

639⁰⁰

A MODEL OF RESPONSE-~~REINFORCER~~ CONTINGENCY

A MODEL OF RESPONSE-REINFORCER CONTINGENCY

By

GREGORY KENNETH SCOTT, B.Sc.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

August, 1982

DOCTOR OF PHILOSOPHY (1982)
(Psychology)

McMASTER UNIVERSITY
Hamilton, Ontario

TITLE: A Model of Response-Reinforcer Contingency

AUTHOR: Gregory Kenneth Scott, B.Sc. (University of Toronto)

SUPERVISOR: Professor J.R. Platt

NUMBER OF PAGES: vii, 91

ABSTRACT

The experiments reported in this thesis investigated the effect of response-reinforcer contingency on response differentiation. Since the failure to control response probability in previous studies had led to difficulties, the present experiments employed percentile schedules to control response probability. Response-reinforcer contingency was indexed by the measure of statistical association between two dichotomous variables known as the phi coefficient (ϕ), and a model of this independent variable was developed to permit a systematic investigation of contingency in operant conditioning. This model was tested using rats in a spatial response differentiation paradigm. The results of three experiments revealed that the higher the value of ϕ , the more effective the shaping of response location to a target location. Despite differences between experiments in the way the independent variable was manipulated, across all three experiments there was a very orderly relationship between asymptotic conditioning and ϕ . These experiments demonstrate the importance of response-reinforcer contingency in response differentiation and provide support for a model of contingency in operant conditioning based on ϕ .

ACKNOWLEDGEMENTS

To my supervisor and friend, John Platt, I express the deepest gratitude for the countless hours devoted to my scientific education. Throughout my graduate career, Shep Siegel and Kris Kristofferson provided constructive criticism and encouragement, for which I am grateful. I thank Steve Link and Harding Bishop, who contributed much to my academic development. I am indebted to Dodi, who enhanced the quality of my life as a graduate student. Finally, for their unfailing support, I dedicate this thesis to my parents, Russell and Margaret Scott.

TABLE OF CONTENTS

	Page
LIST OF TABLES	(vi)
LIST OF FIGURES	(vii)
CHAPTER 1: REVIEW OF EMPIRICAL AND THEORETICAL WORK	1
CHAPTER 2: THE FIRST EXPERIMENT	28
CHAPTER 3: QUANTITATIVE FRAMEWORK FOR THE STUDY OF CONTINGENCY IN OPERANT CONDITIONING	40
CHAPTER 4: THE SECOND EXPERIMENT	50
CHAPTER 5: THE THIRD EXPERIMENT	61
CHAPTER 6: GENERAL DISCUSSION	67
REFERENCES	83
APPENDIX: ASYMPTOTES OF INDIVIDUAL SUBJECTS	88

LIST OF TABLES

	Page
TABLE 1. Phi framework parameter values for the six groups in the first experiment	49
TABLE 2. Parameter values for the three groups in the second experiment	51
TABLE 3. Analysis of variance of target components for the second experiment	55
TABLE 4. Tukey's HSD test of differences in asymptotic target components for the second experiment	57
TABLE 5. Parameter values for the four groups in the third experiment	62
TABLE 6. Group asymptotes for the third experiment	65
TABLE 7. Standard errors of the mean asymptote for all groups in the three experiments	76

LIST OF FIGURES

	Page
FIGURE 1. The joystick apparatus	6
FIGURE 2. The contingency square	12
FIGURE 3. The 2x2 frequency table and the 2x2 joint probability table	24
FIGURE 4. The mean <u>vector</u> of a distribution, determined by the endpoint (X,Y). The length of the mean vector is denoted by R and the target component of the mean vector is denoted by TC	34
FIGURE 5. Group learning curves for the first experiment	37
FIGURE 6. Group asymptotes as a function of q in the first experiment. The curve of the function $TC_{\infty} =$ $q^{1/2}$ is shown	38
FIGURE 7. The 2x2 table determined by the experimentally controlled parameters u, v, and w	45
FIGURE 8. Group learning curves for the second experiment	54
FIGURE 9. Group learning curves for the third experiment	64
FIGURE 10. Group asymptotes as a function of ϕ for all groups in the three experiments. The curve of the function $TC_{\infty} = \phi^{1/2}$ is shown	71

Chapter 1

REVIEW OF EMPIRICAL AND THEORETICAL WORK

Identification of response-reinforcer relationships which lead to increases in response frequency is an essential task in operant conditioning. For many years, learning theorists (Thorndike, 1911; Skinner, 1938) have specified temporal contiguity of responses and reinforcers as the critical relationship. Recently, Rescorla's (1967, 1968, 1969) studies of the role of contingency between conditioned stimuli and unconditioned stimuli in classical conditioning have drawn attention to contingency between responses and reinforcers as a potentially important variable in operant conditioning (Gibbon, Berryman and Thompson, 1974; Hammond, 1980). This thesis evaluates one model of response-reinforcer contingency.

Temporal contiguity of responses and reinforcers

According to temporal contiguity accounts of operant conditioning, when a reinforcer follows a response closely in time, it automatically increases the likelihood that the response will recur. In his demonstration of "superstition" in pigeons, Skinner (1948) arranged for response-independent food delivery every 15 sec. He described the development of obvious, idiosyncratic, stereotyped behavior in 6 of his 8 subjects. Skinner offered a temporal contiguity account of his observations. According to this account, some behavior would be adventitiously followed by food and the resulting increase in frequency of this behavior would increase the likelihood of its recurrence just before another food delivery. Through repeated applications of this

process the behavior would progressively increase in frequency and be maintained by the food presentations, even though food did not depend on the pigeon's behavior.

Herrnstein (1966) elaborated Skinner's temporal contiguity argument. He first trained a pigeon to peck an illuminated disc by delivering food for the first peck after at least 11 sec. had elapsed since the last food delivery (i.e., an FI 11 sec. schedule). Once pecking was frequent, Herrnstein arranged food delivery every 11 sec. without regard to the pigeon's behavior. He reported that the rate of pecking declined, but that it "remained at a substantial level". Herrnstein argued that since pecking was the dominant form of behavior once the response-independent food delivery was arranged, it was very likely to be adventitiously followed by food and maintained at a substantial level. The decline in rate of pecking was explained as a consequence of the less immediate contiguity between pecking and food resulting from the response-independent schedule. Of course, the decline may also be explained as the result of weakening the contingency between pecking and food in the response-independent phase.

Skinner's (1948) observations are not sound support for the temporal contiguity hypothesis because Staddon and Simmelhag (1971) failed to replicate Skinner's demonstration of superstition. They used six pigeons in a "careful study of the superstition situation". Staddon and Simmelhag described the appearance of two kinds of activity-- idiosyncratic interim activities in the early part of the interfood interval, and terminal pecking later in the interval and up to food delivery. They explained the appearance of terminal pecking as temporal classical conditioning. The interim behavior was attributed to an

adaptive mechanism by which organisms "budget their time efficiently". The important point here is that Staddon and Simmelhag not only failed to replicate Skinner's demonstration, but provided evidence against his temporal contiguity account. They noted that behaviors other than pecking were more frequently contiguous with food early in training, yet pecking emerged as the dominant response. It appears that mere temporal contiguity was not sufficient to strengthen those other behaviors early in training.

A contiguity account has been offered to explain the results of several spatial stereotyping studies. Antonitis (1951) compared response variability under continuous reinforcement and extinction. In the continuous reinforcement condition, he arranged for food to be delivered when a rat poked its nose anywhere in a 20-in. slot in one wall of an operant chamber. Antonitis found that over trials an initial bias to the center of the slot became stronger in all his subjects. He argued that the mere temporal contiguity of responding at the center and food increased the relative frequency of responding at the center. A problem with this explanation is that the initial bias was not idiosyncratic. Since all subjects came to respond at the center of the slot, it is likely that some feature of the apparatus made responding at the center preferable. In fact, the foodcup was located on the wall opposite the slot and was closest to the center of the slot. Responding at the center may have required less effort in running between the slot and the foodcup, or reduced the delay of reinforcement. These possibilities mean that the stereotypy in response location cannot be unambiguously attributed to the adventitious temporal contiguity of responses at the center and food.

Eckerman and Lanson (1969) compared response variability under continuous reinforcement, intermittent reinforcement and extinction. They used pigeons and an apparatus designed to be analogous to that used by Antonitis. The operant chamber had a 10-in. response key. The major departure from the design of the apparatus used by Antonitis was the location of the foodcup directly below the response key. Eckerman and Lanson observed a bias to respond to the center of the key in all their birds under the continuous reinforcement condition. This bias did not become stronger with continued training, as would be expected if temporal contiguity were effective.

Herrnstein (1961) used a pigeon apparatus similar to Eckerman and Lanson's. The chamber had a 10-in. rubber response strip and the foodcup was on the opposite wall of the chamber. Under continuous reinforcement, Herrnstein reported a preference for ends of the strip in all three of his pigeons. Only terminal data were presented, so it is not possible to determine if the stereotypy was progressive. Again, the stereotyped responding was not idiosyncratic. Eckerman and Lanson (1969) discuss how the spatial arrangements in the chamber could have given rise to the end-preference observed in Herrnstein's study. The studies reviewed above are usually cited as the strongest evidence for the operation of a temporal contiguity mechanism, yet it seems fair to conclude that none of them provides convincing support for the temporal contiguity hypothesis.


Davis and Platt (1978) designed a joystick apparatus so they could study spatial response differentiation in a situation which minimized response biases of the sort encountered by investigators using a spatial dimension with distinctive endpoints. The apparatus was

designed for rats and will be described in detail because it was used in the experiments to be reported in this thesis. A weighted steel rod, hanging from the ceiling of an operant chamber, could be displaced from the vertical in any direction. When the rod was displaced 7.5 cm. from the vertical it contacted a metal ring at the top of the chamber, the houselight went out, and an electromagnet held the stick in the vertical position for a 10-sec. intertrial interval. A diagram of the joystick apparatus is shown in Figure 1. The locus of points described by the end of the joystick at maximum excursion was a circle of radius 7.5 cm. When the joystick was displaced to its maximum excursion, the end of the joystick determined the endpoint of a particular radius of the circle. The response location of any displacement of the joystick was specified by the clockwise angle between this radius and an arbitrary zero-degree reference radius determined by the end of the joystick when the joystick was displaced toward the center of the foodcup.

Statistics appropriate for describing central tendency and variability of response location will be explained in Chapter 2. Davis and Platt used such statistics to compare distributions of response locations under conditions of nondifferential and differential reinforcement. Under the nondifferential conditions of continuous and intermittent reinforcement, Davis and Platt found no progressive stereotyping of response location, despite initial idiosyncratic biases. The failure of reinforcement to strengthen these initial biases calls into question the efficacy of response-reinforcer contiguity. Under differential reinforcement of response location, where the displacement responses with locations closest to a designated target location on the

FIGURE 1. The joystick apparatus

6a



circle were reinforced, there was a rapid progressive concentration of responding around the target.

The fact that differential reinforcement produced rapid shaping to a target in this situation where temporal contiguity was not sufficient to strengthen initial biases suggested that a contingency of reinforcers on response location might be necessary to produce increases in the relative frequency of particular response locations.

Contingency between responses and reinforcers

Contingency as a possible controlling variable was first studied systematically in classical conditioning by Rescorla (1967, 1968, 1969). Rescorla's notion of contingency was that there must be a difference between the probability of the unconditioned stimulus (US) in the presence of the conditioned stimulus (CS) and in the absence of the CS. If the absence of the CS is denoted by \overline{CS} , then Rescorla's notion of contingency was that there be a difference between $P(US|CS)$ and $P(US|\overline{CS})$. When these two conditional probabilities are equal, there is no contingency of US's on CS's. Rescorla called such a condition a 'truly random control'.

Rescorla (1968) used rats in a conditioned emotional response (CER) paradigm to assess the role of CS-US contingency in conditioning. All rats were first trained to lever-press to obtain food delivered according to a variable-interval (VI) schedule. A VI schedule arranges for the reinforcer to be delivered following the first response t_1 or more seconds since the last reinforcer. Various t_1 values are used by the VI schedule, and the value t appearing in the designation VI t -sec. schedule is simply the mean of the individual t_1 's used in

the construction of the schedule. In the classical conditioning phase of the CER paradigm, the lever was withdrawn and each group was exposed to a different programmed sequence of CS's (tone) and US's (electric shock). Following the classical conditioning, the lever was extended and leverpressing was maintained by food delivered on a VI 120-sec., schedule. Once the rats were lever-pressing at a stable rate, the test phase began. During the test sessions, the lever was always present, and the CS was presented occasionally.

