JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1989, 51, 145-162 NUMBER 1 (JANUARY)

RESPONSE-REINFORCER CONTINGENCY AND SPATIALLY DEFINED OPERANTS: TESTING AN INVARIANCE PROPERTY OF PHI

GREGORY GALBICKA AND JOHN R. PLATT

McMASTER UNIVERSITY

A chamber containing 72 response keys defining the circumference of a circle 1 m in diameter was used to examine the relation between differentiation of response location and a measure of response-reinforcer contingency known as the phi coefficient. A different target key was specified in each successive phase, and response location was differentiated with respect to the target. Criterional and noncriterional responses (i.e., responses "near" and "far" from the target) were defined using targeted percentile schedules to control the overall probability of each response class. By manipulating criterional (and, hence, noncriterional) response probability and the reinforcement probabilities conditional on each, a mathematical invariance property peculiar to phi in contingency analysis was examined. Specifically, diagonally interchanging cell frequencies in a 2 × 2 table relating criterional/noncriterional responses to reinforcement/nonreinforcement leaves phi unchanged. Hence, the degree of response differentiation predicted by phi remains unchanged under the four permutations implied by the various diagonal interchanges. This predicted invariance was examined under values of phi equal to .33, .58, and .82. Increasing phi generally increased the stereotypy of response location. Three of the permutations generated almost interchangeable performance at different phi values. The remaining permutation, however, generated functions relating response concentration to phi with slopes shallower than those obtained under the other permutations. This resulted from relatively higher levels of differentiation, compared to the other permutations, at low phi values. These data strongly suggest boundary conditions on the ability of phi to reflect completely the local processes that are indexed by phi at a molar level.

Key words: response-reinforcer contingency, response differentiation, shaping, differential reinforcement, percentile schedules, response location, pigeons

The term *contingency* suffers from use in multiple contexts in the analysis of behavior, most particularly in operant conditioning. In a general sense, contingency refers to some relation between a response and a consequence, as in the phrases *response-contingent* or *contingencies of reinforcement.* The latter often substitutes for *schedule of reinforcement,* and as such, changes in schedule parameters need not change the contingency as long as the type of schedule remains the same. The statistical meaning of *contingency,* alternatively, refers to the degree of correlation between two or more events. That is, contingency specifies the degree to which the occurrence of one event predicts the occurrence of a subsequent event. In this usage, behavioral contingency depends not only on the probability of reinforcers following responses, but also on the probability of reinforcement given no response. According to a statistical use of the term *contingency,* changing schedule parameters does indeed change the degree the correlation, as does adding response-independent reinforcement, because both of these will change the frequency of the joint events upon which contingency depends.

The quantitative evaluation of contingency in operant behavior is problematic, as Gibbon, Berryman, and Thompson (1974) and Scott and Platt (1985) discuss in detail. In respondent conditioning, presentation of both the unconditional and the conditional stimuli is under experimental control. Hence, adequate exposure to the joint events that specify the degree of contingency can be ensured. The

---

This work was supported by research grant R03 MH40964 from the National Institute of Mental Health (USA) to G. Galbicka and by Natural Sciences and Engineering Research Council grant A8269 (Canada) to J. R. Platt. The research described herein, and the graphic presentation of the results, would not have been possible without the programming and technical expertise of Gary A. Weatherill. Thanks are extended to T. F. Elsmore for many helpful discussions. Address correspondence and reprint requests to Gregory Galbicka, now at the Department of Medical Neurosciences, Walter Reed Army Institute of Research, Washington, D.C. 20307-5100.

transposition of this tactic to operant behavior is difficult because the frequency of the relevant events (i.e., responses and reinforcers) is generally a dependent, rather than an independent, variable. Under traditional schedules, the experimenter cannot guarantee exposure to all response–reinforcer conjunctions, because the frequency of responding cannot be experimentally manipulated with appropriate degrees of precision.

The analysis of contingency in operant behavior has been curtailed further by the predisposition to oppose responding with nonresponding, and to view contingency as a function of reinforcement for these two response classes. Unfortunately, the two response classes are not strictly opposable. The major constraint on this juxtaposition is the inability to measure nonresponding, except by arbitrarily defining temporal windows and time sampling for responses (e.g., Hammond, 1980). Meaningful comparison of responses and nonresponses requires that they both be measured in the same fashion, but the window defining the presence or absence of one may not necessarily be equal to that for the other. That is, if a response takes x s to execute, x-s intervals can contain at most one response. There is no assurance whatsoever that an equal interval without a response contains exactly one nonresponse.

Herrnstein (1970) suggested that relative response frequency, as a dependent measure, solves the second of these problems. Any response differentiation procedure, in this case one differentiating response location, provides a situation in which two or more response classes can be defined and measured using the same metric. Contingency can be studied by specifying one class as criterional and the other as noncriterional, and then varying reinforcement parameters for each. Such variation results in orderly changes in behavior, which can be described by the generalized matching law (Baum, 1974) or by signal-detection theory (see Davison & Tustin, 1978, for the relation between the two), depending on whether response strengthening or stimulus control is emphasized.

Response differentiation procedures that segregate responses according to a fixed criterion (e.g., left- vs. right-key responses) still hinder the analysis of contingency by their inability to control criterional-response probability. That is, as long as the basis for segregating responses into criterional and noncriterional classes is a fixed absolute value, response differentiation requires, by definition, that the probability of criterional responses varies. This simply repeats that criterional-response probability is a dependent variable used to define the success of traditional differentiation procedures. As such, analyses of operant contingency remain constrained by the limitations of procedure first noted by Gibbon, et al. (1974). To the extent that criterional-response probability under these procedures is unspecifiable a priori, contingency analyses must remain post-hoc enterprises.

This constraint can be removed by arranging a percentile schedule (Alleman & Platt, 1973; Galbicka & Platt, 1986; Kuch & Platt, 1976) to control the probability or rate of criterional responses a priori. Briefly, criterional-response probability is controlled by adjusting the absolute value defining criterional responses such that it always demarcates some percentile of the subject's current response distribution. In this manner, as the absolute values of responses are differentiated, the percentile value for the criterion range remains unchanged, as does the expected probability of a criterional response. Because the experimenter can specify the probability of criterional (and hence the complementary probability of noncriterional) responses as independent variables, a quantitative analysis of contingency can be accomplished that is immune to the deficiencies noted by Gibbon et al. (1974).

Galbicka and Platt (1986) and Scott and Platt (1985) used percentile schedules to control criterional-response probability and thereby allow a statistical measure of correlation known as the phi statistic to be manipulated as an independent variable. Galbicka and Platt differentiated long interresponse times of pigeons; Scott and Platt differentiated rats' joystick displacement locations. In both cases, increasing the contingency, as measured by phi, between reinforcement and the differentiated response dimension increased the control over responding by that dimension. The present experiment provides a further test of phi as a viable measure of reinforcement contingency in an operant paradigm by examining an invariance property exhibited by the phi coefficient.

For behavioral procedures in which there are only two possible response classes and two possible events following responding, phi can

