ON RESISTANCE TO EXTINCTION IN A FREE-OPERANT: AN ANALYSIS OF THE EFFECTS OF ALTERNATIVE REINFORCEMENT

by

John C. Damron

B.A., California State University, Fullerton, 1968

M.A., Simon Fraser University, 1975

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
in the Department
of
Psychology

SIMON FRASER UNIVERSITY

May 1981

Approval

Name: John Charles Damron

Degree: Doctor of Philosophy

Title of thesis: On resistance to extinction in a free-operant: An analysis of the effects of alternative reinforcement

Examining Committee:

Chairperson: Dennis Krebs

Vito Modigliani
Senior Supervisor

Bruce Alexander

Chris Davis

William Krane

Ellen Kimmel
External Examiner
College of Education
University of South Florida

Date Approved: [illegible]

PARTIAL COPYRIGHT LICENSE

Title of Thesis/Dissertation:

On resistance to extinction in a free-operant: An analysis
of the effects of alternative reinforcement

Author: ______________________
(signature)

John Charles Damron
(name)

Aug 14, 1981
(date)

ABSTRACT

According to the matching principle for asymptotic operant behavior, responses are expected to be distributed between or among available alternatives in the same proportion as the reinforcements yielded therefrom. In the work below, the matching principle is outlined and instances of supporting data are discussed. It is shown that the matching principle has formal implications for the concept and empirical phenomena of free-operant extinction. In particular, it is argued that operant conditioning and extinction are not, from the perspective of the matching principle, distinctive processes requiring separate explanation but are rather two aspects of the same behavioral process, response allocation and re-allocation as guided by the matching principle. It is, moreover, suggested that the matching principle has formal implications for the generality and robustness of resistance to extinction effects. In particular, it is argued that the matching principle imposes motivational constraints on the generality of the so-called partial-reinforcement effect (PRE), with robust PREs being confined to settings in which the level of alternative or "background" reinforcement is relatively meager.

Three experiments were conducted to test the effects of the level of background reinforcement on resistance to extinction in intermittently reinforced operant responses. In experiment I, 45 subjects (hungry pigeons) were trained in a two-alternative operant conditioning apparatus to peck at the two concurrently available response-keys for occasional access to food-reinforcement. They were then assigned to one of nine groups differing in 1) the level of reinforcement yielded by pecks to the right or "alternative" response-key and/or 2) the ratio of responses to reinforcements arranged for pecks to the left or "target" response-key. Following eight days of acquisition training the reinforcement yielded by the target response was omitted and that response was allowed to undergo extinction. The conditions of reinforcement associated with the alternative response continued uninterrupted during extinction of the target response. It was found that resistance to extinction (or the number of target responses emitted in extinction) was <u>inversely</u> related to the level of prevailing background reinforcement - food yielded by the alternative response, a finding which is consistent with the matching-based account of extinction (i.e., response re-allocation theory) advanced herein.

In order to test for the possibility that the extinction-related response decrements observed in experiment I were produced by interference originating in the repeated occurrence of the alternative response rather than by the motivational decrements posited by response re-allocation theory, a second experiment was conducted. Experiment II was a replication of experiment I with the added provision that the reinforcements yielded by the alternative response were <u>signalled</u> in order to substantially reduce the frequency of that response, the presumed source of potential interference. In contrast to the results of experiment I, the presence of a source of background reinforcement, when signalled, markedly <u>enhanced</u> the resistance to extinction evidenced by the target response in experiment II.

Possible explanations of the results of experiment II were discussed and the hypothesis was advanced that those results are confined to "extinction-naive" subjects. A preliminary investigation of that hypothesis, in which subjects (hungry pigeons) were run through five conditioning-extinction cycles, showed that the response enhancing effects of signalled alternative reinforcement disappeared with repeated exposure to the extinction operation, thus supporting the foregoing hypothesis.

Finally, the theoretical and empirical work reported herein are briefly summarized. Theoretical implications of the present work are delineated and prospects for future research are discussed.

DEDICATION

This thesis is dedicated to Teri, Tracy and Joanna Damron, my family.