Fear conditioned to the CS during the classical conditioning phase--or in more neutral terms, the strength of the CS-US association--was indexed by the decrease in rate of lever-pressing during the CS. In the first experiment there was one contingent condition and two noncontingent conditions. In the contingent procedure US's occurred only during CS periods, and the rats showed conditioned suppression to the CS. In one noncontingent control, shocks were given during both CS and $\overline{\text{CS}}$ periods at the same rate as in CS periods in the contingent group. In the other noncontingent group, the total number of shocks was the same as in the contingent group, but shocks occurred randomly with respect to CS and $\overline{\text{CS}}$ --so here too $P(\text{US}|\text{CS}) = P(\text{US}|\overline{\text{CS}})$: The noncontingent groups did not show suppression to the CS. This experiment demonstrated that mere temporal contiguity of CS's and US's does not produce conditioning, and that the absence of conditioning does not depend on the rate of shocks. As Jenkins, Barnes, and Barrera (1981) pointed out, this experiment alone does not demonstrate the importance of CS-US

contingency because the strength of the contingency, as measured by the difference between $P(US|CS)$ and $P(US|\overline{CS})$, was not varied.

Rescorla's (1968) second experiment examined conditioning as a function of the difference between $P(US|CS)$ and $P(US|\overline{CS})$. He found that conditioned suppression increased with the size of the difference between the two conditional probabilities. It is this result which has frequently been taken as proof that contingency is the determinant of classical conditioning. However, Jenkins, Barnes, and Barrera (1981) argued that this conclusion is premature since Rescorla's data do not discriminate between a contingency hypothesis and a hypothesis first proposed to account for the results of some autoshaping studies.

Gibbon's (1977) scalar expectancy theory has been applied to autoshaping (Gibbon, 1981; Gibbon and Balsam, 1981), and this work has produced the generalization that acquisition of autoshaped keypecking depends on the ratio of the average waiting time per feeding in the trial stimulus to the average waiting time per feeding overall. Jenkins et al. (1981) divorced this empirical generalization from Gibbon's particular theoretical development and called it the 'relative waiting time hypothesis'. Jenkins has adapted the relative waiting time hypothesis to the general classical conditioning situation and shown how it can account for Rescorla's data. Whatever the final resolution of the role of contingency in classical conditioning, Rescorla's work has shown that contiguity is not sufficient to produce conditioning and that a thorough analysis of CS-US relations leads to valuable insights into the essential ingredients of a successful conditioning procedure.

Most of the operant literature is based on procedures involving response-dependent reinforcer delivery. Reinforcer delivery may depend

on the simple occurrence of a discrete response like a lever press, or it may depend on the occurrence of a response having a designated property, like a lever press with a duration longer than t sec. The term 'contingency' is often used to express this dependence of reinforcer delivery on criterion responses. Usually the actual dependency is specified, as for example by a description of a reinforcement schedule, without the benefit of a general concept of contingency that delimits the possible range and meaning of degrees of dependence.

While the term 'contingency' is rife in the operant literature, often it refers to little more than an arrangement in which the reinforcer was delivered if and only if a criterion response occurred. This is the special case of a perfect contingency of reinforcers on criterion responses. A noncontingent procedure usually means that the reinforcer was delivered without regard for the subject's behavior. Of course, since the subject controls its behavior, there may be an actual association between the occurrence of a criterion response and reinforcer delivery -- i.e., a nonzero contingency.


In order to demonstrate that contingency is an important variable in operant conditioning, it is necessary to show that the dependent variable is systematically related to a measure of contingency. Choice of a measure of contingency depends on factors such as the use to which the measure is to be put and preconceived models of the nature of the relationship between the variables (see Costner, 1965; Weisberg, 1974). With the important exception of Gibbon, Berryman, and Thompson (1974), psychologists have not confronted the issue of what properties a measure of contingency should have if it is to be useful in

animal conditioning paradigms. Instead, most investigators have simply followed Rescorla in using the difference between two conditional probabilities as a measure of contingency. Denoting reinforcer delivery by Z , and the occurrence of a criterion response by W , this measure of contingency in operant conditioning is $P(Z|W) - P(Z|\bar{W})$.

One difficulty with this definition of contingency in operant conditioning is that failure to control the probability of a criterion response means that the subject may not be exposed to reasonable samples of both conditional probabilities. For example, the programmed probability of the reinforcer given no response is irrelevant if the subject makes a response on every trial. This problem does not arise in classical conditioning because the analogue of the probability of a response, namely $P(CS)$, is under experimental control. The experimenter can choose $P(CS)$ to ensure that the subject is exposed to reasonable samples of both $P(US|CS)$ and $P(US|\bar{CS})$.

All operant contingency studies using the conditional probability definition of contingency have made reinforcer delivery contingent on the simple occurrence of a discrete response, rather than on the occurrence of a response meeting a criterion on some measurable property of that response. These studies can be located on the contingency square shown in Figure 2. The vertical axis is the probability of the reinforcer given a response, denoted by u ; the horizontal axis is the probability of the reinforcer given no response, denoted by v . The degree of contingency, as measured by the difference between the two conditional probabilities, $(u-v)$, increases in magnitude with increasing distance from the positive diagonal of the square. The positive and negative extremes of this contingency variable are

FIGURE 2. The contingency square



represented by the upper left and lower right corners of the square, respectively. Noncontingent conditions, for which contingency is 0, are represented by points on the positive diagonal of the square.

Most studies of the role of contingency in operant conditioning have used shock as the reinforcer, or more accurately, the punisher. Gibbon (1967) used rats in a discrete-trial procedure to study the two endpoints of the upper edge of the square ($u=1, v=0$ and $u=1, v=1$). Each condition was associated with a distinctive stimulus, and the two conditions were alternated within a session. Lever pressing was supported by a VI schedule of food delivery during the intertrial interval. Since $u=1$ in both conditions, the first response after stimulus onset terminated the stimulus and produced a shock. Failure to respond resulted in shock under the noncontingent condition, since $v=1$; but no shock under the strict punishment condition, since $v=0$. Gibbon found that responding was more frequent in the noncontingent condition.

Neffinger and Gibbon (1975) studied avoidance contingencies represented by points on the bottom edge, the right edge, and the noncontingent diagonal of the square. They used rats in a discrete-trial procedure with no intertrial interval. The response was a lever press. In the first experiment subjects were exposed to the maximal avoidance contingency ($u=0, v=1$), and then were assigned to one of three noncontingent conditions-- $u=v=0$, .5 or 1. Eight of the ten rats stopped responding under the noncontingent condition. In the second experiment, rats were first trained on the maximal avoidance contingency ($u=0, v=1$). Two groups of six rats each were then used to study the right edge and the lower edge of the square. For one group v was reduced while u was held at zero--this corresponds to going from right to left on the bottom

edge. Neffinger and Gibbon found that responding decreased to zero as v approached zero. For the other group, u was increased while v was held at one--this corresponds to moving up the right edge of the square. For three of the six rats, responding decreased to zero with increasing u . The other three rats showed little change as u increased, and they maintained responding at the noncontingent point $u=v=1$. When subjects in the second experiment were exposed to the noncontingent conditions $u=v=0$, .5 and 1, two patterns of responding were found. Half the subjects stopped responding under all the noncontingent conditions, but half the subjects maintained some responding as long as shock density was above zero. Neffinger and Gibbon argued that for the latter subjects, responding was under the control of both contingency and shock density. This is supported by the observation that responding increased as shock frequency increased, even though the contingency was the same--zero--for all noncontingent conditions. Since shock frequency may compete with contingency under some conditions, this points out a drawback to using an aversive conditioning procedure to study contingency, at least when the subjects are rats.

Kop (1974) and Kadden, Schoenfeld and Snapper (1974) studied the four edges of the square and the negative diagonal using rhesus monkeys in a discrete-trial paradigm with no intertrial intervals. The subject was seated in a restraint chair and electric shock was delivered to the tail. In both studies, once the subjects were shaped to press a lever to avoid shock, they were trained on the maximal avoidance contingency ($u=0$, $v=1$). Subjects were then assigned to different groups. In Kop's (1974) study, there were three groups, each with four monkeys. For one group, u was increased while v was held at one. This amounts to moving

up the right edge of the square. Response rates declined monotonically to low levels. No subject showed high levels of responding under the noncontingent condition $u=v=1$, so responding did not come under control of shock density as it did for some of Neffinger and Gibbon's subjects. When the subjects in Kop's first group were exposed to conditions on the top edge of the square in a right-to-left sequence, response rates rapidly dropped to zero. Another group was moved from right to left along the negative diagonal of the square. Denoting a condition by (u,v) , the points examined were $(0,1)$, $(.33,.67)$, $(.67,.33)$ and $(1,0)$. Responding decreased monotonically as u increased, and three of the four subjects reached near zero levels at $(1,0)$. The remaining group was taken from right to left along the bottom edge of the square, and then up the left edge. When subjects were taken from right to left along the bottom edge, the results were not simple. Response rates were higher under $(0,.67)$ and $(0,.33)$ than $(0,1)$. Kop does not state whether these differences were statistically significant. When the extinction condition $(0,0)$ was in effect, response rates dropped below the $(0,1)$ level. It is noteworthy that response rates were well above zero after 20 daily sessions of the extinction condition.

Kadden, Schoenfeld and Snapper (1974) used a procedure very similar to Kop's. They obtained much the same result for a group exposed to conditions on the right edge and top edge of the square, and for a group exposed to conditions on the negative diagonal of the square. For a group taken from right to left along the bottom edge, Kadden et al. did not replicate Kop's finding of an inverted-U relation between response rate and v . Responding did not increase initially, but simply remained at substantial levels as v was reduced. Responding fell

to zero as soon as the subjects were exposed to shock following responses, rather than showing the gradual decline observed by Kop. The finding of these two studies that responding did not decline as v was reduced to near zero agrees with Neffinger and Gibbon's (1975) observations. However, the rhesus monkeys maintained a substantial level of responding even under the extinction condition, in contrast to the cessation of responding by the rats in Neffinger and Gibbon's study.

As noted earlier, a serious problem with any conditional probability definition of contingency in operant conditioning is that there is no way to ensure that the subject will receive substantial exposure to both conditional probabilities. If the experimenter arranges two conditions with equal values of u but different values of v , it would appear that there are two different degrees of contingency. But if a subject were to respond on every trial, he would not be exposed to v at all. So what the experimenter programs may not make sufficient contact with the subject. While the failure to control the subject's response probability prevents the conditional probability specification from guaranteeing exposure to different degrees of contingency, this is not to say that the experimenter's intention is never realized. With a reasonable sample of both conditional probabilities, the subject will be exposed to different contingencies, but not in a well-controlled manner.

Flye and Gibbon (1979) attempted to overcome this problem of controlling the subject's actual experience through control of the joint probability of a response and a shock (p_{11}) or the joint probability of no response and no shock (p_{22}). This is referred to as a joint probability manipulation of contingency. Flye and Gibbon used rats in a paradigm similar to Neffinger and Gibbon's. All rats were first trained

on the maximal avoidance contingency ($p_{11}=p_{22}=0$). For one group of six rats the omission probability p_{22} was manipulated while p_{11} was held at zero. For the other group of six rats the punishment probability p_{11} was varied while p_{22} was held at zero. Note that there are only two noncontingent procedures in this joint probability scheme, one where no shock occurs and one where shock is delivered on every trial. The subjects in the omission group all showed graded decrements in responding with increasing probability of omission. Subjects in the punishment group were idiosyncratic in their reactions to increasing punishment of responding. Some subjects decreased responding to near zero as p_{11} increased, while others continued to respond at substantial levels even when all responses were followed by shock. When subjects in the omission group were subsequently exposed to the noncontingent condition $p_{11}=1, p_{22}=0$, a wide range of response levels was observed. Flye and Gibbon noted that their results paralleled Neffinger and Gibbon's (1975), with some subjects sensitive to contingency alone and some subjects sensitive to both contingency and shock density.

One problem with the joint probability manipulation of contingency is that the joint probability can be held constant only so long as response probability is within a particular range. For example, the value of p_{11} can never exceed the probability of a response. As programmed p_{11} increases, it may well happen that response probability decreases to the point where larger values of p_{11} cannot be achieved. When this occurs, the obtained value of p_{11} will be less than the programmed value. Flye and Gibbon actually obtained this result for most subjects when p_{11} was .5 or larger.