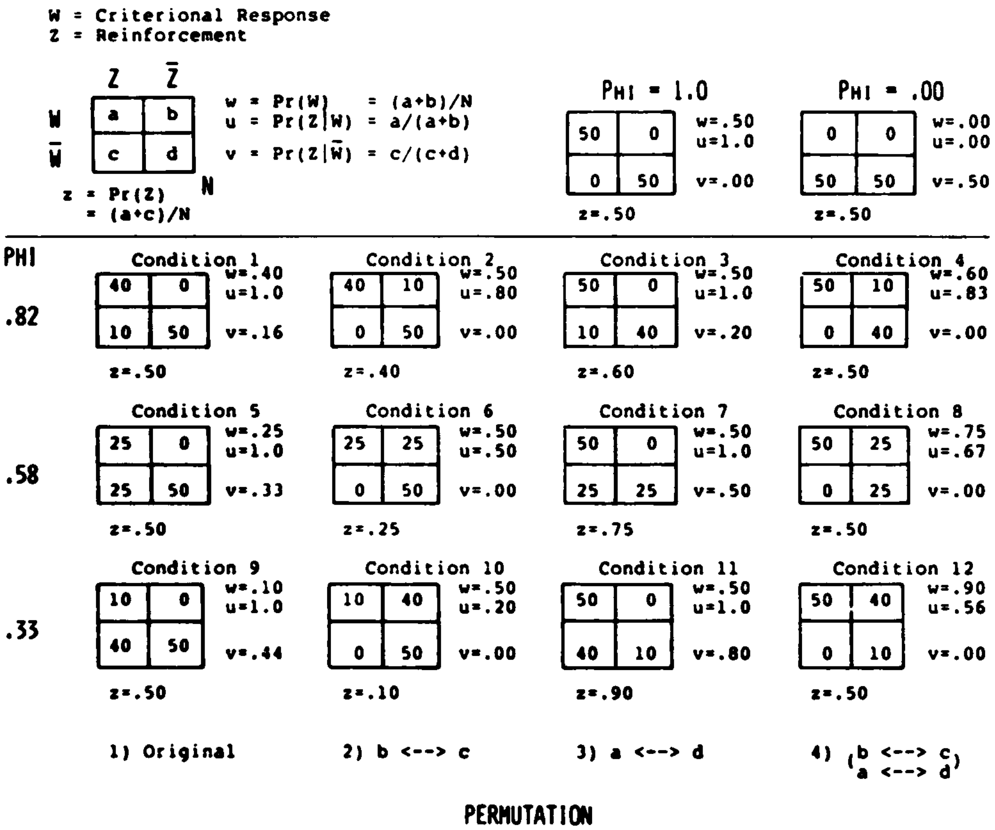

Fig. 1. 2 × 2 frequency tables illustrating the generic case (top left) and all experimental conditions. The letters *a, b, c,* and *d* refer to the frequency of each entry. The expressions surrounding the table relate the frequencies in each cell to the probabilities denoted by *w, u, v,* and *z* (see text for further description of these probabilities). The lower three rows illustrate frequency tables used to generate parameters for the 12 experimental conditions. Each row represents a different phi value, each column a different permutation. Instructions for generating each permutation from the original are at the bottom of the column.

be calculated from the frequencies of the four resulting response–event conjunctions. The generic case is illustrated in the 2 × 2 frequency table in the top left of Figure 1. The two responses represented in Figure 1 are criterional and noncriterional responses; the two possible events are reinforcement or nonreinforcement. The values *a, b, c,* and *d* represent the frequencies of observing the joint occurrence of the four possible response–event pairs. Hence, *a* and *c* represent the frequency of reinforced-criterional and -noncriterional responses, respectively, and *b* and *d* represent the frequency with which nonreinforced-criterional and -noncriterional responses occurred. Phi, expressed as a function of these cell frequencies, is

$$\phi = \frac{ad - cb}{[(a + b)(c + d)\ (a + c)(b + d)]^{1/2}}. \qquad (1)$$

In terms of cell frequencies, phi involves taking the difference of the diagonal cross products and dividing by the square root of the product of the marginal sums. This form of the phi coefficient reveals an invariance property of phi that is less obvious in other forms of the statistic. Diagonally interchanging the frequencies in the 2 × 2 table (i.e., swapping *a* for *d*, *c* for *b*, or both) leaves phi unchanged. The numerator of Equation 1 remains unchanged because multiplication is commutative. Interchanging the diagonals also does not quantitatively change the denominator of the fraction, but simply changes whether a particular marginal term appears as a row or column sum. Because the product of the marginal sums forms the divisor, the commutative property of multiplication again maintains the integrity of the expression.

This invariance property, then, specifies a formal similarity between reinforcement of criterional responses and nonreinforcement of noncriterional responses, as well as between reinforcement of noncriterional responses and nonreinforcement of criterional ones. Note that the diagonal interchange, although specifying the above symmetry, requires that a differing portion of responses be classified as criterional, to the extent that the interchanged cell frequencies differ. That is, if w denotes the probability of a criterional response, then w equals $(a + b)/N$, where N is the total cell entries. Diagonally switching a with d or b with c must change w whenever the swapped cell frequencies differ.

The diagonal interchange will also affect the conditional probabilities of reinforcement given criterional and noncriterional responses, denoted u and v, respectively, as well as the overall reinforcement frequency, denoted z. Expressions for each of these in terms of the cell frequencies are given in Figure 1. Because each is based on a row or column sum, diagonal interchanges may modify each probability, yet the molar relation (phi) remains constant. If phi reliably predicts the degree of response differentiation, the invariance of phi across diagonal interchanges suggests invariance in the degree of response differentiation generated. Should this result be obtained, it would restrict candidates for measures of behavioral contingency to those sharing this invariance. Other than phi, the only commonly used metrics of contingency sharing this invariance are the various monotone functions of the cross-product ratio, ad/bc. Clearly, the commutative property of multiplication will protect the integrity of these ratios across diagonal interchanges.

Phi can be distinguished easily from ratios of the cross products by setting either b or c at zero, such that cross-product ratios, but not phi, become extreme for all values of ad. With b set at zero, phi reduces to

$$\begin{aligned}\phi &= ad/[ad(c + d)(a + c)]^{1/2}\\ &= \{ad/[(c + d)(a + c)]\}^{1/2},\end{aligned}$$

and when c equals zero,

$$\phi = \{ad/[(a + b)(b + d)]\}^{1/2}.$$

Thus, variation in the nonzero parameters continues to produce systematic variation in phi, even with one cell set at zero.

The relative merits of the two metrics, then, can be ascertained by comparing responding under a set of conditions in which either b or c is set at zero. With this constraint, the relation between criterional responding and reinforcement, as measured by any of the cross-product ratios, will be equally maximal under all conditions. That is, if a function of the cross-product ratio is a viable metric of behavioral contingency, responding should not systematically vary under these conditions. Alternatively, varying the relative proportion of a and c in determining the marginal $a + c$ generates different values of phi. Thus, values of .33, .58, and .82 were arranged in the experiment reported here. Within each of these phi values, the four possible permutations of the contingency table generate four isophi conditions. Thus, if response differentiation varies as a function of phi, it would be expected to change across phi values, but not within permutations that give the same value. In order to generate complete functions across phi in a new response–species combination, conditions of phi = 0 and 1.0 were also examined.

The response dimension of interest was not the rate of key pecking, but rather the location of pecking. The manipulandum was a strip of 72 response keys arranged so as to form a circle 1 m in diameter. The circular arrangement eliminated major differences in discriminative stimuli correlated with end versus middle keys in a linear arrangement (see Eckerman & Lanson, 1969). Pecks on these keys were reinforced differentially with respect to their location relative to a target key. No explicit discriminative stimuli were used to mark the target location. To decrease further the probability of location preferences, this apparatus used a centrally located feeder, and trials were started by pecks into the feeder. This ensured that the subjects always were approximately equidistant from all keys at trial onset.

METHOD

Subjects

Subjects were 4 adult male White Carneau (*Columba livia*) pigeons (S01, S02, S03, and S04). They were maintained throughout the experiment at 80% of their free-feeding weights by restricted postsession access to mixed grain. Between experimental sessions, subjects were housed individually in a colony room maintained on a 15:9 hr light/dark cycle and were given free access to grit and water.

Apparatus

The experimental chamber had a cylindrical working area 103 cm in diameter and 35 cm deep. The walls of the cylinder were sheet metal painted flat gray, the top was a sheet of clear acrylic, and the floor was 0.6-cm hardware cloth. Seventy-two response keys, each 4 cm wide by 4.5 cm high, formed a strip around the cylinder 24 cm above the floor. Each key thus subtended 5° of arc around the perimeter of the cylinder. Any force on the key in excess of 0.15 N, which moved the base across a distance of 2 mm, registered a response. To ensure against simultaneous pecks on multiple keys, black acrylic bars 4.5 cm long, 0.5 cm wide, and extending 1.2 cm into the working area were mounted between neighboring keys. The keys could be transilluminated by a strip of 36 1.1-W white lights mounted behind and above the keys. When necessary during preliminary training, electrical tape placed on the back of some keys prevented transillumination. Centered in the floor of the chamber was a cylindrical feeder of black acrylic, 7.6 cm in diameter and rising 2.5 cm above the floor. A 2.5-cm diameter circular aperture centered in the top of the feeder cylinder allowed access to mixed grain dependent on the operation of a food magazine mounted underneath the cylinder. Feeder operation was signaled by illuminating the inside of the cylinder with four 1.1-W white lights mounted underneath the top of the feeder cylinder and by the sound of a motor used to raise and lower the food magazine on a worm gear. Two photoresistors mounted at right angles within the plane of the top of the feeder cylinder transduced feeder entries. A 60-W light bulb mounted above the center of the cylinder, approximately 2 m above the chamber, provided general illumination throughout experimental sessions except during the feeder cycle.

