contingency, a second CS (**B**) was included. **B** possessed one of two contingencies. In two groups **B** was a perfect positive predictor of the US; in that it preceded all of the US deliveries and the US never occurred in its absence (p(US|**B**)=1; p(US|-**B**)=0; Δp_B=1). Alternatively B was uncorrelated with the US (Δp_B=0). Therefore, four treatments crossed two levels of **A** contingency (Δp_A= 0.5 or 0) and two levels of **B** contingency (Δp_B= 1 or 0). The exact trial frequencies, conditional probabilities and contingencies are shown in Table 5. The four treatments are labelled by their joint contingencies for **A/B** (.5/1, .5/0, 0/1, 0/0).

Both the normative analysis involving relative contingencies and the associative

analysis predict that the cue with the highest contingency will acquire a conditioned response at the expense of weaker predictors. Based simply on **A**'s absolute contingency, responding to **A** should be higher in treatments .5/1 and .5/0 than in treatments 0/1 or 0/0. However, if the predictive value of cue **B** is compared with **A**, it is clear that **B** is a more valid predictor, of the US in treatment .5/1 because it perfectly signals the occurrence of the US. Assuming that animals are sensitive to, and learn about the most likely predictor of the US then this analysis predicts that **B**'s strong relationship will reduce the ability of **A** to elicit a conditioned response in group .5/1. The Rescorla-Wagner model makes the same prediction although this prediction is based on the outcome of associative competition. The precise predictions of the model will be described in more detail in the discussion. Unlike previous relative validity experiments with multiple CSs (i.e., Wagner et al., 1968; Cole, Barnet & Miller, 1995), in this experiment the goal is to define explicitly the Δp contingency of the CSs. All four cells of the contingency table will be defined in terms of programmable combinations of events.

Method

Subjects: Thirty-six rats obtained from the same breeder and maintained under the same conditions as Experiment 1 were used in this experiment.

Apparatus: In addition to the lever and houselight used in Experiment 1, a 2000 Hz square wave Tone was also used (approximately 80 dB). The speaker (35 ohm, 7 cm in diameter) was located on the ceiling of the chamber adjacent to the houselight.

Procedure: The details of the pre-training are the same as those described in Experiment 1. The conditioning procedure involved the successive trial training procedure described previously. Each session consisted of 64 ten second trials presented on a variable time 45 second (range 12 s - 90 s) schedule. The trials were signalled by the entrance into the chamber of the retractable lever (**X**). On half the trials (32), the retraction of the lever at the end of the ten second interval coincided with the delivery of a single food pellet, while on

the other half no pellets were delivered. Two other discrete cues (**A**, **B**) were each presented on 32 trials. The relationships between these two stimuli and the US presentations are described in Table 5. Essentially the four treatments were designed to cross two levels of **A** contingency (Δp_A= .5 and 0) with two levels of **B** contingency (Δp_B= 1 and 0) producing four treatments each receiving simultaneous training with one of four **A/B** contingency pairs (.5/1, .5/0, 0/1, 0/0). The first value refers to the contingency of cue **A** and the second to the contingency of cue **B**. There were four different types of trial: **ABX**, **AX**, **BX** compounds as well as **X** alone trials. In the two groups in which **A** had a moderate positive contingency with the US, **A** was present during 24 of the 32 trials (75%) which ended with the US and 8 of 32 the trials (25%) which did not (p(US|**A**)=.75, p(US|no**A**)=.25; Δp_A=.50). In the two other groups, **A** had a zero contingency with the US and was present during half of the US and noUS trials (p(US|**A**)=.50, p(US|no**A**)=.50; Δp_A=0). The second CS (**B**) was either perfectly predictive or nonpredictive of US occurrence. In the two treatments in which **B** was perfectly predictive of the US it only appeared during the US trials (p(US|**B**)=1, p(US|no**B**)=0; Δp_B=1). For the zero contingency, **B** appeared during half of US and noUS trials (p(US|**B**)=.50, p(US|no**B**)=.50; Δp_B=0).

Whether cues **A** and **B** were the tone or the light was counterbalanced in each group of 9 subjects. Five received training with **A** as a tone and **B** as the light and four received the **A** as a light and **B** as the tone.

Following 10 conditioning sessions, the target cue **A** was presented for 20 trials alone in the absence of **X**, **B** or the US.

Results and Discussion

As predicted by the relative contingency learning hypothesis and the RWM, the conditioned response controlled by **A** was influenced both by its absolute contingency and by its contingency relative to **B**. In general, when **A** was a moderate predictor of the US,

Table 5. Trial types and frequencies, conditional probabilities and Δp calculation for **A** and **B** in relation to trial context **X** for treatments in Experiment 3.

	Moderate Positive Contingency (Δp_A=.5)		Zero contingency (Δp_A=0)	
$\Delta p_A/\Delta p_B$	**.5/1**	**.5/0**	**0/1**	**0/0**
ABX--->US	24	16	16	8
ABX--->no US	0	8	0	8
BX--->US	8	0	16	8
BX--->no US	0	8	0	8
AX-->US	0	8	0	8
AX-->no US	8	0	16	8
X-->US	0	8	0	8
X-->no US	24	16	16	8
Total US	32	32	32	32
Total no US	32	32	32	32
P(US|**B**)	1	.5	1	.5
P(US|-**B**)	0	.5	0	.5
Δp_B	1	0	1	0
P(US|**A**)	.75	.75	.5	.5
P(US|-**A**)	.25	.25	.5	.5
Δp_A	.5	.5	0	0

there was a greater increase in tray entries than when **A** was uncorrelated (.5/1, .5/0). However, **A** controlled a weaker tray entry response, regardless of its contingency if **B** was a perfect predictor (.5/1). Both **A**'s and **B**'s contingency also influenced whether the trial marker (lever) would elicit lever pressing. There were no stimulus counterbalancing effects during training and so the analysis is presented collapsed for this factor.

Tray Entries

Pre-CS intervals. There were strong differences in Pre-CS tray entries suggesting, that in spite of the use of the discrete trial marker and that none of the USs were presented in the absence of **A**, **B** or **X**, the static contextual cues controlled different levels of tray entry behaviour. Tray entries were lower in the two groups in which **B** was a perfect predictor of the US. Figure 11 presents the mean log tray entries in blocks of three sessions (omitting session 1) during the 10 second Pre-CS interval. It is clear that baseline tray entries were fairly consistent across training in the two groups in which **B** was uncorrelated with the US (i.e., treatments .5/0 and 0/0). However, baseline tray entries were much lower in the two treatments in which **B** was a perfect predictor of the US (treatments .5/1, 0/1). An ANOVA found a reliable difference in levels of responding dependent on **B**'s contingency (Δp_B=1 or 0) as well as the interaction between **B**'s contingency and Blocks [$F(1,32)=24.1$ and $F(2,64)=6.60$ respectively]. There was no significant change across training in the two groups receiving zero contingency training of **B** [$F(2,32)<1$] but there was a significant decrease in responding in the two groups in which **B** was perfectly correlated with the US [$F(2,32)=10.4$]. There were no reliable effects on Pre-CS responding involving **A**'s contingency.

CS trials. The general pattern of tray entries during the discrete trials shows increased tray entries when either **A** or **B**'s relationship with the US was positive. The groups in which **B** was a perfect predictor of the US (.5/1 and 0/1), showed more tray entries whenever **B** was present (**ABX** and **BX** trials). In the two treatments in which **B**

Figure 11.The Mean natural log tray entries per minute during the 10 second Pre-CS interval in three session blocks during training in Experiment 3. Four groups crossed two levels of contingency for cue **A** (0.5- squares; 0-triangles) and two levels of contingency for cue **B** (1- filled symbols; 0-open symbols).

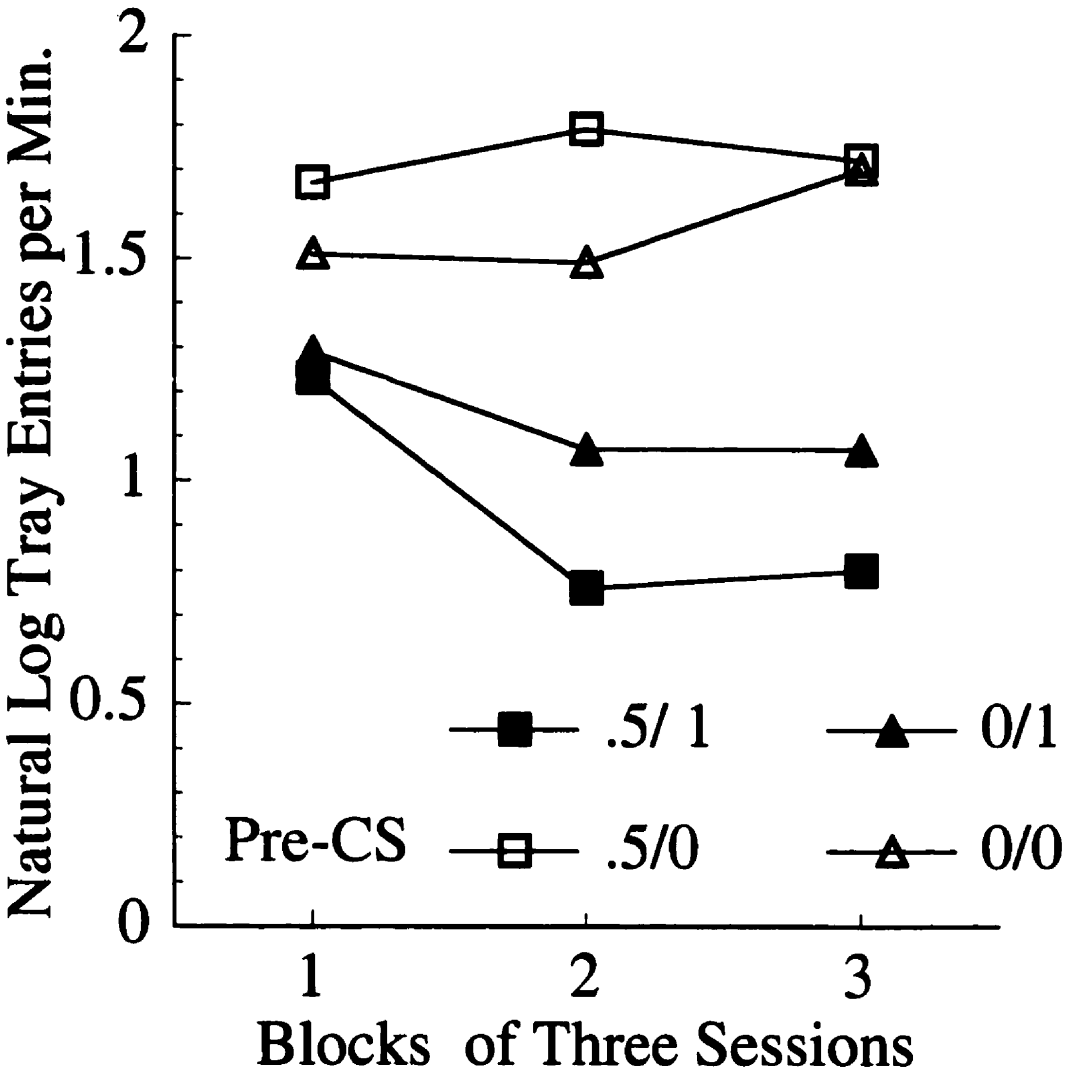
Natural Log Tray Entries per Min.
2
1.5
1
0.5
0
1
2
3
Blocks of Three Sessions
Pre-CS
.5/ 1
0/1
.5/0
0/0

had a zero contingency with the US (.5/0 and 0/0) **A** controlled tray behaviour when it was moderately positive. Finally, in treatment 0/0 when both **A** and **B**'s relationships with the US were zero there were fewer tray entries on all four types of trial. As we will see, in this group (0/0) the trial context **X** elicited increased lever pressing.

Figure 12 presents the mean natural log rates of tray entries separated by the four trial types. It is clear that there were more tray entries in the groups in which **B** was a perfect predictor (the filled symbols) on **ABX** and **BX** trials (top panels) compared to **AX** or **X** trials (bottom panels). This indicates, somewhat unsurprisingly, that more tray entries occurred on trials which were always followed by the US. An overall four-way analysis of variance for **A**-contingency (.5 or 0) , **B**-Contingency (1, 0) trial type (**ABX**, **AX**, **BX**, **X**) and Blocks showed that the four types of trial elicited different levels of responding [$F(3,96)=183$] and that these levels changed over the three blocks of trials. The main effect of blocks and the interaction between blocks and trial type were reliable [$F(2,64)=10.6$ and $F(6,192)=11.6$]. In addition, the pattern of responding to the four types of trial was influenced by the particular combination of **A/B** contingency. The interactions between trial type and both **B**'s contingency [$F(3,96)=155$] and **A**'s contingency [$F(3,96)=9.64$] were reliable. By the end of training on the last block of trials, tray responding was controlled by both **B**'s and **A**'s contingency in the different groups.

On **ABX** trials (the upper left hand panel of Figure 12) responding was highest in groups .5/1 and 0/1. These are the two groups in which **B** had a perfect positive contingency with the US. Responding was intermediate in .5/0 and lowest in 0/0. A two factor analysis of variance on the last block of **ABX** trials found a reliable effect for **B** and **A** contingency [$F(1,32)=14.9$ and $F(1,32)=7.09$] and an interaction between these factors [$F(1,32)=6.52$]. There were no differences between the three treatments containing at least one positive predictor, .5/1, 0/1 and .5/0 but tray entries were reliably lower in group 0/0 [$F(1,32)= 27.5$]. On **BX** trials (the upper right hand panel) tray entries rates were again

Figure 12. The Mean natural log tray entries per minute during trial presentation in three session blocks during training in Experiment 3. Four groups crossed two levels of contingency for cue **A** (0.5-squares; 0-triangles) and two levels of contingency for cue **B** (1-filled symbols; 0-open symbols). Response rates are separated by the four trial types presented (**ABX**, **BX**, **AX**, **X**).

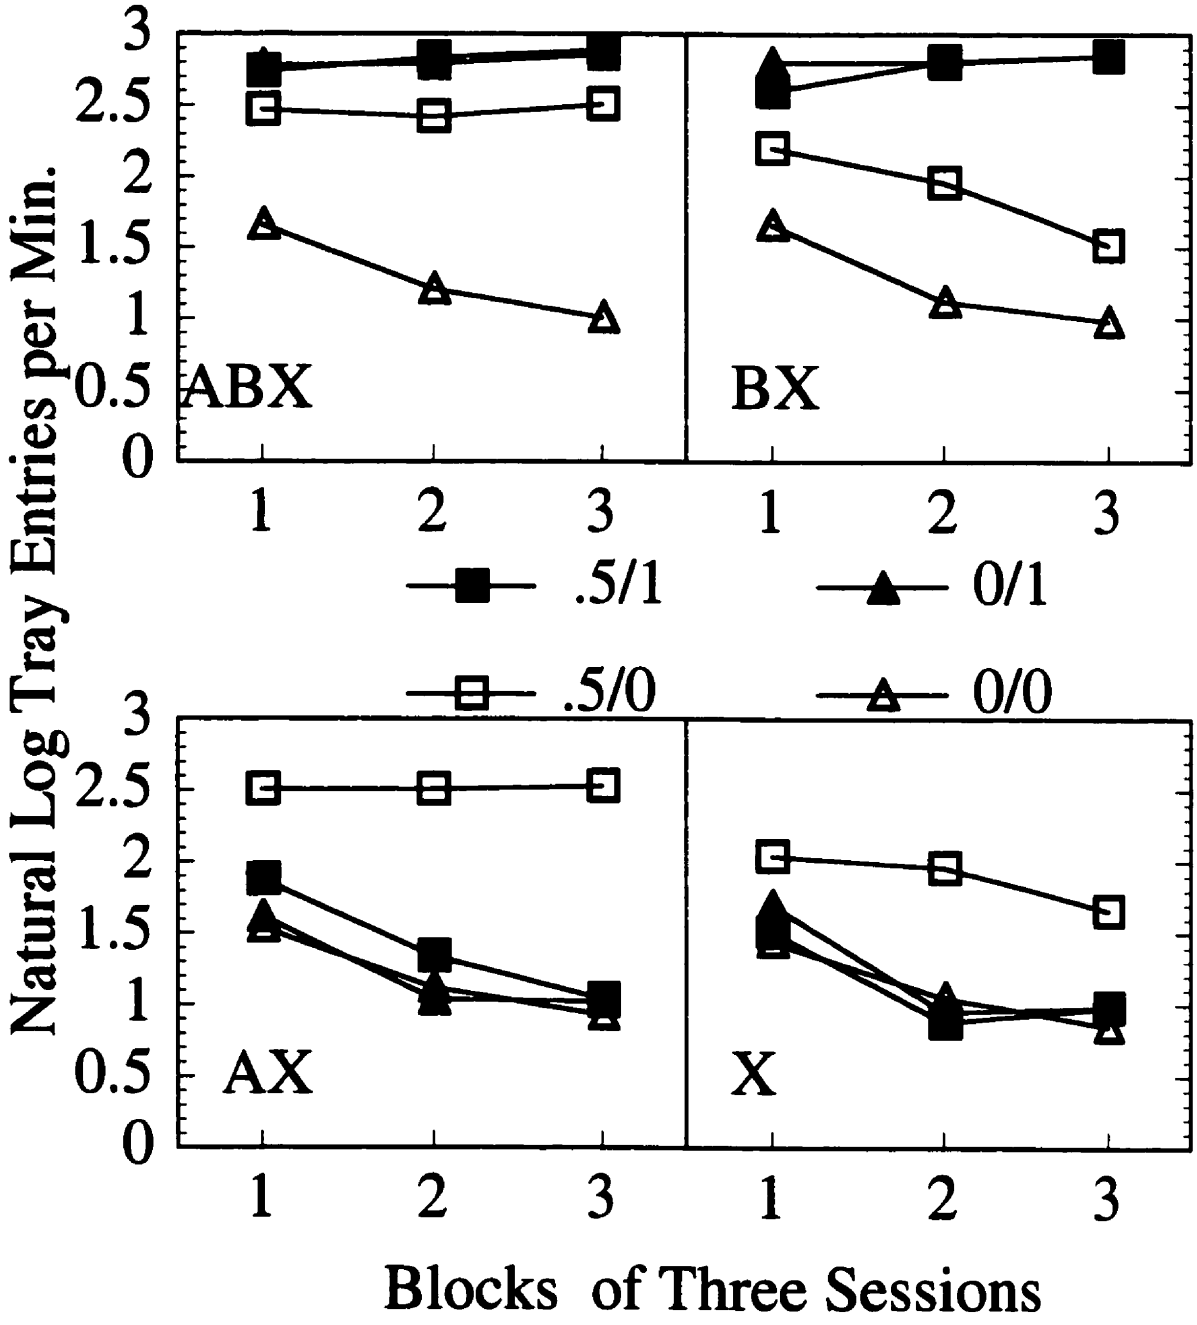
Natural Log Tray Entries per Min.
3
2.5
2
1.5
1
0.5
0
ABX
BX
AX
X
1
2
3
.5/1
0/1
.5/0
0/0
Blocks of Three Sessions

greater in the two treatments in which **B** had a perfect positive contingency. Only the main effect for **B**'s contingency was reliable [$F(1,32)=30.0$]. On **AX** trials (the lower left hand panel) responding was reliably greater in group .5/0 and lower in the other three groups . Both main effects for **A** and **B** contingency [$F(1,32)=5.11$ and $F(1,32)=6.87$] and the interaction of these two factors [$F(1,32)=6.5$] were reliable. A contrast confirmed that levels of responding in group .5/0 were higher than in the other three groups [[$F(1,32)=18.41$]. There were no reliable differences between any of the four treatments on tray entries during **X** trials [max $F(1,32)=2.03$], although group 0/0 did seem to respond somewhat more than the other three groups, the effect was not reliable.

Lever Pressing

CS trials. Figure 13 presents the mean log lever presses in blocks of three sessions. The overall pattern suggests an inversion of the tray entry data. There was more lever pressing in groups in which **B**'s relationship with the US was lower, (i.e., .5/0, 0/0), and there was almost no lever pressing in the other groups (.5/1 and 0/1). This pattern is consistent with the hypothesis that when **A** and **B** were of insufficient predictive value to control a conditioned tray entry response, the discrete trial marker **X** (the lever) which was a relatively more reliable predictor, became more strongly conditioned.

The statistical analyses of the lever pressing data support these observations. The ANOVA with **A**-Contingency (.5, 0), **B**-Contingency (1, 0) Trial Type (**ABX**, **AX**, **BX** and **X**) and Blocks of Sessions found reliable main effects for **B**-contingency, Trial Type [$F(1,32)=15.1$ and $F(3,96)=15.3$ respectively] and the Trial Type by Blocks interaction [$F(6,192)=5.84$].