There has been very little investigation of contingency in appetitive situations. Hammond's (1980) study resembles Neffinger and Gibbon's (1975), with the major difference being that Hammond used a free-operant procedure and food as the reinforcer. The entire session was divided into one-second, unsignalled units. For any second in which the rat made at least one lever press, food was delivered with probability u . For any second in which no response was made, food was delivered with probability v . Hammond also defined contingency as the difference between these two conditional probabilities. In the first experiment, the rats were initially exposed to a positive contingency ($u=.05$, $v=0$) and then shifted to a zero contingency ($u=v=.05$). All subjects showed a large decline in responding when the zero contingency was introduced, but responding was not eliminated. In the second experiment, two groups of rats were first trained on a positive contingency ($u=.05$, $v=0$). One group was shifted to a zero contingency ($u=v=.05$), while the other group was shifted to a negative contingency ($u=0$, $v=.05$). Hammond found that the negative contingency was much more effective in reducing the level of responding. Two other groups were first trained on a positive contingency ($u=.12$, $v=0$). One group was shifted to a weaker positive contingency ($u=.12$, $v=.08$) and the other group was shifted to a zero contingency ($u=v=.12$). There was considerable overlap in response rates between the two groups, but the zero contingency did produce a significantly larger decline in responding than did the intermediate contingency. The earlier criticisms of a conditional probability definition of contingency apply here, but Hammond's study does provide some evidence for the importance of a contingency variable in an appetitive conditioning paradigm. A more

serious problem, noted by Hammond, is the arbitrary definition of a nonresponse. As Hammond asks, is it reasonable to interpret a five-second period with no responding as five consecutive nonresponses? Of course, the discrete-trial shock studies are just as arbitrary in defining a nonresponse as a period of nonresponding equal to the trial duration.

In both the conditional and joint probability contingency studies, failure to control the probability of a response can prevent the programmed contingency from being achieved. In general it would be advantageous to be able to control response probability, though it would make little sense in studies using response rate as the dependent variable. There is, however, a fundamental problem with contingency studies that make reinforcer delivery contingent on the simple occurrence of a discrete response. The subject determines when a response will occur, and the experimenter has no control over the occurrence of a discrete response. This precludes the experimental control of the probability of a discrete response. Of course, studies such as Neffinger and Gibbon (1975) and Flye and Gibbon (1979) can obtain useful results because much of the time the subject's response probability is within a range that permits a close approximation to the programmed contingency to be achieved.

Fortunately, it is possible to avoid problems stemming from the failure to control response probability if reinforcer delivery is made contingent upon the occurrence of a criterion response property rather than on simple occurrence of a discrete response. If the response property used to identify criterion response can be measured on at least an ordinal scale, then the probability of a criterion response

can be controlled by a percentile schedule (Platt, 1973). The probability of a criterion response can then be chosen to ensure adequate exposure to both conditional probabilities in $P(Z|W)-P(Z|\bar{W})$, or to allow particular values of a joint probability to be realized. In addition, with the probability of a criterion response controlled, a possible source of between-subject variance is removed.

Percentile schedules define a criterion response by reference to a distribution of the subject's recent response values. The notion underlying the percentile procedure is that response values can be thought of as coming from a theoretical population distribution, so that a criterion response can be defined as one with a value coming from a predetermined part of the population distribution--for example, the upper quarter. A percentile schedule uses an ordered list of the subject's m most recent response values to estimate percentile points of the population distribution, and identifies criterion responses as those with values either above or below an estimated percentile point. If shaping is effective, the population distribution will gradually change and the value associated with that percentile point will vary, but the probability of a criterion response will remain constant. The rationale for the percentile procedure is that if the m values are a random sample from the population distribution, then the expected proportion of that distribution falling between any two consecutive sample values is $1/(m+1)$ (see Platt, 1973). If a criterion response is defined as one with a value larger than k of the m ordered values, then the proportion of responses which are criterion responses is $1-(k/(m+1))$. In practice, when shaping is effective, the sample values do not all come from the same population distribution. Further, the assumption of random sampling may be violated by sequential dependencies between successive

response values. However, Platt (1973) has pointed out that there is a wide range of moderate values of m for which violations of the assumptions are not severe enough to be of practical significance.

A slightly more complicated type of percentile schedule is the targeted percentile, which defines a criterion response as one with a value coming from a predetermined proportion of the population distribution closest to a target value. The procedure uses two simple percentile schedules, one for the most recent responses with values below the target, and one for the most recent responses with values above the target. Webster (1976) has discussed the details, but for the present purpose it is sufficient to note that the targeted percentile procedure identifies criterion responses as those with values in the upper tail of the distribution of values below the target and those with values in the lower tail of the distribution of values above the target. A targeted percentile schedule maintains a constant probability of a criterion response, defined as one with a value relatively close to a target value.

Consider Davis and Platt's (1978) joystick paradigm for the study of spatial response differentiation. It is not possible to control the probability of a joystick-displacement response. However, a targeted percentile schedule will control the probability of a criterion response, defined as one with a location closer to the target than a predetermined proportion of the population distribution. The first experiment to be reported in this thesis investigated the role of a novel response location-reinforcer contingency variable in the joystick paradigm, and employed targeted percentile schedules to avoid possible

problems resulting from failure to control the probability of a criterion response.

Davis and Platt (1978) had already studied the extremes of response location-reinforcer contingency. A nondifferential, intermittent reinforcement schedule corresponds to zero contingency because the reinforcer is delivered without regard to response location. A targeted percentile schedule establishes a perfect contingency between response location and reinforcers since criterion responses are always followed by the reinforcer and noncriterion responses are never followed by the reinforcer. A trial on which the nondifferential schedule is in effect may be termed a noncontingent trial, and a trial on which the targeted percentile schedule is in effect may be termed a contingent trial. The innovation of the first experiment of this thesis was to arrange intermediate degrees of contingency by randomly mixing noncontingent and contingent trials in various proportions within a daily session. The index of response location-reinforcer contingency was the proportion of trials in a session which were contingent trials, and was denoted by q . The measure of shaping to a target was the same as that used by Davis and Platt (1978), and will be described in Chapter 2. The purpose of the first experiment was to manipulate q between groups to determine whether shaping to a target was related in an orderly way to this index of contingency. The experiment revealed that the larger the degree of contingency, the more effective was the shaping to a target.

The evidence to this point indicated that contingency, whether indexed by $(u-v)$ or q , was a potent variable in operant conditioning. The next step in the study of contingency seemed to require a broader definition of contingency or a more general measure of contingency than

(u-v) or q. One approach to finding a reasonable measure of the dependence of reinforcer delivery on the criterion response is to find a suitable measure of the statistical association between reinforcer delivery and the criterion response. This notion of contingency dates back to Pearson (1904), who asserted that "...the greater the contingency, the greater must be the amount of association or correlation between the two attributes."

Any measure of association between two dichotomous variables can be expressed in terms of the four cell frequencies of the 2x2 frequency table shown in Figure 3. The occurrence of a criterion response is denoted by W, and reinforcer delivery is denoted by Z. Cell entries are absolute frequencies of the four joint events. Note that the relative frequency of the criterion response is $(a+b)/N$, and the conditional relative frequency of the reinforcer given a criterion response is $a/(a+b)$. Since probabilities are approximated by relative frequencies, the terms are often used synonymously. For example, the 'probability of a criterion response' is used to denote the relative frequency of a criterion response. Note also that the joint probability table shown in Figure 3 can be obtained from the frequency table by dividing cell frequencies by the sample size, N. The important difference between the two tables of Figure 3 is that the joint probabilities must sum to 1.

The choice of a statistical measure of association for use as a measure of contingency in animal conditioning paradigms was the subject of a trenchant analysis by Gibbon, Berryman, and Thompson (1974). They considered several traditional measures of association as measures of CS-US contingency and eliminated all but one of the candidates by pointing out an invariance property they possess which makes them

FIGURE 3. The 2x2 frequency table and the 2x2 joint probability table

**FREQUENCY
TABLE**

		Reinforcer		
		Z	\bar{Z}	
Response	W	a	b	a + b
	\bar{W}	c	d	c + d
		a + c	b + d	N

**JOINT PROBABILITY
TABLE**

		Reinforcer		
		Z	\bar{Z}	
	W	p_{11}	p_{12}	p_{1+}
	\bar{W}	p_{21}	p_{22}	p_{2+}
		p_{+1}	p_{+2}	1.0

NOTATION:

$$w = P(W) = p_{11} + p_{12}$$

$$u = P(Z|W) = p_{11} / (p_{11} + p_{12})$$

$$v = P(Z|\bar{W}) = p_{21} / (p_{21} + p_{22})$$

incapable of describing partial reinforcement effects. This invariance argument will be presented in Chapter 3, along with a more detailed discussion of measures of association.

Gibbon, Berryman, and Thompson (1974) discussed the merits of the one candidate, the phi coefficient (ϕ), not excluded by their invariance argument. They observed that ϕ is sensitive to changes in the probability of the CS, $P(\text{CS})$, and that Stein, Sidman, and Brady's (1958) results suggest that $P(\text{CS})$ may be an important variable. Gibbon et al. (1974) discussed problems with conventional trial and probability definitions in classical conditioning, and noted that ϕ values for Stein, Sidman, and Brady's (1958) study cannot be calculated unambiguously as a result of these problems. It is noteworthy that while ϕ is sensitive to changes in $P(\text{CS})$, the usual measure of CS-US contingency, $P(\text{US}|\text{CS}) - P(\text{US}|\overline{\text{CS}})$, is not. The data from the Stein et al. (1958) study are not conclusive, but they do favor choice of ϕ .

Gibbon, Berryman, and Thompson (1974) stated that "... ϕ , or some monotone function of ϕ , appears a reasonable measure of the power of the training contingencies at asymptote in the conditioned suppression paradigm." However, they noted that there is a fundamental difference between classical conditioning and operant conditioning which undermines the usefulness of ϕ in operant conditioning. In classical conditioning, ϕ is sensitive to changes in $P(\text{CS})$, which is under experimental control. But in operant conditioning, ϕ is sensitive to changes in the probability of a response, and this parameter is not under experimental control. This means that:

"... ϕ is calculable only on a post hoc basis. One cannot then argue that partial contingencies control behavior when behavior in turn defines the degree of contingency."

This points out a further problem with making reinforcers contingent on the simple occurrence of a discrete response: the probability of a response cannot be experimentally controlled, so that utility of ϕ as a measure of contingency for operant conditioning can never be tested. But this problem can be sidestepped if reinforcers are made contingent on the occurrence of responses with a criterion value of some property. For then, as described earlier, a percentile schedule can be used to control the probability of a criterion response; and hence, the value of ϕ can be fixed prior to the experiment.

The use of a percentile schedule to control the probability of a criterion response removes the obstacle to applying ϕ in operant conditioning. Since Gibbon, Berryman, and Thompson (1974) selected ϕ partly because the other candidates they considered cannot account for partial reinforcement effects in classical conditioning, it is important to show that the partial reinforcement effect obtains in the operant conditioning paradigm used to test ϕ . The second and third experiments to be reported in this thesis were designed to test the usefulness of ϕ as a measure of response location-reinforcer contingency in the joystick paradigm.

Outline of thesis

Chapter 2 describes the statistics used to summarize response distributions on a circular response dimension, and presents the first experiment. In this experiment, the proportion of trials in a session which were contingent trials was manipulated between groups. The result was that the higher the contingency of reinforcers on response location, the more effective was the shaping of responding to a target.

Chapter 3 provides a discussion of measures of association, including arguments for choosing ϕ as a measure of contingency in animal conditioning paradigms. Chapter 3 develops a quantitative framework, based on ϕ , for studying the influence of response location-reinforcer contingency on spatial response differentiation. The first experiment is reinterpreted as evidence for the usefulness of the framework.

Chapters 4 and 5 present the second and third experiments as additional evidence for the usefulness of the ϕ framework. The three parameters in the expression for ϕ were manipulated in these experiments. Regardless of how the value of ϕ was varied, the higher the response location-reinforcer contingency as indexed by ϕ , the more effective was the shaping of responding to a target location.

Chapter 6 summarizes the empirical work, discusses problems and unresolved issues, and outlines directions for future research.

Chapter 2

THE FIRST EXPERIMENT

The purpose of the first experiment was to determine the relationship between a measure of shaping to a target location in the joystick situation, and a measure of the contingency of reinforcers on response location. The measure of shaping indicates the degree to which subjects respond toward the target location, and will be described later. The index of contingency is the proportion of trials in a daily session which were contingent trials, and is denoted by q . On contingent trials, only criterion responses, those responses with locations closest to the target, were followed by food delivery. A targeted percentile schedule identified criterion responses on contingent trials. On noncontingent trials, responses were followed by food with the same overall probability as on contingent trials, but irrespective of response location. Contingent and noncontingent trials were mixed randomly in a daily session, with the proportion of trials which were contingent trials experimentally controlled.

METHOD

Subjects

Forty male Wistar rats obtained from Canadian Breeding Farms served as subjects. Their weights at the beginning of the experiment ranged from 242 to 330 gm. For the entire experiment each rat was maintained at 85% of its free feeding weight by supplemental rations of Purina Rat Chow given in the home cage immediately after the experimental session.

Apparatus

Four Lehigh Valley Electronics (Model #143-20-215) operant conditioning chambers in sound-attenuating enclosures were used. Each chamber was 30.5-cm. long by 24.3-cm. wide by 24.5-cm. high. A 3.8-cm. circular hole in the front wall provided access to a foodcup mounted behind the wall. A BRS-Foringer pellet dispenser delivered .045-gm. Noyes pellets into the foodcup. The chamber was illuminated by a 2.8-W incandescent houselight mounted 17.5 cm. above the foodcup. The front and rear walls of the chamber were aluminum; the side walls of the chamber were plexiglass. A plexiglass floor was mounted over the standard grid floor.