ACKNOWLEDGEMENTS

I'd like to express my gratitude to the members of my supervisory committee for their many helpful comments and suggestions and for their moral support during the preparation of this manuscript. In particular, I'd like to thank Dr.s B. K. Alexander and C. M. Davis for their contributions to the development of the final form of the present manuscript. I moreover thank Dr. Alexander for our numerous intellectual encounters. I have been gratified by such exchanges and have benefited greatly from them. Most especially, I'd like to express my gratitude to Dr.s K. R. Burstein and V. Modigliani, my teachers and advisors, for their support and encouragement throughout my graduate studies and for a level of ongoing intellectual stimulation that has without question had a distinctive, pervasive and enduring effect on my thinking.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTORY REMARKS

> Neither Pavlovian reinforcement theory nor the law of effect, taken by themselves, contain any set of statements sufficient to account for extinction: although both attribute increases in response probability to the presentation of a reinforcing event, it is not clear how either can explain, without further assumptions, why the omission of reinforcement should result in a decrease in the probability of responding. Hence, apparently, the need for theories of extinction.
>
> N. J. Mackintosh, 1974, p. 405

One of the more important aims of the scientific enterprise is the isolation and careful description of the regularities in the events or processes under experimental scrutiny. A second and more formidable aim of science is the formulation of theoretical explanations of observed regularities in the form of general laws, principles or, in the case of a large number of related observations, abstract paradigms. Not infrequently, theories and paradigms provide scientists with a vantage point from which the subject matter of an entire discipline may be viewed and tentatively understood (cf. Kuhn, 1970). Thus, Newton's theory of universal gravitation and Darwin's theory of evolution by natural selection provided physics and biology with paradigmatic frameworks within which the various phenomena of those disciplines could be ordered and explained and which, as a consequence, shaped the course of research in these respective enterprises.

As the virtual cornerstone of behavioral psychology, the principle of reinforcement (or, the law of effect) has achieved something of the same status as the foregoing paradigms insofar as it has 1) provided behavioral psychologists with model problems, experiments and theoretical solutions (cf. Kuhn, 1970, p. 3), and 2) is routinely invoked to explain a wide variety of human and infrahuman behavior (Skinner, 1953, 1974; Honig and Staddon, 1977).[1]

Two recent developments add considerable substance to the foregoing contention. First, reinforcement as a behavioral process is regarded by a number of theoreticians as the ontogenetic analogue of natural selection (see, e.g., Skinner, 1966; Gilbert, 1973; Staddon and Simmelhag, 1971; and Catania, 1978). Indeed, operant conditioning and evolution have been characterized as simply two examples of adaptation by blind variation and environmental selection (Skinner, 1974). Secondly, recent work by Herrnstein and his colleagues has extended the generality of the law of effect and considerably sharpened the quantitative precision of reinforcement-based analyses of behavior (Herrnstein, 1970, 1979; de Villiers, 1977).

Bearing this in mind, the introductory quotation by Mackintosh may be regarded as a source of genuine puzzlement, for the upshot of that quote, which I believe represents the "conventional view," is that there is <u>nothing</u> inherent in the principle of reinforcement which explains why the omission of a reinforcing event should result in a decrement in the probability of responding. Which is to say, there is a considerable segment of instrumental behavior for which the law of effect is, according to Mackintosh, wholly irrelevant. Hence, the apparent need for separate theories of extinction.

Mackintosh's statement to the contrary notwithstanding, a major thesis of what follows is that a sophisticated principle of reinforcement has substantive relevance for the data of experimental extinction and indeed, bears centrally on questions about the very nature of extinction. And further, I suggest that the foregoing follows fairly straightforwardly from the logic of what is perhaps the most highly developed quantitative model of reinforcement, the so-called matching principle (Herrnstein, 1970, 1979).

Curiously, although the law of effect has existed in one or another form for nearly eighty years, the matching principle began unfolding only twenty years ago (Herrnstein, 1961, 1979) and gained status as a general quantitative principle of behavior only in the 1970's (Herrnstein, 1970; de Villiers, 1977). Why? There are, I suggest, at least three assumptions psychologists have traditionally made about the nature of the reinforcement process which have impeded the development of quantitative models of reinforcement such as the matching principle. The first of these is the familiar assumption that a reinforcer must be virtually contiguous with the response that produced it before conditioning can occur. Implicit in this view is the presumption that complex molar performances can, in some sense, be reduced to and explained at the level of discrete and virtually simultaneously occurring behavioral and environmental events (Spence, 1947). Such a view ignores the possibility that molar behavior may be considerably more sensitive to molar relationships such as, for example, the rate at which a response yields a reinforcer (Baum, 1973). Research in classical (e.g., Revusky and Garcia, 1971; Rescorla and Wagner, 1972) and instrumental conditioning (e.g., Herrnstein, 1969; Baum, 1973) supports the idea that contingencies or correlations, <u>not</u> contiguities, are the major determinants of conditioned behavior. Insofar as it caused learning theorists to search for regularities at an unproductive level of analysis, adherence to the age-old contiguity perspective probably impeded the development of modern correlation-based models of the conditioning process.