The central concern in this experiment was whether the differences between the groups following training suggested that they were sensitive to the relative contingencies. Analyses of lever pressing during **ABX**, **AX**, **BX** and **X** trials separately on the last block of trials found the same overall pattern. There was more lever pressing on all four types of

Figure 13.The Mean natural log lever presses per minute during trial presentation in three session blocks during training in Experiment 3. The lever was stimulus **X**. Four groups crossed two levels of contingency for cue **A** (0.5- squares; 0- triangles) and two levels of contingency for cue **B** (1-filled symbols; 0-open symbols). Response rates are separated by the four trial types presented (**ABX**, **BX**, **AX**, **X**).

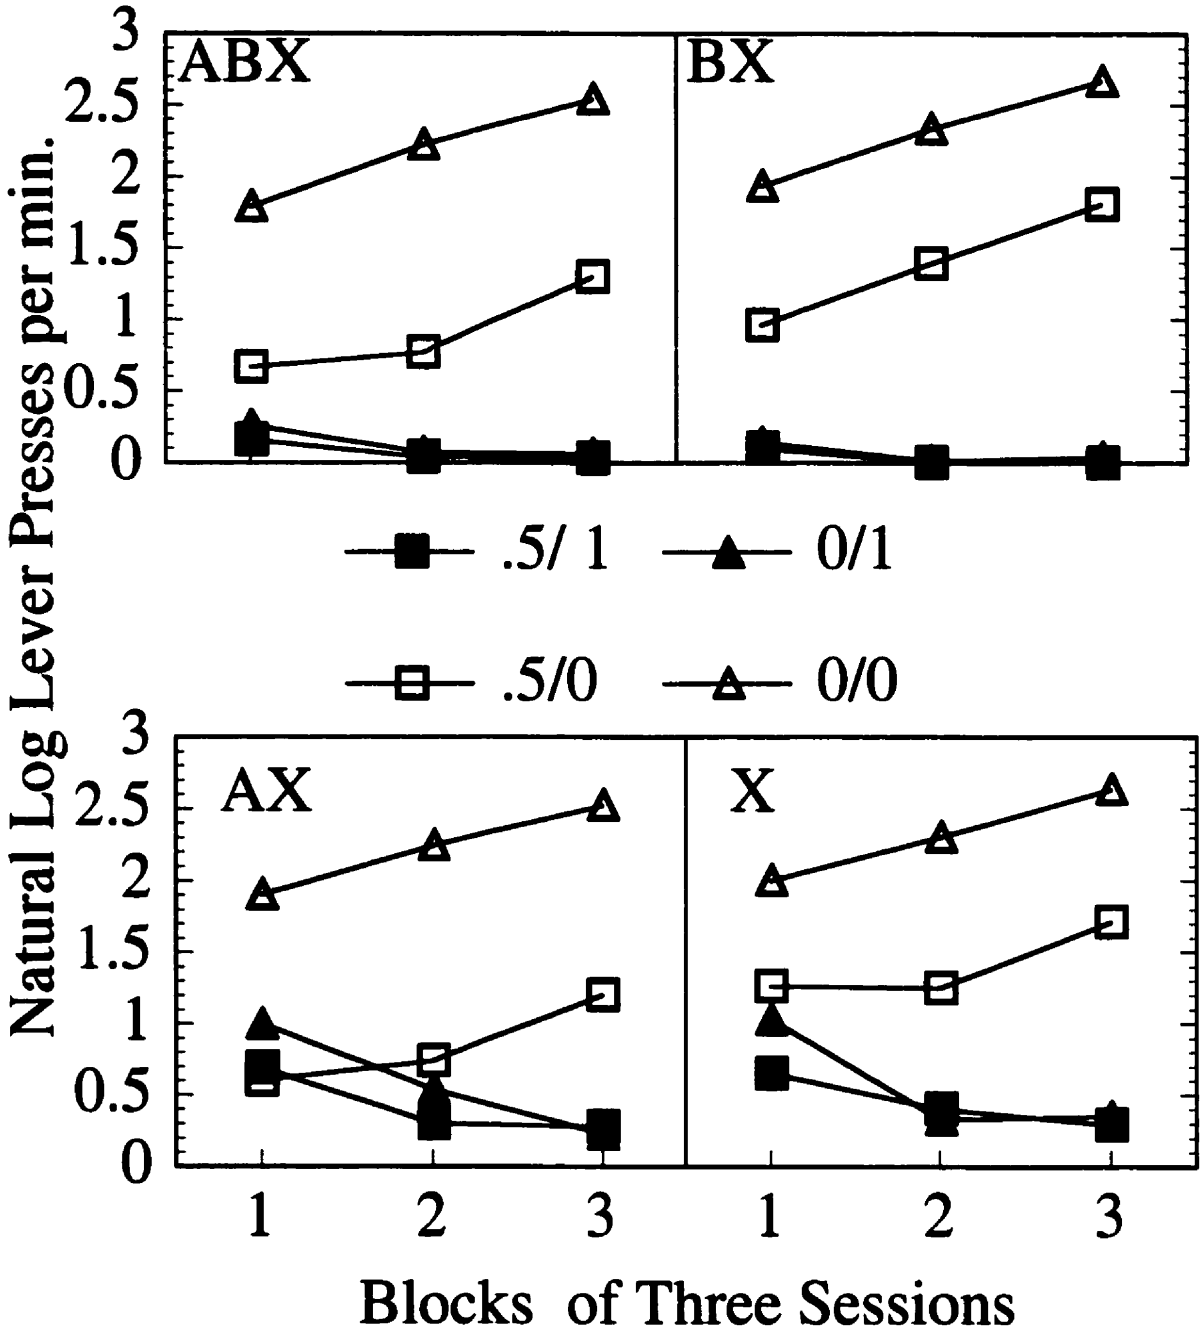
ABX
BX
AX
X
Natural Log Lever Presses per min.
3
2.5
2
1.5
1
0.5
0
.5/ 1
0/1
.5/0
0/0
1
2
3
Blocks of Three Sessions

trial in the two groups when **B** was relatively less predictive (.5/0 and 0/0). Only the main effect for **B**'s contingency was reliable [$F(1.32)=14.9$]. Generally lever pressing was blocked only when **B** was a perfect predictor of the US.

Test

Of primary interest was how the rats would respond to cue **A** when it was presented without the trial marker (ie., the lever) during the extinction test. These rates and each group's corresponding Pre-CS rates during the test are presented in Figure 14. Tray entries during the test of the moderately correlated or uncorrelated cue **A** were determined by two factors, **A**'s absolute contingency with the US and by its contingency with the US relative to **B**'s contingency. Consistent with the relative contingency hypothesis there was a greater elevation of tray entries when **A** was positively related to the US (.5/1 and .5/0) and **A** elicited a greater increase when **B** had a zero contingency, regardless of **A**'s contingency (.5/0 and 0/0). In addition, the analysis of the counterbalancing manipulation found that as a CS the tone elicited more overall tray entries than the light. The overall analysis of responding to **A** as a function of **A**-Contingency (.5, 0), **B**-Contingency (1,0), Stimulus (Light or Tone) and Response Interval (Pre-CS, CS) found reliable main effects for **B**-Contingency and Response Interval [$F(1,28)=10.62$ and $F(1,28)=19.5$ respectively]. These effects support the observation that tray entries were greater when B was uncorrelated and during the CS rather than the Pre-CS. In addition the interactions between Response Interval and **B**-Contingency [$F(1,28)=6.62$] and Response interval with **A**-Contingency [$F(1,28)=4.05$] were both reliable, suggesting that the conditioning response interval effect was influenced by both **A** and **B** contingencies independently. The only effect involving the stimulus counterbalancing manipulation was the **A**-Contingency by Stimulus interaction [$F(1,28)=10.8$]. Separate tests of the **A**-Contingency effect indicated that it was reliable for both stimuli [minimum $F(1,16)=6.06$].

The two interactions of Stimulus contingency and response interval confirm that the

Figure 14. Mean natural log rates of tray entries in the four treatments during the Pre-CS and stimulus **A** during the extinction test of Experiment 3.

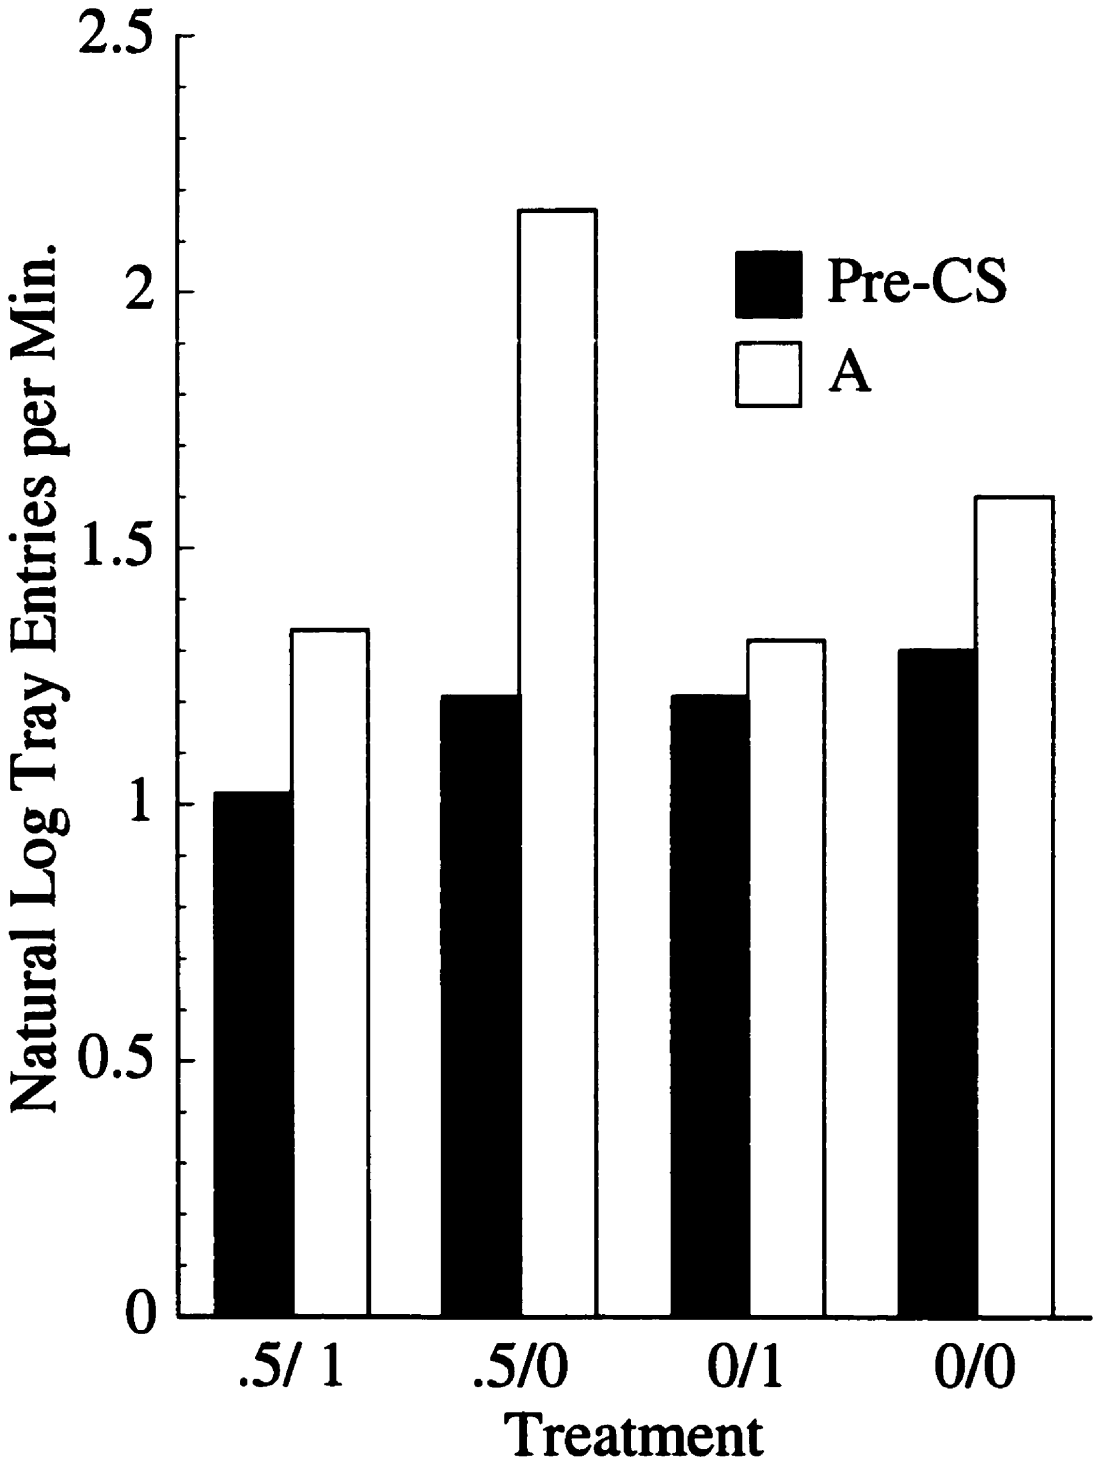
2.5
2
1.5
1
0.5
0
Natural Log Tray Entries per Min.
Pre-CS
A
.5/ 1
.5/0
0/1
0/0
Treatment

conditioned response elicited by **A** was influenced by both **A**'s absolute contingency and in relation to **B**'s contingency. Individual tests of the elevation in tray behaviour elicited by **A** in comparison with the Pre-CS in each of the four treatments found no evidence of this conditioning to **A** in treatment 0/1 [$F(1,7)<1$]. **A** was blocked by **B**. There were reliable levels of tray entries in both groups .5/0 and 0/0 [Minimum $F(1,7)=8.21$]. It is interesting to note that although the overall level of responding in the .5/1 group was low both during **A** and the Pre-CS that the conditioning effect was just reliable suggesting that there was still a reliable level of conditioned responding to **A** in this treatment [$F(1,7)=5.75$]. Therefore, the presence of a perfect predictor did not completely block the moderately positive cue **A**'s ability to control tray entries.

These results are consistent with the hypothesis that rats are sensitive to the relative validity of the cues with validity defined as CS-US contingency. The Δp contingency influenced the response controlled by **A**. **A** produced a greater elevation in tray behaviour above that observed during the Pre-CS when **A** was positively correlated with the US. Furthermore, the observed elevation was mediated by **B**'s contingency. The question remains whether this type of effect is clearly predicted by any associative theory or whether it requires that the animals calculate the conditionalized contingencies between all the cues and the US (Cheng, 1997). The Rescorla-Wagner model was used to simulate the four treatments of this experiment. Figure 15 displays the asymptotic predictions for cue **A** from the model using the parameters for the simulations in Experiment 1. A visual comparison suggests that the ordinal relation of the means from the data match the simulations. In both conditions in which **B** was a perfect predictor of the US, the model predicts that most of the associative strength is acquired by **B**, however, before **B** has acquired this strength **A** is able to acquire some small amount of associative strength. This strength is eventually lost but a residue is present even after 500 trials which produces the small difference in strength between groups .5/1 and 0/1. When **B** is uncorrelated with the US, **A** is able to compete

with the context for associative strength and even more so when it is positively correlated with the US. When both cues are uncorrelated then the context is free to acquire most of the associative strength leaving almost none for **A** to acquire (although there was still a small conditioning effect for **A**).

There are two measures of context conditioning that can be used to test this prediction about the strength of the context. Firstly, the differences in Pre-CS tray entries during the test provide one measure. Examination of Figure 14 and the statistics suggested that Pre-CS tray entries rates were greatest in group 0/0 but the statistics showed that this group did not differ from .5/0. A second measure comes from the level of conditioning elicited by the trial marker **X**. The lever pressing data does support the prediction of the model in relation to the contextual conditioning. Although it looks like group 0/0 elicited more lever pressing than any of the other three treatments, a post hoc analysis of the lever pressing response to **X** on the last block of training trials showed that group 0/0 showed reliably more lever pressing than either group .5/1 or 0/1 [minimum $F(1,16)=13.9$], there were no reliable differences involving group .5/0. Therefore in general the data support the associative model's predictions, specifically that contextual conditioning, as measured by lever pressing to the trial context, should be greater when none of the cues are perfect predictors.

Previous work on the relative validity effect has shown how behaviour controlled by a CS with positive predictive relationships with the US can be reduced by concurrent training with stronger predictors of the US (Wagner et al., 1968; Cole, Barnet, Miller, 1997). In this experiment a relative validity effect was also demonstrated with a CS trained with a zero contingency. This experiment is the first demonstration of blocking of a zero contingency by a positive predictor. The results suggest that the zero contingency training generated a level of association with the US that could be blocked. This result has also been found with human judgments of contingency (Baker, Mercier, Vallée-Tourangeau, Frank &

Figure 15.Predictions of the cumulative final associative strength for cue **A** in Experiment 1 using the Rescorla-Wagner (1972) model. The simulations involved presenting the same ratio of trial types for the four treatments described in Table 5 and using the same learning parameters as those used previously. Cue **A** was presented with either a positive or zero CS-US contingency (Δp=.5 or 0). The trial context had the same overall relationship with food in each treatment. Predictions for **A** following the two contingency treatments are shown as a function of **B**'s contingency.

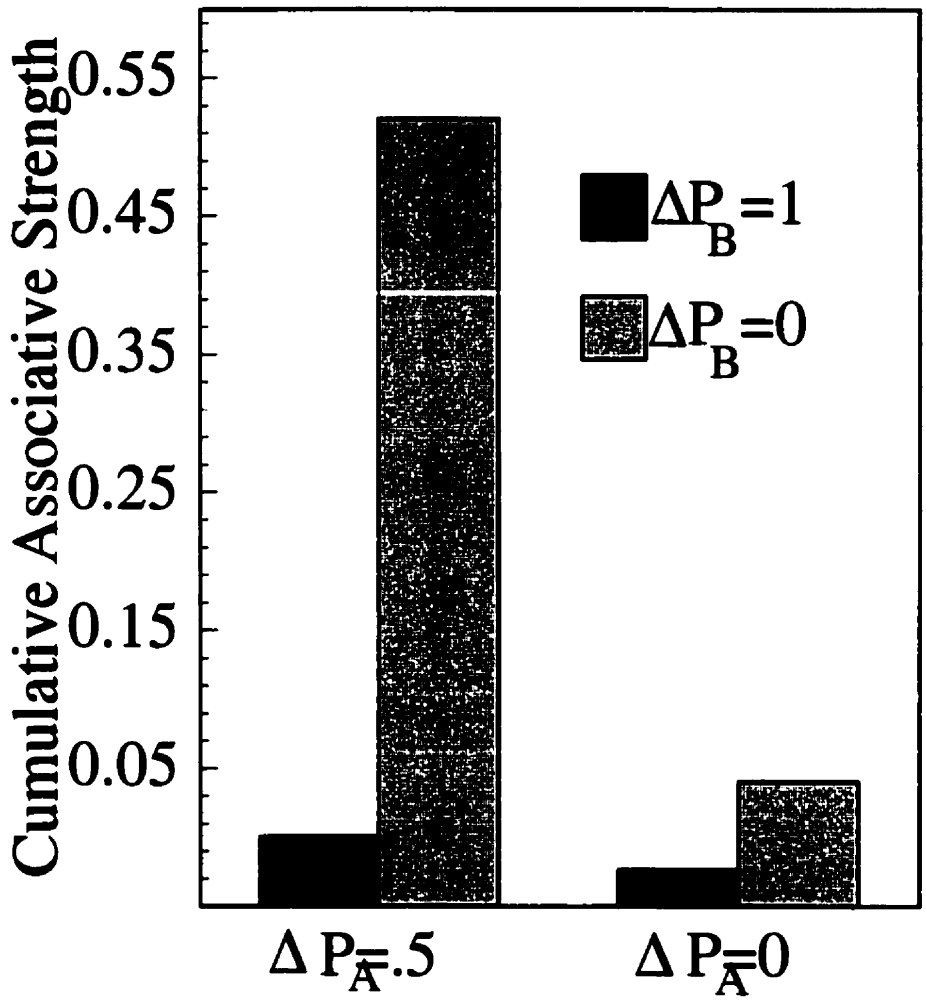
Cumulative Associative Strength
0.55
0.45
0.35
0.25
0.15
0.05
ΔP_B=1
ΔP_B=0
ΔP_A=.5
ΔP_A=0

Pan, 1993). Human subjects will consistently judge the causal effectiveness of a cue as being lower if it has been trained with a more valid predictor.