A weighted stainless-steel rod, the joystick; was suspended from the center of the ceiling of each chamber (see Figure 1). A minimum force of approximately .2 N was required to displace the joystick from its vertical position until it contacted a metal ring mounted on the ceiling. The locus of points described by the end of the joystick at maximum excursion was a circle of radius 7.5 cm., so that one degree arc of the circle subtended 1.31 mm. at the circumference.

The vertical plane passing through the center of the foodcup determined a diameter of the circle. This diameter was taken as the X-axis of a rectangular coordinate system with origin at the center of the circle. The joystick was connected to two variable resistors wired as voltage dividers and mounted at right angles to each other above the ceiling. The voltage drop across one resistor was a linear function of the displacement of the joystick along the X-axis; the voltage drop across the other was a linear function of the displacement of the joystick along the Y-axis. For each location of the joystick at maximum

excursion there was a unique ordered pair of voltage values, and the computer transformed these voltages into Cartesian coordinates X_1 and Y_1 . Hence, the location of any displacement response was recorded as a point (X_1, Y_1) on the circumference of the unit circle in Cartesian space.

For descriptive purposes, a location was specified by the clockwise angle between the radius with endpoint (X_1, Y_1) on the unit circle and the X-axis. For example, a response recorded as $(0, 1)$ would be described as having a location of 270° .

A Digital Equipment Corporation PDP8/E computer, located in a separate room, generated experimental conditions and recorded response locations.

Procedure

The experiment was run in two replications because of time and equipment limitations. There were 16 subjects in the first replication and 24 subjects in the second replication. The procedure described below applies to both replications, except where a procedural difference is explicitly noted.

After magazine training, all rats were hand shaped to approach and then operate the joystick. Once the rat was pushing the joystick reliably, an automated fixed-ratio 1 (FR 1) schedule was implemented. Every operation of the joystick through its maximum excursion, regardless of its location, produced the immediate delivery of a food pellet into the foodcup. At the same time as food delivery, the houselight went out and the electromagnet in the floor was activated for 10 seconds. The joystick was held in its vertical position during this intertrial interval and could not be operated until the houselight came

on at the beginning of the next trial. Two sessions of FR 1 training were given, each ending after 50 responses had been made. In all subsequent conditions, a session ended when each rat had made 128 responses or one hour had elapsed, whichever came first. The experiment was conducted 6 days a week.

Next, a single session of a random-ratio 2 (RR 2) schedule was in effect. Under this schedule, 50% of the rat's responses, determined by a pseudo-random number generator, were followed by food delivery. Following this, a RR 5 schedule which reinforced a randomly determined 20% of the rat's responses was in effect for 20 sessions in the first replication, but as no progressive change in responding was observed, the RR 5 schedule was in effect for only one session in the second replication.

Subjects were then randomly assigned to groups. Each group received a different value of q , the proportion of trials in a daily session which were contingent trials. On contingent trials only criterion responses were followed by food delivery. A targeted percentile schedule defined a criterion response as one with a location among the 20% currently closest to the target. The schedule used the subject's 24 most recent response locations to estimate the population distribution. On noncontingent trials, a randomly determined 20% of the subject's responses were followed by food delivery. Thus, 20% of the subject's responses were followed by food delivery on both contingent and noncontingent trials.

The groups differed in the value of q , the proportion of trials in a session which were contingent trials; and hence, in the degree to which reinforcers were contingent on response location. In the first

replication there were four groups of four subjects each, with $q = .2, .4, .6$, and $.8$. In the second replication there were six groups of four subjects each, with $q = 0, .2, .4, .6, .8$, and 1 . This phase of the experiment lasted for 34 sessions. A subsequent phase exposed all subjects to a perfect contingency ($q=1$) for 14 sessions. The purpose of this phase was to show that the group differences obtained in the previous phase were not due to some groups having rats that simply could not be shaped to a target.

RESULTS

Description of the dependent variable:

A subject made 128 joystick-displacement responses in a daily session, and to summarize its performance it is useful to have measures of the central tendency and concentration of the distribution of response locations. Also, since the present experiment involved differential reinforcement of response locations, it is desirable to have a measure of shaping, or the relationship between the subject's distribution of response locations and the target angle.


Statistics appropriate for summarizing distributions of responses on a circular dimension have been applied in the study of animal homing and navigation (Batchelet, 1965), and were employed in the experiments reported in this thesis. The location of each response was recorded as a point (X_i, Y_i) on the circumference of the unit circle with center $(0,0)$. Each response can also be represented by a unit vector with starting point $(0,0)$ and endpoint (X_i, Y_i) , and response location can be specified by the clockwise angle between this vector and the X-axis. The mean vector of the distribution of individual response

vectors provides a measure of the central tendency and concentration of the distribution. The mean vector has the starting point (0,0), and is determined by the endpoint (\bar{X}, \bar{Y}) . The coordinates of the endpoint are computed from the formulas $\bar{X} = \sum X_i / N$ and $\bar{Y} = \sum Y_i / N$, where N is the number of individual response vectors (X_i, Y_i) in the subject's distribution. A diagram showing the mean vector of a distribution is given in Figure 4. While the mean vector is determined most conveniently by calculation of its endpoint (\bar{X}, \bar{Y}) within the rectangular coordinate system, its polar coordinates (A, R) are more useful for descriptive purposes. The mean angle A is the clockwise angle between the X-axis and the mean vector, and is a measure of the central tendency of the distribution of response locations. R is the length of the mean vector, and is a measure of the concentration of the distribution of response locations about the mean angle.

When the quadrant in which the mean vector lies is known, the mean angle A is determined by the acute angle α between the mean vector and the X-axis. The quadrant in which the mean vector lies is identified by the signs of \bar{X} and \bar{Y} . The acute angle α is computed by $\alpha = \arctan (|\bar{Y}|/|\bar{X}|)$ if $|\bar{X}| > |\bar{Y}|$ and $\alpha = 90^\circ - \arctan (|\bar{X}|/|\bar{Y}|)$ if $|\bar{X}| < |\bar{Y}|$. The mean vector length R is computed by $R = (\bar{X}^2 + \bar{Y}^2)^{1/2}$. If the mean vector length (R) is zero, no mean angle is defined; otherwise $0 \leq A < 360$. In Figure 4, the mean angle is between 180 and 270 degrees. The mean vector length ranges from 0 to 1, with higher values corresponding to greater degrees of concentration. A value of 1 would be produced by a subject who always responded to the same location.

The dependent variable in the experiments reported in this thesis is a measure of shaping, or the degree to which the subject came

FIGURE 4. The mean vector of a distribution, determined by the endpoint (\bar{X}, \bar{Y}) . The length of the mean vector is denoted by R and the target component of the mean vector is denoted by TC



to respond toward the target. It is called the target component of the mean vector, and is the projection of the mean vector on the reference line determined by the origin and the point on the circle corresponding to the target location. The target component of the mean vector is also shown in Figure 4, where it is denoted by TC. Since the target (270^0) in the experiments of this thesis was on an axis of the coordinate system used to specify response locations, the calculation of the target component was greatly simplified; in the present case the formula was $TC = \bar{Y}$, the ordinate of the mean vector.

The value of the target component ranges from -1 to +1. If the response distribution were highly concentrated and centered on the target, the target component would take a value close to +1. A target component of zero could result in two different ways. If the mean angle were 90 degrees away from the target, the projection on the target line would be zero, regardless of the mean vector length R. If the distribution had the property that every response had a counterpart 180 degrees away, the mean vector length would be zero and its projection would be zero. Note that a special case of this property occurs when responses are distributed uniformly about the circle. Finally, if the distribution were highly concentrated 180 degrees away from the target, the target component would be close to -1. A decrease in the concentration of the distribution, reflected by a lower value of R, results in a decrease in the magnitude of the target component.

The data:

The target components for all the subjects in a group were averaged after each session, and the mean target component for each group was plotted over sessions to produce a learning curve for each

group. The group learning curves are presented in Figure 5. The asymptote is defined as the mean target component over the last five sessions. Group asymptotes and their standard errors are shown to the right of the learning curves. Note that the $q=.4$ group had only seven rats because one rat died of a respiratory infection.

The $q=0$ group showed no systematic change in performance over sessions. Note that the average target component for this group was below zero for all 34 sessions, and the asymptote of $-.21$ was below the expected value of zero. One of the four subjects in this group was stereotyped away from the target: its asymptote was $-.75$. This stereotypy was present from the beginning and did not change over sessions. With this subject excluded, the group asymptote was $-.02$.

The learning curves of the other five groups were negatively accelerated with different asymptotes. The asymptotic level was reached in about 30 sessions by all five groups, and the group asymptotes increased with increasing values of q . The relation between group asymptotes (TC_{∞}) and q is shown in Figure 6. Most of the variance in group asymptotes can be accounted for by a negatively accelerated function of q , such as $TC_{\infty}=(q)^{1/2}$. In fact, this simple function accounts for 98% of the variance in group asymptotes. When the deviant subject in the $q=0$ group is included, the group asymptote is the bracketed point, but the function $TC_{\infty}=(q)^{1/2}$ still accounts for 94% of the variance in group asymptotes.

On the last session of the perfect contingency ($q=1$) phase, the average target component was .80 or higher for all six groups. This was

FIGURE 5. Group learning curves for the first experiment





FIGURE 6. Group asymptotes as a function of q in the first experiment. The curve of the function $TC_{\infty} = q^{1/2}$ is shown



a substantial increase for the groups that had low q values in the preceding phase, and their performance was comparable to that of the groups with high q values in the preceding phase.

DISCUSSION

These results provide good evidence that the contingency of reinforcers on response location, as indexed by q , is an important determinant of shaping to a target location. While these results demonstrate the importance of an intuitively appealing contingency variable in an operant conditioning paradigm, the manipulation of q was only one among many ways of varying the degree of contingency between reinforcers and response location. The next chapter presents arguments for the phi coefficient (ϕ), a measure of association in a 2×2 table, as a useful measure of response-reinforcer contingency. A quantitative framework based on ϕ is shown to subsume the contingency variable denoted by q .

Chapter 3

QUANTITATIVE FRAMEWORK FOR THE STUDY OF CONTINGENCY

IN OPERANT CONDITIONING

Both CS-US contingency and response-reinforcer contingency express a dependence or association between two dichotomous nominal variables, so any measure of the statistical association between nominal variables in a 2x2 table is a potential measure of contingency in conditioning paradigms. The statistical literature on measures of association is extensive (Goodman and Kruskal, 1979; Edwards, 1963; Costner, 1965; Weisberg, 1974) and there are many useful textbook discussions of these measures (Kendall and Stuart, 1973; Bishop, Fienberg, and Holland, 1975; Reynolds, 1977). The task is to select from this literature a measure of association with properties suitable for use in animal conditioning. While there is a large variety of different measures of association in larger $I \times J$ tables, many of these are equivalent to each other in 2x2 tables. Referring to the frequency table shown in Figure 3, most measures of association in 2x2 tables are functions of either the cross-product ratio ad/bc , or the phi coefficient $\phi = (ad-bc)/((a+b)(a+c)(b+d)(c+d))^{1/2}$ (Bishop et al., 1975).

One measure that is not a function of either the cross-product ratio or the phi coefficient is the difference in row proportions $a/(a+b) - c/(c+d)$. This is a difference between two conditional probabilities, and the problems with using this difference as an index of response-reinforcer contingency have been discussed in Chapter 1. In classical conditioning, this difference is Rescorla's index of contingency $P(US|CS) - P(US|\overline{CS})$. Gibbon, Berryman, and Thompson (1974)

have drawn attention to the possible importance of the probability of the CS, $P(\text{CS})$, and Rescorla's index is not sensitive to changes in $P(\text{CS})$. These considerations weaken the case for the difference in conditional probabilities as a measure of contingency for conditioning applications.

A potential measure of contingency should not change when all cell entries of the 2×2 table are multiplied by the same constant. This simply represents continued exposure to the same contingency, and it is the relation between asymptotic conditioning and the degree of contingency that is of interest. The cross-product ratio, ϕ , and the difference in conditional probabilities all have this desirable property.

Gibbon et al. (1974) noted that some measures of association have an invariance property that disqualifies them as measures of contingency: they are not changed when any row or column of the 2×2 table is multiplied by a constant. This invariance property means that the measure takes on one of its extreme values whenever one cell of the table is 0, and that the measure has the same extreme value for any two tables with 0 in the same cell. The 2×2 tables for any two procedures represented by intermediate points on the same edge of the contingency square shown in Figure 2 have 0 in the same cell, and hence produce the same value for any such measure. Consider procedures represented as points on the left edge of the square; the corresponding 2×2 tables all have 0 in the lower left cell, and hence the same value of any measure with this invariance property. Yet Gibbon et al. (1974) noted that Rescorla (1968) found partial reinforcement effects along this edge. A measure with this invariance property cannot describe graded effects

along any edge of the square. The cross-product ratio has this invariance property, so all functions of the cross-product ratio are disqualified as contingency measures.