The second of the aforementioned assumptions has to do with the notion that one or another absolute property of reinforcement is the critical determining factor in instrumental conditioning. To be sure, this is usually not a highly articulated assumption, but it is clearly implicit in the way in which researchers routinely program, manipulate, report the effects of and discuss, reinforcement rates and magnitudes. Until fairly recently, explicit manipulation of relative reinforcement values by researchers has been notably uncommon. And indeed, researchers have routinely employed experimental procedures in which response frequencies are plotted as functions of absolute rates or magnitudes of reinforcement. Unfortunately, response frequencies in such procedures are usually highly variable and surprisingly insensitive to changes in absolute amount (i.e., rate or magnitude) of reinforcement (Catania, 1963, 1966). Clearly, such lack of order makes successful quantification of the behavior-reinforcement relationship virtually impossible. In contrast, when reinforcement is viewed in relative terms, its orderly relationship to behavior becomes (as we shall soon see) readily apparent. Moreover, the reported insensitivity of behavior to absolute reinforcement levels becomes entirely comprehensible when one considers that, at almost all absolute values of reinforcement in a single-alternative procedure, the target response yields nearly 100% of the available reinforcement. That response's relative reinforcement therefore remains essentially unchanged across all nonzero values of the programmed reinforcer.

The third and final assumption, which is also something of a methodological convention, relates intimately to the second. It is the assumption that researchers may be expected to find simple and orderly relationships between response frequencies and rates of reinforcement in single-alternative experimental settings. This assumption is troublesome not because there are no systematic relationships to be found at this level, but because researchers are apt not to discover them. The reason is straightforward: try as we may to eliminate them, subjects will always have unprogrammed sources of reinforcement available to them in such situations, even, as Herrnstein (1970, p. 388) has suggested, these amount to little more than the subjects' own "itches, irritations and other calls for service." According to Herrnstein (1970, p. 389), the problem of identifying and measuring these alternatives may be insoluble, which, if unavoidably true, renders the task of precisely measuring the effects of a quantity of programmed reinforcement on behavior very difficult indeed, for the nature and extent of the effects of unprogrammed reinforcement on behavior cannot be readily assessed.

How then are the effects of reinforcement on behavior to be measured? According to one prominent view (i.e., Herrnstein, 1970, 1971), the relative effects of reinforcement on behavior may be assessed by providing subjects with two or more alternatives, between or among which it may freely and continuously choose (Herrnstein, 1970, 1971, 1979). Such a procedure, effectively overcomes the problem of unprogrammed or "extraneous" reinforcement by equalizing, as far as possible, its effects on the programmed alternatives thereby yielding an orderly, quantifiable relationship between relative response frequencies and relative rates of reinforcement (cf. Herrnstein, 1970).

The remainder of this work will concern itself with an empirically well-validated perspective on the reinforcement process which has either expressly or tacitly rejected the traditional assumptions discussed above. That perspective, the matching principle, posits that organisms tend to distribute their behavior across available alternatives in proportion to the relative molar value of the consequences yielded therefrom. To the extent that the matching principle is silent about the particular mechanisms underlying reinforcement, it is not a theory of reinforcement. It is rather a theory <u>about</u> reinforcement, or about the action of reinforcement on behavior. In a related sense, the matching principle is a theory of choice (or alternatively, a choice theory of instrumental behavior), for it implies that whenever an organism behaves in a particular way, it has made an implicit choice between alternatives based on their values relative to each other.

The remainder of this work is concerned with the formal implications of the matching principle for free-operant extinction and with an empirical examination of one of those implications.[2] The plan of what follows is to introduce the reader in chapter two to the methods and major theoretical issues characterizing work on the matching principle. In chapter three, the formal or logical properties of the matching principle will be outlined followed by a discussion of the implications of the matching principle for extinction and resistance to extinction. Chapter four will report on a line of research conducted for the purpose of testing one of the implications for resistance to extinction discussed in chapter three. Finally, chapter five will provide a brief summary of the work reported herein and a general discussion.

FOOTNOTES

[1] For example, the principle of reinforcement has been invoked in explanations of social and nonsocial behavior (e.g., Homans, 1974; Skinner, 1953), language acquisition (e.g., Skinner, 1957; Staats, 1968), personality (e.g., Lundin, 1974), psychopathology (e.g., Ullmann and Krasner, 1969), child development (e.g., Bijou and Baer, 1961) and organizational behavior (e.g., Luthans and Kreitner, 1975).