An LSI® 11/1 computer in an adjacent room arranged experimental conditions and collected data. Sessions were conducted 5 days per week and were monitored routinely through a closed-circuit television camera mounted in the upper corner of the room housing the experimental apparatus. Every attempt was made to maintain the location of the chamber within the room, as well as the location of other equipment in the apparatus room, constant throughout the course of the experiment. Otherwise, no explicit stimuli were provided to orient the subjects in the chamber. The chamber was cleaned daily of large debris and dusted approximately weekly.

Procedure

Pretraining. Subjects all had prior histories of pecking in standard pigeon operant conditioning units, and hence the amount of pretraining required was minimal. Pecking in the current apparatus was established using an autoshaping procedure involving 8-s illuminations of a single key followed by 4-s presentations of food, timed from feeder entry. Key illuminations were separated by varying intervals during which all keys were dark, lasting on the average 30 s. Any peck on the lighted key resulted in immediate food presentation, whereas pecks on any darkened key had no consequence. Once pecking was established, subjects were exposed to a single session during which four keys separated by 90° were illuminated, followed by a single session during which all keys were illuminated, with the autoshaping procedure still in effect, and pecks on any lighted key immediately producing food. Also at this point, the feeder-cycle duration was decreased to, and subsequently maintained at, 1.5 s timed from feeder entry. During these and all subsequent phases of the experiment, trials were initiated by an interruption of the feeder photoresistor circuit. This occurred either directly following eating, if the subject's beak remained in the feeder, or after the subject interrupted the photoresistor circuit by pecking into (most typically) or standing on the feeder.

Trial initiation was signaled by illumination of all response keys. They remained illuminated until a key was pecked (i.e., all trials contained one and only one peck). If the response was reinforced, the keylights and houselight were darkened, and the magazine was raised inside the feeder cylinder and illuminated. After feeding, the magazine was lowered inside the cylinder, the feeder lights darkened, and the houselight reilluminated. Unreinforced responses darkened the keylights but had no other consequence. Interrupting the feeder photoresistor circuit when the houselight was on initiated the next trial and illuminated the keys. Interruptions of the circuit during trials had no scheduled consequences.

The probability of food per peck, still irrespective of location, was next reduced. Subjects S02 and S03 were exposed to 20 sessions with a probability of food equal to .5, S01 and

S04 were exposed to probabilities of .5 and .33 for single sessions each, then to .2 for 15 sessions. After a pilot condition involving a targeted percentile schedule (see below), these 2 subjects were also exposed for 20 sessions to a probability of food per peck equal to .5. Following this condition, the experiment proper began.

Percentile schedules. All other phases used targeted percentile schedules (e.g., Davis & Platt, 1983) to control criterional-response probability while differentiating responding around a specified absolute target location. Locations were specified relative to a standard, arbitrary zero. Because targeted percentiles are actually two simultaneously operating nontargeted percentiles, it is necessary first to describe the operation of the latter. Galbicka (1988) and Platt (1973) provide extensive discussion of percentile procedures.

Percentile schedules segregate responses into criterional/noncriterional classes according to whether they exceed a certain percentile of the current response distribution. Any distribution of m observations can be ranked ordinally from lowest to highest, resulting in rank orders 1 through m. A subsequent observation must be contained in one of $m + 1$ intervals defined by the m ranks; below the first rank, between consecutive ranks, or above the highest rank. A fundamental theorem of nonparametric statistics (see Smith, 1953) states that given random and independent sampling, an observation sampled from the same population as the distribution has an equal probability of being contained in any interval, or a probability of $1/(m + 1)$. This is true independent of the values comprising the distribution, because it is based on their ranks and not their absolute values.

The probability that the next observation will be contained in one of k of the $m + 1$ intervals equals the probability of falling within a single interval multiplied by the number of intervals involved, or $k/(m + 1)$. This also represents the probability of a response below the kth rank, because k intervals are contained in this range. The probability the sample will exceed the kth rank is the complement, or

$$w = 1 - [k/(m + 1)].$$

Solved for k, any probability w of a criterional response can be generated by specifying the rank that must be exceeded for a response to be considered criterional:

$$k = (m + 1)(1 - w). \qquad (2)$$

Equation 2 specifies, given a sample of the most recent m responses (the *control memory*), the rank (k) the current response must exceed to observe criterional responses with the specified probability (w). The required rank (k) will remain constant for any w and m combination. Because the control memory is continuously updated to include only the most recent m responses, the cardinal value of the kth rank is subject to continuous flux, whereas the probability of observing a response ranked greater than k always equals w. Responses exactly tying the cardinal value of the kth rank are considered criterional with a random probability w.

All this follows from the assumption that the current response is sampled from the same population as that from which the distribution was derived, and that sampling occurs in a random and independent manner. To ensure that the current distribution most closely estimates that from which the current response was selected, the control memory is updated to include only the most recent m responses. As control memory size (m) is increased, the control memory will more likely include responses no longer reflecting the current population of responses. Hence, extremely large memories are to be avoided. Alternatively, requiring that sampling be random and independent means the effects of sequential dependencies in responding should be minimized. Increasing m will decrease the effect of any first or higher order sequentials. The latter are seldom observed at substantial levels in operant behavior, even when explicitly reinforced (see Shimp, 1973; Weiss, Laties, Siegel, & Goldstein, 1966). Hence, m must be large enough to attenuate the effects of sequential dependencies but small enough to reflect the current response distribution. The value in the present study was 24, typical of values used previously.

Percentile schedules, then, always define the most extreme w proportion of the responses as criterional, selecting more extreme cardinal values as differentiation progresses, in order that the probability of criterional responses remains constant. This arrangement cannot produce response values clustered around a discrete target value; rather, it can specify only a directional change from current values until differentiation reaches asymptote.

Defining as criterional a fixed proportion of responses closer to a discrete target than the rest of the distribution requires a targeted percentile schedule. These schedules first divide the differentiated dimension at the target value. This is particularly appropriate in the present apparatus, where there are equal response opportunities on either side of the target, but such symmetry is not prerequisite. Responses to each side of the target can now be treated independently from those to the other side, running two percentile schedules simultaneously so as to differentiate responding towards the target. Thus, responses on either side of the target are considered criterional if they are closer than k of the m responses on the same side, even though "being closer" requires moving in different directions on opposite sides of the target.

The only problem unique to targeted percentiles is how to define the control memory. Should the memory consist of the most recent m responses, independent of which side of the target they fall, or should individual, equal-sized memories be kept for responses on each side of the target? By including only recent responses, the first alternative tracks changes in behavior more closely. The price paid for this property is variation across trials in the relative proportion of memory values to the right and left of the target. The alternative strategy keeps constant the number of memory values on each side of the target, but runs the risk of not keeping pace with behavior. For example, the distribution of "left responses" would not change if responding consistently favored the right. A subsequent response to the left of the target would be ranked relative to a distribution no longer characteristic of the current response population. Although both symmetric numbers of memory values and response recency are desirable characteristics of the control memory, criterional-response probability remains controlled across a wide range of memory sizes. Because variation in memory size is less likely to disrupt differentiation than is use of responses that may lie considerably outside the current distribution, the memory was structured to include the most recent m responses, independent of their location. Thus, the number of memory values on each side of the target varied from trial to trial.

An attempt was made, however, to increase the symmetry in the number of responses on each side of the target as differentiation progressed. The criterional-response probability was modified depending on how responses were distributed in the control memory. If the memory were balanced, $m/2$ responses would fall on each side of the target. If s represents the number of responses in the memory on the same side as the current response, multiplying w in Equation 2 by the quantity $m/2s$ would, given a symmetric memory, leave w unaffected. However, as s approaches zero (that is, as more memory values are on the other side of the target), $wm/2s$ approaches infinity, increasing the proportion of responses on the (nonfavored) side considered criterional. For $wm/2s$ greater than or equal to one, any response on the nonpreferred side was criterional. Conversely, $wm/2s$ decreases criterional-response probability if the current response is on the preferred side (i.e., $m/2s < 1$ for $s > m/2$, and, hence, $w > wm/2s$). This strategy, however, is increasingly self-defeating as memory values predominate to one side. With this correction, exclusive preference (i.e., $s = m$) leads to half as many responses being criterional as programmed. Hence, this correction was used only while the adjusted criterional-response probability for responses on the nonpreferred side was less than or equal to unity (i.e., only if the number of memory values on the nonpreferred side equaled or exceeded $mw/2$, and, hence, $s \leq m - mw/2$).

If memory values predominated further to one side of the target such that $s > m - mw/2$, the quantity $(1 - w)$ in Equation 2 was multiplied by m/s. At the point of transition between the two corrections, each weights w by a factor of $1/(2-w)$. As s approaches m beyond the transition point, the latter correction approaches unity, and the probability of a criterional response is once again as specified by w.