Two common measures of association that are functions of the cross-product ratio are Yule's coefficient of association, Q , and Yule's coefficient of colligation, Y . It was mentioned earlier that many different measures of association for $I \times J$ tables are identical in 2×2 tables. Goodman and Kruskal's ordinal measure of association γ is identical to Yule's Q in 2×2 tables (Goodman and Kruskal, 1979, p. 20). If the 2×2 table is transformed by appropriate multiplications of rows and columns so that all marginal totals become .5, the value of Yule's Y is not changed. Goodman and Kruskal's measures of predictive association λ_a , λ_b , and λ all become equal in the transformed table, and all three are equal to the absolute value of Yule's Y (see Goodman and Kruskal, 1979, p. 37; Kendall and Stuart, 1973, p. 567).

The difference in conditional probabilities and functions of ϕ change when a row or column of the 2×2 table is multiplied by a constant, and they attain their extreme values only when both cell entries on the same diagonal of the 2×2 table are 0. While no conclusive evidence against the use of the difference in conditional probabilities has been presented, some potential problems resulting from its use have been discussed. Hence, the most promising candidate for a measure of contingency is a function of ϕ , and the simplest candidate is ϕ itself. The phi coefficient is related to the Pearson chi square statistic (χ^2) by the equation $\phi = (\chi^2/N)^{1/2}$, where N is the sum of the four cell frequencies in the 2×2 table--i.e., the sample size. The chi square statistic is a measure of the departure of the two variables from

statistical independence, so the variables are thought to be associated to the degree that they are not independent.

The chi square statistic itself is not suitable as a measure of dependence or association because its magnitude depends on the sample size. One approach to devising an interpretable measure of dependence is to divide χ^2 by the maximum value it may attain. In an $I \times J$ table, the maximum value of χ^2 is $N(t-1)$, where t is the smaller of I and J (Cramér, 1971, p. 282). The resulting measure of dependence or association, $\chi^2/N(t-1)$, is due to Cramér and measures the relative departure from independence. For a 2×2 table, the maximum value of χ^2 is N , and this special case of Cramér's statistic is phi squared: $\phi^2 = \chi^2/N$. Phi squared is sometimes called mean square contingency, and ϕ is called root mean square contingency. Values of phi squared range from 0 for no association to 1 for perfect association. In 2×2 tables it may be useful to distinguish between positive association and negative association. The phi coefficient is often used because it varies from -1 to $+1$, and can distinguish positive association from negative association. The value 0 corresponds to no association or statistical independence and the absolute value of ϕ increases with increasing strength of association.

The expression $\phi = (\chi^2/N)^{1/2}$ emphasizes the fact that ϕ reflects the degree to which there is non-independence between the two variables. But ϕ has another interpretation: it is a special case of the Pearson product moment correlation coefficient r . If the two categories of each variable are assigned arbitrary numbers and r is calculated for the resulting N pairs of numbers, the result is equal to ϕ (Conover, 1980,

p. 189). The phi coefficient is sometimes called the fourfold point correlation to emphasize the equivalence of ϕ and r in 2x2 tables.

Gibbon, Berryman, and Thompson (1974) argued the case for ϕ as a measure of CS-US contingency, but they did not derive an expression for ϕ in terms of operant conditioning parameters because they believed that one of the parameters could not be experimentally controlled. However, in Chapter 1 it was argued that this parameter, the probability of a criterion response, can be controlled by a percentile schedule. Hence, it is useful to have an expression for ϕ in terms of operant conditioning parameters.

The textbook formula for ϕ in terms of the 2x2 joint probability table shown in Figure 3 is $\phi = (p_{11}p_{22} - p_{21}p_{12}) / (p_{1+}p_{2+}p_{+1}p_{+2})^{1/2}$. An experimentally convenient way to determine all of the terms in this formula is to define $u = P(Z|W)$, $v = P(Z|\bar{W})$, and $w = P(W)$, where Z is reinforcer delivery and W is the occurrence of a criterion response. All the joint probabilities are determined once u , v , and w are specified; the 2x2 table determined in this way is shown in Figure 7. Note that z , the overall probability of the reinforcer, is determined by the formula $z = uw + v(1-w)$ once u , v , and w are specified. Applying the ϕ formula to the 2x2 table in Figure 7 gives:

$$\begin{aligned}\phi &= \frac{(uw(1-v)(1-w) - v(1-w)(1-u)w)}{(w(1-w)z(1-z))^{1/2}} \\ &= \frac{(u-v)w(1-w)}{(w(1-w)z(1-z))^{1/2}} \\ &= (u-v) \left(\frac{w(1-w)}{z(1-z)} \right)^{1/2}\end{aligned}$$

Since z is determined once u , v , and w are specified, it is a secondary

FIGURE 7. The 2x2 table determined by the experimentally controlled parameters u , v , and w

		Reinforcer		
		Z	\bar{Z}	
Response	W	$u \cdot w$	$(1-u) \cdot w$	w
	\bar{W}	$v(1-w)$	$(1-v)(1-w)$	$1-w$
		z	$1-z$	1.0

$$w = P(W)$$

$$v = P(Z|\bar{W})$$

$$u = P(Z|W)$$

$$z = u \cdot w + v(1-w)$$

or derivative variable in the ϕ framework. Writing z in terms of u , v , and w gives an expression for ϕ in terms of the primary, experimentally controlled variables u , v , and w :

$$\phi = \frac{(u-v)(w(1-w))^{1/2}}{([uw + v(1-w)][1-(uw + v(1-w))])^{1/2}}$$

This expression for ϕ provides the basis of a quantitative framework for experimentation. If ϕ is a useful measure of an effective contingency variable, then manipulation of ϕ through systematic variations in the three component parameters should have orderly effects on asymptotic conditioning levels.

If this quantitative framework is to organize experimental work on contingency, it should be able to accomodate the conditional probability studies reviewed in Chapter 1 and the first experiment reported in this thesis. From the perspective of the ϕ framework, the $(u-v)$ studies are incomplete because the probability of a response, w , was not controlled. However, for large differences in $(u-v)$, no variation in w would make the rank order of ϕ values different from the rank order of $(u-v)$ values. Further, it could happen that the variation in w was fortuitously small, so that ϕ values were ordered the same as $(u-v)$ values, even for conditions with moderate differences in $(u-v)$.

Since the q manipulation of the first experiment is an intuitively appealing contingency variable that does not appear in the ϕ formula, it would strengthen the case for the ϕ framework if it could subsume the first experiment. It is easy to show that it can. For the most general q manipulation, let the probability of the reinforcer on a contingent trial be u' given a criterion response and v' given a noncriterion response, and let the probability of the reinforcer on a

noncontingent trial be h for both types of response. Then the probability of the reinforcer given a criterion response is $u = u'q + h(1-q)$, and the probability of the reinforcer given a noncriterion response is $v = v'q + h(1-q)$. It is reasonable to combine contingent and noncontingent trials to get these conditional probabilities, since the two types of trial are not signalled, and hence, cannot be distinguished by the subject. What the subject actually experiences are the two conditional probabilities u and v . Substituting these expressions for u and v into the ϕ formula gives:

$$\phi = q(u' - v') \left(\frac{w(1-w)}{z(1-z)} \right)^{1/2}$$

where $z = (u'q + h(1-q))w + (v'q + h(1-q))(1-w)$.

Thus the ϕ framework readily accomodates the q -manipulation. The first experiment is a very simple special case of the q -manipulation. The probability of the reinforcer on noncontingent trials was set equal to that on contingent trials, so $h = u'w + v'(1-w)$. This makes $z = h = u'w + v'(1-w)$. In addition, in the first experiment $u'=1$ and $v'=0$, so $z=w$ and the square root terms cancel, leaving

$$\phi = q(1-0)=q$$

This means that in manipulating q in the first experiment, we were in fact manipulating ϕ .

Further, note that solving for q in terms of the other variables gives $q=(u-v)/(u'-v')$; and so $q=(u-v)$ in the first experiment. Thus, we were manipulating ϕ in the first experiment through variations in $(u-v)$ alone. The values of u and v are easily calculated using the formulas $u = u'q + h(1-q)$ and $v = v'q + h(1-q)$. The ϕ framework parameters for the first experiment are shown in Table 1.

This reinterpretation of the first experiment reveals it to be a traditional manipulation of $(u-v)$ with the important difference that the probability of a criterion response was controlled: $w=.2$. The results of the first experiment show that the higher was the value of ϕ , the more effective was the shaping to the target location. Asymptotic conditioning level (TC_{∞}) was related to q by a negatively accelerated function like $TC_{\infty} = (q)^{1/2}$ and it has been shown that $\phi=q$. Therefore, a simple negatively accelerated function like $TC_{\infty} = (\phi)^{1/2}$ describes the relation between asymptotic conditioning and ϕ . So in retrospect, the first experiment provided evidence for the usefulness of ϕ as a measure of an effective response location-reinforcer contingency variable. Since ϕ was manipulated through variations in $(u-v)$ alone in the first experiment, the next step was to provide additional evidence for the usefulness of the ϕ framework.

Chapter 4

THE SECOND EXPERIMENT

The results of the first experiment were consistent with the predictions of the ϕ framework: the larger the value of ϕ as determined by $(u-v)$, the larger was the measure of shaping at asymptote. The second experiment was designed to test two predictions derived from the ϕ framework. The first prediction is that when the probability of the reinforcer is 1 given a criterion response and 0 given a noncriterion response, the probability of a criterion response will have no effect on the asymptotic conditioning level. In the notation of the ϕ framework, the prediction is that when $u = 1$ and $v = 0$, the value of w is not important. The reason is that when the values $u = 1$ and $v = 0$ are substituted into the ϕ formula, the result is $\phi = 1$ regardless of the value of w . Other work with percentile schedules (Alleman and Platt, 1973; Yap, 1979) has suggested that w may exert a strong influence on shaping, but the ϕ framework implies that w will have no effect on shaping so long as $u = 1$ and $v = 0$. To allow a test of this prediction, the present experiment included two groups with $u = 1$ and $v = 0$ but different values of w . These are Groups I and III in Table 2.

A second prediction of the ϕ framework is that a group which is exposed to intermittent reinforcement of the criterion response ($u < 1$) will not be shaped as effectively as a group which is exposed to continuous reinforcement of the criterion response ($u = 1$). This is because ϕ decreases when u is reduced.

The present experiment had a secondary purpose. In the ϕ framework, the overall probability of the reinforcer is $z = uw + v(1-w)$.

Table 2

Parameter Values for the Three Groups in the Second Experiment

	Group		
	I	II	III
u	1.0	.4	1.0
w	.2	.5	.5
z	.2	.2	.5
φ	1.0	.5	1.0

$v=0$ for all groups

The value of z is determined once u , v , and w are specified, and z is merely a derivative variable in the ϕ framework. However, since z is an important parameter in other frameworks, the second experiment was designed to assess the importance of z relative to ϕ . Referring to Table 2, note that Groups I and II have the same value of z but different values of ϕ , while Groups I and III have different values of z but the same value of ϕ .

METHOD

Subjects

Twenty-four male Wistar rats obtained from Canadian Breeding Farms served as subjects. Their weights at the beginning of the experiment ranged from 230 to 280 gm. As in the first experiment, the rats were maintained at 85% of their free feeding weights.

Apparatus

The apparatus for this experiment was identical to that used in the first experiment.

Procedure

As in the first experiment, the rats were trained to operate the joystick, and were then exposed to two sessions on a FR 1 schedule and one session on a RR 2 schedule. Next, a RR 5 schedule was in effect for two sessions. A daily session ended when a rat had made 128 responses or 1 hour had elapsed, whichever came first. The experiment was conducted 6 days a week.

Subjects were assigned randomly to 3 groups. The target location was 270^0 and a targeted percentile schedule identified criterion responses. The ϕ framework parameters of the 3 groups are presented in Table 2. Noncriterion responses were never followed by food delivery, so $v = 0$ for all groups. The probability of the reinforcer given a criterion response (u), and the probability of a criterion response (w), were manipulated between groups. This phase of the experiment was 34 sessions long. Following this phase, Group II was exposed to a perfect contingency ($\phi = 1$) for 22 sessions. The only change in the parameters of this group was the increase in the value of u from .4 to 1.

RESULTS

The group learning curves are presented in Figure 8. All three learning curves were negatively accelerated and reached their asymptotic level in about 28 sessions. The results were entirely consonant with the ϕ framework: Groups I and III reached the same asymptotic level, while Group II was below this level. The asymptote is defined as the mean target component over the last five sessions. In the first and third experiments reported in this thesis, only the functional relation between ϕ and group asymptote was investigated. No predictions about differences between particular group asymptotes were made. However, the present experiment made predictions about differences between group means, and so inferential statistical analysis was necessary. A groups by sessions analysis of variance of target components is shown in Table 3. All three groups showed an increase in target component over sessions; $F(33,693) = 85.86$, $p < .001$; and there was a significant difference among the groups; $F(2,21) = 12.76$, $p < .001$. The mean of the

FIGURE 8. Group learning curves for the second experiment




Table 3

Analysis of variance of target components for the second experiment.