One interpretation of these results which is consistent with both the associative and normative models is that the rats were sensitive to the relative contingencies of the cues. Reductions to both A's absolute and relative contingency reduced the strength of the conditioned tray response. This suggests the possibility that both have a similar effect on a cue's validity as a predictor of the US in support of the ΔP definition of validity. With respect to the sufficiency of the simple unconditional CS-US contingency to elicit a conditioned response, Papini and Bitterman (1990) are correct to point out that absolute contingency can not be the sole determinant of the CR. At the computational level, in addition to the absolute unconditional contingency the animals in these experiments seem to be sensitive to the more complex conditional contingencies in which a cue's contingency is calculated conditional upon the presence of other cues (Baker, Murphy &Vallée-Tourangeau, 1996; Cheng & Holyoak, 1995).

Experiment 4

One problem with the design used in Experiment 3 is that it may insufficiently control for an alternative interpretation. The target cue **A** was presented during two types of trial **ABX** and **AX**, and due to the nature of the contingencies, the four treatments received varying number of these trial types. It may be possible that the strength of the conditioned response elicited by **A** arose not from its absolute and relative contingency but from the amount of generalised excitation from **AX**. Simply by comparing experience with **AX** in the four groups one could predict the ordinal relation between the four groups. In the .5/0 group all the **AX** trials were paired with the US while in the 0/0 group only 50% were paired with the US and in the two blocking groups (.5/1, 0/1) there was no experience of **AX** and the US. In .5/1 and 0/1 there were no pairings with the US but 8 and 16 extinction trials each session respectively (see Table 5). Therefore, simply comparing the four groups

on the basis of experience with the compound which is most similar to the test cue **A**, in this case **AX**, one might anticipate the present pattern of responding to **A** alone on the test (i.e., .5/0>0/0>.5/1>0/1). This ordering matches the conditioning data in Figure 15, therefore contingency learning could emerge from an associative network that makes use of the notion of stimulus configurations and generalised associative strength (e.g., Pearce, 1987).

The relative number of **AX**-US alone would not explain why there was a reliable elevation in tray entries to **A** in the .5/1 group since this group received no pairings of **AX** with the US. However, it is still possible that it was experience with **AX** that contributed to the present pattern of results. This interpretation also requires that **AX** trials contributed more to the response controlled by **A** than **ABX** trials. At least one theory of conditioning makes this same prediction. Pearce's (1987; 1994) configural model of associative learning predicts that **AX** trials should have a stronger impact on the response elicited by **A** than **ABX** trials. The model, described in more detail in the introduction, predicts that the response controlled by **A** will be determined by a weighted product of the associative strengths of configurations similar to **A**. While the Pearce model does not specify precisely how similarity might be instantiated, it does suggest that as a first approximation similarity be calculated on the basis of the ratio of the proportion of common elements to unique elements in the test stimulus and the configural stimulus. Since half of the stimuli present in **AX** are present in the test stimulus **A**, presentation of **A** alone should elicit half of the associative strength of the compound **AX**. However, because **A** only makes up one third of the stimuli present in the triple compound **ABX** only one third of the associative strength of compound **ABX** with the US would be elicited by the presentation of **A**. Therefore, other things being equal this theory suggests that **AX**'s association with the US will be generalize more strongly by presentation of **A** than by **ABX**'s.

One way to test this alternative account of the data in Experiment 3 is to employ

training in which the target stimulus **A** is always embedded in the same number of alternate predictors. The previous explanation arises because **A** is sometimes presented with two cues **B** and **X** on **ABX** trials, but other times presented with only one other cue, **X** on **AX** trials. However, if **A** was always embedded in the same number of cues then there would be no difference in experience with compounds between the different contingency treatments except with respect to **A**'s overall contingency. Experiment 4 involved a systematic extension of the previous findings but differed in that **A** was always trained in compound with two other cues either **BX** or **CX**. If the response to **A** is still a function of its relative contingency then the argument that conditioning to **A** was simply a function of differential experience with the training compounds would be less tenable.

Experiment 4 included a systematic replication of two treatments from Experiment 3 (.5/1 and .5/0). These two treatments were also run to test another possible explanation of the data reported in Experiment 3. There are many different combinations of trials which will produce the same nominal contingency for two cues and a US. For example, for a given combination of contingencies, the actual frequency of each trial type, and the distribution of USs can differ without altering the overall unconditional contingencies between each CS and the US. The experimenter's decision on how to instantiate contingencies in trial frequencies then may influence the outcome of the experiment. In Experiment 3 the control treatment, .5/0, was designed to equate the frequencies of each trial type with that in the blocking treatment .5/1 (Compare the trial frequencies in Table 5 ignoring whether the US was present). One by-product of this choice is that in the control treatment .5/0 there is an increase in **A**'s contingency conditional upon the presence and absence of **B**. **A**'s contingency can be estimated unconditionally in which case all trials containing a given cue are entered into calculations of its contingency with the US. Alternatively **A**'s contingency can be calculated conditional upon the presence of a set of local contextual cues (Cheng & Novick, 1992; Baker, Murphy & Vallée-Tourangeau, 1996;

see also Figure 2)

Table 6 presents the trial frequencies only for the control treatment .5/0 used in Experiments 3 and 4. The calculation of **A**'s unconditional contingency indicates that it possesses a moderate positive contingency with the US Δp_A=.5. This relationship is the same as the one programmed in the comparison experimental .5/1 blocking treatment. However, calculating **A**'s contingency in either the presence ($\Delta p_{A|B}$) or absence of **B** ($\Delta p_{A|-B}$) indicates that **A**'s conditional contingency in treatment .5/0 is .67, whereas in treatment .5/1 the conditional contingency is .5 regardless whether the calculation is unconditional or conditional upon the presence of **B**. Therefore **A**'s conditional contingency with respect to **B** and the US is different in these two treatments in Experiment 3. Having two relationships with the US following a single training regime is similar to Simpson's paradox (Spellman, 1996) in which a stimulus possesses more than one predictive relationship with an outcome depending upon the choice of comparison cue. This is of course very similar to the problem that Experiment 1 was meant to address. Use of the discrete trial marker was meant to define the subjective CS-US contingencies.

In Experiment 3 the high level of responding in the .5/0 group may reflect the animal's sensitivity to the increase in **A**'s contingency conditional upon the presence of **B** rather than its contingency relative to **B**. There has been little previous research designed to test whether animals are sensitive to these type of relationships. Holland (1989) has shown that animals are sensitive to simple conditional relationships, but there is no research that systematically tests the type of conditional statistical relationships described here. One goal of Experiment 4 was to demonstrate the relative validity effect from Experiment 3 independent of any changes in conditional contingencies.

Experiment 4 involved a .5/1 experimental treatment along with a .5/0 control which eliminated this problem. The trial frequencies used in Experiment 4 are presented in Table 6 for comparison. Note that **A**'s unconditional contingency of .5 matches its conditional

Table 6. Trial types and frequencies, conditional probabilities and Δp calculation for treatment .5/0 in Experiments 3 and 4 comparing the unconditional Δp_A and the conditional contingencies in the presence of **B** ($\Delta p_{A|B}$) and its absence ($\Delta p_{A|\text{-}B}$).

Unconditional Contingency Calculations

Expt 3	US	noUS	p(US)
A	24	8	.75
-A	8	24	.25
		Δp_A	.5

Expt 4	US	noUS	
	30	10	.75
	10	30	.25
		Δp_A	.5

Conditional Contingency Calculations

Expt 3	US	noUS	
A.B	16	8	.67
-A.B	0	8	0
		$\Delta p_{A\|B}$	.67
A.-B	8	0	1
-A.-B	8	16	.33
		$\Delta p_{A\|\text{-}B}$	.67

Expt 4	US	noUS	
	15	5	.75
	5	15	.25
		$\Delta p_{A\|B}$	.5
	15	5	.75
	5	15	.25
		$\Delta p_{A\|\text{-}B}$	.5

contingencies in the presence and absence of **B**.

In Experiment 4 a moderate positive cue (Δp_A=.5) was paired with an either more or less valid cue. Two treatments using the new control for conditional contingencies from Experiment 3 (.5/1 and .5/0) involved training with **ABX**, **BX**, **AX** and **X** trials. Two new groups were used which received the same training except that on trials without **B** (**AX** and **X** trials) a fourth stimulus **C** was added (**ACX** and **CX**). The two groups with this extra stimulus received the same moderate positive **A** and either perfect or uncorrelated **B**. Now however, stimulus **C** was included and was either a perfect predictor of the absence of the US or was uncorrelated with the US. These new treatment groups .5/1/-1 and .5/0/0 are very similar to the groups used by Wagner, Logan, Haberlandt & Price (1968) in their demonstration of the relative validity effect. In their design, the common element (**X**) was always paired with either of two cues (**BA**, **CA**) but **A**'s CS-US contingency was not defined because there were never any trials in the absence of **A** and therefore it is difficult to tell whether their results reflect sensitivity to **A**'s relative contingency.

Although exclusion of the trial context during the extinction test in Experiment 3 did not eliminate the contingency effects described in Experiment 2, in the present experiment two tests of conditioning to **A** were performed. The first included the trial context. Following this test the animals received additional training followed by a second test of **A** in the absence of the trial context. These two tests were conducted to compare the effects of testing in the presence and absence of the discrete trial context, to ensure that responding to cue A was not a function of a conditioned response controlled by the trial context.

Method

Subjects: Forty-eight rats were used in this experiment, all other details of the subjects are the same as described in the previous experiments.

Apparatus: In addition to the auditory stimulus used in Experiment 3 (i.e., the square

wave tone) subjects in .5/1/-1 and .5/0/0 also received training with an auditory click stimulus. The clicker was a 20 Hz string of clicks (approximately 80 Db) presented from the same speaker as the tone.

Procedure: The rats were first trained to retrieve food from the food tray using the method described in Experiment 1. Following this session the rats were divided into four treatment groups. Two treatments (.5/1 and .5/0) were similar to the two treatments from Experiment 3 including only two CSs (**A** & **B**) and the trial context (**X**). The other two treatments included another CS (**C**). The notation for labelling these two treatments refers to the Δp contingency of **A** and **B** and **C** ($\Delta p_A/\Delta p_B/\Delta p_C$). Groups .5/1/-1 and .5/0/0 received training with **A**, **B** and **C**. Cue **A** was moderately related to the US and **B** and **C** were either perfect predictors (one positive and the other negative) or uncorrelated with the US.

A total of 80 trials were presented during each hour long session. Trials were presented on the basis of an irregular 25 second schedule (2-48 s). Each trial was ten seconds in duration. The start of each trial was marked by the entrance of the lever (**X**). In addition, on most trials combinations of the other cues were also present. In all treatments cue **A** had a moderate positive contingency with the US [p(US|**A**)=.75; p(US|-**A**)=.25; Δp=.50]. This was accomplished by pairing 75 % of the trials involving **A** (**ABX** and **AX** or **ACX**) with the US but only pairing 25 % of the trials on which **A** was absent with the US. In treatments .5/1, .5/0 the four types of trial were **ABX**, **BX**, **AX** and **X**. For these two treatments, the stimuli in the role of **A** and **B** were counterbalanced so that for half of each treatment **A** was the light and for the other half it was the tone. In the other two treatments, the same trial types were presented except that on all the trials in which **B** was not presented, the extra stimulus **C** was present. The two treatments .5/1/-1 and .5/0/0 therefore received training with **ABX**, **BX**, **ACX** and **CX**. For these treatments **A** was always the light, **B** was always the tone and **C** was always the clicker. Table 7 presents the

frequencies for each trial type during training as well as the frequency with which each type of trial was paired with the US.

Training continued for 14 sessions followed by a single extinction test session in which **AX** and **X** were each presented 20 times in the absence of **B** and **C** and the US. Following this test the conditioning training continued during phase 2 for a further 8 sessions followed by another test of **A** by itself.

Results and Discussion

Overall the findings of this experiment were consistent with the findings of Experiment 3. In spite of the two controls introduced, the fourth cue **C** and equating the conditional with the unconditional contingency of **A**, rats still showed sensitivity to the relative contingency of the cues. Comparison of the training and test data in groups .5/1 and .5/0 with .5/1/-1 and .5/0/0 suggests that there was little difference in the pattern of responding as a result of the inclusion of cue **C**. As in Experiment 3 the two treatments with a perfect cue **B** (.5/1 and .5/1/-1) show concentrated tray entries on trials that included **B**. This discrimination was not found in the two groups in which **B** was uncorrelated with food (.5/0 and .5/0/0).

Tray Entries

Pre-CS intervals. There were reliable differences in rates of Pre-CS tray entries. These are shown in the top left panel of Figure 16. The two groups with a perfectly correlated CS (.5/1 and .5/1/-1) resulted in reliably weaker rates of tray entries during the Pre-CS. In each of the experiments of this thesis Pre-CS tray entry rates have been shown to be sensitive to the relative validity of the cues. During the first 3 training blocks the overall ANOVA for **B**-Contingency [1 or 0], whether the fourth stimulus **C** was included during training, and blocks (1-3) found that only the main effect for **B**'s contingency was reliable [$F(1,44)=9.01$]. A similar analysis of the second phase of training during Blocks 4 and 5 following the first test found the same reliable main effect [$F(1,44)=12.7$].

Table 7. Trial types, frequencies, conditional probabilities and Δp calculations for treatments in Experiment 4.

	Single Blocker		Two Blockers	
$\Delta p_A/\Delta p_B/\Delta p_C$	**.5/1**	**.5/0**	**.5/1/-1**	**.5/0/0**
ABX--->US	30	15	30	15
ABX--->no US	0	5	0	5
BX--->US	10	5	10	5
BX--->no US	0	15	0	15
AX-->US	0	15		
AX-->no US	10	5		
X-->US	0	5		
X-->no US	30	15		
ACX-->US			0	15
ACX-->no US			10	5
CX-->US			0	5
CX-->no US			30	15
Total US	40	40	40	40
Total no US	40	40	40	40
p(US|**B**)	1	.5	1	.5
p(US|-**B**)	0	.5	0	.5
Δp_B	1	0	1	0
p(US|**C**)	-	-	0	.5
p(US|**C**)	-	-	1	.5
Δp_C	-	-	-1	0
p(US|**A**)	.75	.75	.75	.75
p(US|no **A**)	.25	.25	.25	.25
Δp_A	.5	.5	0	0

CS trials. The other four panels in Figure 16 show the mean log rates of tray entries during the different trial types. In general, the rats showed discrimination between the four types of trial in accordance with prediction based upon the relative contingencies. In the two groups in which **B** was a perfect predictor of food (.5/1 and .5/1/-1) tray entries were greater on trials with **B**. Tray behaviour in the other two groups was moderately high on all trial types. The statistical analysis of the two between-group factors **B**-contingency (1 or 0), whether **C** was included in the discrimination, as well as the two within factors, the trial types (**ABX**, **BX**, **AX** or **ACX**, **X** or **CX**), and blocks (1-3) found a reliable main effect for trial type and the trial type by **B** contingency interaction [$F(3,132)=103$ and $F(3,132)=76.3$]. There were also reliable block effects on some trial types. The two-way trial types by blocks and the three way interaction of trial types, **B**-contingency and blocks were also reliable [$F(6,264)=8.30$ and $F(6,264)=4.01$]. None of the effects involving the presence and absence of was reliable. The effects for trial types and the interaction between trial types and **B** contingency were also reliable during phase two following the first test (blocks 4 and 5 minimum $F(3,132)=74.3$]. None of the effects involving blocks was reliable in phase 2. The absence of a blocks effect suggests that learning was at asymptote.

Separate analyses of tray entry rates on each of the four trial types was justified by the reliable interactions.There were no reliable effects for the inclusion of **C**. During **ABX** trials responding was high in all four treatments, there were no reliable differences during blocks 1-3 [max $F(1,44)=2.7$]. During Blocks 4-5 there was a small effect for **B** contingency suggesting that responding on **ABX** trials was greater for the two groups in which **B** was a perfect predictor of the US [$F(1,44)=4.37$]. On **BX** trials there was more reliable evidence for this difference. The effect for **B** contingency on Blocks 1-3 was reliable [$F(1,44)=38.8$] as was the main effect for Blocks and the interaction of these two factors [$F(2,88)=9.67$ and $F(2,88)=3.69$]. On Blocks 4-5 the effect for **B** contingency was also reliable [$F(1,44)=6.47$]. During the **AX** or **ACX** trials it is clear that tray entries in the

Figure 16. The Mean natural log tray entries per minute during trial presentation in four session blocks during training in phase 1 and phase 2 of Experiment 4. Four treatments involved a moderately positive cue **A** ($\Delta p=0.5$) with either perfectly predictive or correlated alternate cues). Response rates are separated by the four trial types presented (**ABX**, **BX**, **AX or ACX**, **X or CX**). Two treatments received training with stimulus **C** which was present on trials in which **B** was not present.

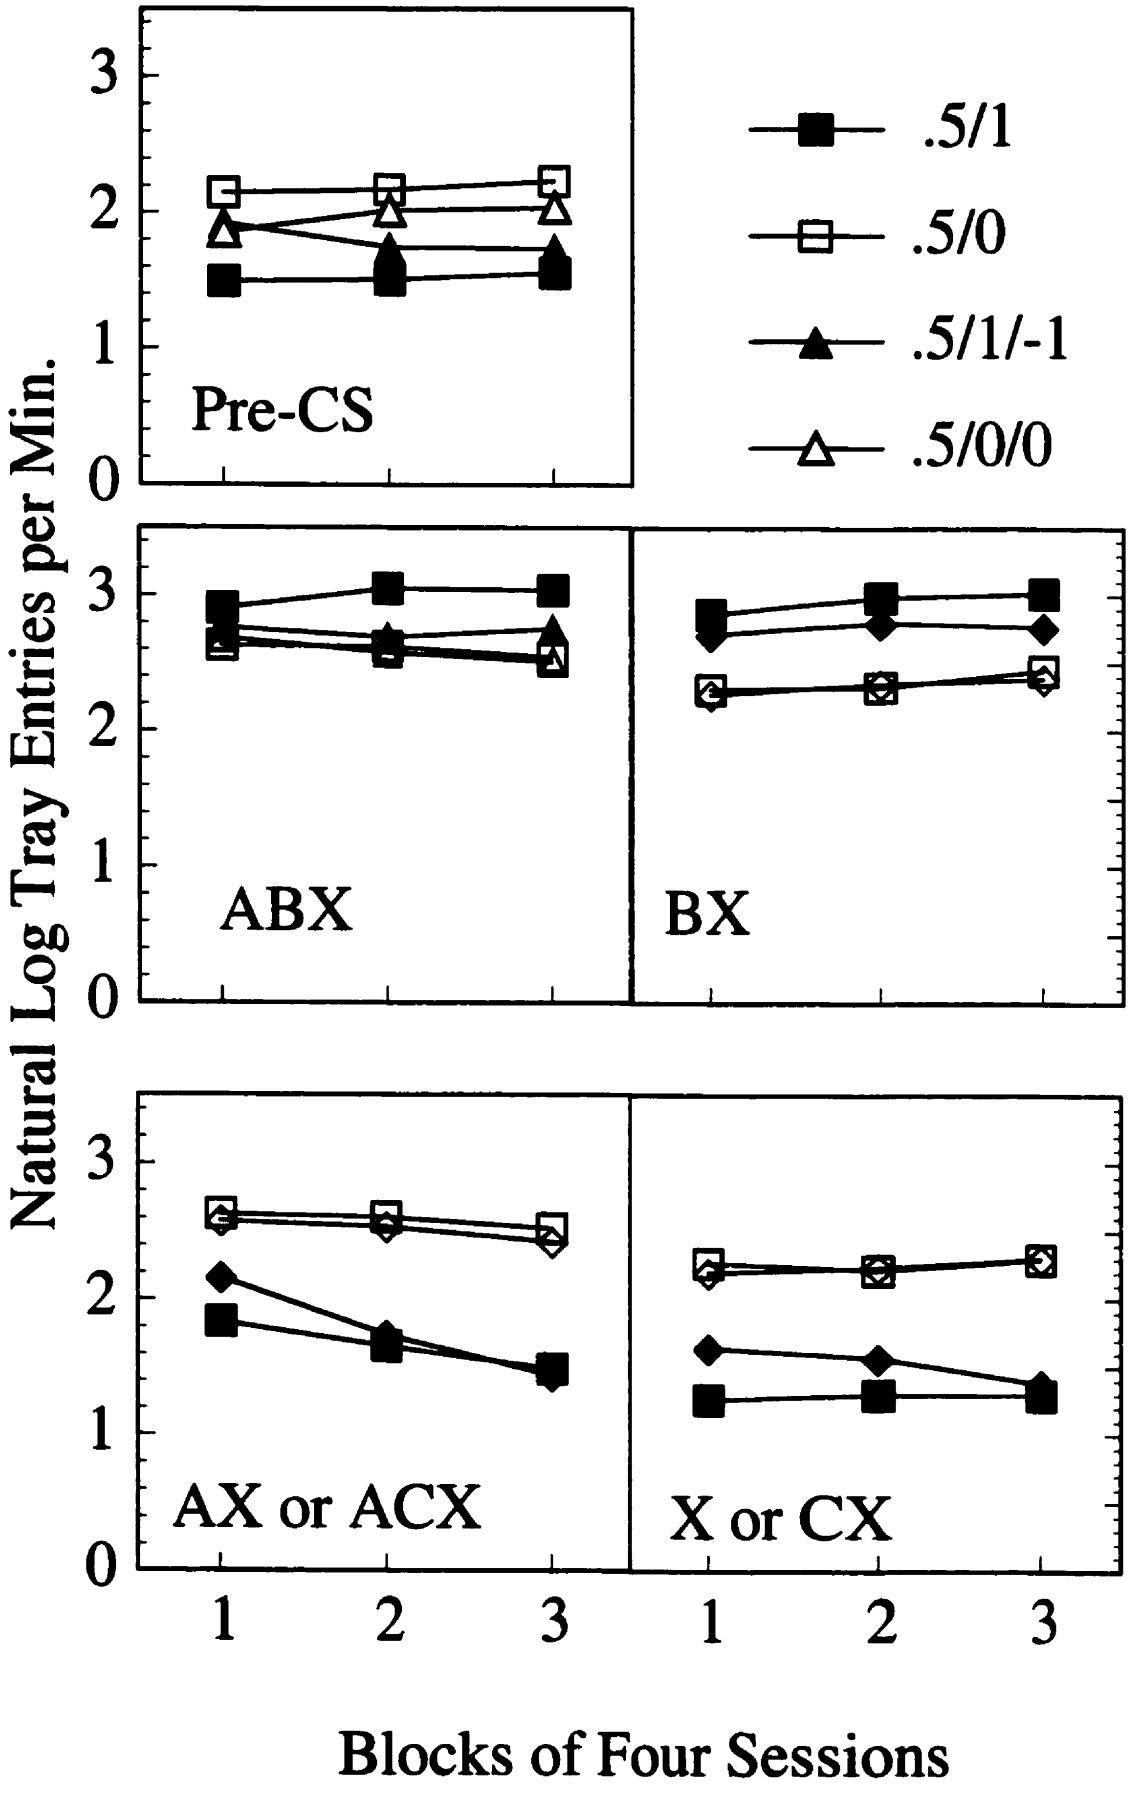
.5/1
.5/0
.5/1/-1
.5/0/0
Pre-CS
ABX
BX
AX or ACX
X or CX
Natural Log Tray Entries per Min.
Blocks of Four Sessions
0
1
2
3

two treatments in which **B** was a perfect predictor decreased. The main effects for **B** contingency and Blocks, as well as the interaction between the two factors, were all reliable and the **B** contingency effect was maintained on Blocks 4 and 5 [minimum $F(2,88)=3.69$]. The pattern observed on **X** or **CX** trials was similar except that the effects involving Blocks were not reliable. However, the main effect for **B** contingency was reliable during both the initial training phase and the second phase [minimum $F(1,44)=42.1$]. Overall this pattern of data is similar to that found in Experiment 3 and suggests that tray behaviour was not substantially influenced by either the inclusion of cue **C** or the change in conditional contingency between A and the US.