Experimental conditions. The experiment proper began with the exposure of all subjects to 20 sessions of a targeted percentile with a maximal contingency (i.e., phi = 1, see top of Figure 1) between response location and reinforcement. Target locations were always 90° from the mean location during the last 5 days of the preceding condition except during phi = 0 (nondifferential reinforcement), for which targets are meaningless. Use of the mean location ensured equal shifts between conditions, independent of how effectively responding was differentiated in the immediately preceding condition. Target shifts between phases were always in the same direction for any subject, counterclockwise for S01 and S02, clockwise for S03 and S04.

The experimental manipulation involved examining behavior under the 12 additional conditions shown in the bottom three rows of Figure 1. These conditions were obtained by generating 2 × 2 frequency tables for three values of phi (.33, .58, and .82, see different rows), and then generating for each of these the four possible permutations resulting from diagonally interchanging pairs of cells within the tables, as shown in the different columns. The phi values were obtained by arbitrarily setting the overall reinforcement probability to .5 during the first permutation ("Original") under each phi value. Under phi = .58, these reinforcers were equally allocated to both criterional and noncriterional responses (i.e., $a = c$). Under phi = .82, 80% of the reinforcers followed criterional responses, and under phi = .33, 20% of the reinforced responses were criterional responses. All unreinforced responses were noncriterional responses (i.e., $b = 0$) in this permutation.

Figure 1 also presents for each condition the probability of a criterional response (w) and the conditional probabilities of reinforcement for criterional and noncriterional responses (u and v, respectively) used to generate the frequencies of joint events dictated by the various conditions. As a concrete demonstration of the process used to generate the experimental parameters for the four permutations, consider Conditions 5 through 8. Cell frequencies for Condition 5 were, given the constraints above, 25, 0, 25, and 50 for Cells a, b, c, and d, respectively. This yields, from the equations presented at the top left of Figure 1, values of w, u, and v equal to .25, 1.0, and .33, respectively. Permutation 2 (Condition 6) always involved interchanging Cells b and c, such that w now equaled .50, u equaled .50, and v equaled 0. Permutation 3 (Condition 7) interchanged Cells a and d of the original cells, generating frequencies of 50, 0, 25, and 25 for Cells a through d, respectively. Hence w again equaled .50, but u increased to 1.0 and v to .50. Finally, Permutation 4 (Condition 8) interchanged Cells a and d and Cells b and c of the original matrix simultaneously, generating cell frequencies of 50, 25, 0, and 25. These frequencies dictated programmed values of w equal to .75, u equal to .67, and v equal to 0. Under all of these permutations, w, u, and v varied across rather wide ranges, but the obtained value of phi was always .58. A similar process was used to obtain w, u, and v values for the other eight conditions of the experiment.

Figure 1 also lists the overall expected reinforcement frequency (z) for each condition. Conditions were presented in the order indicated for each subject in Table 1. Originally, only Conditions 1 through 8 were planned, and subjects were exposed to the eight conditions in a random fashion with the constraint that there be as many transitions between different conditions across as within phi values. All subjects were subsequently exposed to phi = 0 (see top of Figure 1). When aspects of the data suggested it might be fruitful to examine a value of phi lower than .58, Conditions 9 through 12 were generated and presented according to a Latin square. Subsequent to these, Condition 10 was replicated with new targets, and subjects then were again exposed to a phi = 0 condition. All conditions remained in effect for 20 sessions of 100 trials each, except the first two to which S01 and S04 were exposed. These remained in effect for 35 sessions, but responding showed little change after 15 sessions. Hence, all data presented are from Sessions 15 to 20 of a condition.

RESULTS

Figure 2 presents response distributions for each subject during the final five sessions of each experimental condition. First consider the three distributions in the rightmost column for each subject. These were obtained under phi = 0 (both exposures) and phi = 1.0 conditions. Phi = 0 distributions for S01, S02, and S04 were relatively flat, indicating little bias towards any particular key in the absence of differential reinforcement. Further, because these data were obtained after differentiation to a particular target (indicated by the dashed arrow), these distributions provide little evidence that previous differentiation was maintained in the absence of differential reinforcement. The phi = 0 distributions for Subject S03 showed some control by the previous target but a much more striking key preference at 200° under nondifferential reinforcement.

This location preference was effectively

Table 1

Condition orders, targets, mean and modal angles, mean resultants and target components of mean resultants for each subject by condition.

Condition	P[a]	Order	Target	Mean	Mode	Result	TC[b]	Order	Target	Mean	Mode	Result	TC
			Subject S01							Subject S02			
Phi = 1	—	1	345	347	345	.991	.990	1	25	30	30	.936	.932
1	1	2	75	76	90	.906	.906	3	205	235	200	.535	.462
2	2	5	355	358	360	.966	.965	4	325	325	320,330	.933	.933
3	3	6	90	97	105	.981	.973	7	225	225	220	.757	.757
4	4	9	25	28	30	.974	.973	8	315	319	320	.967	.964
5	1	3	165	176	175	.662	.650	2	120	114	125	.816	812
6	2	4	265	268	270	.928	.926	5	55	60	75	.832	.828
7	3	7	185	184	185	.687	.687	6	150	133	115,135	.648	.621
8	4	8	275	303	200	.191	.168	9	50	61	45	.790	.774
8[c]	4	10	120	127	125	.856	.850	—	—	—	—	—	—
9	1	12	250	282	25,265	.067	.057	12	190	225	265	.220	.180
10	2	13	290	299	290	.949	.937	11	280	282	290	.812	.812
10[c]	2	16	340	343	345	.854	.853	15	80	90	70,105	.619	.609
11	3	14	30	53	165	.147	.135	14	95	164	165	.311	.112
12	4	15	200	249	200	.199	.131	13	125	193	165	.280	.104
Phi = 0	—	11	120	123	195	.156	.156	10	50	117	290	.006	.002
Phi = 0[c]	—	17	340	322	345	.222	.211	16	80	148	235	.164	.062
			Subject S03							Subject S04			
Phi = 1	—	1	125	125	125	.919	.919	1	340	345	340	.863	.859
1	1	9	160	166	200	.895	.890	9	25	28	30	.942	.941
2	2	5	110	117	115	.902	.896	8	110	113	105	.893	.892
3	3	4	235	202	165	.351	.293	5	5	17	5	.805	.787
4	4	2	35	41	50	.957	.952	4	95	99	90	.887	.884
5	1	7	300	322	320	.787	.729	10	295	311	295,305	.126	.121
5[c]	1	—	—	—	—	—	—	11	150	155	125	.585	.582
6	2	6	25	33	30	.959	.950	7	200	199	205	.869	.869
7	3	3	310	323	295	.800	.779	3	160	171	175	.654	.642
8	4	8	230	253	200	.552	.508	2	250	254	230	.811	.809
8[c]	4	—	—	—	—	—	—	6	290	293	290	.762	.761
9	1	13	65	186	200	.643	−.336	16	235	200	220	.246	.201
10	2	14	95	101	85	.814	.810	15	315	325	305	.459	.452
10[c]	2	15	10	11	5	.813	.813	17	115	123	150	.592	.587
11	3	11	85	170	200	.819	.067	14	45	47	20, 45, 55	.453	.453
12	4	12	80	173	200	.795	−.048	13	110	133	175	.615	.566
Phi = 0	—	10	160	176	200	.854	.821	12	150	200	195	.297	.192
Phi = 0[c]	—	16	10	230	200	.209	−.160	18	115	141	150	.319	.287

[a] P, permutation.

[b] TC, target component.

[c] Replication. Condition 8 was replicated because of bimodal performance in the case of S01, as was Condition 5 for S04. Condition 8 was replicated due to technical problems in the case of S04.

abolished under phi = 1.0, and responding was located around the area of the target. All subjects demonstrated more location stereotypy under phi = 1.0, with a modal response value within 10° of the target. It should be remembered that the apparatus allowed only 5° resolution of response locations. Mean response locations over the last five sessions of this phase deviated from the target by 2°, 5°, 0°, and 5° for S01, S02, S03, and S04, respectively. Note that transposing the circular location dimension to a linear distribution produces artificial bimodality when the target is near either extreme (see S04 under phi = 1.0, for an example).