	<u>SS</u>	<u>df</u>	<u>MS</u>	<u>F</u>	
<u>Between Subjects</u>					
Groups	73098.33	2	36549.16	12.76	p<.001
Subjects within Groups	60144.62	21	2864.03		
<u>Within Subjects</u>					
Sessions	256920.10	33	7785.46	85.86	p<.001
Groups x Sessions	10025.26	66	151.90	1.68	p<.001
<u>Sessions x Subjects</u>					
within groups	62839.26	693	90.68		

(Note that all data points were multiplied by 100 prior to the analysis).

34 target components was virtually identical for Group I and Group III-- .753 and .757 respectively--but was only .554 for Group II. Thus, the main effect of groups reflects the consistently less effective shaping of the subjects in Group II. The learning curve of Group II diverges from the other two, and this is reflected in the significant interaction of Groups and Sessions; $F(66,693) = 1.68, p < .001$.

Table 4 shows the asymptotes, defined as the mean of the last five target components, and the results of a Tukey HSD test. There was no difference between the asymptotes of the two $\phi = 1$ groups, while the asymptote of the $\phi = .5$ group was significantly lower than the other two asymptotes.

When the probability of the reinforcer given a criterion response (u) for Group II was increased to 1 so that subjects were exposed to a perfect contingency, the average target component rose to .91 on the last of the 22 sessions. This was as high as the asymptotes of the other two groups in the preceding phase.

DISCUSSION

The ϕ framework predicts that the probability of a criterion response (w) will have no effect on shaping when $u = 1$ and $v = 0$; that is, when the probability of the reinforcer is 1 given a criterion response and 0 given a noncriterion response. This prediction was confirmed by the finding that Group I and Group III did not differ, since they both had $u = 1$ and $v = 0$ but different values of w (.2 and .5 respectively). The second prediction, that a group with $u < 1$ will not be shaped to the target as effectively as will a group with $u = 1$, was confirmed by the result that the asymptote of Group II was lower than that of Group III. The only difference between these two groups was

Table 4

Tukey's HSD test of differences in asymptotic target components for the second experiment.

	<u>Group</u>		
	I	II	III
Asymptotic target component (sessions 30-34 inclusive)	0.912	0.736	0.874

Denominator of ratio used in Tukey's HSD test:

$$D = (MS_{\text{subj. within groups}}/n \times j)^{1/2} = (290.19/8 \times 5)^{1/2} = 2.693$$

where n = number of subjects in each group

j = number of sessions on which each group mean is based

<u>Comparison</u>	<u>Tukey ratio</u>
I and III	1.41
I and II	6.54
III and II	5.12

With 3 and 21 degrees of freedom, the critical value of the Tukey ratio is 3.57 at the .05 level of significance and 4.62 at the .01 level. So the first difference is not significant, while the other two are significant at the .01 level.

(Note that all data points were multiplied by 100 prior to the analysis).

that Group II had a value of u less than one. The average target component of Group II increased when the value of u was increased from .4 to 1, so u is an important variable in this differentiation paradigm. In their search for a measure of contingency, Gibbon, Berryman, and Thompson (1974) eliminated all their candidates except for ϕ on the grounds that none of them could account for the partial reinforcement effect in classical conditioning. In order to favor choice of ϕ as a measure of contingency for operant conditioning, it is important to show that a partial reinforcement effect obtains in an operant paradigm. The present result that reducing u impairs shaping to the target indicates that there is a partial reinforcement effect in this spatial differentiation paradigm, and so the other measures of association considered by Gibbon, Berryman, and Thompson (1974) could not be used.

Although confirmation of the first prediction was surprising, it does not actually conflict with results from Alleman and Platt (1973) or Yap (1979), since these studies did not include groups with $u = 1$ and $v = 0$ but different values of w . Rather, they concluded that the probability of a criterion response, w , was always an important variable because variations in w had large effects in their studies. Confirmation of the second prediction does conflict with the assumption made by Brandon (1969) and Yap (1979) that intermittent reinforcement of the criterion response does not hinder shaping. Yap (1979) confounded changes in w with changes in u in order to keep the overall probability of the reinforcer (z) the same for all groups. Her conclusion that group differences were the result of differences in w alone is challenged by the present finding that changes in u alone can have a large effect on shaping.

Groups I and II had the same value of z , the overall probability of reinforcer delivery, yet their performance was different. Groups I and III had different values of z , yet their performance was the same. This finding is consistent with the role of z as a derivative variable in the ϕ framework, but is inconsistent with frameworks which give this variable a primary role in response differentiation. It also suggests that Yap was misguided in her efforts to keep z constant by confounding changes in w with changes in u . The present results suggest that it would be better to investigate the effects of w by keeping u constant and letting z covary with w . The failure of z to account for the present findings is to be contrasted with the success of ϕ in providing a complete account of the variance between the group asymptotes.

The ϕ framework implies that groups with the same value of ϕ will exhibit equally effective shaping, even if they do not have identical values of all three parameters which determine ϕ . This isocontingency prediction received some support from the present results: Groups I and III had identical ϕ values but did not have identical values of all three component parameters, yet their performance was the same.

The experimental evidence to this point strongly suggested that ϕ was a useful measure of an effective response location-reinforcer contingency variable. The ϕ framework made correct ordinal predictions about asymptotic conditioning levels when ϕ was manipulated through variations in $(u-v)$ alone, and through variations in both u and w . The only parameter of the ϕ framework that had not been manipulated in either the first or the second experiment was v , the probability of the reinforcer given a noncriterion response. The purpose of the third

experiment was to determine whether or not the ϕ framework made correct predictions about the relation between asymptotic performance and ϕ when ϕ was manipulated through variations in v alone.

Chapter 5

THE THIRD EXPERIMENT

The value of ϕ decreases as v , the probability of the reinforcer given a noncriterion response, increases. Hence, the ϕ framework implies that the effectiveness of shaping to a target will decrease with increasing v . The third experiment employed four groups differing only in the value of v to test this implication. For all groups, a criterion response was designated as one of the 20% closest to the 270^0 target ($w = .2$) and all criterion responses were followed by food delivery ($u = 1$). Table 5 presents the four values of v and the resulting ϕ values.

METHOD

Subjects

Twenty-four male Wistar rats obtained from Canadian Breeding Farms served as subjects. Their weights at the beginning of the experiment ranged from 203 to 319 gm. The rats were maintained at 85% of their free feeding weights.

Apparatus

The apparatus for this experiment was identical to that used in the first and second experiments.

Procedure

As in the first two experiments, the rats were trained to operate the joystick, and were then exposed to two sessions on a FR 1 schedule and one session on a RR 2 schedule. This was followed by three sessions on a RR 5 schedule. Subjects were then assigned randomly to four groups for the differential reinforcement phase in which a targeted

Table 5

Parameter Values for the Four Groups in the Third Experiment

v	.10	.20	.40	.60
ϕ	.80	.67	.48	.34

$w = .2$ and $u = 1$ for all groups.


percentile schedule identified criterion responses. The parameter values of the four groups are shown in Table 5; these conditions were in effect for 50 sessions. Following this phase, the v value of the $v = .60$ group was changed to $.10$, so that this group had the same parameter values as the $v = .10$ group in the preceding phase. This change in the value of v increased the value of ϕ from $.34$ to $.80$. This phase lasted for 18 sessions, and was intended to show that the poor performance of subjects when $v = .60$ did not mean that they could not be shaped to a target, but was the result of the low ϕ value.

RESULTS

The group learning curves are presented in Figure 9. The learning curves were negatively accelerated, and all groups reached their asymptotic level in about 30 sessions. At the 30th session it appeared that some groups might be on the increase, so subjects were run for 50 sessions to ensure that the asymptotes were stable. There were no systematic increases after the 30th session. The asymptote is defined as the mean target component over the last five sessions. Group asymptotes are shown in Table 6. The top two groups, $v = .10$ and $v = .20$, did not differ significantly ($t_{10} = 1.25$, $p > .10$). Shaping was less effective in the $v = .40$ group than in the top two groups, and was less effective in the $v = .60$ group than in the $v = .40$ group.

During the phase in which the $v = .60$ group was exposed to the new value $v = .10$, one of the subjects was in very poor health. Its fur was much duller than usual and its weight kept dropping despite an increasing ration of food in the home cage. The obvious poor health of this subject, R-22, led to its data being excluded from the calculation of the group asymptote for the phase in which $v = .10$. Based on the

FIGURE 9. Group learning curves for the third experiment



SESSIONS

Table 6

Group Asymptotes for the Third Experiment

<u>v</u>	<u>Asymptotic Target Component (Sessions 46-50 inclusive)</u>
.10	.80
.20	.85
.40	.61
.60	.37

other five subjects, the group asymptote was .78. This was as high as the asymptote of the $v = .10$ group in the preceding phase.

DISCUSSION

The insignificant reversal in the asymptotes of the top two groups was the only departure from the prediction that the asymptotes would be ordered in the same way as the ϕ values. The lack of difference between the $v = .10$ and $v = .20$ groups is not too surprising when it is considered that neither the v values nor the corresponding ϕ values (.80 and .67 respectively) for these two groups were very different, and that there were only six subjects in each group. It seems fair to conclude that the results provide support for the ϕ framework.

It should be mentioned that there was one anomalous subject in the experiment: R-12 in the $v = .60$ group. R-12 was the only subject that did not show any tendency to respond toward the target. Its responding was highly variable; the only regularity in its performance was a tendency to respond within 60^0 on either side of the foodcup--this occurred on 35 of the 50 sessions. Also, R-12 was the only subject with a negative asymptote (-.24). Among the other 23 subjects, the smallest asymptote was .17. With R-12 excluded, the $v = .60$ group asymptote increased from .37 to .48, but was still well below the $v = .40$ group asymptote of .61.

Chapter 6

GENERAL DISCUSSION

Any experiment designed to test the usefulness of ϕ as a measure of contingency in an operant paradigm must control the probability of a criterion response, since this probability influences the value of ϕ . Percentile schedules are ideally suited to this task. The experiments reported in this thesis employed percentile schedules, and the results strongly suggest that ϕ is a useful measure of an effective contingency variable. In some paradigms, different responses necessarily differ in effort, execution time, or delay of reinforcement, so that the study of response differentiation is complicated by the presence of a counter-vailing process. For example, lever presses of different durations cannot involve the same effort and execution time, so duration differentiation may be affected by a minimizing process (see Platt, 1979). Differentiation of response location in the joystick paradigm is relatively uncontaminated by such opposing processes. While other paradigms may not be as well designed for the study of shaping by response-reinforcer contingency as the joystick paradigm, the generality of the present findings should be explored, since the ϕ framework can be applied to any dimension on which response values can be rank ordered.

An incidental advantage of the joystick paradigm is that negative contingencies and positive contingencies only differ semantically, so that negative contingencies do not have to be studied in a parallel series of experiments. If a criterion response is defined as one of the 100w% closest to a target A, then a negative contingency

can be arranged by reinforcing criterion responses with probability $u=a$, noncriterion responses with probability $v=b$, and choosing $a < b$. Now consider a condition in which a target, B, is 180° away from target A, and criterion responses are defined as the $100(1-w)\%$ closest to the target B. If criterion responses are reinforced with probability $u=b$ and noncriterion responses are reinforced with probability $v=a$, this is a positive contingency, because $a < b$. But this condition is indistinguishable from the negative contingency defined in terms of target A. Since for every negative contingency there is a corresponding positive contingency, it is not necessary to map out negative contingency functions in the joystick paradigm.

The procedure used to arrange contingencies in the present experiments only allows good approximations to the nominal contingencies, because the procedure makes u , v , and w random variables with fixed expected values. For example, the probability of reinforcer delivery following a criterion response is u , but in any particular session the relative frequency of criterion responses which are followed by reinforcer delivery will only approximate u . A little variance in the independent variable is a concomitant of the procedure, but there is no obviously superior procedure. The variance is slight and the results are orderly, so this problem does not appear to be of much practical significance.

One shortcoming of the joystick paradigm is the artificiality of the task. Displacement of the joystick is not easy to shape: it takes from 20 minutes to 1-1/2 hours. Learning to push the stick in a particular direction does not seem to make contact with a rat's natural repertoire. It may be argued that the learning observed in this

artificial situation has little to do with learning contingencies in the natural environment. The standard response to this assertion is that rats exhibit great behavioral plasticity and that the best way to study fundamental learning processes is in a highly controlled laboratory situation, however artificial the setup may appear. Another shortcoming of the joystick paradigm is that some subjects show a weak bias to respond toward the foodcup. The bias is always weak enough to be overridden by a strong response location-reinforcer contingency, but it can interfere with weak contingencies, as shown by the responding of R-12 in the third experiment.