Lever Pressing

CS trials. Levels of conditioned lever pressing were also a function of **B**'s contingency. Figure 17 shows how the pattern for this behaviour is reversed in comparison with the tray entries. Lever pressing was greater in the two groups in which **B** was uncorrelated. The overall ANOVA for Trial Type, **B** contingency, whether **C** was included and Blocks found reliable effects for Trial Type, **B** contingency, Blocks and three two way interactions, Trial Type by Blocks, **B** contingency by Blocks and Trial Type by **B** contingency [minimum $F(2,88)=3.62$]. The main effects for Trial Type and **B** contingency conducted on the phase two training from Blocks 4-5 were also reliable, but none of the effects involving blocks was significant.

Individual analyses of the rates of lever pressing during **ABX**, **BX**, **AX** (or **ACX**), **X** (or **CX**) trials found that lever pressing was always stronger in the two groups in which **B** was uncorrelated and **A** and **X** were relatively more reliable predictors of the US. The main effects for **B** contingency were reliable on each type of trial [minimum $F(1,44)=10.8$]. The main effects for Blocks were only reliable on trials containing **A**, **ABX** and **AX** or **ACX** trials [minimum $F(2,88)=5.45$]. Reliable two way interactions between **B**-contingency and Blocks confirm the observation that lever pressing showed an

Figure 17.The Mean natural log lever presses per minute during trial presentation in four session blocks during training in Experiment 4. Four treatments involved a moderately positive cue **A** ($\Delta p=0.5$) with either perfectly predictive or correlated alternate cues. Response rates are separated by the four trial types presented (**ABX**, **BX**, **AX or ACX**, **X or CX**). Two treatments received training with stimulus **C** which was present on trials in which **B** was not present.

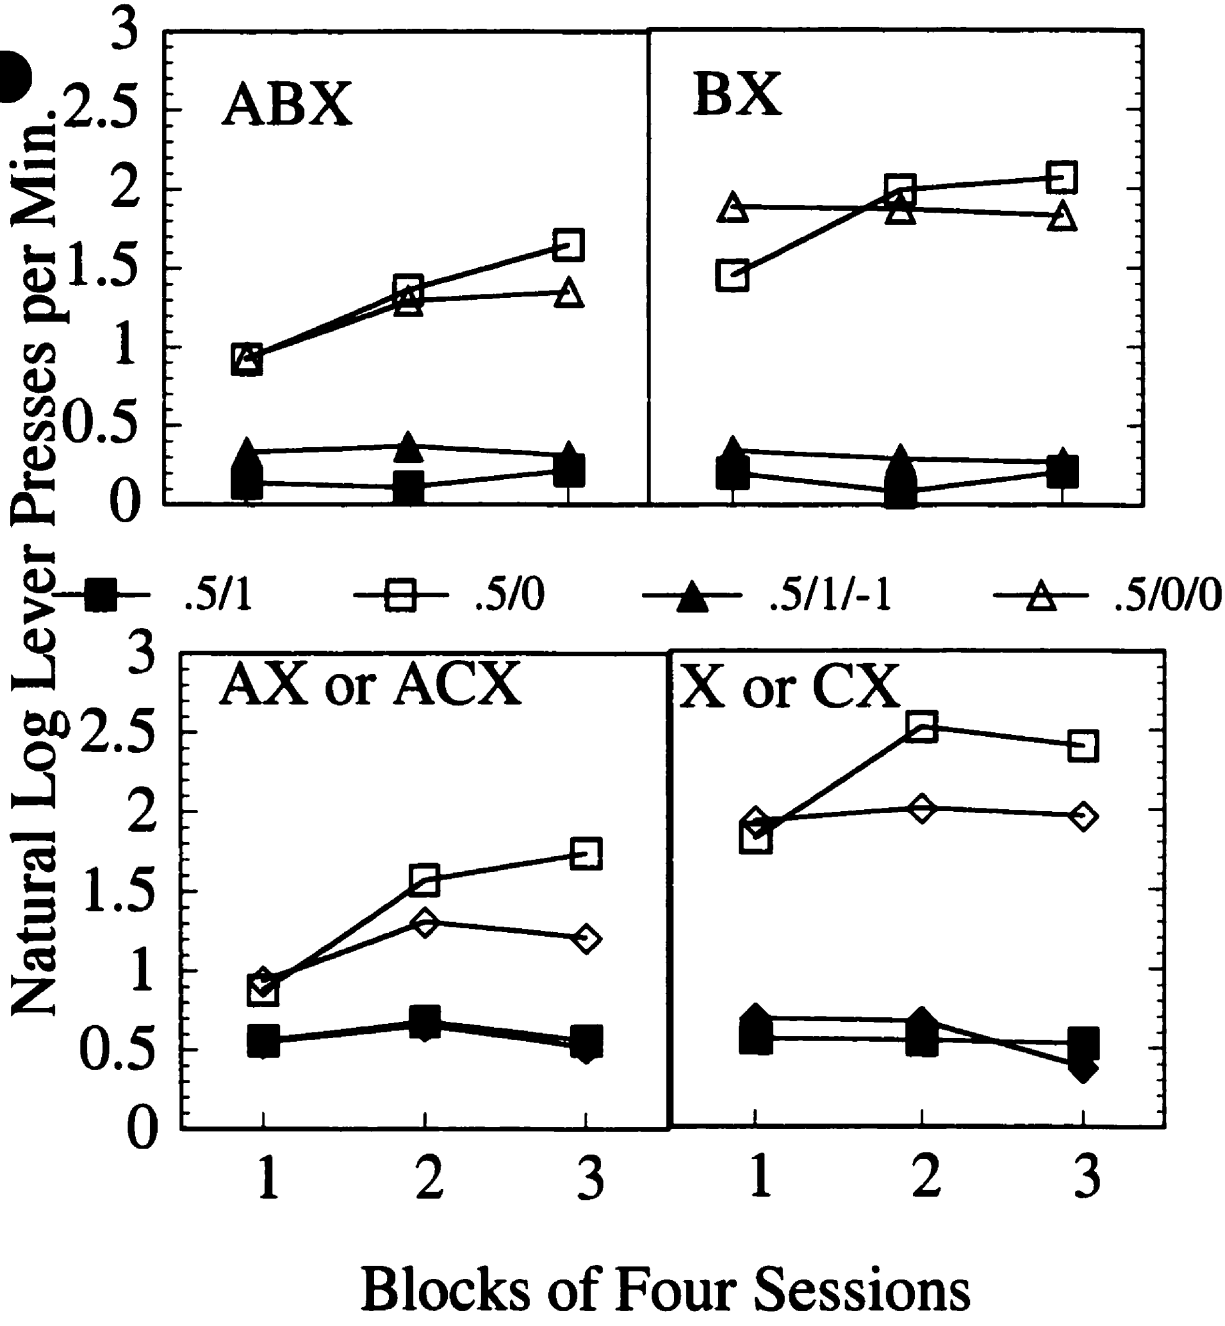
ABX
BX
AX or ACX
X or CX
Natural Log Lever Presses per Min.
Blocks of Four Sessions
.5/1
.5/0
.5/1/-1
.5/0/0
0
0.5
1
1.5
2
2.5
3

increase over blocks in the two groups in which **B** was uncorrelated, but that this increase was not observed in the other two groups [$F(2,88)=4.64$, $F(2,88)=4.13$ and $F(2,88)=3.58$ for lever pressing on **ABX**, **AX** (or **ACX**) and **X** (or **CX**) trials respectively]. Consistent with the analysis of tray entries, none of the analyses comparing responding in the groups that received training with the fourth CS, **C**, was reliable.

Test

Tray Entries Figure 18 displays the tray behaviour during the Pre-CS interval and in the presence of the compound **AX** during the first test and in the presence of **A** by itself during the second extinction test. During the first test, **AX** elicited elevated tray behaviour above that observed during the Pre-CS only in the two groups in which **B** was uncorrelated (.5/0 and .5/0/0). The analysis of the effect of **B**'s contingency ($\Delta p_B=1$ or $\Delta p_B=0$), **C**-Presence (no **C**, **C**) and Response Interval (Pre-CS, CS) found reliable effects for **B**'s contingency, Response Interval and the interaction between these two factors [minimum $F(1,44)=6.96$, $p<.01$]. A separate analysis showed a reliable elevation of tray behaviour during **AX** from that observed during the Pre-CS only in the two groups in which **B** was uncorrelated with the US [$F(1,22)=12.0$, $p<.01$], while the two groups in which **B** was contingent failed to show any difference between responding during the Pre-CS and **AX** [$F<1$]. There was no effect of the inclusion of cue **C** during the test [maximum $F<1$].

The second test of **A** without **X** found generally the same pattern of responding except lower Pre-CS response rates overall and in particular for the .5/1 group. In this treatment there was a reliable conditioning effect for **A**. The overall analysis found a reliable main effect for **B**'s contingency, and Response Interval [minimum $F(1,44)=4.07$] but not the interaction between the two [$F(1,44)=2.60$ $p<.11$]. The interaction of Response Interval and **C**-Presence was reliable in this analysis [$F(1,44)=5.86$] which supports the observation that the elevation in responding during **A**, above that recorded during the Pre-CS, was greater in the groups that did not receive training with **C**. These results generally

Figure 18.Mean rates of tray entries during the 10 second Pre-CS and during the stimulus presented during the two tests in Experiment 4. Test 1 involved nonreinforced exposure to the **AX** compound. The second test involved only stimulus **A**.

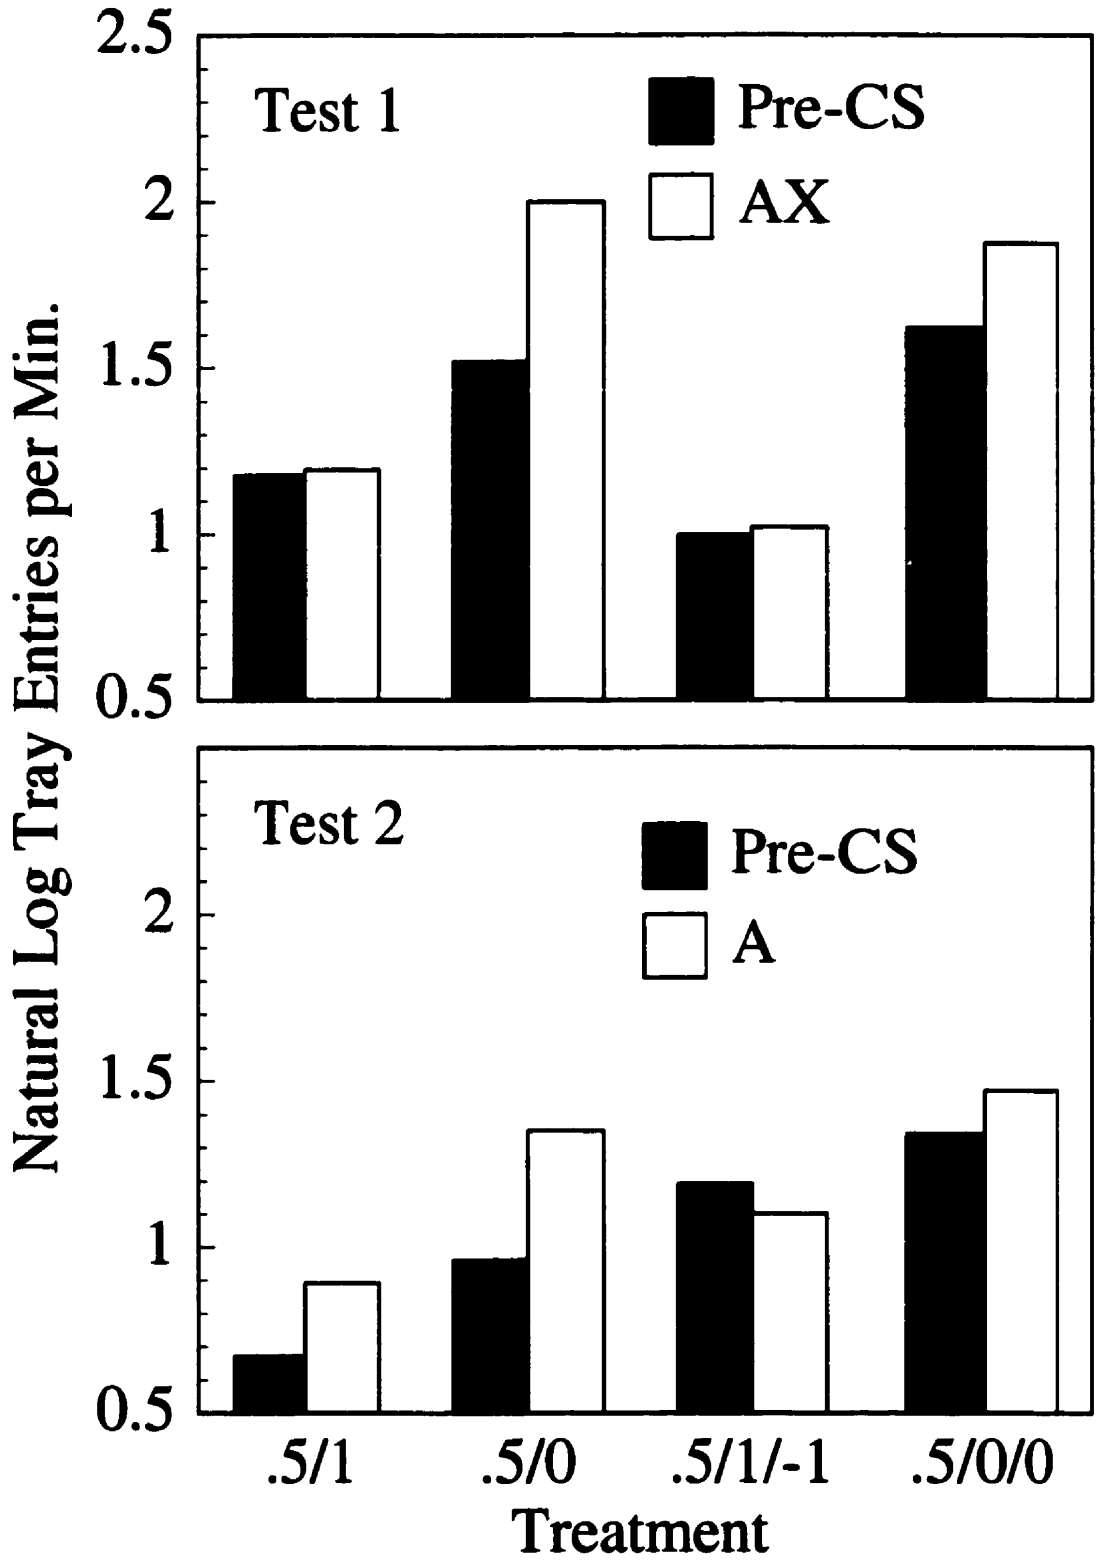
Test 1
Pre-CS
AX
Test 2
Pre-CS
A
2.5
2
1.5
1
0.5
Natural Log Tray Entries per Min.
.5/1
.5/0
.5/1/-1
.5/0/0
Treatment

confirms the findings from test 1, a perfect predictor of the US reduced responding to a less accurate predictor. The effects and overall response rates found in Test 2 appear smaller, however it must be remembered that these tests were not counterbalanced and so the test of **A** alone always followed the initial extinction test of **AX** and therefore the weaker effect may reflect a weaker **A**-US association or experience with the nonreinforced test situation. Overall levels of responding were lower, suggesting that animals may have learned to withhold responding during test sessions.

Lever Pressing. The mean rates of lever pressing during the first test of responding to **AX** suggest, as with the previous experiments, that the trial marker acquired a stronger association with the US when **B** was uncorrelated with the US. The mean rates of lever pressing during the test in groups .5/1 and .5/1/-1 were .789 and .717 while in the two groups without perfect predictors of the US .5/0 and .5/0/0 they were 1.97 and 1.33 respectively. An overall analysis of **B**'s contingency (Δp_B=1 or Δp_B=0), **C**-Presence (no **C**, **C**) only found reliable effect for **B** contingency, $F(1,44)$= 13.2.

This experiment demonstrates that the conditioned response elicited by a moderate predictor for the US was weaker when training involved a better predictor of the US. These relative validity effects like Wagner, et al.'s (1968) original experiments were not simply a function of the number of pairings between **A** the US in the different compounds, rather the effect is related to the overall predictive relation between **A** and the US relative to the other trained cues. This effect was also not a function of any difference between the conditional and unconditional contingency between **A** and the US which may have contributed to the effect observed in the .5/0 treatment of Experiment 3. While this experiment does not rule out the possibility that rats can learn conditional contingent relationships it does show that the present relative validity effect was not an artifact of an increase in the conditional contingency.

There was additional support for the role of the trial context in mediating

contingency effects. The trial context acquired greater control over lever pressing as the relative validity of the other two discrete cues were lower.

Experiment 5

The experiments described so far have examined how learning a CS-US contingency can be influenced by concurrent learning of other stimuli (discrete CSs or trial markers). In general the CSs used in these experiments are considered predictors of the US. These predictors acquire greater control over conditioned responses as their contingency increases relative to other possible predictors. In Experiments 1 and 2 the discrete CS was compared with a trial marker cue while in Experiments 3 and 4 multiple discrete cues and a trial marker cue were each trained with different positive relationships with the outcome. A computational account based on calculations of normative contingencies might explain and predict when animals will respond to one stimulus over another. One interesting question raised by this research is whether these relative validity effects are dependent on the valence of a cue's predictiveness.

There are at least two factors that might contribute to the relative validity effect. The relative strengths of the relationships between each CS and US is one factor. So far we have considered how relatively more contingent predictors may reduce the ability of weaker predictors to enter into association with the US. From the perspective of the normative theory involving sensitivity to causal relationships this result may be seen to make some intuitive sense. A stimulus fails to elicit a preparatory or consummatory response because there is a better predictor in the environment to which an animal might attend and learn. The RWM is able to predict this result simply by using the notion that temporal contiguity between CS and US favours more accurate predictors.