Figure 2 also presents, for each subject, response distributions from Conditions 1 through 12 in the leftmost three columns of each matrix

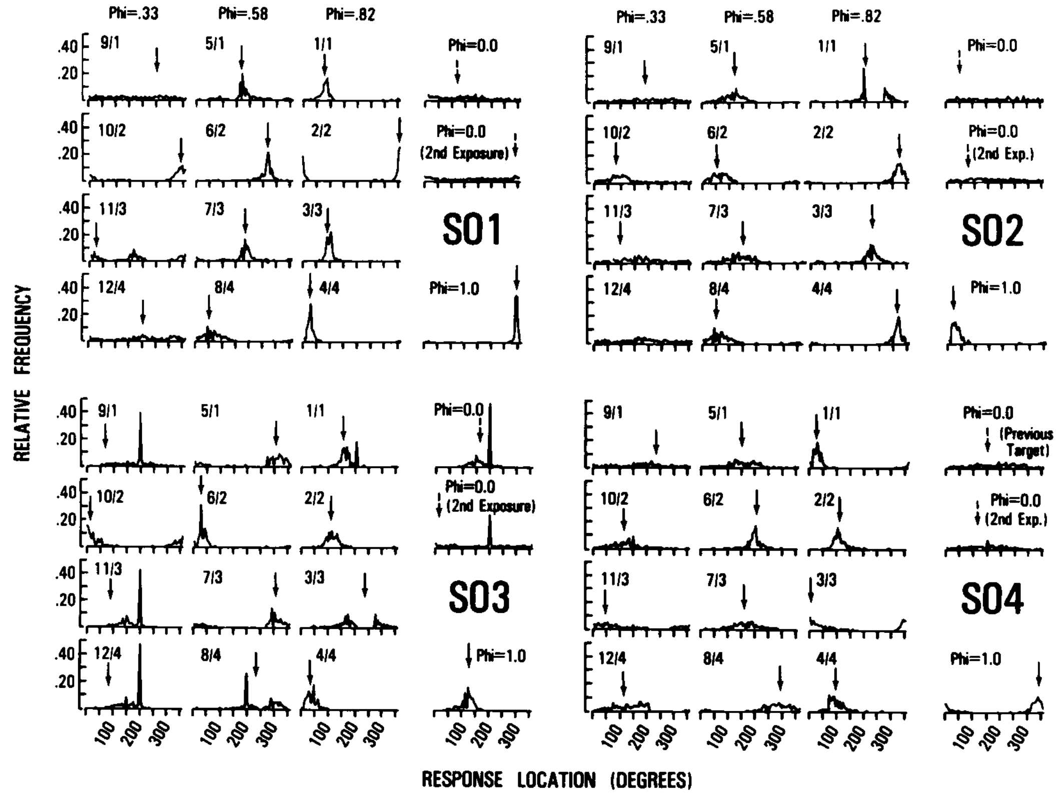

Fig. 2. Relative frequencies of pecks at each of the 72 response locations during the last five sessions of each condition. Each subject's data are shown in a 4 × 4 matrix. The rightmost column shows data obtained under both exposures to phi = 0 (top two) and phi = 1.0 (bottom right). The remaining three columns show distributions obtained under phi = .33, .58, and .82 (left to right) under each permutation (different rows). The numbers in these distributions refer to the condition and the permutation, respectively. The arrow shows the target. For phi = 0, the dashed arrow shows the target from the previous condition. The distributions show cumulative responses at each location divided by the total responses during the same period. Response location was determined from an arbitrary zero that was the same for all subjects.

of distributions. Rows represent data from the same procedural permutation, columns from identical phi values. Under phi = .82, responding tracked changes in the target location rather well, with the mode of the distribution within 20° of the target in 14 of the 16 cases. One exception involved a mode at 200° for S03 in Condition 1; however, the secondary mode was within 20° of the target. The other exception was Condition 3 for the same subject, which resulted in a highly bimodal distribution. This bimodality was especially peculiar, given that the target was 235°, near the 200° preference shown by this subject under phi = 0 as well as under Conditions 1, 8, 9, 11, and 12. Conditions showing pronounced preference for this location occurred sequentially (see Table 1) and after Condition 3. With phi = .82, there appeared to be little systematic relation between permutations and the response-location distributions.

Under phi = .58, responding again generally peaked within 20° of the target, with the exception of the previously mentioned data from Condition 8 for S03 showing the strong position preference, and the data from Condition 5 for S04. Distributions generally were less dispersed under the second than under the other permutations at this phi value, although the effect was small for S02. They were generally more dispersed than under the same permutation at phi = .82 (but compare Conditions 6 and 2 for S03, for example).

When phi = .33, responding was even more dispersed than when phi = .58 in the same permutation with all subjects except S03. This

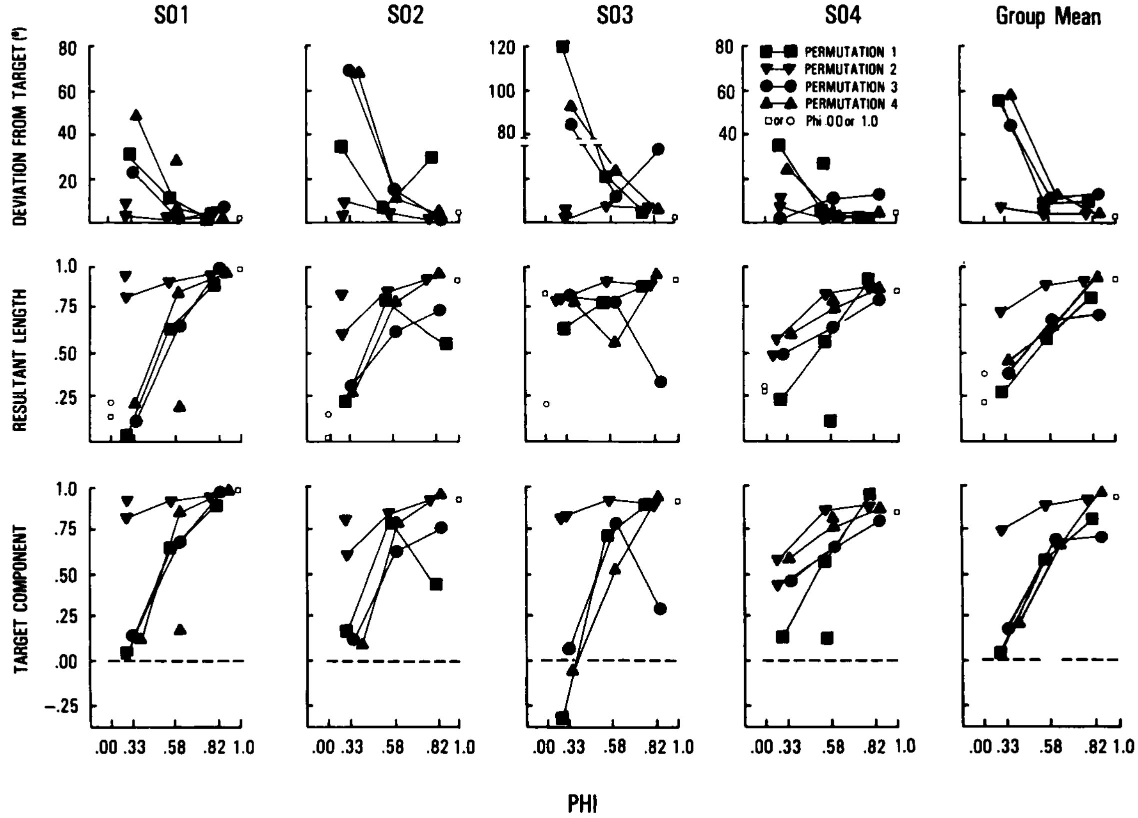

Fig. 3. Summary measures for each subject (leftmost four columns) and for the group of 4 subjects (right column) as a function of phi under the last five sessions of each condition. Different permutations are indicated by different symbols and are connected by solid lines. Unconnected points are from first exposures to the conditions that were replicated. Data from phi = 1.0 and, where appropriate, phi = 0, are shown by open symbols (square—first exposure, circle—replication). The top row shows the deviation between the mean response angle and the target, irrespective of sign. Note the break in the vertical axis for S03. The middle row shows the length of the resultant obtained through vector addition, and the bottom row shows the target component of the resultant (see text for further details). Values have been slightly displaced along the *x* axis where necessary to increase clarity.

subject showed increased pecking at 200°, with over 40% of all pecks occurring on this key during Conditions 9, 11, and 12. The remaining permutation generated concentrated responding around the target. The second permutation also produced slightly more concentrated responding than the others at phi = .33 for S01 and S02.