All conditions in the experiments reported in this thesis had overall reinforcer probabilities of at least .20. The reason for this is that pilot work had shown that rats cease responding when the overall reinforcement probability is below about .15. Hence, a strong positive contingency cannot produce good shaping if the overall reinforcer probability is much below .20, so that the ϕ framework prediction of good shaping will not be confirmed. But it is not unusual for many excellent predictive frameworks to break down when an incidental variable attains an extreme value, and the ϕ framework makes correct predictions about asymptotic conditioning for a wide range of parameter values. Perhaps this problem could be circumvented by using subjects of a species that will tolerate lower overall reinforcement probabilities.

In the experiments reported here, the 2x2 table of each condition was determined by specifying the probability of a criterion response, w , and the two conditional probabilities $P(Z|W)$ and $P(Z|\bar{W})$. This is the only sensible way to use the three degrees of freedom in the

table, because there is a fundamental asymmetry in the operant conditioning procedure to which the table has reference. Effective shaping can be produced by controlling the occurrence of a reinforcer following a criterion response, but not by controlling the occurrence of a criterion response following a reinforcer. Referring to Figure 7, it can be seen that the present experiments determined the 2x2 table by controlling w and the two conditional probabilities u and v . It might appear that another way to determine the table would be to control z (the overall probability of the reinforcer) and two other conditional probabilities, for example $P(W|Z)$ and $P(W|\bar{Z})$. But controlling these conditional probabilities is a very different matter from controlling $P(Z|W)$ and $P(Z|\bar{W})$. For example, $P(Z|W)$ can be controlled because the experimenter can ensure that a reinforcer is delivered following a criterion response, but $P(W|Z)$ cannot be controlled because the experimenter cannot ensure that a criterion response will occur following a reinforcer delivery.

The ϕ framework made correct predictions in all three experiments reported in this thesis. If ϕ is a generally useful measure of contingency, then there should be an orderly relation between asymptotic conditioning and ϕ across the entire range of ϕ values, regardless of how these values were produced. One way to address this issue is to determine the relation between group asymptote and ϕ across all three experiments. This relation is presented in Figure 10. The bracketed points are group asymptotes calculated with the deviant subjects' data included; these points were not used in the calculation

FIGURE 10. Group asymptotes as a function of ϕ for all groups in the three experiments. The curve of the function $TC_{\infty} = \phi^{1/2}$ is shown



of ω^2 . A negatively accelerated function of ϕ could account for most of the variance in group asymptotes; in fact, the simple function $TC_{\infty} = \phi^{1/2}$ accounts for 93% of the variance in group asymptotes. Clearly, this is not the best fitting function; however, the important point is that the relation between group asymptote and ϕ is orderly. The negatively accelerated relation between asymptotic target component and ϕ is not conclusive evidence for a fundamentally nonlinear relation between concentration of responding about the target and ϕ : it may reflect a negatively accelerated relation between concentration and the measure of concentration, target component. This is suggested by the observation that a moderate degree of concentration of responding about the target will produce a target component of about .90, and a very high degree of concentration can produce a target component no more than .10 higher than this. Perhaps another measure would reflect concentration of responding in a more linear fashion; and hence, produce a linear relation between concentration and ϕ .

The very orderly relation between asymptotic conditioning and ϕ depicted in Figure 10 provides strong evidence that ϕ is a suitable measure of contingency in the joystick paradigm. Additional evidence might be produced by testing the isocontingency property that all groups with the same ϕ value will exhibit the same asymptotic conditioning, regardless of how that value was generated. Unfortunately, it is difficult in practice to perform a convincing experimental test of isocontingency. The major problem is that an attempt to show that several groups with the same ϕ value do not differ in asymptotic conditioning is merely a test of the null hypothesis. A better approach to testing isocontingency would start with at least two sets of groups. All groups within the same set would have the same ϕ value, but no two

groups in the same set would have the same value of any of the parameters, u , v , and w . Each group in a set would have a counterpart in the other set that differed in the value of only one of the three parameters, but of course had a different ϕ value. The isocontingency property would be supported by finding that the asymptotic conditioning of each group was closer to that of the other groups in the same set than to that of its counterpart in the other set. A drawback to this type of experiment is that a large difference in the asymptotic conditioning of two groups in the same set could not be explained as the effect of a difference in the value of any one parameter. Certainly there are other ways to test the ϕ framework that are superior to tests of the isocontingency property because they will provide useful information about the effects of u , v , and w even if the framework is not supported.

The ϕ framework makes predictions about the final level of acquisition, asymptotic performance, but it does not address the issue of the conditions necessary for maintenance of asymptotic performance. Once a group has reached asymptote, it is quite possible that the contingency could be lowered without changing the performance. Davis and Platt (1978) found that some subjects maintained the asymptote attained on a perfect contingency for 30 days after a zero contingency was implemented.

It has been argued that ϕ is the most suitable measure of contingency for operant conditioning, but the difference in conditional probabilities, $(u-v)$, deserves some attention because it has been used in operant conditioning with some success. The present data allow a post hoc comparison of ϕ and $(u-v)$. All three experiments had a group

with $(u-v) = .4$; the resulting mean asymptotic target components were .58, .73, and .48 with standard errors of .09, .04, and .15 respectively. Certainly the group with an asymptote of .73 is different from the other two groups with the same $(u-v)$ value. By contrast, the corresponding ϕ values are .4, .5, and .34; so the asymptotes are perfectly ordered by ϕ . The first and third experiments each had a group with $(u-v) = .6$; the asymptotes were .82 and .61 with standard errors of .04 and .08 respectively. Again the asymptotes are quite different despite identical $(u-v)$ values, and again the corresponding ϕ values, .6 and .48, account for the difference in the asymptotes.

Groups with the same ϕ value but different $(u-v)$ values can be compared similarly. The second and third experiments each had a group with $\phi = .5$ (actually, .5 and .48). The asymptotes were .73 and .61, a difference not much larger than the standard error of the second group, .08. Further, the $(u-v)$ values of the two groups are .4 and .6 respectively, so that $(u-v)$ is much farther from the mark than is ϕ . The same pattern holds for the $\phi = .8$ groups from the first and third experiments. The asymptotes were .87 and .80, and the standard errors were both .03. While the difference in asymptotes is large relative to the standard errors, the result is more damaging to $(u-v)$ than to ϕ , since the $(u-v)$ values of .8 and .9 are in the wrong order to account for the difference in the asymptotes.

While the post hoc comparison of $(u-v)$ and ϕ strongly favors ϕ as the measure of contingency, it is desirable to test their relative merits experimentally. This can be done by comparing two groups with the same values of u and v but with the values of w chosen so as to produce very different values of ϕ . Only the ϕ framework identifies w

as an important variable so as to predict that the two groups will have different asymptotes.

When the data from all three experiments is examined, it is clear that there is a tendency for variability of individual subject asymptotes about their group mean to decrease with increasing ϕ . Table 7 shows all the ϕ values and the corresponding standard errors of the mean asymptote. Note that the deviant subjects have been excluded from the $\phi = 0$ and the $\phi = .34$ groups. A manipulation of response location-reinforcer contingency is in a sense a manipulation of the control over response location by the degree to which food delivery depends on response location. The weaker the contingency, the less stringent the control of response location by food delivery. It is reasonable to expect that the weaker the control over response location, the less similar will be the performance of individual subjects in a group, and hence, the larger will be the standard error of the group asymptote. One implication of this negative correlation between ϕ and standard error is that the model underlying the analysis of variance is not appropriate for the data obtained in the present experiments. The homogeneity of variance assumption is violated by the correlation between the independent variable and within-group variance. In practice, analysis of variance of target components may be used when the groups being compared have ϕ values in a range over which the standard errors are of comparable size. But generally, it is not wise to ignore the implicit relation between control and within-group variance.

The ϕ framework provides a very cogent account of the behavior observed in the experiments reported here, but the framework is not a model of behavior. It is simply a model of an independent variable. It

Table 7

Standard errors of the mean asymptote for all groups
in the three experiments.

ϕ	<u>Standard error</u> <u>of the mean asymptote</u>
0	.19
.2	.06
.34	.15
.4	.09
.48	.08
.5	.04
.6	.04
.67	.02
.8	.03
.8	.03
1.0	.02
1.0	.01
1.0	.02

specifies a set of significant parameters and a formula which may be used to systematize experimental procedures. Such a model is useful to the extent that it exposes orderly relationships between behavior and experimental parameters. The framework serves to guide research and theory, and its usefulness is to be judged by its success in organizing procedures and results, relative to competing models.

The major shortcoming of the ϕ framework is that it does not suggest what basic processes are responsible for the simple relationship between shaping and ϕ . It is implausible that rats are doing anything like computing ϕ values over sessions. The ϕ value provides a good account of behavior at the level of an entire session, but it remains to be determined what trial-by-trial mechanisms generate this order. It may well be that a trial-by-trial contiguity mechanism produces the regularities so well described by the contingency variable ϕ . Since ϕ is a function of the four joint probabilities in the 2x2 table, perhaps ϕ simply expresses the long term effects of the subject's exposure to the trial events which give rise to the joint probabilities. For example, the joint probability of a criterion response and food delivery is the relative frequency of trials on which a criterion response occurred and was followed by food. By simple temporal contiguity, such trial events could increase the relative frequency of responses at the same location. Similarly, the other three joint events could influence the distribution of response locations. The four joint probabilities indicate the relative frequencies of the four joint events, and ϕ combines these probabilities in a manner which accurately predicts the long term effects of exposure to the four joint events in those proportions.

Another relation at the level of the session that could be the result of a trial-by-trial contiguity mechanism is the correlation-based law of effect (Baum, 1973), which states that "...behavior increases in frequency if the increase is correlated with an increase in rate of reinforcement...". Baum (1973) noted that the basis for the correlational law of effect could be a trial-by-trial contiguity mechanism, but that focussing on correlation over a session was to be preferred because the "...concept of correlation has the additional advantage that it draws together apparently diverse procedures into a single conceptual framework." This comment applies with equal force to the ϕ framework. The correlation-based law of effect and the ϕ framework both identify correlations over a session between behavior and reinforcement as critical to understanding behavioral change, but the similarity of the two models is merely superficial, for they point to correlations between different entities.

The relative waiting time hypothesis (Jenkins, Barnes & Barrera, 1981) has been found to provide a good account of autoshaping results in situations where CS-US contingency cannot be invoked to explain the results. But the contingency was indexed by the difference between $P(US|CS)$ and $P(US|\overline{CS})$, and it may be suspected that a contingency account would be successful if contingency were measured by ϕ instead of the difference between the two conditional probabilities. Unfortunately, while the relative waiting time hypothesis is stated in terms of time, the ϕ framework is stated in terms of probabilities; and the autoshaping paradigm has temporal properties for which the ϕ framework is not suited. The artifice of constructing probabilities

from ratios of times is certainly questionable, but the only way to come up with a $P(US)$ in the autoshaping paradigm is to divide the number of times food was delivered by the total time in a session. The construction of a probability in this way is peculiar, and mixing probabilities derived in fundamentally different ways in the same model is factitious. The inevitable conclusion is that the ϕ framework cannot be applied to the autoshaping paradigm because the framework is based on a correlation between punctate events and the autoshaping paradigm cannot be reduced to punctate events. Whether the ϕ framework will be successful in classical conditioning situations which can be reduced to relations among punctate events awaits further experimentation. Gibbon, Berryman and Thompson (1974) have discussed some improvements in the way trials and probabilities are specified that must be made to permit the ϕ framework to be tested in classical conditioning.

There are no studies using human subjects that examine response differentiation as a function of response-reinforcer contingency, but there is a small literature on the judgment of contingency. In the typical experiment, subjects are presented with 2x2 tables, or with samples from bivariate distributions, and are then asked to estimate the strength of association between the two variables by choosing a number between 0 (no association) and 100 (perfect association). Nisbett and Ross (1980) reviewed studies which showed that when subjects have preconceived notions about the relationship between the variables, judgment of covariation reflected "...true covariation far less than it reflected theories or preconceptions of the nature of the associations that 'ought' to exist." Results from studies using relatively neutral stimulus materials have been varied. Jennings, Amabile, and Ross (1980)

asked subjects to judge the correlation between two continuous variables. In one condition, for example, the stimulus materials pictured men holding walking sticks. For each set of 10 pictures, the subject was required to estimate the strength of association between the man's height and the length of his walking stick. Jennings, Amabile, and Ross (1980) found that the function relating subjective covariation estimate to objective correlation was positively accelerated. This finding may be compared to the negatively accelerated function relating asymptotic target components to ϕ in the experiments reported in this thesis, but the many differences between the studies vitiate interpretations of the difference between the functions. The control of ad hoc verbal responses by contingency may have little to do with the control of responses actually involved in the contingency. The motivation to detect a contingency must be very different in the two paradigms, and it is moot whether or not animals detect a contingency even when their responding is systematically related to contingency. Further, the independent variable in the Jennings, Amabile, and Ross (1980) study was not the correlation between two dichotomous variables (ϕ), but the correlation between two continuous variables (r). Estimation of the magnitude of each variable is necessary for judging correlation when the variables are continuous, but not when they are dichotomous. Finally, even if the studies had used the same independent variable, different forms of the function relating dependent to independent variable could reflect little more than different properties of the dependent variables used in the two kinds of study.