The focus so far has been to examine how two positive predictors of the US interact. Consider how the absence of a stimulus can also be a very accurate predictor of the

US. In conditioning terminology this is the notion of a conditioned inhibitor, which accurately predicts the absence of a US, and the CS *absence* predicts the presence of the outcome (e.g., Rescorla, 1969a). A strong negative predictor of an outcome can provide perfect information about when the outcome will and will not occur, without ever actually being paired with the US. The question arises whether the informativeness of a perfect negative predictors can reduce responding controlled by positive predictors that may be less accurate but at least sometimes paired with the US.

There is research showing how stimuli which themselves signal no US can interfere with associative learning. For example, a well known effect is the CS pre-exposure effect (or latent inhibition) in which acquisition of excitatory conditioning to a CS is impeded by pre-exposure of the CS (Lubow & Moore, 1959). Furthermore Reed (1995) has examined the CS pre-exposure effect with compounds and has shown that pre-exposure to a compound **AB**$^-$ will also interfere with subsequent conditioning of the elements of a compound. This effect can be reduced with prior nonreinforced exposure to one element of the compound. Pre-exposure to **A**$^-$ interferes with the CS preexposure effect found with **B** produced by **AB**$^-$ training. The implication is that a signal for no US can interfere with another signal for no US. Reed suggests that the associative rules governing blocking with a positively correlated stimulus are the same as those governing negatively correlated stimulus. However, there is little research to show whether relative validity effects can occur with negatively correlated cues, and whether these effects are a function of the negative Δp contingency. From the perspective of normative contingencies, the valence of a cue's predictive relationship is irrelevant. From the perspective of traditional associative theories cues are learned because they are involved in pairings between the CS and US. The following experiment will examine whether a perfect negative predictor (i.e., one that is never paired with the US) can interfere with responding controlled by a moderate positive predictor for the US.

The experiment that follows reports the effect of a perfectly positive predictor, a perfectly negative predictor and a nonpredictive cue (Δp_B=1, -1 or 0) on responding to a moderately positive predictor (Δp_A=.5) . If the relative validity effect is purely a function of relative informativeness about when the US will occur, animals may be able to represent the informativeness of both the presence and absence of the stimulus. Theoretically, this might require a more complex associative structure possibly involving associations with noUS events. If, on the other hand, the relative validity effect arises out of a competition between the two predictors for direct association with the US then the relative validity effect may only occur in group .5/1.

Although so far I have described negative continency blocking as possibly involving blocking by knowledge of the absence of the US, the RWM predicts negative contingency blocking without involving representations of the absence of the US. The model predicts that a perfect predictor of the absence of the US will cause a reduction in associative strength to a moderate positive predictor. To understand the model's prediction requires examining the model's behaviour during this type of experimental design in more detail.

The predictions for the final associative strength in the three groups described here are presented in Figure 19. The finding that a cue should acquire more associative strength following .5/0 treatment than .5/1 contingency treatment has been dealt with in the discussion of Experiment 3. The model correctly predicts that the moderately positive cue (open circle) **A** acquires more associative strength in treatment .5/0 (middle panel) than in treatment .5/1 group (top panel). In simulations of treatment .5/1, the relatively more valid perfect predictor **B** (Δp_B=1), successfully competes with the less valid **A** for the limited amount of associative strength (λ). Note all three cues **A** , **B** and the context **X** initially acquire similar amounts of associative strength, but that over time **B** acquires all the strength and **A** and **X** lose associative strength as a result of nonreinforced exposure. **X**

Figure 19. Cumulative associative strength predicted by the Rescorla-Wagner model (1972) for the three treatments in Experiment 5. These strengths represent the asymptotic weights predicted by the model. The perfectly predictive **B** acquires all the associative strength in treatment .5/1, **A** and **X** share strength in treatment .5/0 and **X** acquires all the strength in treatment .5/-1.

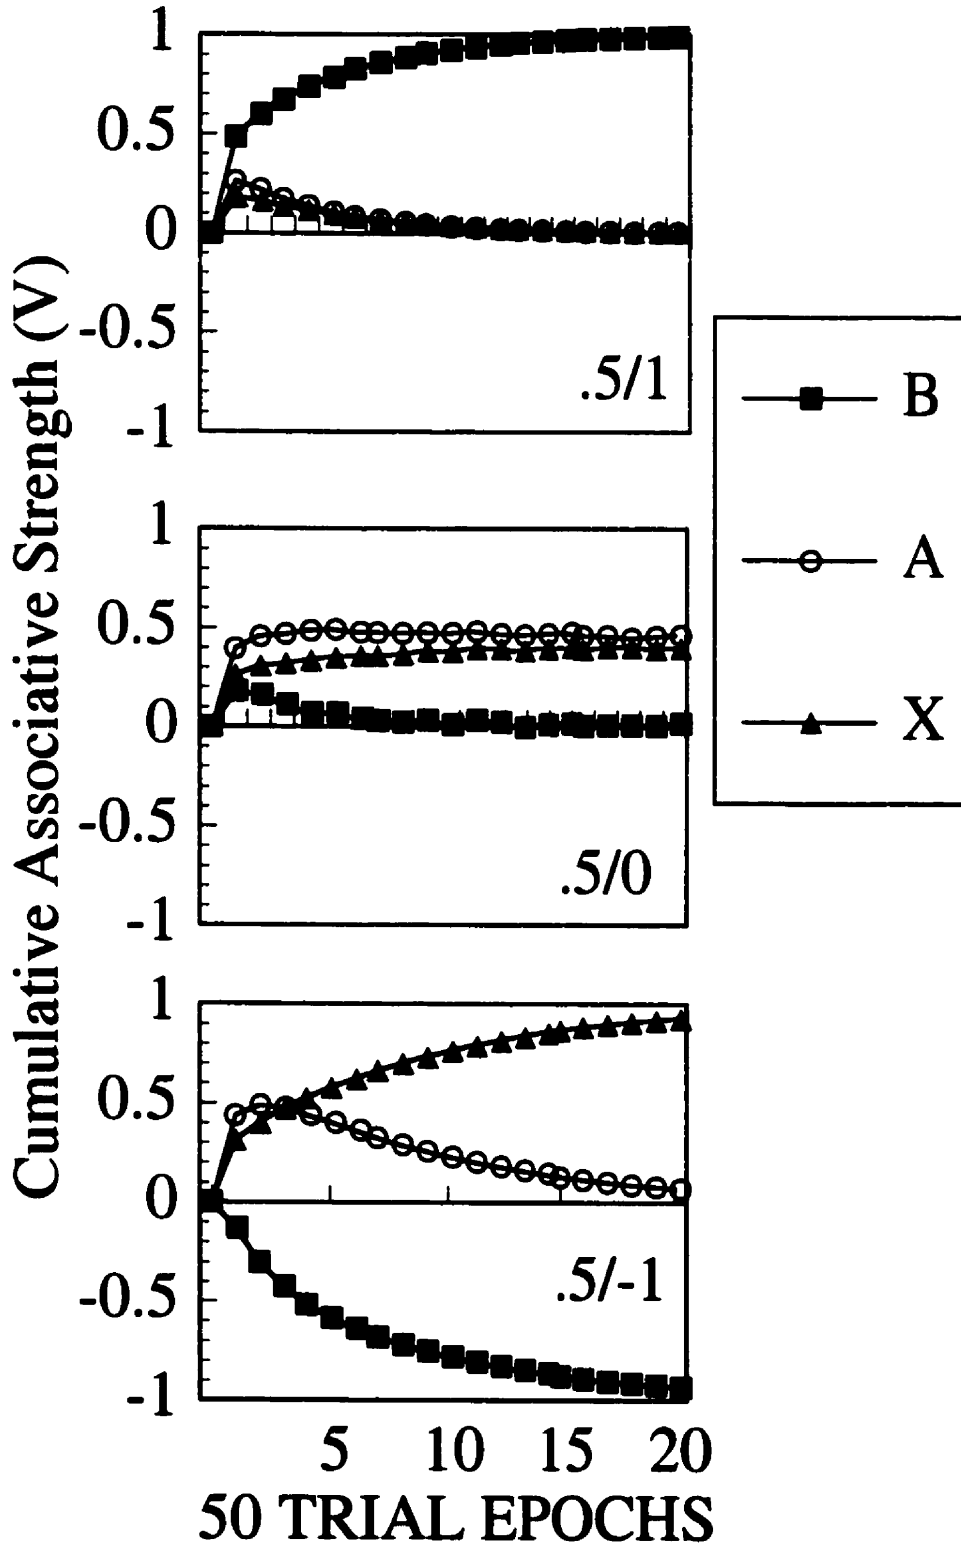
Cumulative Associative Strength (V)
1
0.5
0
-0.5
-1
.5/1
.5/0
.5/-1
B
A
X
5
10
15
20
50 TRIAL EPOCHS

loses strength faster and more completely because it receives the greatest nonreinforced experience. Eventually **A** and **X** have lost most of any initial strength acquired during the early stages of training.

In contrast, during the course of training with the negative contingency treatment .5/-1 (bottom panel), only **A** and **X** are able to acquire any positive associative strength. **B** is never present on trials with the US so it can not compete directly with **A** for association with the US. However, as **B** acquires inhibitory strength following pairing with the two excitatory cues **A** and **X**, **B** slowly increases its ability to *protect* **A** and **X** from losing associative strength on extinction trials. Cue **B** successfully competes for all the negative associative strength, however, unlike the .5/1 treatment, **A** and **X** are able to acquire and initially maintain gains in positive associative strength. The surprising prediction of the model is that **X** receives the most protection from the loss of positive strength on nonreinforced trials. Eventually, the partially reinforced context cue **X** is able to compete with **A** for associative strength until it ultimately acquires all of it. The bottom panel of Figure 19 shows how initially **A** acquires more associative strength than **X** on blocks 2 and 3, but that **X** eventually acquires **A**'s associative strength. At asymptote, the model predicts that a perfect negative predictor can reduce responding to a moderately positive cue. According to the model the negative contingency blocking effect is mediated by the strength of the context **X**. In summary, the model predicts that experience with either a perfectly positive (top panel) or perfectly negative predictor (bottom panel) of the US will reduce the acquisition of associative strength by the less valid predictor **A**. Both of these predictions involve competition between two positively related cues, **A** and **B** in the first case and **A** and **X** in the latter.

In addition to this prediction of the RWM, this experiment will test the generalisability of the basic multi-cue relative validity effect reported in Experiments 3 and 4. Experiments 3 and 4 did not involve complete counterbalancing. Experiment 5 involves

two treatments (.5/1 and .5/0) that have been used in both Experiments 3 and 4. To increase the generality of these findings Experiment 5 used a diffuse auditory stimulus in the role of the trial context, rather than the lever.

Although the use of two CSs that elicit different responses has some advantages, this might expose these experiments to a possible criticism.The associative and nonassociative explanations of contingency learning discussed so far explain the effect in terms of some type of cognitive learning. Either associations are formed and combined by specific rules, or alternatively various mental calculations are performed. More importantly, during training the rats were assumed to be exposed to all the stimuli and that the conditioning mechanism, whatever that may be, selected which stimuli to respond to. It is possible, however, that the results reported may reflect a peripheral attention deficit. This account suggests that in the current preparation while rats had their heads in the tray area, they did not or could not attend to the lever and, hence, in some groups failed to learn about it. This potential alternative explanation might imply that the present contingency effects are an artifact of the spatial location of the cues. However, this does not challenge the computational description of Pavlovian conditioning. Rats are still sensitive to absolute and relative CS-US contingency, but rather the algorithm that supports this computation may be even simpler than associative competition or comparison.

Experiment 5 uses a diffuse auditory stimulus rather than the localized lever. The auditory stimulus should be less likely to elicit a competing CR or peripheral response. Although it is never possible to entirely eliminate such arguments, this will weaken the response or peripheral attentional competition argument which holds that the relative validity effects reported here is a function of these peripheral mechanisms.

Method

Subjects and Apparatus: Thirty six rats were used for this experiment. In addition to the stimulus light (**A**) and tone (**B**) from previous experiments a mechanical relay placed on the

outside wall of the training chamber was made to open and close at 7 cycles per second. This produced an audible clicking (**X**) and vibrations to the box. All other equipment unless otherwise stated was the same as that used in previous experiments.

Procedure: Initially rats were trained to retrieve food pellets from the food tray using the same treatment regime described in earlier experiments. Following this, twelve animals were randomly assigned to each of the three treatments (.5/1, .5/0 and .5/-1). Each session consisted of a successive discrete trial discrimination procedure in which (**X**) signalled each of 64 ten second trials presented on an irregular 25 (2-58) second schedule. The end of a trial was always signalled by the cessation of the relay clicker which coincided with the delivery of the US on half the trials (32). Each treatment involved the same four trial types **ABX**, **BX**, **AX** and **X** alone. The frequency of these trial types in the three treatments and how the USs were distributed is described in Table 8. 75% of the trials with **A** were paired with the US and 25% of the trials without **A** were paired with the US. Thus, **A** was a moderate positive cue for the delivery of the US [p(US|**A**)=.75 p(US|-**A**)=.25, Δp=0.5]. The three treatments differed with respect to how well **B** signalled the US. In treatment .5/1, all of the US deliveries were signalled by a trial with **B** and no USs occurred on trials without **B**. In treatment .5/-1 none of the US deliveries were signalled by trials containing **B** and all trials without B were paired with the US. In treatment .5/0, half of both US and NoUS trials contained **B**. Conditioning training continued for 14 sessions followed by a single extinction test containing 32 ten second nonreinforced presentations of **A** in the absence of any other discrete cue. All test trials were presented using the same variable interval as during training.

Results and Discussion

Training

Pre-CS intervals. By the end of the first phase of training, tray entries during the 10 second Pre-CS period was equivalent in all three groups. The top panel of Figure 20

Table 8. Trial types, frequencies, conditional probabilities and Δp calculations for treatments in Experiment 5.

	(Δp_A=.5)		
$\Delta p_A/\Delta p_B$	**.5/1**	**.5/0**	**.5/-1**
ABX--->US	24	12	0
ABX--->no US	0	4	8
BX--->US	8	4	0
BX--->no US	0	12	24
AX-->US	0	4	24
AX-->no US	8	12	0
X-->US	0	4	8
X-->no US	24	12	0
Total US	32	32	32
Total no US	32	32	32
p(US|**B**)	1	.5	0
p(US|-**B**)	0	.5	1
Δp_B	1	0	-1
p(US|**A**)	.75	.75	. 75
p(US|-**A**)	.25	.25	.25
Δp_A	.5	.5	.5

shows the mean natural log tray entry rates from the Pre-CS interval in four session blocks (omitting sessions 1 and 2). A Groups by Blocks analysis of the Pre-CS data found a reliable main effect for Blocks and reliable Groups by Blocks interaction [$F(2,66)=23.8$ and $F(4,66)=2.82$]. The interaction reflects the stronger Pre-CS responding in group .5/-1 on Blocks 1 and 2. Post Hoc comparisons of each block found that .5/-1 differed from .5/1 on block 1 and .5/-1 differed from both .5/1 and .5/0 on Block 2 but these differences were not reliable by block 3.

CS trials. The rates of tray entering during the four trial types (**ABX**, **BX**, **AX** and **X**) are presented in the lower four panels of Figure 20. These rates generally reflect the fact that the .5/-1 group responded less on the two trial types containing the negatively correlated stimulus **B** (**ABX** and **BX**), that group .5/1 responded more on these two types of trial and that group .5/0 responded more on trials containing **A** (**ABX** and **AX**). Individual analyses of the data support these observations. During **ABX** trials, responding was lower for the .5/-1 treatment reflecting that **ABX** was never paired with the US. The Treatment by Blocks analysis found a reliable main effect for Treatment and a Treatment by Blocks interaction [$F(2,33)=11.2$ and $F(4,66)=3.17$ respectively]. Post hoc analyses showed that .5/-1 differed from the other two groups on Blocks 2, 3. Responding during **BX** trials similarly discriminated the treatments. Only the main effect of Blocks was reliable [$F(2,33)=3.11$]. A post hoc analysis of the data from the final block found a reliable main effect for treatment [$F(2,33)= 6.45$]. Individual post hoc comparisons found only that tray entry rates were higher in .5/1 than in the .5/-1. The Treatment by Blocks analysis on **AX** trials found a reliable main effect for Treatment and a reliable Treatments by Block interaction [$F(2,33)=25.9$ and $F(4,66)=2.74$]. The interaction reflects the low responding to **AX** trials in the .5/1 group. Post hoc analyses revealed that the Treatment effects were present on Blocks 2, 3 of training. Finally, the Treatments by Blocks analysis of responding to the trial marker alone (**X**) revealed reliable main effects for Treatment and

Figure 20.The Mean natural log tray entries per minute during the trial presentations in four session blocks during Experiment 5. Three treatments involved a moderately positive cue **A** (Δp=0.5) with either a perfect positive, perfect negative or uncorrelated alternate cue. Response rates are separated by the four trial types presented (**ABX**, **BX**, **AX**, **X**).

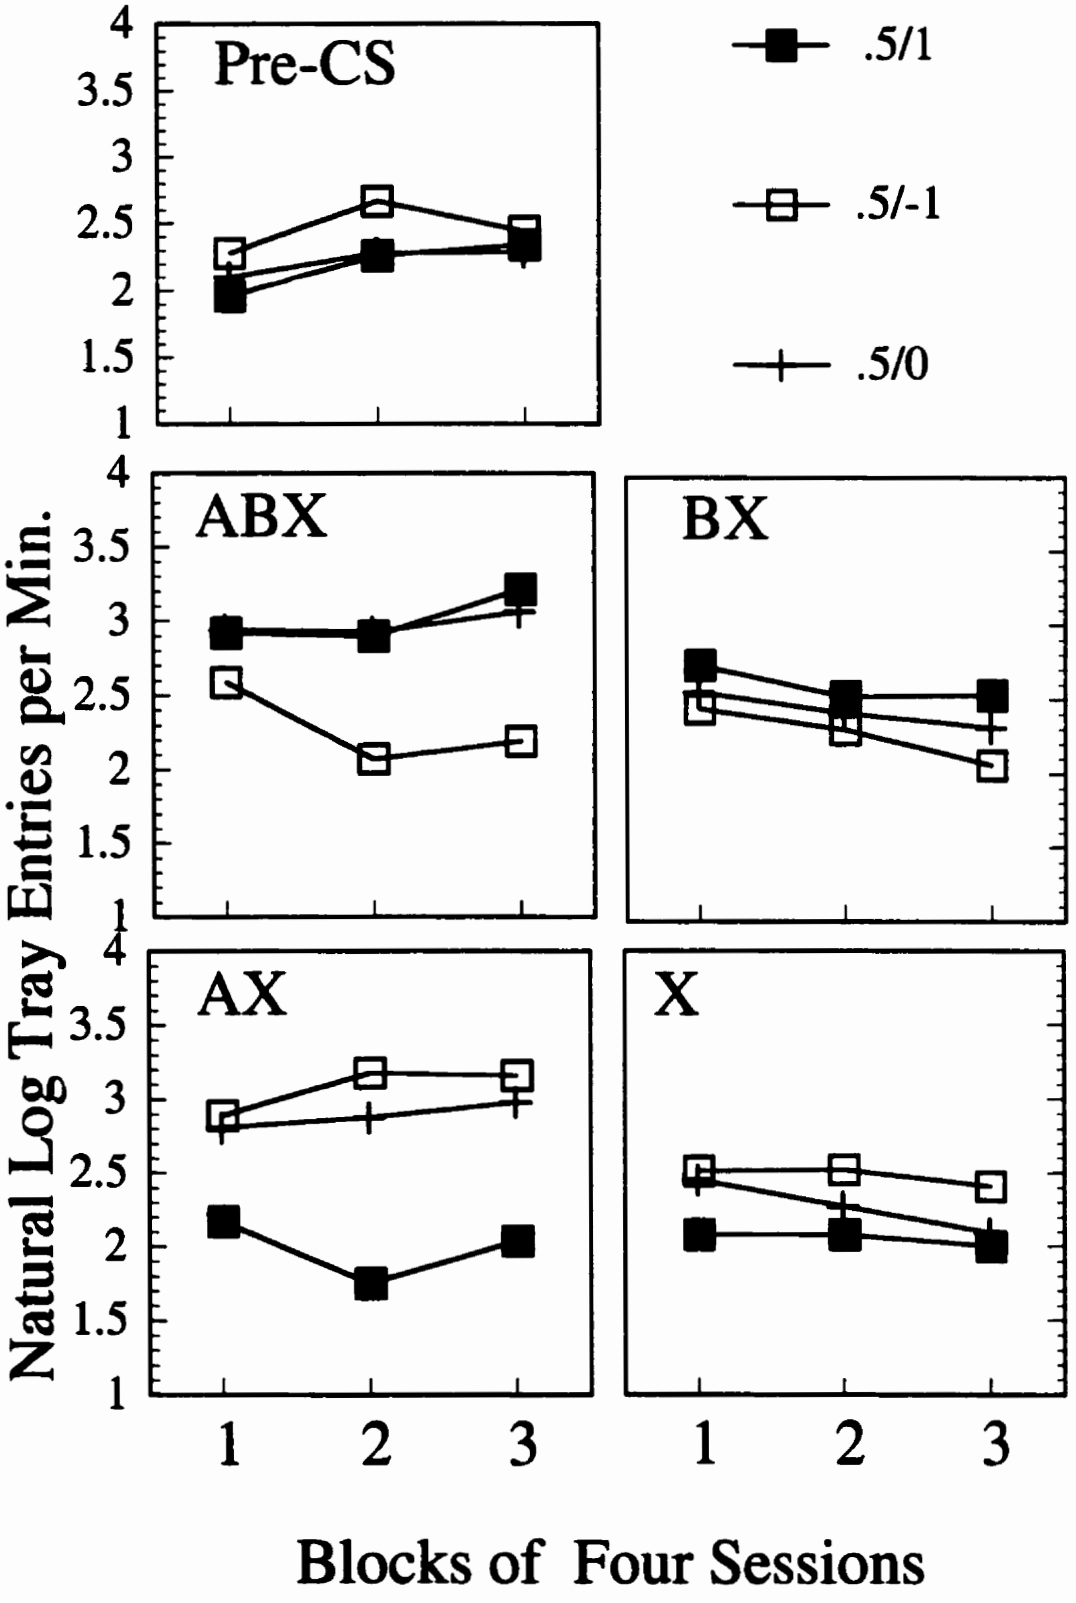
Pre-CS
ABX
BX
AX
X
.5/1
.5/-1
.5/0
Natural Log Tray Entries per Min.
Blocks of Four Sessions
4
3.5
3
2.5
2
1.5
1
1
2
3

Blocks [$F(2,33)=5.46$ and $F(2,66)=6.70$ respectively] but not the interaction. However the figure suggests that the interaction in this analysis may be confounded by the two main effects. An analysis of the data from the third block showed a main effect for treatment [$F(2,33)=3.77$]. A planned comparison of responding to **X** in treatment .5/-1 compared with the other two treatments found a reliable difference supporting the observation from the figure that in this treatment the trial context controlled a higher rate of tray entries than in either the other two treatments [$F(1,33)=7.16$]. The difference between treatment .5/1 and .5/0 was not reliable[$F<1$]. In summary, these training differences generally reflect the sensitivity of the three groups to the presence of elements of the compounds that signalled an increase in the likelihood of the US and support the prediction of the RWM that conditioning of the trial context should be greatest in treatment .5/-1.