Figure 3 presents three measures summarizing aspects of the distributions shown in Figure 2. Different symbols denote different permutations. Lines connect the same permutation under the various phi values. Each column except the rightmost presents data from an individual subject; the rightmost shows group means. In cases in which a condition was replicated, the value obtained during initial exposure is shown as an unconnected point. Data from phi = 1.0 and, where appropriate, phi = 0 are shown as unconnected unfilled symbols. Group means include data from all exposures to each condition. Values used in constructing these figures are also presented in Table 1.

The top row in Figure 3 shows that the difference between the mean and target locations, irrespective of sign, generally decreased as phi increased. This was true of all but the second permutation, which more often was characterized by small deviations at all phi values. Subjects S02, S03, and S04 each showed one instance (during the first, third, and third permutations, respectively) under which deviations from target did not monotonically decrease with increases in phi. In the first two cases, the response distributions under the phi = .82 condition were strongly bimodal (see Figure 2). In the last, the mean deviation never exceeded 12°.

The mean deviation reflects the central tendency of the response distribution. Unfortunately, because the dimension of interest is a circular and not a linear one, standard deviation is invalidated as a measure of dispersion. Instead, each response was treated as a unit vector with a direction given by the response location, and vector addition was used to obtain mean, or resultant, vectors for each condition. The length of such resultants ranges from 0 to 1.0. Highly concentrated responding generates resultants approaching 1.0, whereas the resultant approaches 0 either as responding becomes increasingly dispersed or if responding becomes multimodal with radially symmetric modes. Because determination of resultant lengths involves trigonometric transformations of response location, length is a nonlinear function of response concentration. Increasing response concentration is reflected in rapidly increasing resultants that exponentially approach a limit of 1.0.

The second row of Figure 3 shows resultants for each subject under the various conditions. For S01, S02, and S04, increasing phi generally increased resultant lengths, indicating increasing concentration of responding. The exceptional function for S02 had a resultant length at phi = .82 almost equal that obtained under phi = .33, well below the resultants at phi = .82 for the other permutations. This resulted from the previously noted bimodality in this subject's response distribution under Condition 1. Resultants for S03 were not clearly related to phi. Only under the first permutation was a monotonic function obtained. Under all permutations, resultants under phi = .33 were longer than any obtained with the other subjects. The reason for the concentration was, in the case of the first, third, and fourth permutations, the location bias noted previously. Responding during the second permutation showed no sign of this bias; instead, it was concentrated in the vicinity of the target. At phi = 0, S03's responding was highly concentrated during initial exposure, but less so when reexposed. In both cases, the mode was 200°, accounting for 45% and 25% of all responses, respectively. All other subjects showed little concentrated responding under phi = 0. Under phi = 1.0, resultants ranged from .86 to .99.

Permutation appeared to affect resultant lengths only at the lowest phi value, if at all. As was the case with deviations from target location, the second permutation generated relatively more concentrated responding (i.e., longer resultants) than did the other permutations. This effect was particularly pronounced with S01 and S02. Although S03 generated long resultants under all permutations at phi = .33, only under the second permutation was this not the result of this subject's strong location bias. But even including these data, the resultants for the group (rightmost panel) showed a clear separation between the second and the other permutations as phi decreased. The resultants obtained under the first, third, and fourth permutations were tightly clustered at phi = .33 and .58. The disparate resultants at phi = .82 were due primarily to the bimodal distributions obtained under Conditions 1 and 3 for S02 and S03, respectively. Without these values, the group resultants for Permutations 1 through 4 were .914, .924, .848, and .946, respectively. Excluding data from conditions in which S03 demonstrated substantial bias towards 200° (Conditions 8, 9, 11, and 12) would further increase the difference between responding under Permutation 2 and the others. With these data excluded, resultants for the group under Permutations 1 through 4 at phi = .33 were .178, .739, .304, and .365, respectively, and under Condition 8 the group mean increased to .676.

Mean deviation and resultant vector length each only partially characterize the differentiation of response location. Responding must simultaneously be concentrated and be in or around the target area to show differentiation. Vectors provide a simple means of determining the concurrence between a response distribution and a target location. The resultants discussed above were determined relative to an arbitrary zero. The correspondence between each and its associated target may be scaled by projecting the resultant onto a diameter containing the target. That is, the resultant can form the hypotenuse of a right triangle, the other two sides of which are contained on a diameter through the target angle and a perpendicular of this diameter. The side on the diameter is the *target component* of the resultant vector. The length of the target component relative to the resultant increases as the mean deviation from the target reduces to zero, in which case the two segments are equivalent and the target component equals the resultant. As responding deviates further from the target, the length of the target component goes to zero at a mean angle ±90° from the target, then becomes increasingly negative as the mean angle approaches the antitarget. Because resultant lengths range from 0 to 1.0, target components range from −1 to +1. Only when responding is highly concentrated *and* around the target will target components approaching +1 be obtained. Increasing dispersion decreases the resultant, and hence the associated target component, even though the mean angle might accurately track the target. Alternatively, highly concentrated responding around nontarget areas will project only partially onto the target angle, thus also decreasing the target component.

The bottom of Figure 3 shows target components for the different subjects and for the group. Increasing phi increased target components in a monotonic manner in all but two of the functions presented. The two discrepancies involve the two bimodal distributions previously noted. Bimodality decreases target component by increasing the degree of dispersion, thus decreasing the length of the resultant, and by shifting the mean deviation away from the target, further decreasing the length of the projection (see panels directly above for S02 and S03). The only systematic effect of permutation again involved Permutation 2. Except for S04, the slope of these functions was considerably shallower than slopes from the other permutations. Subjects S01 and S03 showed practically identical differentiation at phi = .33 under this permutation as under phi = .82 in general.

In addition to examining individual responses with respect to location, both absolutely and relative to a target location, consecutive responses on a circular dimension can be analyzed with respect to the deviation between them, without any bias arising from the preceding location. That is, the angular deviation between responses on consecutive trials is not subject to artificial limitation arising from the response location on any trial, because subjects can always deviate ±180° from any location. One interesting comparison is between deviations obtained following reinforced responses relative to those obtained following nonreinforced responses. Figures 4 and 5 plot the relative frequencies of these two measures. In general, subjects returned to an area close to that of the previously reinforced response; that is, for any particular condition, deviations after a reinforced response were smaller than after a nonreinforced one. The degree of dispersion in the deviations following nonreinforced trials also tended to be smaller during Permutations 2 and 4 relative to Permutations 1 and 3.

DISCUSSION

Response location on a circular dimension was controlled in direct relation to location–reinforcer contingency as measured by phi. That is, deviations from target generally decreased as phi was increased, resultant vector lengths (indicating concentration of responding) increased as phi was increased, and as a result of these the target components also increased. Thus, at least under the constraints involved here (i.e., either *b* or *c* equal to zero), no commonly used function of the cross-product ratio adequately predicts the variation in behavior obtained under different 2 × 2 tables.

The only exceptions to the above statements occurred in conditions that generated bimodal responding, or in those that for Subject S03 resulted in responding around the 200° position. Once this position preference developed during Condition 8, the only conditions generating responding controlled by target location were Condition 1 and Condition 10. Responding during Condition 10 demonstrates that the position preference did not simply predominate during the latter half of the experiment for this subject, because this was the penultimate condition to which this subject was exposed. Responding occurred predominantly at 200° during conditions preceding and following Condition 10, yet no responding was observed at this location under the latter condition. Rather, responding was concentrated around the targets assigned with each exposure to Condition 10, those being 95° and 10°.