Smedslund (1963) and Jenkins and Ward (1965) found a lack of correspondence between judgments and the contingency between dichotomous

variables, both when the information was presented serially and in the form of a 2x2 table. Jenkins and Ward (1965) used a two response-two outcome task, and subjects were asked to judge the degree of control their choice of response had over the outcomes. Denoting the response alternatives by R_1 and R_2 , the actual contingency was indexed by the difference between the probability of the positive outcome given R_1 and the probability of the positive outcome given R_2 . This difference in conditional probabilities is just $(u-v)$ in the notation of the ϕ framework. Jenkins and Ward (1965) found little correspondence between judgments and actual contingency, but noted that judged control increased with the frequency of the positive outcome. They suggested that when there are two active response alternatives, the systematic error in judgments of control may arise because subjects tend to assume that no positive outcomes would occur if no response was made.

Allan and Jenkins (1980) tested this hypothesis by including a condition in which one alternative was an active response and the other alternative was to make no response. Subjects in this condition and the conditions with two active response alternatives were presented with five contingent and five noncontingent problems. The hypothesis about the source of the systematic error was supported by the results from the noncontingent problems, but not by the results from the contingent problems. For noncontingent problems, judged influence or control increased with the probability of the positive outcome in the conditions with two active responses, but not in the condition where one alternative was to make no response. But for contingent problems, there was a tendency for judgments of influence or control to increase with probability of the positive outcome in the conditions with two active

responses and in the condition with only one active response. If judgment of the influence the response exerts over the outcome is taken as a judgment of the contingency between the response and the outcome, then it is reasonable to conclude that Allan and Jenkins' (1980) subjects are not good judges of the contingency as indexed by $(u-v)$. For example, in the condition with only one active response, the judgment was as high for a problem with $(u-v) = .2$ as for two problems with $(u-v) = .4$, and was higher than for another problem with $(u-v) = .2$ but different values of u and v .

It is unlikely that the results of the human studies on judgment of contingency have much bearing on the results from animal studies. In the first place, ad hoc verbal reports are not necessarily consistent with responding entering the contingency, and response differentiation is not measured in the human studies. It is not even clear that there is a reinforcer capable of producing response differentiation in the human studies. The lack of a systematic relation between judgments of influence and actual contingency may reflect inappropriate measures for one or both variables. Possible drawbacks to the use of $(u-v)$ as a measure of contingency have been pointed out repeatedly in this thesis. Whatever the case may be, there are too many significant differences between the human and animal contingency paradigms to permit a sensible integration of the results.

REFERENCES

- Allan, L.G. & Jenkins, H.M. The judgment of contingency and the nature of the response alternatives. Canadian Journal of Psychology, 1980, 34, 1-11.
- Alleman, H.D. & Platt, J.R. Differential reinforcement of interresponse times with controlled probability of reinforcement per response. Learning and Motivation, 1973, 4, 40-73.
- Antonitis, J.J. Response variability in the white rat during conditioning, extinction, and reconditioning. Journal of Experimental Psychology, 1951, 42, 273-281.
- Batchelet, E. Statistical methods for the analysis of problems in animal orientation and certain biological rhythms. American Institute of Biological Sciences Monograph, January, 1965.
- Baum, W.M. The correlation-based law of effect. Journal of the Experimental Analysis of Behavior, 1973, 20, 137-153.
- Bishop, Y., Fienberg, S. & Holland, P. Discrete multivariate analysis: Theory and practice. Cambridge, Mass.: The MIT Press, 1975.
- Brandon, P.K. Effects of error-contingent time-out on counting behavior in rats. Unpublished doctoral dissertation, University of Michigan, 1969.
- Conover, W.J. Practical nonparametric statistics. Toronto: John Wiley & Sons, 1980.
- Costner, H.L. Criteria for measures of association. American Sociological Review, 1965, 30, 341-353.
- Cramér, H. Mathematical methods of statistics. Princeton: Princeton University Press, 1971.

- Davis, E.R. & Platt, J.R. Response shaping on a continuous spatial dimension. Psychonomic Society, 1978, San Antonio, Texas.
- Eckerman, D.A. & Lanson, R.N. Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. Journal of the Experimental Analysis of Behavior, 1969, 12, 73-80.
- Edwards, A.W.F. The measure of association in a 2x2 table. Journal of the Royal Statistical Society, 1963, A, 126, 109-114.
- Flye, B.L. & Gibbon, J. Partial avoidance contingencies: Absolute omission and punishment probabilities. Journal of the Experimental Analysis of Behavior, 1979, 31, 351-371.
- Gibbon, J. Discriminated punishment: Avoidable and unavoidable shock. Journal of the Experimental Analysis of Behavior, 1967, 10, 451-460.
- Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. Psychological Review, 1977, 84, 279-325.
- Gibbon, J. The contingency problem in autoshaping. In C.M. Locurto, H.S. Terrace & J. Gibbon (Eds.), Autoshaping and conditioning theory. New York: Academic Press, 1981.
- Gibbon, J. & Balsam, P. Spreading association in time. In C.M. Locurto, H.S. Terrace, & J. Gibbon (Eds.), Autoshaping and conditioning theory. New York: Academic Press, 1981.
- Gibbon, J., Berryman, R. & Thompson, R.L. Contingency spaces and measures in classical and instrumental conditioning. Journal of the Experimental Analysis of Behavior, 1974, 21, 585-605.
- Goodman, L.A. & Kruskal, W.H. Measures of association for cross classifications. New York: Springer-Verlag, 1979.

- Hammond, L.J. The effect of contingency upon the appetitive conditioning of free-operant behavior. Journal of the Experimental Analysis of Behavior, 1980, 34, 297-304.
- Herrnstein, R.J. Stereotypy and intermittent reinforcement. Science, 1961, 133, 2067-2069.
- Herrnstein, R.J. Superstition: A corollary of the principles of operant conditioning. In W.K. Honig (Ed.), Operant behavior: Areas of research and application. New York: Appleton-Century-Crofts, 1966.
- Jenkins, H.M., Barnes, R.A. & Barrera, F.J. Why autoshaping depends on trial spacing. In C.M. Locurto, H.S. Terrace & J. Gibbon (Eds.), Autoshaping and conditioning theory. New York: Academic Press, 1981.
- Jenkins, H.M. & Ward, W.C. Judgment of contingency between responses and outcomes. Psychological Monographs, 1965, 79, No. 1 (Whole No. 594).
- Jennings, D., Amabile, T.M. & Ross, L. Informal covariation assessment: Data-based vs. theory-based judgments. In A. Tversky, D. Kahneman & P. Slovic (Eds.), Judgment about uncertainty: Heuristics and biases. New York: Cambridge University Press, 1980.
- Kadden, R.M., Schoenfeld, W.N. & Snapper, A.G. Aversive schedules with independent probabilities of reinforcement for responding and not responding by rhesus monkeys: II. Without signal. Journal of Comparative and Physiological Psychology, 1974, 87, 1189-1197.
- Kendall, M.G. & Stuart, A. The advanced theory of statistics, (Vol. 2), Inference and relationship. London: Charles Griffin & Co., 1973.

- Kop, P.F.M. Aversive schedules with independent probabilities of reinforcement for responding and not responding by rhesus monkeys: I. With signal. Journal of Comparative and Physiological Psychology, 1974, 87, 1180-1188.
- Neffinger, G.G. & Gibbon, J. Partial avoidance contingencies. Journal of the Experimental Analysis of Behavior, 1975, 23, 437-450.
- Nisbett, R. & Ross, L. Human inference: Strategies and shortcomings of social judgment. New Jersey: Prentice-Hall Inc., 1980.
- Pearson, K. Mathematical contributions to the Theory of Evolution - XIII. On the theory of contingency and its relation to association and normal correlation. Draper's Company Research Memoirs: Biometric Series I, 1904.
- Platt, J.R. Percentile reinforcement: Paradigms for experimental analysis of response shaping. In G.H. Bower (Ed.), The psychology of learning and motivation: Advances in theory and research (Vol. 7). New York: Academic Press, 1973.
- Platt, J.R. Temporal differentiation and the psychophysics of time. In M.D. Zeiler & P. Harzem (Eds.), Advances in analysis of behavior: Reinforcement and the organization of behavior (Vol. 1). New York: Wiley, 1979.
- Rescorla, R.A. Pavlovian conditioning and its proper control procedures. Psychological Review, 1967, 74, 71-80.
- Rescorla, R.A. Probability of shock in the presence and absence of CS in fear conditioning. Journal of Comparative and Physiological Psychology, 1968, 66, 1-5.

- Rescorla, R.A. Conditioned inhibition of fear resulting from negative CS-US contingencies. Journal of Comparative and Physiological Psychology, 1969, 67, 504-509.
- Reynolds, H.T. The analysis of cross classifications. New York: The Free Press, 1977.
- Skinner, B.F. The behavior of organisms. New York: Appleton-Century-Crofts, 1938.
- Skinner, B.F. "Superstition" in the pigeon. Journal of Experimental Psychology, 1948, 38, 168-172.
- Smedslund, J. The concept of correlation in adults. Scandinavian Journal of Psychology, 1963, 4, 165-173.
- Staddon, J.E.R. & Simmelhag, V.L. The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. Psychological Review, 1971, 78, 3-43.
- Stein, L., Sidman, M. & Brady, J.V. Some effects of two temporal variables on conditioned suppression. Journal of the Experimental Analysis of Behavior, 1958, 1, 153-162.
- Thorndike, E.L. Animal intelligence: Experimental studies. New York: Macmillan, 1911.
- Webster, J.B. An experimental analysis of response differentiation. Unpublished doctoral dissertation, McMaster University, 1976.
- Weisberg, H.F. Models of statistical relationship. American Political Science Review, 1974, 68, 1138-1155.
- Yap, P.S. Differential reinforcement of response location in the white rat. Unpublished honors thesis, McMaster University, 1979.

APPENDIX

Asymptotes of individual subjects

Asymptotic target components of individual subjects in the first experiment.

Subjects in the first replication are indicated by an asterisk.

<u>Group</u>	<u>Subject</u>	<u>Asymptote</u>
$\phi = 0$	R-1	-.33
	R-2	.32
	R-3	-.06
	R-4	-.75
$\phi = .2$	R-13*	.42
	R-14*	.56
	R-15*	.54
	R-16*	.37
	R-5	.34
	R-6	.62
	R-7	.15
	R-8	.23
$\phi = .4$	R-9*	.27
	R-11*	.80
	R-12*	.79
	R-9	.83
	R-10	.55
	R-11	.35
	R-12	.48
$\phi = .6$	R-5*	.93
	R-6*	.79
	R-7*	.78
	R-8*	.85
	R-13	.82
	R-14	.88
	R-15	.54
$\phi = .8$	R-16	.93
	R-1*	.91
	R-2*	.73
	R-3*	.81
	R-4*	.79
	R-17	.93
	R-18	.91
	R-19	.88
$\phi = 1$	R-20	.97
	R-21	.96
	R-22	.85
	R-23	.91
	R-24	.94

Asymptotic target components of individual subjects in the second experiment.

<u>Group</u>	<u>Subject</u>	<u>Asymptote</u>
I ($\phi = 1$)	R-1	.88
	R-2	.95
	R-3	.95
	R-4	.92
	R-13	.87
	R-14	.87
	R-15	.91
	R-16	.93
II ($\phi = .5$)	R-5	.88
	R-6	.79
	R-7	.75
	R-8	.82
	R-17	.73
	R-18	.75
	R-19	.65
	R-20	.50
III ($\phi = 1$)	R-9	.90
	R-10	.87
	R-11	.83
	R-12	.84
	R-21	.95
	R-22	.96
	R-23	.83
	R-24	.81

Asymptotic target components of individual subjects in the third experiment.

<u>Group</u>	<u>Subject</u>	<u>Asymptote</u>
v = .10 (ϕ = .80)	R-1	.83
	R-2	.82
	R-3	.89
	R-13	.70
	R-14	.84
	R-15	.73
v = .20 (ϕ = .67)	R-4	.81
	R-5	.92
	R-6	.93
	R-16	.82
	R-17	.83
	R-18	.79
v = .40 (ϕ = .48)	R-7	.37
	R-8	.92
	R-9	.71
	R-19	.51
	R-20	.65
	R-21	.49
v = .60 (ϕ = .34)	R-10	.74
	R-11	.91
	R-12	-.24
	R-22	.17
	R-23	.41
	R-24	.19