Test

The crucial data from the test of tray entry rates during the Pre-CS and during exposure to stimulus **A** when presented by itself are presented in Figure 21. This was the first time that the subjects received stimulus **A** without the accompanying stimuli **B** and/or **X**. To assess whether each group showed increased rates of tray entering in the presence of **A** Figure 21 compare tray behaviour during the Pre-CS with that during the **A**. An analysis of Treatments (.5/1, .5/-1 and .5/0) by Response Interval (Pre-CS, **A**) found a reliable Treatment by Interval interaction [$F(2,33)=3.32$]. Individual tests of conditioning were conducted using single degrees of freedom F-tests for each group. Responding was elevated during presentation of **A** only for the .5/0 group [$F(1,11)=7.40$; $F(1,11)<1$ for groups .5/-1 and .5/1]. A direct post hoc test of responding during the Pre-CS comparing the .5/-1 group with the other two groups found a reliable difference [$F(1,33)=6.02$]. This shows that Pre-CS response rates were higher in the Negative contingency group. A test of responding during **A** comparing .5/1 with the other two groups was also reliable [$F(1,33)=5.89$]. The rate of responding during **A** was lowest in the positive contingency

Figure 21.Mean rates of tray entries during the 10 second Pre-CS and during the stimulus presented during the test in Experiment 5.The test involved nonreinforced exposure to A.

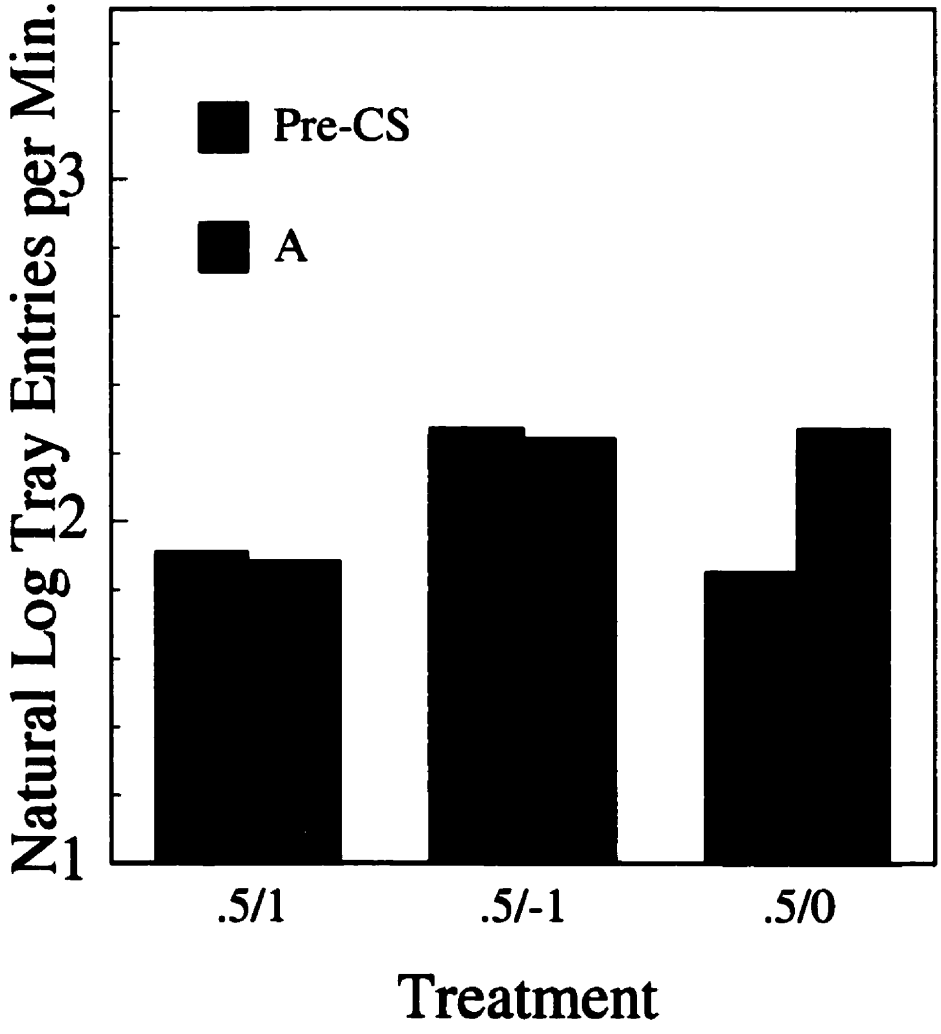
Natural Log Tray Entries per Min.
3
2
1
Pre-CS
A
.5/1
.5/-1
.5/0
Treatment

blocking treatment (.5/1). These results are consistent with the observation from Figure 21 that only the control group (.5/0) showed evidence of a conditioned response to **A**. The absence of a conditioned response in the .5/1 group replicates the relative validity effect from Experiments 3 and 4 in which the lever interfered with the elevation in tray behaviour observed to a moderately correlated **A**. This present experiment which used an auditory cue in the role of trial marker, suggests that the previous results were not dependent on incompatible responses generated by the lever. In spite of the relatively high rate of responding during **A** in the negative contingency group responding in this group was elevated during both the Pre-CS and during the CS suggesting that the moderately positive cue did not elevate tray entry rates from that observed during the Pre-CS.

The findings described here are consistent with the prediction that informativeness of a stimulus determines whether a CS will block or reduce responding to a cue. In group .5/1, a strong positive CS and in group .5/-1 a strong negative correlation between a CS and the food reduced the strength of a conditioned response controlled by a less accurate positive predictor of the US. Experiments 3 and 4 demonstrated the attenuating effect of strong positive predictors while Experiment 5 showed that strong negative predictors also reduce responding to moderate positive predictors.

From the perspective of theories of associative competition (Rescorla & Wagner, 1972), relative validity effects emerge as a function of direct competition between CSs for association with the US. It might not be surprising if negatively contingent CSs that, themselves never paired with the US were unable to mediate relative validity effects, if these effects emerge out of competition for association with the US. However, the results clearly show that negative predictors of a US can decrease the conditioned response of a positive but less correlated cue. One of the features of the Rescorla-Wagner model is that inhibitory strength is only acquired by a stimulus as a result of pairing with other cues that have themselves been paired with the US and have excitatory strength. This asymmetry

between the process describing acquisition of excitatory and inhibitory strength is a feature of the model (Baker, 1974). During nonreinforced stimulus training the elements of compound CSs will compete for negative associative strength. Like strong positive predictors, strong negative predictors will successfully block less accurate negative predictors from acquiring negative associative strength. This means that a cue can receive protection from the effects of extinction. In the case of the .5/-1 treatment in Experiment 5 the model predicts that cue **X**, that was itself reinforced 50%, would receive protection from the loss of excitatory strength when paired with **B**. The partially reinforced stimulus **X** becomes 'super'-excitatory, in that it acquires all the available strength (see Figure 19). This is much more than would be expected in the absence of the inhibitory cue. The empirical support for this prediction is derived from the rates of responding to stimulus **X** during training and during Pre-CS rates of tray entering. Pre-CS rates of responding, a measure of contextual conditioning, were greatest in the .5/-1 treatment. In addition the analysis of tray entries to the trial marker **X** also suggests that responding to **X** was greater in treatment .5/-1 than either .5/1 or .5/0. This finding provides direct support for the theoretical explanation of the relative validity effect as described by the RWM. The context, **X** is predicted to acquire most of the associative strength and block acquisition of associative strength by the moderate cue, **A**. The finding of significant Pre-CS rates of responding in some treatments suggests perhaps that the discrete cue trial marker was not entirely successful at blocking contextual conditioning.

In summary, the Rescorla-Wagner model provides very specific predictions about how cue **X** should be protected from the normal loss of associative strength that it would experience during nonreinforced training, and this allows **X** to compete with the moderate positive predictor **A** resulting in the attenuated control of responding by **A**. This experiment provides a unique demonstration that a partially reinforced stimulus with a 50% reinforcement schedule (the trial context) is able to successfully compete with a stimulus

which is paired with a greater probability with the US (**A** was reinforced 75% of the time it was presented). This result emerges from the competition for associative strength predicted by RWM. The Rescorla-Wagner model provides a plausible associative account of sensitivity to relative CS-US contingency, both positive and negatively correlated cues.

The associative comparator theories described earlier can also account for these results if it is assumed that the comparator stimulus for **A** in treatment .5/-1 is **X** and the comparator stimulus for **A** in .5/1 is **B** (Miller & Schachtman, 1985). While comparator theory is explicit as to how associations are formed, and is clear in stating that the strongest relative association will control behaviour, these models are not explicit as to how to identify the comparison stimuli. One might expect that **A** should be the comparator stimulus in both .5/0 and .5/-1 groups since it is reinforced on 75% of the trials. However, as we have just seen **X** would have to be the comparator stimulus in group .5/-1 in order to predict the blocking even though it is reinforced on only 50% of the trials. Other comparison theories also fail to accurately predict the blocking effect of X. For example, the average waiting time for the US (Gibbon & Balsam, 1981; Jenkins, Barnes & Barrrera, 1981) is greater during **X** than **A** in treatment .5/1, yet X still seems blocks **A**'s ability to control responding. Without a mechanism for selecting comparison stimuli, comparator theories are unable to make accurate predictions of some selective association effects.

General Discussion

The findings of these experiments contribute to a body of literature which draws an analogy between learning during Pavlovian conditioning and sensitivity to a measure of overall predictiveness, CS-US contingency. Conditioning preparations involve instituting a statistical relation or contingency between the presentation of the CS and the US. This is a relationship that can be described by various contingency metrics (Allan, 1980; Hammond & Paynter, 1980). It has been suggested that learning occurring during conditioning requires sensitivity to this relationship while Papini and Bitterman (1990) have called for

the abandonment of this theory. The present results are quite consistent with the former interpretation. It is difficult to imagine a simple explanation of these data that does not involves the use of the notion of the overall CS-US contingency.

In addition, contingency learning is similar to the type of learning required to perform causal induction embodied by Mill's Joint Method of Agreement and Difference (Mill, 1843/1973). Identifying causes of events requires assessing how consistently an outcome is signalled by an event [i.e., p(CS|US)] as well as how likely the outcome will occur in the absence of the event [i.e.,p(CS|-US]. Similarly, Pavlovian conditioning may have evolved to help animals distinguish relevant predictors from irrelevant or redundant predictors of important biological outcomes (Mackintosh, 1977).

Rather than suggesting that animals are engaged in causal reasoning in the sense commonly used to describe human reasoning, it is possible that evolution took advantage of the fact that the true cause of an event is often also the most contingent candidate. Humans have since developed a philosophical and mathematical framework for describing causal relationships, but much of this is presaged by the way learning systems work. Associative processes like those described by Rescorla and Wagner (1972; or Pearce, 1987) can mimic this form of causal reasoning using a competitive mechanism for accumulating associative strength which dynamically updates associative strength on the basis of the relative likelihood of the US in the presence and absence of the CS. Sensitivity to basic aspects of the causal structure of the environment allows the rat to behave as if it expects causes to produce their outcomes. This can be described by simple network models of learning involving associations between events (Gluck & Bower, 1988; Pearce, 1987; Rescorla & Wagner, 1972; Sutton & Barto, 1981). Similarly, the Rescorla-Wagner model has been used as a model of human causal reasoning (Baker, Murphy & Vallée-Tourangeau, 1996).

The possibility that has been tested in this thesis is that Pavlovian conditioning may reflect a simple evolved causal reasoning mechanism to deal with a fundamental

environmental problem. Previous research has had difficulty testing this hypothesis given the nature of contingent relationships and the associated problems of defining them in an experimental situation. In traditional conditioning procedures it is difficult to separate the effects of CS-US contiguity from CS-US contingency. The use of the discrete trial marker in this thesis allowed control over the signalling information normally provided by environmental contextual cues. This is important for several theories because most theories posit that learning about the context mediates contingency sensitivity. The present procedure allowed direct control over the trial context and provided a measure of contextual conditioning. The results from all 5 experiments generally supported the theory that rats were sensitive to changes in the Δp relationship between CS and US. In addition, this sensitivity depended upon the Δp of accompanying cues. It was proposed that whether a stimulus becomes conditioned is determined at least in part on the basis of its relative Δp contingency with respect to all other CSs and contextual cues.

Generally, conditioned responses were greater with a CS that signalled an increased likelihood of the delivery of food from that experienced in its absence [$p(US|CS)>p(US|\text{-}CS)$]. A stimulus' absolute contingency resulted in more conditioning. In Experiments 1 and 2, conditioned tray entries or lever pressing, increased appropriately as Δp increased. This is similar to the basic contingency sensitivity effect first demonstrated by Rescorla (1968, 1969a, 1969b). However, unlike Rescorla's, experiments the contingency was defined relative to a second CS, a trial context, rather than simply to the background cues and measurements of conditioning to this cue supported the choice of theory for the effect. The trial context was of similar salience and of the same duration as the target CS and elicited its own conditioned response (Experiments 1, 2, 3 and 4) or the same conditioned response as the CS (Experiment 5). Conditioned tray entries elicited by the CS were stronger with relatively stronger CS-US contingencies (i.e., as they approach -1 or +1) in relation to the discrete trial marker.

In the associative learning literature, Context -US associations are almost universally the intervening variable invoked to account for sensitivity to CS-US contingency (Baker, Murphy, Vallée-Tourangeau, 1996). Experimental support for this idea was provided by Odling-Smee (1975a; 1975b; 1978), who demonstrated how decreases in CS-US contingency increased a conditioned response directed at contextual cues. This was particularly important because it supported the associative explanation of contingency learning devised by Rescorla and Wagner (1972). Their theory suggests that the CS and the Context compete for association (Baker, Singh & Bindra, 1985; Durlach, 1983; Durlach, 1989b; Tomie, 1976). An important idea that follows from the RWM is that contexts and CSs have similar roles during conditioning. They both acquire associations with the US. But, in the standard conditioning preparation in comparison to a 10 second CS the general cues of the context are quite different (i.e., that they are longer in duration, multi-modal etc.). It might be difficult to control the amount of exposure to these cues. Standard contextual cues may also play a second role in a standard conditioning procedure, in addition to their role in defining the contingency. Contextual cues have also been implicated in fundamentally different associative structures. For example, contexts have been described in terms of complex spatial cues (i.e., spatial maps; O'Keefe & Nadel, 1978) or in terms of their ability to signal other associations (i.e., occasion setting; e.g., Holland, 1992). The RWM suggests that contexts and CSs, however, share a common role in spite of any other differences.

The RWM proposes that contexts like the CS can act like a single configural signal for the US and thereby acquire association with the US much like a regular discrete CS. This similarity was the source of the idea for the trial marker. A discrete CS could be used in the place of the general contextual cues (e.g., Baker, Singh & Bindra, 1985; Tanner, Rawlins & Mellanby, 1989). In this thesis a trial marker was adopted that logically maintained the contingency information role normally played by the general chamber cues.

The RWM is not alone in ascribing an important role to contextual cues. Other associative theories make similar predictions, but the mathematical formalisation and relatively modest number of internal variables makes the Rescorla-Wagner model an appealing candidate. In general, the RWM predicted the 1) sensitivity to $\Delta p=.5$ and 0 contingencies in Experiments 1 and 2 as well as the 2) effect of US density on conditioning to the same Δp contingency and 3) the sensitivity to relative contingencies (Experiment 3, 4 and 5) including the predicted role for the context in the explanation of the negative contingency blocking effect (Experiment 5).

Nonassociative accounts also propose a role for the context (Cheng, 1997). One normative computational theory of causal reasoning attributes causal reasoning to a conflict between the information provided by the putative cause and the base rate of the effect. Rather than viewing this similarity as a conflict between an associative and normative causal reasoning account, I have argued that the similarity is actually an important similarity between a representational theory of Pavlovian conditioning and a computational account of causal reasoning.

Pavlovian conditioning at a computational level mirrors the type of human causal reasoning performed by scientists. Both are governed by the goal of identifying the likely cause of an event and both have settled on a similar solution, which is to use covariation information. There are however certain mismatches between animal behaviour and this contingency norm. For example, in many of the experiments reported here there was evidence that some cues that were not the *best* or most valid predictor still attracted some level of conditioned responding (e.g, the Pre-CS rates in most experiments). The general contextual cues of the chamber were always less valid then any of the discrete CSs and therefore should never have elicited any behavioural change. This relative validity effect on the general contextual cues is not anticipated by the normative account, it should be added the RWM also fails to predict this difference (see also, Williams, Frame & LoLordo, 1992

for a failure to find blocking of contextual associations). The RWM predicts that the associative strength of the general contextual cues should be close to zero given the relatively stronger discrete cues (see Murphy, McDonald & Baker, 1998).

The different conditioning theories available are successful in varying degrees at predicting the data from these experiments. The many attempts to incorporate these notions in simple associative terms have proven quite successful. Rescorla and Wagner used the idea that a US has only limited abilities to strengthen an association resulting in cue competition for association. Wagner (1980) has proposed an account based on a limited memory capacity. Pearce & Hall (1980; and Mackintosh, 1974) proposed that stimuli are differentially associable as a result of changes in CS-US contingency. They claimed that associability was determined to the extent to which a CSs did not predict its consequences. Finally, Pearce has recently attempted to describe some selective association effects from the perspective of processes of configuration and generalization (e.g., Pearce, 1997). Each of these attempts to explain Pavlovian conditioning have come to the conclusion that conditioning was more complex and multi-determined than initial conditioning-extinction theories credited (Rescorla, 1988). However, what has been lacking in previous attempts to discriminate the different approaches is the conceptual framework for describing the environment in which the animals are learning.

An appreciation of environmental contingencies and an optimal assessment of contingencies can be used to aid our understanding of the mechanism that underlies conditioning. The fact that contingencies are determined subjectively evaluated on the basis of relative validity comparisons provides an important framework for understanding conditioning. This in turn provides testable predictions about when conditioned behaviour might occur. The results of experiments designed to test these predictions provides a method for discriminating among the different associative algorithms that have been proposed to account for conditioning.