Responding during Condition 10 differed from that obtained under the other permutations of phi = .33 for Subjects S01 and S02 as well. For both birds, responding during each exposure to the second permutation demon-

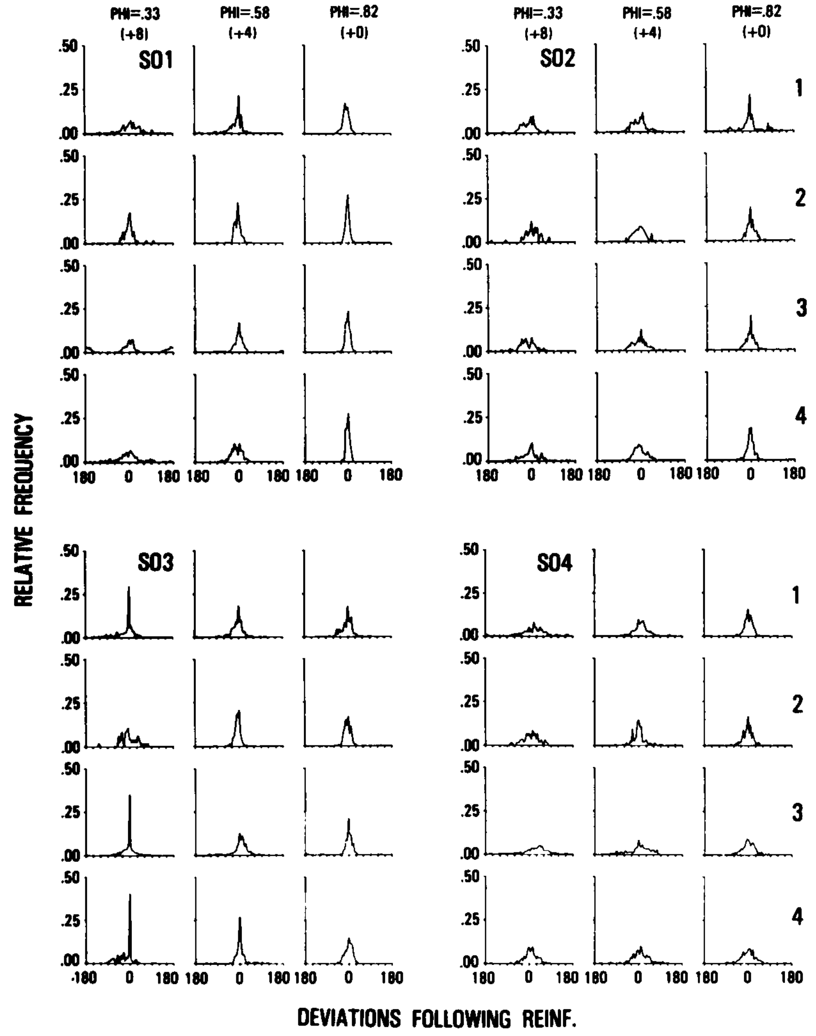

Fig. 4. Relative frequencies of response deviations (the angular displacement between successive responses) following reinforced responses for each subject under the 12 experimental conditions. Deviations following reinforcement were segregated into 5° bins, summed across the last five sessions, and divided by the total number of reinforcers delivered during this period to obtain the relative frequencies. Permutations are given by the numbers at the right of the figure. Condition numbers can be obtained by adding the permutation to the number in parentheses under the phi value at the top of each column.

strated much greater control by response location than did the other permutations. This was also generally true under phi = .58; however, the difference was smaller due to the greater differentiation obtained with the other permutations. Hence, although response concentration increased monotonically under all permutations as a function of phi, the slope of the function for the second permutation was clearly lower than those from the other three. The diagonal interchange thus produced only partially interchangeable results. The first, third, and fourth permutations produced roughly comparable target components for the group (see Figure 3, bottom left), particularly after noting that the differences observed at .82 were almost entirely due to the bimodal performance obtained under Conditions 1 and 3 for S02 and S03, respectively. That the remaining permutation differs from the others, however, is equally clear.

Do these data demonstrate sufficient invar-

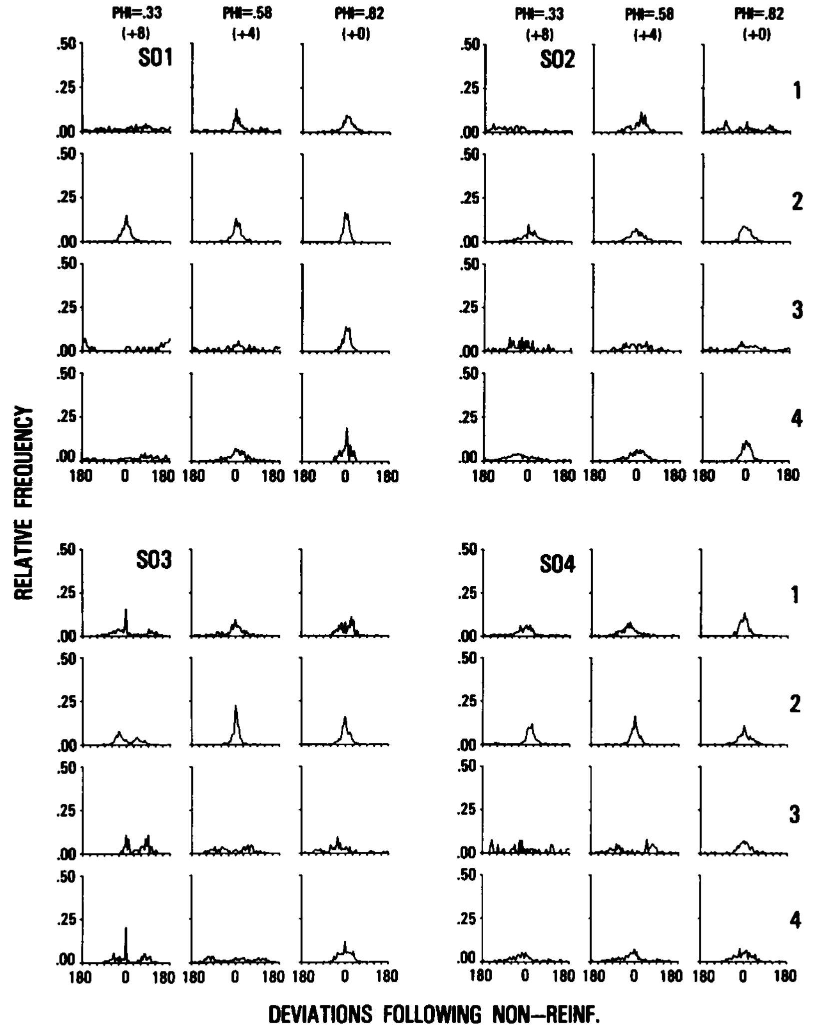

Fig. 5. Relative frequencies of response deviations following nonreinforced responses for each subject under the 12 experimental conditions. Deviations following nonreinforced responses were segregated into 5° bins, summed across the last five sessions of a condition, and divided by the total number of nonreinforced responses during this period. All other plotting conventions are as in Figure 4.

iance to provide support for phi as a metric of contingency in operant differentiation? Three of the four permutations generated roughly interchangeable target components. Previous data on location differentiation from Scott and Platt (1985) and on interresponse-time differentiation from Galbicka and Platt (1986) add weight to the conclusion that phi is a viable contingency metric. Although all permutations were not equally subsumed by phi in the present study, other contingency metrics fare no better, and usually less well than phi. The difference in conditional probabilities ($u - v$ in the present nomenclature), extensively used in respondent conditioning (e.g., Rescorla, 1968), generates values for Conditions 9 through 12 of .56, .20, .20, and .56. Thus it actually predicts a decrease in response-reinforcer contingency under Condition 10 relative to that in Conditions 9 and 12, and the same contingency as Condition 11. Responding was only substantially differentiated under the former. Another measure of contingency recently offered would in the present nomenclature be $(u - v)/u$ (see Thomie & Khouri, 1984; Thomie & Loukas, 1983). This measure specifies that the contingency during the second permutation would be perfect, but it also says the same for the fourth permutation. Hence its use solves one problem while creating another.

We do not mean to argue that phi should be adopted unconditionally on the basis that it provides the least poor fit to the data. Rather, we note that no measure of contingency, when directly tested against phi, has provided as good a fit. At the same time, the present data argue strongly that the use of phi be tempered until further experimentation can delineate the variables responsible for the difference in differentiation obtained under the second versus the other permutations.

The most striking difference seen when comparing the permutations is the low overall reinforcement probability under the second permutation relative to the others. Perhaps the relative scarcity of reinforcement leads to greater effectiveness of that reinforcer when delivered, generating greater response differentiation as a result. This sort of interpretation is consistent with marginal utility notions of reinforcement, in which the context is critically important in determining the effectiveness of a single reinforcement at a very molar level (e.g., Brandon, 1969). Such marginal utility might interact with the processes normally indexed by phi to generate the deviation observed.

Absolute or relative reinforcement probability cannot by themselves easily account for the present data. Increasing absolute reinforcement rates has been suggested to increase discriminability under signal-detection-like procedures (Nevin, Jenkins, Whittaker, & Yarensky, 1982; but cf. McCarthy & Davison, 1982, for conflicting results). To the extent that stimulus discrimination is an analogue of response differentiation, similar effects might be expected in the present situation. That clearly was not true for overall reinforcement probability, because the second permutation had the lowest reinforcement frequency of all permutations, not the highest. More to the point, however, the same reinforcement frequency was provided under all conditions comprising Permutations 1 and 4, yet differentiation varied drastically. Conversely, although relative reinforcement frequency for criterional responses was maximal during Permutations 2 and 4, only the former generated consistent response differentiation. Condition 12 stands out, both for providing a relatively high overall reinforcement probability and for reinforcing only criterional responses. Yet, with all subjects, this condition generated little differentiation.