Papini and Bitterman (1990) rejected the notion of contingency in Pavlovian conditioning. However, it would be difficult to reconcile the data presented here with their thesis. It seems possible that Papini and Bitterman have simply confused two different and separable levels of analysis. The findings of these experiments do not demand that rats learn contingencies by calculation of a sophisticated statistical concept such as Δp. Rather, a consideration of the various theoretical accounts demonstrates how there are many simpler possibilities available. Δp was used as a theoretical tool to guide experimental design because it suggested normative calculations to which the animals behaviour may or may not conform. Given a framework that suggested that animals learn about the most valid relationship in the environment the question that follows is; by what criteria are relationships evaluated? Is it by the number of CS-US pairings, the proportion of reinforced CS trials, predictive valence of the CS or some variable. In these experiments animals almost always behaved as if they were sensitive to the Δp contingency but, as Marr (1982) has shown, computational level analyses only describe the informational content of a cognitive problem without providing a representational or algorithmic analysis. In these experiments the Rescorla-Wagner model's algorithm provides a good account of the CS and trial context data as well as the relative validity effects.

The second important theoretical advance suggested by these experiments is the idea that the contingency between any CS and a US is evaluated with respect to the general context but also with respect to any other discrete CS. Experiments 3, 4 and 5 explored multiple CS conditioning and sensitivity to relative contingencies. In all experiments animals developed the strongest conditioned response to the CS with the strongest relative contingency.

The basic relative validity effect described by Wagner, Logan, Haberlandt & Price (1968) along with other selective association effects (Kamin 1969; Rescorla, 1968) have been pivotal in shaping much experimental inquiry in learning since their initial description.

These experiments showed that unconditional CS-US contiguity history can not always be used to predict whether an animal will develop a conditioned response to a CS. Relative validity effects have been interpreted as demonstrating a number of different psychological processes. None, however, have been very specific. One interpretation suggests that CS processing is mediated by informativeness (Egger & Miller, 1962). However, it has never been clear what information the CS was supplying, nor how it was evaluated. For example, Kamin suggested that the blocking effect demonstrated that surprise mediates learning. However, the conditions that elicited surprise were almost entirely defined by the blocking procedure that he was trying to explain in the first place. Similarly others have suggested that these effects reflect sensitivity to redundancy (Egger & Miller, 1962; Rickert, Lorden, Dawson, Smyly, and Callahan, 1979), but have not described exactly what makes a stimulus redundant. The closest that any writer has come to making an explicit statement about this issue is Wagner himself who proposed that validity or informativeness of cues is defined by the probability that it will be paired with the US (p(US|CS). Wagner's reliance on this conditional probability to explain his results might be appropriate for his experiments which only involved a partially reinforced test cue. With the design of that experiment the p(US|CS) is highly correlated with Δp. The relative contingency hypothesis proposed here suggests rather that validity is determined by the relative contingency defined approximately by Δp [p(US|CS)- p(US|noCS)].

Causal Knowledge

One question that emerges from this animal research is related to causal reasoning in humans. To what extent does Pavlovian conditioning actually involve learning of causal relations in the way that humans may understand cause? The notion of causation in human reasoning seems to involve at least two distinct ideas. Firstly, the acquisition of causal knowledge seems likely to involve both an appreciation of covariation (Baker, Murphy & Vallée-Tourangeau, 1996) and possibly an appreciation of a causal model in which causes

and their outcomes are linked by generative power (Cheng, 1997; Dickinson & Shanks, 1995; Kant, 1871/1965; White, 1991). Cheng for example, argues that true causal knowledge involves this second form of understanding that there is a causal force that produces an event. This idea is similar to the notion that causes are perceived as transmitting a form of energy to produce their effect. This knowledge is referred to as causal power (Cheng, 1997). One might argue that humans have an understanding of causation at both these levels, but that animals may only be able to learn about causes at the level of covariation. If at least part of causal knowledge is acquired by an appreciation of covariation, then the experimental evidence described here demonstrates that, at least at the level of performance, rats have a relatively sophisticated appreciation of relative contingencies and therefore of the causal structure of their environments. This appreciation goes beyond simply learning about the probability of reinforcement and is more akin to the methods of agreement and difference used to describe the acquisition of causal knowledge. From this perspective rats could be argued to possess at least one form of causal knowledge.

Others are more conservative and have suggested that only the second of these forms is true causal knowledge. Dickinson & Shanks (1995) have proposed that true causal knowledge emerges only from an appreciation of the causal energy supplied by a cause (see also, Kumer, 1995). Their data from an instrumental conditioning procedure with rats does not rule out the idea that animals can learn this second form of true causal relationships. These experiments involve rats acquiring instrumental behaviour, such as lever pressing leading to food. Their interpretation is that the animals have learned that their behaviour produces the food. In contrast, they suggest that Pavlovian conditioning experiments of the form described in this thesis have only ever showed that an animal learns covariation, not that they actually learn a causal relation between a CS and US. Animals might show increased tray entries, but we will never know whether they have any knowledge about

whether they expect food. Of course they can not expect that the animal would learn in what physical way the appearance of the CS generatively produces the US; since in most experiments the causal relation is only instantiated in the computer program written by the experimenter (Einhorn & Hogarth, 1986). But they argue that the animal responds to the CS without any knowledge that the CS predicts the occurrence of a specific US.

The idea that the animals have a truer causal understanding because the behaviour involves an instrumental action may simply reflect a biased understanding of human causal reasoning. Philosophers have argued that there is *only* contingency knowledge for causal relations (Mill, 1843/1973), and that all causal knowledge, regardless of how sophisticated, is acquired from contingency information. Differences in causal knowledge may only be quantitative rather than qualitative.

If this is true then any system that represents this type of relationship has true causal knowledge. The problem in the past has been to explain how both humans and non-humans are able to learn about an event and its outcome and to ignore many other potential causes. The relative contingency hypothesis predicts that contingencies are evaluated with respect to all possible cues and that the most contingent cue acquires strength at the expense of a less correlated cue. A similar idea was explored by Cheng (1997) in human causal reasoning. She argued that people acquire causal information by performing multiple conditional contingency calculations, and that the cue with the strongest conditional contingency with an outcome would be identified as the cause. This thesis explored this idea with rats. An associative interpretation was shown to provide a plausible description of the underlying psychological processes and offer an explanation for why certain cues are able to interfere with learning other cues.

Implications

The basic notion that Pavlovian conditioning involves sensitivity to contingencies was established by Rescorla's first experiment. However, the experimental design and

procedure is open to criticism (Papini & Bitterman, 1990). Furthermore, even those that accept the experimental results still resist the basic finding. I will now consider one associative learning experiment that demonstrates as an example the failure of current research to incorporate the notion of contingencies into experimental design.

In the field of research in Pavlovian conditioning there is still a belief that conditioning involves learning about simple CS-US pairings. To test this hypothesis researchers compare a CS-US paired group of animals with an unpaired CS-US group. To test simple conditioning-extinction theories of conditioning this may be an appropriate strategy. However from a Δp contingency theory and in the light of the experiments reported in this thesis this is like comparing a perfect positive contingency $\Delta p=1$ with a perfect negative contingency $\Delta p=-1$. Any differences in conditioned behaviour between these two treatments could be due to excitatory conditioning in the paired group, but equally could be due to inhibitory learning in the unpaired group or a combination of both. The point is that these two groups do not provide a test for Pavlovian conditioning (Rescorla, 1967).

A very recent experiment on associative learning, reported in *Animal Learning and Behavior* ignores the contingency issue entirely. Ungless (1998) reports a test of Pavlovian conditioning in the snail. The experimental procedure involved pairing an apple odor (CS) with carrot food (US). He trained a paired group that received 10 minute exposure to the CS and US and compared them with an unpaired group that received separate experience with the two stimuli. When he tested the CS for its ability to elicit a conditioned tentacle extension response he reported more responses in the paired than the unpaired group. The problem with this research is in the conclusion which states "...exposure to the stimuli unpaired did not promote learning" (p. 17). This interpretation while it may be correct, does not even consider the idea that the snails are sensitive to the negative contingency between the CS and US. This example points to a failure to acknowledge conditioning mechanisms

that are sensitive to more than simple CS-US pairings.

The contingency hypothesis and its more complex version, the relative contingency hypothesis, may have important implications for current research in animal learning which includes human sensitivity to event-outcome relations. Like the Pavlovian conditioning experiments with rats and other animals, humans also show good learning about contingency information. For example, Vallée-Tourangeau, Murphy and Baker (1998) presented human subjects with information about six sets of fictitious patients and whether they had contracted a virus. The task was to try and learn whether a virus was a good predictor of a disease. Each set of patients represented a set of trials for a different overall virus-disease contingency. Subjects were presented with 40 trials successively just like the CS-US pairings in a Pavlovian conditioning experiment. Subjects were told whether a patient had contracted a fictitious virus (e.g., Threbaggia) and whether they had a fictitious disease (e.g., Ork's complex). In any given treatment the virus could have a positive, negative or zero contingency with the disease diagnosis. At the end of the 40 trials subjects were asked to rate the relationships between the virus and the disease. The data showed that subjects were quickly able to learn the Virus-Disease contingency and they were able to discriminate the different contingencies. Other research has shown similar effects using different scenarios (Price & Yates, 1995; Wasserman, Chatlosh, Elek and Baker, 1993). In addition, the selective association effects with multiple predictors have also been demonstrated (Baker, Mercier, Vallée-Tourangeau, Frank & Pan, 1993). These experiments are often cast in terms of associative learning and the same type of associative models including RWM (Shanks & Dickinson, 1987) and Pearce provide a good fit to the data.

Future research in associative learning may examine the departures from normative contingency learning in both humans and animals. The appreciation of the complexity of environmental contingencies is accompanied by the realization that the real environment in

which animals have evolved is richly more complex than the 2 and 3 cue experiments described in this thesis. One suggestion for future research is to compare normative predictions with behaviour as the complexity of the CS-US relationships is increased. In addition to this ecological argument for increased contingency complexity, associative models like RWM which were designed to account for the relatively simple validity experiments of the 60s may not be equipped to account for data in more complicated scenarios. There is at least pilot work conducted on this issue that suggests that the RWM may not provide an adequate account of contingency learning with more than 3 cues (Murphy & Baker, 1995).

Statement of Original Contribution

Rescorla (1968) provided empirical evidence that the likelihood of the US in the absence of a CS, in addition to the likelihood of the US in the presence of a CS, contributed to the strength of a conditioned response. The experiments in this thesis provided new empirical support for the predictions of the RWM for this effect. In addition his contingency effect and the theoretical account was extended to the more general case of contingency sensitivity with multiple CSs. Once the associative theory that could account for Rescorla's basic effect emerged (i.e., Rescorla-Wagner Model), little additional research was conducted to test the implications of the model (although see Odling-Smee, 1975a; 1978). In all the experiments of this thesis the trial marker provided a method of defining CS-US contingencies using marked temporal intervals. In Experiments 1 and 2 the strength of conditioned tray entries to the CS (i.e., light) and the separate lever pressing response controlled by the trial marker together provided a direct test of the reciprocal conditioning predicted by the model. Conditioned tray entries to the CS were higher when it was a positive predictor of the US than when it was uncorrelated. In contrast the trial marker (the lever) elicited lever pressing that showed the opposite pattern. This reciprocal conditioning pattern is consistent with the theoretical predictions of the effect suggested by the RWM.

This is the first demonstration of this difference with the proper control conditions to rule out alternate explanations of the contingency effect. The three positive and three zero contingency conditions allowed direct tests that ruled out alternate explanations of the contingency effect based on, the number of CS presentations, the number of US presentations and the number of CS-US pairings.

Experiments 3, 4 examined whether the contingency effect found with a single CS was part of a general sensitivity to the contingencies of multiple CSs. Both experiments showed how the strength of the conditioned response controlled by either a moderately positively correlated cue or an uncorrelated cue was a function of both its absolute Δp contingency and its contingency relative to the strength of the contingency of other cues in the environment. Experiments 3 and 4 extended the range of selective association phenomenon to demonstrate blocking effects with multiple contingencies. There was support for the RWM as a general model for these effects.

Finally, Experiment 5 tested an interesting prediction of the competitive model that had never been tested in the Pavlovian conditioning literature. The model makes the surprising prediction that a training context might successfully compete with a moderate predictor of the US if training included a perfect negative predictor of the US. This prediction and the generally supporting results were compared with various associative and non-associative theories. The Rescorla-Wagner model was shown to provide accurate ordinal predictions in all treatments. Other models either failed to make specific predictions about some of the treatments or failed to make the correct predictions (e.g., Miller & Schachtman, 1985). These experiments provided new support for the RWM as the best description of contingency learning in Pavlovian conditioning

Conclusion

In five experiments rats were trained with CS-US contingencies and in all experiments animals showed sensitivity to both the absolute and relative CS-US

contingency as defined by Δp. These results are consistent with a computational level analysis of Pavlovian conditioning. The conclusion was that Pavlovian conditioning seems to exhibit some of the same properties as a causal reasoning mechanisms. It is quite good at identifying the most accurate predictor of food. The normative Δp model is a useful theoretical tool because it 1) provided predictions about how behaviour might be organised 2) helped organize data and 3) finally it provides a framework for considering why conditioning may have evolved in the first place. Previous experiments including Rescorla's (1968) original experiments on contingency learning effect failed to properly control the events that constitute a subjective Δp . This thesis used a discrete trial context that signalled all trials. This procedure had two advantages. First it allowed the definition of the complete CS-US contingency and second it provided the substrate for a measurable conditioned response. In spite of equivalent CS-US contiguity in many of the treatments the conditioned response was determined by an appreciation of a more global relative relationship. The Pavlovian conditioning mechanism, whatever that turns out to be, must minimally be sensitive to these relative relationships.

References

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

Ayres, J.J.B., Benedict, J.O., & Witcher, E.S. (1975). Systematic manipulation of individual events in a truly random control in rats. *Journal of Comparative and Physiological Psychology, 88*, 97-103.

Baker, A.G. (1974). Conditioned inhibition is not the symmetrical opposite of conditioned excitation: A test of the Rescorla-Wagner model. *Learning & Motivation, 5*, 369-379.

Baker, A.G. (1977). Conditioned inhibition arising from a between-sessions negative correlation. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 144-155.

Baker, A.G. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning & Behavior, 5*, 315-319.

Baker, A.G., & Mercier, P. (1982). Extinction and the context and latent inhibition. *Learning and Motivation*, 13, 391-416.

Baker, A.G., & Mercier, P. (1989). Attention, retrospective processing and cognitive representations. In S.B. Klein & R.R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 85-116). Hillsdale, N.J. : Lawrence Erlbaum Associates, Inc..

Baker, A.G., Mercier, P., Gabel, J., & Baker, P.A. (1981). Contextual conditioning and the US preexposure effect in conditioned fear. *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 109-128.

Baker, A.G., Mercier, P., Vallée-Tourangeau, F, Frank, R., & Pan, M. (1993). Selective associations and causality judgments: The presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning,*

Memory and Cognition, *19*, 414-432.

Baker, A.G., Murphy, R.A., & Vallée-Tourangeau, F. (1996). Associative and normative models of causal induction: Reacting to versus understanding cause. In D.R.Shanks, D.L.Medin & K.J.Holyoak (Eds.) *The Psychology of Learning and Motivation: Causal Learning*, *34*, 1-45.

Baker, A.G., Singh, M., & Bindra, D. (1985). Some effects of contextual conditioning and US predictability on Pavlovian conditioning. In P.D. Balsam & A. Tomie (Eds.), *Context and learning* (pp. 73-104). Hillsdale, N.J.: Lawrence Erlbaum Associates, Inc..

Baker, A.G., Vallée-Tourangeau, F & Murphy, R.A. (In press) Asymptotic judgment of cause in a relative validity paradigm. *Memory & Cognition.*

Balaz, M.A., Gustin, P., Cacheiro, H., & Miller, R.R. (1982). Blocking as a retrieval failure: Reactivation of associations to a blocked stimulus.*The Quarterly Journal of Experimental Psychology*, *34B*, 99-113.

Balsam, P.D. (1985). The functions of context in learning and performance. In P.D. Balsam & A. Tomie (Eds.), *Context and learning*. Hillsdale, N.J.: Lawrence Erlbaum Associates, Inc..

Benedict, J.O., & Ayres, J.J.B. (1972). Factors affecting conditioning in the truly random control procedure in the rat. *Journal of Comparative and Physiological Psychology*, *78*, 323-330.

Bennett, C.H., Maldonado, A., & Mackintosh, N.J. (1995). Learned irrelevance is not the sum of exposure to CS and US.*The Quarterly Journal of Experimental Psychology*, *48B*, 117-128.

Blanchard, R.J., & Blanchard, C. (1969). Crouching as an index of fear. *Journal of Comparative and Physiological Psychology*, *67*, 370-375.

Bonardi, C. & Hall, G. (1996). Learned irrelevance: No more than the sum of CS and US

preexposure effects? *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 183-191.

Bouton, M.E. (1984). Differential control by context in the inflation and reinstatement paradigms. *Journal of Experimental Psychology: Animal Behavior Processes, 10*, 56-74.

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

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

Bouton, M.E., & King, D.A. (1986). Effect of context on performance to conditioned stimuli with mixed histories of reinforcement and nonreinforcement. *Journal of Experimental Psychology: Animal Behavior Processes,12*, 4-15.

Bouton, M.E., & Peck, C.A. (1989). Context effects on conditioning, extinction, and reinstatement in an appetitive conditioning preparation. *Animal Learning & Behavior, 17*, 188-198.

Bouton, M.E., & Nelson, J.B. (1994). Context-specificity of target versus feature inhibition in a feature-negative discrimination. *Journal of Experimental Psychology: Animal Behavior Processes,20*, 51-65.

Bouton, M.E., Rosengard, C., Achenbach, G.G., Peck, C.A., & Brooks, D.C. (1993). Effects of contextual conditioning and unconditioned stimulus presentation on performance in appetitive conditioning. *The Quarterly Journal of Experimental Psychology, 46B*, 63-95.

Breland, K. & Breland, M. (1965). The misbehavior of organisms. In T.E.McGill (Eds.). *Readings in Animal Behavior* (pp.455-460). New York: Holt, Rinehart and Winston.

Bush, R.R., & Mosteller, F. (1951). A mathematical model for simple learning.

Psychological Review, *58*, 313-323.

Capaldi, E.J. (1966). Partial reinforcement: A hypothesis of sequential effects. *Psychological Review*, *73*, 459-477.

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.

Cheng, P.W., & Holyoak, K.J. (1995). Complex adaptive systems as intuitive statisticians: causality, contingency, and prediction. In H.L. Roitblat & J-A. Meyer (Eds.), *Comparative approaches to cognitive science* (pp.271-302). Cambridge, MA: MIT Press.

Cheng, P.W., & Novick, L.R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.

Church, R.M. (1997). Quantitative models of animal learning and cognition. *Journal of Experimental Psychology: Animal Behavior Processes*, *23*, 379-389.

Cole, R.P., Barnet, R.C., & Miller, R.R. (1995) Effect of relative stimulus validity: Learning or performance deficit?, *Journal of Experimental Psychology: Animal Behavior Processes*, *21*, 293-303.

Cole, R.P., Denniston, J.C., & Miller, R.R. (1996). Reminder-inducted attenuation of the effect of relative stimulus validity. *Animal Learning & Behavior*, *24*, 256-265.

Cole, R.P., Gunther, L.M., & Miller, R.R. (1997). Spontaneous recovery from the effect of relative stimulus validity. *Learning and Motivation*, *28*, 1-19.

Cotton, M.M., Goodall, G. & Mackintosh, N.J. (1982). Inhibitory conditioning resulting from a reduction in the magnitude of reinforcement. *Quarterly Journal of Experimental Psychology*, *34B*, 163-180.

Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge

University Press.

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

Dickinson, A., & Charnock, D.J. (1985). Contingency effects with maintained instrumental reinforcement.*Quarterly Journal of Experimental Psychology, 37B*, 397-416.

Dickinson, A., & Shanks, D.R. (1995). Instrumental action and causal representation. In D. Sperber, D. Premack & A.J.Premack (Eds.) *Causal cognition: A multidisciplinary debate* (pp.5-25). Oxford, UK: Clarendon Press.

Durlach, P.J. (1983). Effect of signalling intertrial unconditioned stimuli in autoshaping. *Journal of Experimental Psychology: Animal Behaviour Processes, 9*, 374-389.

Durlach, P.J. (1989a). Role of signals for unconditioned stimulus absence in the sensitivity of autoshaping to contingency. *Journal of Experimental Psychology: Animal Behavior Processes, 15*, 202-211.

Durlach, P.J. (1989b). Learning and performance in Pavlovian conditioning: Are failures of contiguity failures of learning or performance. In S.B. Klein & R.R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory*. Hillsdale, N.J. : Lawrence Erlbaum Associates, Inc..