An alternative account of the present data may lie in the deviations between consecutive responses. The smaller deviations following reinforced versus nonreinforced trials (compare Figures 4 and 5) can be described as a "win-stay, lose-shift" pattern. This phrase merely describes the greater probability, after reinforcement, of returning to the same general area relative to moving off to a substantially different location. This general pattern was observed under all conditions, suggesting that the effect did not depend on the location and/or frequency of reinforcement. It should be emphasized that both criterional and, in Permutations 1 and 3, noncriterional responses could be reinforced. Hence, these distributions do not describe concentration around a particular location but rather a return to the location reinforced on the last trial, whatever that might have been.

Deviations following nonreinforced responses were smaller under Permutation 2 and, to a lesser degree, Permutation 4, than under Permutations 1 and 3. Again the location of nonreinforced pecks was not correlated with the size of this effect. For example, Conditions 9 and 11 programmed 50 and 10 nonreinforced trials per day. Further, the noncriterional range $(1 - w)$ was substantially different in these two conditions. Despite these two differences, the deviation distributions following nonreinforced responses were not substantially different.

Permutations 2 and 4 each provide no reinforcement for noncriterional responses (i.e., $v = 0$). It is not readily apparent why this should generate smaller deviations following nonreinforced responses. These occur when responses are either noncriterional or unreinforced criterional ones, hence the probability of reinforcement for pecking the same location (i.e., not deviating) would be 0 or u, respectively. The reinforcement probability for deviating would be u or 0, depending on whether deviating generates a criterional or noncriterional response, respectively. Thus, there appears to be little differential consequence of doing either.

Permutations 1 and 3, however, always reinforce criterional responses (i.e., $u = 1$). All unreinforced responses are noncriterional. Thus, pecking a location which was not reinforced on Trial n will be reinforced with probability v on Trial $n + 1$. Deviating from that location will be reinforced with a probability of v or 1, depending on whether the new location is noncriterional or criterional, respectively. Therefore, deviation following nonreinforcement is differentially reinforced under these permutations. Reinforcement probability can only increase, following a nonreinforced peck, if a different location is pecked. Reinforcing deviation differentially relative to Permutations 2 and 4 might be expected to generate more dispersion in the distribution of response locations pecked.

Permutation 4 generated much less concentrated responding than did Permutation 2, even though, according to the above analysis, both provide the same nondifferential reinforcement for deviating. A possible explanation for the difference between the two involves the tendency under all conditions to return to a location if pecking there was reinforced on the preceding trial. Under the fourth permutation, as under the second, only criterional responses were reinforced. Under the former, however, the range of responses considered criterional (i.e., w) was greater, particularly as phi decreased. As a result, reinforcement would occur over a wider range of locations relative to the other permutations, and hence pecking would recur within that widened range.

Rather than attribute behavioral control to the molar measure phi, the above suggests a much simpler, and more local, analysis. Subjects return to a previously reinforced location until a response at that location is not reinforced, and then shift to a different location depending on the degree of differential reinforcement correlated with shifting. Hiller (1987) showed that response deviations, as opposed to locations, can themselves be differentiated, such that the distribution of deviations comes to correspond more closely to that required for reinforcement. Subjects learn to move between two locations with the deviation specified for reinforcement. This leads to dispersed response-location distributions. That such behavior can be differentiated should not be surprising, in that it merely restates optimal foraging theory. Subjects stay in a patch to the extent that remaining is correlated with continued access to food. When moving to a new patch (i.e., shifting to a new location in the present apparatus) is differentially reinforced, deviation increases in frequency.

Locational control then becomes a by-product of the reinforcement for switching between locations. That is, concentrated responding occurs not because of the differential reinforcement for pecking at particular locations, but instead is due to the absence of differential reinforcement for shifting to *other* locations. As shifting is increasingly reinforced (either by making more patches available or increasing patch density), control by a particular response location diminishes. This might occur in the present experiments either because the noncriterional range increased or because reinforcement probability for noncriterional responses increased. A perspective analogous to the one above has been advanced with respect to concurrent-schedule performance (cf. Silberberg & Ziriax, 1982), wherein responding is described not in terms of the reinforcers obtained by pecking each of the two locations, but rather in terms of the differential reinforcement of shifting between the two.

Whatever the relative merits of modifying or abandoning phi, we should reiterate that it is not offered as a model of behavior, but only as a means of characterizing the long-term effects of reinforcement on separate, mutually exclusive operants. The interaction between these operants may most easily be characterized by a molar value such as the phi coefficient, but that coefficient merely scales the relative contiguity between reinforcement and the two response classes. Its value is one of simplifying the description of the interaction, not one of simplifying the interaction itself. The deviation from invariance observed in the present study suggests a need for more data on the boundary conditions within which phi is an appropriate descriptor of a behavioral procedure. Such data may shed new light on the local mechanisms or processes responsible for the shaping of responding through differential reinforcement.

REFERENCES

Alleman, H. D., & Platt, J. R. (1973). Differential reinforcement of interresponse times with controlled probability of reinforcement per response. *Learning and Motivation*, **4**, 40–73.

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior,* **22,** 231-242.

Brandon, P. K. (1969). *The effects of error-contingent time out on counting behavior in rats.* Unpublished doctoral dissertation, University of Michigan.

Davis, E. R., & Platt, J. R. (1983). Contiguity and contingency in the acquisition and maintenance of an operant. *Learning and Motivation,* **14,** 487-512.

Davison, M. C., & Tustin, R. D. (1978). The relation between the generalized matching law and signal-detection theory. *Journal of the Experimental Analysis of Behavior,* **29,** 331-336.

Eckerman, D. A., & Lanson, R. N. (1969). Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior,* **12,** 73-80.

Galbicka, G. (1988). Differentiating *The Behavior of Organisms. Journal of the Experimental Analysis of Behavior,* **50,** 343-354.

Galbicka, G., & Platt, J. R. (1986). Parametric manipulation of interresponse-time contingency independent of reinforcement rate. *Journal of Experimental Psychology: Animal Behavior Processes,* **12,** 371-380.

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

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

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior,* **13,** 243-266.

Hiller, G. (1987). *Stereotypy of a spatial operant in pigeons and the law of effect.* Unpublished doctoral dissertation, McMaster University, Hamilton, Ontario.

Kuch, D. O., & Platt, J. R. (1976). Reinforcement rate and interresponse time differentiation. *Journal of the Experimental Analysis of Behavior,* **26,** 471-486.

McCarthy, D., & Davison, M. (1982). Independence of stimulus discriminability from absolute rate of reinforcement in a signal-detection procedure. *Journal of the Experimental Analysis of Behavior,* **37,** 371-382.

Nevin, J. A., Jenkins, P., Whittaker, S., & Yarensky, P. (1982). Reinforcement contingencies and signal detection. *Journal of the Experimental Analysis of Behavior,* **37,** 65-79.

Platt, J. R. (1973). Percentile reinforcement: Paradigms for experimental analysis of response shaping. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 7, pp. 271-296). New York: Academic Press.

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

Scott, G. K., & Platt, J. R. (1985). Model of response-reinforcer contingency. *Journal of Experimental Psychology: Animal Behavior Processes,* **11,** 152-171.

Shimp, C. P. (1973). Sequential dependencies in free-responding. *Journal of the Experimental Analysis of Behavior,* **19,** 491-497.

Silberberg, A., & Ziriax, J. M. (1982). The interchangeover time as a molecular dependent variable in concurrent schedules. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 131-151). Cambridge, MA: Ballinger.

Smith, K. (1953). Distribution-free statistical methods and the concept of power efficiency. In L. Festinger, & D. Katz (Eds.), *Research methods in the behavioral sciences* (pp. 536-577). New York: Dryden Press.

Thomie, A., & Khouri, P. (1984). The effects of response-reinforcer contingency on time allocation in the open field. *Learning and Motivation,* **15,** 1-11.

Thomie, A., & Loukas, E. (1983). Correlations between rats' spatial location and intracranial stimulation administration affects rate of acquisition and asymptotic level of time allocation preference in the open field. *Learning and Motivation,* **14,** 449-471.

Weiss, B., Laties, V. G., Siegel, L., & Goldstein, D. (1966). A computer analysis of serial interactions in spaced responding. *Journal of the Experimental Analysis of Behavior,* **9,** 619-626.

Received August 4, 1987
Final acceptance September 13, 1988