Durlach, P.J. & Shane, D.O. (1993). The effect of intertrial food presentations on anticipatory goal-tracking in the rat. *Quarterly Journal of Experimental Psychology, 46B*, 289-318.

Dweck, C.S. & Wagner, A.R.(1970). Situational cues and correlation between CS and US as determinants of the conditioned emotional response. *Psychonomic Science, 18*, 145-147.

Einhorn, E.J. & Hogarth, R.M. (1986). Judging probably cause. *Psychological Bulletin,*

99, 3-19.

Egger, M.C., & Miller, N.E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology, 64*, 97-104.

Fitzgerald, R.D. (1963). Effects of partial reinforcement with acid on the classically conditioned salivary response in dogs. *Journal of Comparative and Physiological Psychology, 56*, 1056-1060.

Gallistel, C.R. (1990). *The organization of learning*. Massachusetts: The MIT press.

Gamzu, E.R. & Williams, D.R. (1973). Associative factors underlying the pigeon's key pecking in auto-shaping procedures. *Journal of the Experimental Analysis of Behaviour, 19*, 225-232.

Garcia, J., & Koelling, R.A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science, 4*, 123-124.

Gewirtz, J.C., Brandon, S.E. &, Wagner, A.R. (1998). Modulation of the acquisition of the rabbit eyeblink conditioned response by conditioned contextual stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 24*, 106-117.

Gibbon, J., Berryman, R., & Thompson, R.L. (1974). Contingency spaces and measures of classical and instrumental conditioning. *Journal of the Experimental Analysis of Behaviour, 21*, 585-605.

Gibbon, J., Farrell, L., Locurto,, C., Duncan, H., & Terrace, H. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning & Behavior, 8*, 45-59.

Gibbon, J. & Balsam, P. (1981). Spreading association in time. In C.M. Locurto, H.S. Terrace, & J.Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219-253). New York, N.Y.: Academic Press.

Gibbon, J., Locurto, C. & Terrace, H.S. (1975). Signal-food contingency and signal frequency in a continuous trials auto-shaping paradigm. *Animal Learning and*

Behavior, 3, 317-324.

Gluck, M.A., & Bower, G.H. (1988). From conditioning to category learning: An adaptive network model.*Journal of Experimental Psychology: General*, 117, 227-247.

Goddard, M.J. & Jenkins, H.M.(1987). Effect of signalling extra unconditioned stimuli on autoshaping. *Animal Learning & Behavior, 15*, 40-46.

Gormenzano, I. & Kehoe, E.J. (1981). Classical conditioning and the law of contiguity. In P. Harzem and M.D. Feiler (Eds.), *Predictability, correlation and contiguity* (pp.1-45). New York, N.Y.: John Wily & Sons Ltd..

Granger, R.A.H. Jr. &, Schimmler, J.C. (1986). The computation of contingency in classical conditioning. *The Psychology of Learning and Motivation*, 20,137-192.

Grau, J.W. & Rescorla, R.A. (1984). Role of context in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes,10*, 324-332.

Hall, G. (1991). *Perceptual and associative learning*. Oxford: Clarendon Press.

Hall, G. & Channell, S. (1985). Latent inhibition and conditioning after preexposure to the training context. *Learning and Motivation, 16*, 381-397.

Hall, G., & Honey, R.C. (1989). Contextual effects in conditioning, latent inhibition and habituation: Associative and retrieval functions of contextual cues. *Journal of Experimental Psychology: Animal Behavior Processes,15*, 232-241.

Haliam, S.C., Grahame, N.J. & Miller, R.R. (1992). Exploring the edges of Pavlovian contingency space: An assessment of contingency theory and its various metrics. *Learning and Motivation, 23*, 225-249.

Hammond, L.J. (1980). The effect of contingency upon the appetitive conditioning of free-operant behavior. *Journal of the Experimental Analysis of Behavior, 34*, 297-304.

Hammond, L.J. & Paynter, W.E. (1983). Probabilistic contingency theories of animal conditioning: A critical analysis. *Learning and Motivation, 14*, 527-550.

Hammond, L.J., & Weinberg, M. (1984). Signalling unearned reinforcers removes the suppression produced by a zero correlation in an operant paradigm. *Animal Learning & Behavior, 12*, 371-377.

Hearst, E., & Jenkins, H.M. (1974). Sign tracking: The stimulus-reinforcer relation and directed action. *Monograph of the Psychonomic Society*. Austin, TX.

Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267-272.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press.

Holland, P.C. (1977) Conditioned stimulus as a determinant of the form of the Pavlovian conditioned response. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 77-104.

Holland, P. C. (1992) Occasion setting in Pavlovian conditioning. In D. L. Medin (Ed), *The Psychology of Learning and Motivation*, (Vol 28, pp. 69-125). San Diego CA: Academic Press.

Holyoak, K.J., Koh, K. & Nisbett, R.E. (1989). A theory of conditioning: Inductive learning within rule-based default hierarchies. *Psychological Review, 96*,315-340.

Hull, C. L. (1952). *A behavior system.* New York: John Wiley and Sons.

Hurwitz, H.M., & Davis, H. (1983). The description and analysis of conditioned suppression: A critique of the conventional suppression ratio. *Animal Learning & Behavior, 11*, 383-390.

Jenkins, H.M., Barnes, R.A. & Barrera, F.J. (1981). Why autoshaping depends upon trial spacing. In C.M. Locurto, H.S. Terrace, & J.Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 255-284). New York, N.Y.: Academic Press.

Jenkins, H.M. & Shattuck, D. (1981). Contingency in fear conditioning a reexamination.

Bulletin of the Psychonomic Society, *17*, 159-162.

Jenkins, H.M., & Lambos, W.A. (1983). Tests of two explanations of response elimination by noncontingent reinforcement. *Animal Learning & Behavior*, *11*, 302-308.

Kamin, L.J. (1969). Selective associations and conditioning. In W.K. Honig & N.J. Macintosh (Eds.), *Fundamental issues in associative learning*. (pp. 42-64). Halifax, N.S.: Dalhousie University Press.

Kant, I. (1950). Critique of pure reason (N.K.Smith, trans.). London: Macmillan. (Original work published in 1781).

Kaye, H. & Pearce, J.M. (1984). The strength of the orienting response during Pavlovian conditioning. *Journal of Experimental Psychology*, *10*, 90-109.

Keller, R.J., Ayres, J.J.B. & Mahoney, W.J. (1977). Brief versus extended exposure to truly random control procedures. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*, 53-65.

Killcross, S., & Dickinson, A. (1996). Contextual control of latent inhibition by the reinforcer. *Quarterly Journal of Experimental Psychology*, *49B*, 45-59.

Kremer, E.F. (1971). Truly random and traditional control procedures in CER conditioning in the rat. *Journal of Comparative and Physiological Psychology*, *76*, 441-448.

Kremer, E.F. (1974). The truly random control procedure: Conditioning to the static cues. *Journal of Comparative and Physiological Psychology*, *86*, 700-707.

Kummer, H. (1995). Causal knowledge in animals. In D. Sperber, D. Premack & A.J.Premack (Eds.) *Causal cognition: A multidisciplinary debate* (pp,26-36). Oxford, UK: Clarendon Press.

Lea, S.E.G. (1981). Correlation and contiguity in foraging behaviour. In P.Harzem and M.D.Zeiler (Eds.) *Predictability, correlation, and contiguity* (pp. 355-406). New York, N.Y.: John Wily & Sons Ltd..

Lubow, R.E. & Moore, A.U. (1959). Latent inhibition: The effect of nonreinforced pre-exposure to the conditional stimulus. *Journal of Comparative and Physiological Psychology, 52,* 415-419.

Lindblom, L.L. & Jenkins, H.M. (1981). Responses eliminated by noncontingent or negatively contingent reinforcement recover in extinction. *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 175-190.

Lopez, F.J., Shanks, D.R., Almaraz, J., & Fernandez, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24,* 672-694.

Lovibond, P.F., Preston, G.C., & Mackintosh, N. J. (1984). Context specificity of conditioning, extinction, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes, 10,* 360-375.

Luongo, A. (1976). Stimulus selection in discriminative taste-aversion learning in the rat. *Animal Learning & Behavior, 4,* 225-230.

Mackintosh, N.J. (1974). *The psychology of animal learning.* London: Academic Press

Mackintsoh, N.J. (1977). Conditioning as the perception of causal relations. In Butts and Hintikka (Eds.) *Foundational problems in the special sciences* (pp.241-250). Dordrecht, Holland: D.Reidel Publishing Company.

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

Mackintosh, N.J. (1983). *Conditioning and associative learning.* Oxford: Oxford University Press.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* San Fransisco, CA: W.H.Freeman, (Ch. 1).

Matzel, L.D., Schachtman, T.R. & Miller, R.R. (1988). Learned irrelevance exceeds the

sum of CS-preexposure and US-preexposure deficits.*Journal of Experimental Psychology: Animal Behavior Processes, 14,* 311-319.

McLaren, I.P.L., Kaye, H., & Mackintosh, N.J. (1989). An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition. In R.G.M. Morris (Ed.) *Parallel Distributed Processing: Implications for Psychology and Neurobiology*, Oxford: Clarendon Press.

Mill, J.S. (1973). *A system of logic*. Toronto: University of Toronto Press.(Original work published in 1843).

Miller, R.R., Barnet, R.C. & Grahame, N.J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin, 117,* 363-386.

Miller, R.R., Hallam, S.C., Hong, J.Y. &, Dufore, D.S. (1991). Associative structure of differential inhibition: Implications for models of conditioned inhibition.*Journal of Experimental Psychology: Animal Behavior Processes, 17,* 141-150.

Miller, R.R. & Matzel, L.D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G.H. Bower (Ed.), *The psychology of learning and motivation, Vol. 22* (pp. 51-92). San Diego, CA: Academic Press.

Miller, R.R. & Schachtman, T.R. (1985). The several roles of context at the time of retrieval. In P.D. Balsam & A. Tomie (Eds.), *Context and learning.* (pp. 167-194) Hillsdale, N.J. : Lawrence Erlbaum Associates, Inc..

Murphy, R.A. & Baker, A.G. (1995, June). *With relative validity guiding conditioning, two blockers are as good as one.* Paper presented at the meeting of the Canadian Society for Brain, Behaviour and Cognitive Science, Halifax, NS.

Murphy, R.A. & Baker, A.G. (1996). *Summation of associative strength between a context and discrete stimulus with a selective association procedure.* Poster session presented at the XXVI International Congress of Psychology, Montreal, QC.

Murphy, R.A., McDonald, R.J., & Baker, A.G. (1998). Lesions of the Hippocampus or

fornix do not interfere with the relative validity effect on a discrete stimulus or the context. *Behavioural Brain Research, 92,* 55-66.

Nelson, K.J. (1987). Conditioned inhibition from incomplete reductions in the probability of reinforcement. *Quarterly Journal of Experimental Psychology, 39B,* 365-391.

Odling-Smee, F.J. (1975a). The role of background during Pavlovian conditioning. *Quarterly Journal of Experimental Psychology, 27,* 201-209.

Odling-Smee, F.J. (1975b). Background stimuli and the interstimulus interval during Pavlovian conditioning. *Quarterly Journal of Experimental Psychology, 27,* 387-392.

Odling-Smee, F.J. (1978). The overshadowing of background stimuli by an informative stimulus in aversive Pavlovian conditioning. *Animal Learning and Behavior, 6,* 43-51.

O'Keefe, J. and Nadel, L. (1978). *The hippocampus as cognitive map.* Oxford: Oxford University Press.

Papini, M.R. & Bitterman, M.E. (1990). The role of contingency in classical conditioning. *Psychological Review, 97,* 396-403.

Pavlov, I.P. (1927). *Conditioned Reflexes.* New York: Dover.

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. & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned stimuli. *Psychological Review, 87,* 532-552.

Pearce, J.M. & Redhead, E.S. (1993). The influence of an irrelevant stimulus on two discriminations.*Journal of Experimental Psychology: Animal Behavior Processes, 19,* 180-190.

Peterson, G.B., Ackil, J.E., Frommer, G.P. & Hearst, E.S. (1972). Conditioned approach and contact behavior toward signals for food or brain-stimulation reinforcement. *Science*, 177, 1009-1011.

Price, P.C., & Yates, F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition*, *21*, 561-572.

Prokasy, W.F.Jr. (1965). The acquisition of observing responses in the absence of differential external reinforcement. *Journal of Comparative and Physiological Psychology*, *49*, 131-134.

Quinsey, V.L. (1971). Conditioned suppression with no CS-US contingency in the rat. *Canadian Journal of Psychology*, *25*, 69-82.

Redhead, E.S., & Pearce, J.M. (1995). Stimulus salience and negative patterning. *Quarterly Journal of Experimental Psychology*, *48B*, 67-83.

Reed, P. (1995). Enhanced latent inhibition following compound pre-exposure. *Quarterly Journal of Experimental Psychology*, *48B*, 32-45.

Rescorla, R.A. (1966). Predictability and number of pairings in Pavlovian fear conditioning. *Psychonomic Science*, *4*, 383-384.

Rescorla, R.A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review*, *74*, 71-80.

Rescorla, R.A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 1-5.

Rescorla, R.A. (1969a). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*, 77-94.

Rescorla, R.A. (1969b). Conditioned inhibition of fear resulting from negative CS-US contingencies. *Journal of Comparative and Physiological Psychology*, *67*, 504-509.

Rescorla, R.A. (1969c). Conditioned inhibition of fear. In W.K. Honig & N.J. Macintosh

(Eds.). *Fundamental issues in associative learning.* Halifax: Dalhousie University Press.

Rescorla, R.A. (1971). Variation in the effectiveness of reinforcement and nonreinforcement following prior inhibitory conditioning. *Learning and Motivation*, 2, 113-123.

Rescorla, R.A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43*, 151-160.

Rescorla, R.A. (1997). Summation: Assessment of a configural theory. *Animal Learning & Behavior, 25*, 200-209.

Rescorla, R. A., Durlach, P. J., and Grau, J. W. (1985) Contextual learning in Pavlovian conditioning. In Balsam, P. & Tomie, A. (Eds.)*Context and Learning*.Hillsdale NJ/ Erlbaum.

Rescorla, R.A., & LoLordo, V.M. (1965). Inhibition of avoidance behavior. *Journal of Comparative and Physiological Psychology, 59*, 406-412.

Rescorla, R.A., & Skucy, (1969). Effect of response-independent reinforcers during extinction. *Journal of Comparative and Physiological Psychology, 67*, 381-389.

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

Rickert, E.J., Lorden, J.F., Dawson Jr., R., Smyly, E. and Callahan, M.F. (1979) Stimulus processing and stimulus selection in rats with hippocampal lesions, *Behavioral Neural Biology, 27*, 454-465.

Rizley, R.C. & Rescorla, R.A. (1972). Associations in second-order conditioning and sensory preconditioning.*Journal of Comparative and Physiological Psychology, 81*, 1-11.

Schachtman, T.R., Brown, A.M., Gordon, E.L., Catterson, D.A. & Miller, R.R. (1987). Mechanisms underlying retarded emergence of conditioned responding following inhibitory training: Evidence for the comparator hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes, 13*, 310-322.

Scheffé, H. (1953). A method of judging all contrasts in the analysis of variance. *Biometricka, 40*, 87-104.

Schmajuk, N.A. (1997). *Animal learning and cognition.* Cambridge, UK: Cambridge University Press.

Seligman, M.P. (1992). *Helplessness: On development, depression & death.* New York: Freeman and Company.

Siegel, S. (1977). Morphine tolerance acquisition as an associative process. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 1-13.

Sevenster, P. (1973). Incompatibility of response and reward. In R.A. Hinde & J.Stevenson-Hinde (Eds.). *Constraints on learning: Limitations and predispositions* (pp.265-283). London: Academic Press.

Shanks, D.R. (1995). *The psychology of associative learning.* Cambridge, UK: Cambridge University Press.

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

Spellman, B.A. (1996). Conditionalizing causality. In D.R.Shanks, D.L.Medin & K.J.Holyoak (Eds.) *The Psychology of Learning and Motivation: Causal Learning, 34*, 167-206.

Spence, K. W. (1940) Continuous versus noncontinuous interpretation of discrimination learning. *Psychological Review, 47*, 271-288.

Sosa, E., & Tooley, M. (1993). Introduction. In E.Sosa, M.Tooley (Eds.) *Causation.* Oxford: Oxford University Press.

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

Tanner, J., Rawlins, J.N.P. & Mellanby, J.H. (1987). Manipulation of CS-US conditional probability and of the CS-US trace interval on conditioning to the CS and to background stimulus in a CER situation. *Learning and Motivation, 18,* 371-391.

Thomas, D.R., Burr, D.E.S. &, Eck, K.O. (1970). Stimulus selection in animal discrimination learning: An alternative interpretation. *Journal of Experimental Psychology, 86,* 53-62.

Thorndike, E.L. (1911/1965). *Animal Intelligence: Experimental Studies.* New York: Hafner Publishing Company.

Tolman, E.C., & Brunswik, E. (1935). The organism and the causal texture of the environment. *Psychological Review, 42,* 43-47.

Tomie, A. (1976). Interference with autoshaping by prior context conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 2,* 323-324.

Turner, C. & Mackintosh, N.J. (1972). Stimulus selection and irrelevant stimuli in discrimination learning by pigeons. *Journal of Comparative and Physiological Psychology, 78,* 1-90.

Udell, H. & Rescorla, R.A. (1979). Conditioning of simultaneous and successive common elements in a discrimination and pseudo discrimination, *Bulletin of the Psychonomic Society, 14,* 453-456.

Ungless, M.A. (1998). A Pavlovian analysis of food-attraction conditioning in the snail (Helix aspersa). *Animal Learning & Behavior, 26,* 15-19.

Vallée-Tourangeau, F., Murphy, R.A. & Baker, A.G. (1998). Causal induction in the presence of a perfect negative cue: Contrasting predictions from associative and statistical models.*Quarterly Journal of Experimental Psychology, 51B,* 173-191.

Vallée-Tourangeau, F., Murphy, R.A., Drew, S., & Baker, A.G. (1998). Judging the

importance of constant and variable candidate causes: A test of the Power PC theory. *Quarterly Journal of Experimental Psychology, 51A*, 65-84.

Van Hamme, L.J., & Wasserman, E.A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements.*Learning and Motivation, 25*, 127-151.

Wagner, A.R. (1969). Stimulus validity and stimulus selection in associative learning. In W.K. Honig & N.J. Macintosh (Eds.). *Fundamental issues in associative learning.* Halifax: Dalhousie University Press.

Wagner, A.R. (1981). SOP: A model of automatic memory processing in animal behavior. In N.E. Spear and R.R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). New Jersey: Erlbaum.

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

Wagner, A.R. & Rescorla, R.A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R.A. Boakes & M.S. Halliday (Eds.), *Inhibition and learning* (pp. 301-336). London: Academic Press.

Wagner, A.R., Siegel, S., Thomas, E. & Ellison, G.D. (1964). Reinforcement history and the extinction of a conditioned salivary response. *Journal of Comparative Physiological Psychology, 58*, 354-358.

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

Wasserman, E.A. (1974). Stimulus-reinforcer predictiveness and selective discrimination learning in pigeons. *Journal of Experimental Psychology, 103*, 284-297.

Wasserman, E.A., Elek, S.M., Chatlosh, D.L. & Baker, A.G. (1992). Rating causal relations: Role of probability in judgments of response-outcome contingency.

Journal of Experimental Psychology: Learning, Memory and Cognition, 19, 174-188.

Weiss, S.J. (1972). Stimulus compounding in free-operant and classical conditioning. *Psychological Review, 78*, 189-208.

Weisman, R.G. & Litner, J.S. (1969). Positive conditioned reinforcement of Sidman avoidance behavior in rats.*Journal of Comparative Physiological Psychology, 68*, 597-603.

White, P.A. (1990). Ideas about causation in philosophy and psychology. *Psychological Review, 108*, 3-18.

Williams, D.A, Frame, K.A. & LoLordo, V.M (1992). Discrete Signals for the unconditioned stimulus fail to overshadow contextual or temporal conditioning. *Journal of Experimental Psychology Animal Behavior Processes, 18*, 41-55.

Wilson, P.N. & Pearce, J.M. (1989). A role for stimulus generalization in conditional discrimination learning. *Quarterly Journal of Experimental Psychology, 41B*, 243-273.

Witcher, E.S. & Ayres, J.J. (1980). Systematic manipulation of CS-US pairings in negative CS-US correlation procedure in rats. *Animal Learning & Behavior, 8*, 67-74.

Zimmer-Hart, C.L. & Rescorla, R.A. (1974). Extinction of pavlovian conditioned inhibition. *Journal of Comparative and Physiological Psychology, 86*, 837-845.