[2] If one is called for, the rationale for the present exercise (to quote Herrnstein in a slightly different context) is that the matching principle has "... a degree of verification that justifies an efforts to explore its logical properties at this point" (Herrnstein, 1974, p. 159) and that, moreover, "few now question its approximate accuracy" (Herrnstein, 1979, p. 486).

CHAPTER II

THE MATCHING PRINCIPLE

> The problem is to make sense of our intuition that action is affected by its consequences to the actor, with reward strengthening, or punishment weakening, the behavior that gives rise to them. Adaptiveness is elegantly accounted for if the pigeon does whatever is pleasurable (and therefore rewarding), avoids whatever is painful (and therefore punishing), and its pleasures and pains are so constituted that doing what comes naturally is by and large doing the right thing (pragmatically, if not morally). This is, of course, just the venerable hedonistic doctrine, which modern psychology has taken over notwithstanding its long and stormy history in philosophy. But if the doctrine is not to slip back into the circularity of which it was both accused and guilty, it must seek the concrete, functional relationships between reward and punishment, on the one hand, and the strength of behavior, on the other. Or, to use the parlance of psychology, the principle of reinforcement must either be stated objectively or not at all.
>
> R. J. Herrnstein, 1971, p. 399

Perhaps because of the yet-to-be articulated reasoning in the quotation above. R. J. Herrnstein (1961) designed and executed what has come to be regarded by behavior theorists as a fairly important experiment. Herrnstein trained hungry pigeons to peck at either of two lighted disks or "keys" on the wall of an operant conditioning apparatus. Key-pecks occasionally produced brief access to a tray of pigeon food. Variable-interval schedules of reinforcement made response-contingent food intermittently available by arranging a series of unpredictable minimum interreinforcement times ranging from several seconds to several minutes. Insofar as responses on one key had no effect on the alternate schedule, the two schedules were mutually independent. The pigeons had continuous access to the two alternatives and were free to distribute their pecks between them in any way they chose.

Fixing total within-apparatus reinforcement at no more than 40 per hour, the aim of Herrnstein's experiment was to examine how various distributions of reinforcement affected the allocation of key-pecks to the two alternatives. The results of the experiment were exceedingly noteworthy: although absolute response output varied considerably, Herrnstein found that all subjects allocated their responses to the two alternatives in the same proportion as the rates of reinforcement yielded by them. Hence, while the effect of a fixed amount of reinforcement on absolute frequency of response was variable, its effect on relative frequency of response was virtually *invariant*. Numerous replications of Herrnstein's groundbreaking experiment have been conducted using different experimental organisms (including human beings) (e.g., Schroeder and Holland, 1969; Shull and Pliskoff, 1967), reinforcers other than food (e.g., Shull and Pliskoff, 1967), aversive rather than appetitive reinforcers (e.g., de Villiers, 1974), more than two alternatives (e.g., Reynolds, 1963), reinforcement parameters other than rate (e.g., Catania, 1963) and alternative procedures for arranging the schedules (e.g., Stubbs, and Pliskoff, 1969). Virtually all of the foregoing confirmed Herrnstein's initial finding. Hence, its generality seems to be well established.

Methodological considerations

With rare exception, two experimental arrangements have been employed in matching research. In the first, the so-called two-key procedures, the subject is free to switch back and forth between two spatially disparate manipulanda each of which is associated with a schedule of reinforcement (de Villiers, 1977, p. 234). In the second, the so-called CO-key procedure, the subject may alternate between two reinforcement schedules programmed on the same key by emitting a changeover response on a second key, with each reinforcement schedule associated with a distinctive stimulus (de Villiers, 1977, p. 234). Many of the differences between these two procedures are of little technical importance but in research requiring precise measures of *local* response frequencies, or of the *time* allocated to a particular alternative, the CO-key procedures is clearly preferable. (See fig. 1.)

The changeover delay

A methodological contrivance in concurrent schedule programming that may be of theoretical interest is the so-called changeover delay (COD). Herrnstein (1961) found that, when given a choice between two variable-interval schedules, pigeons sometimes tended to distribute their responses between the alternatives indiscriminately or, that is to say, independently of the reinforcement rates associated with those alternatives. However, when a short delay (e.g., 1.5-2.0 seconds) was imposed between switching from one to a second alternative and delivery of reinforcement for responding to that alternative, the pigeons' distribution of responses matched, or were in proportion to, the relative rates of reinforcement arising from the alternatives. That delay, the COD, would therefore seem to be a precondition for demonstrating the matching relation.

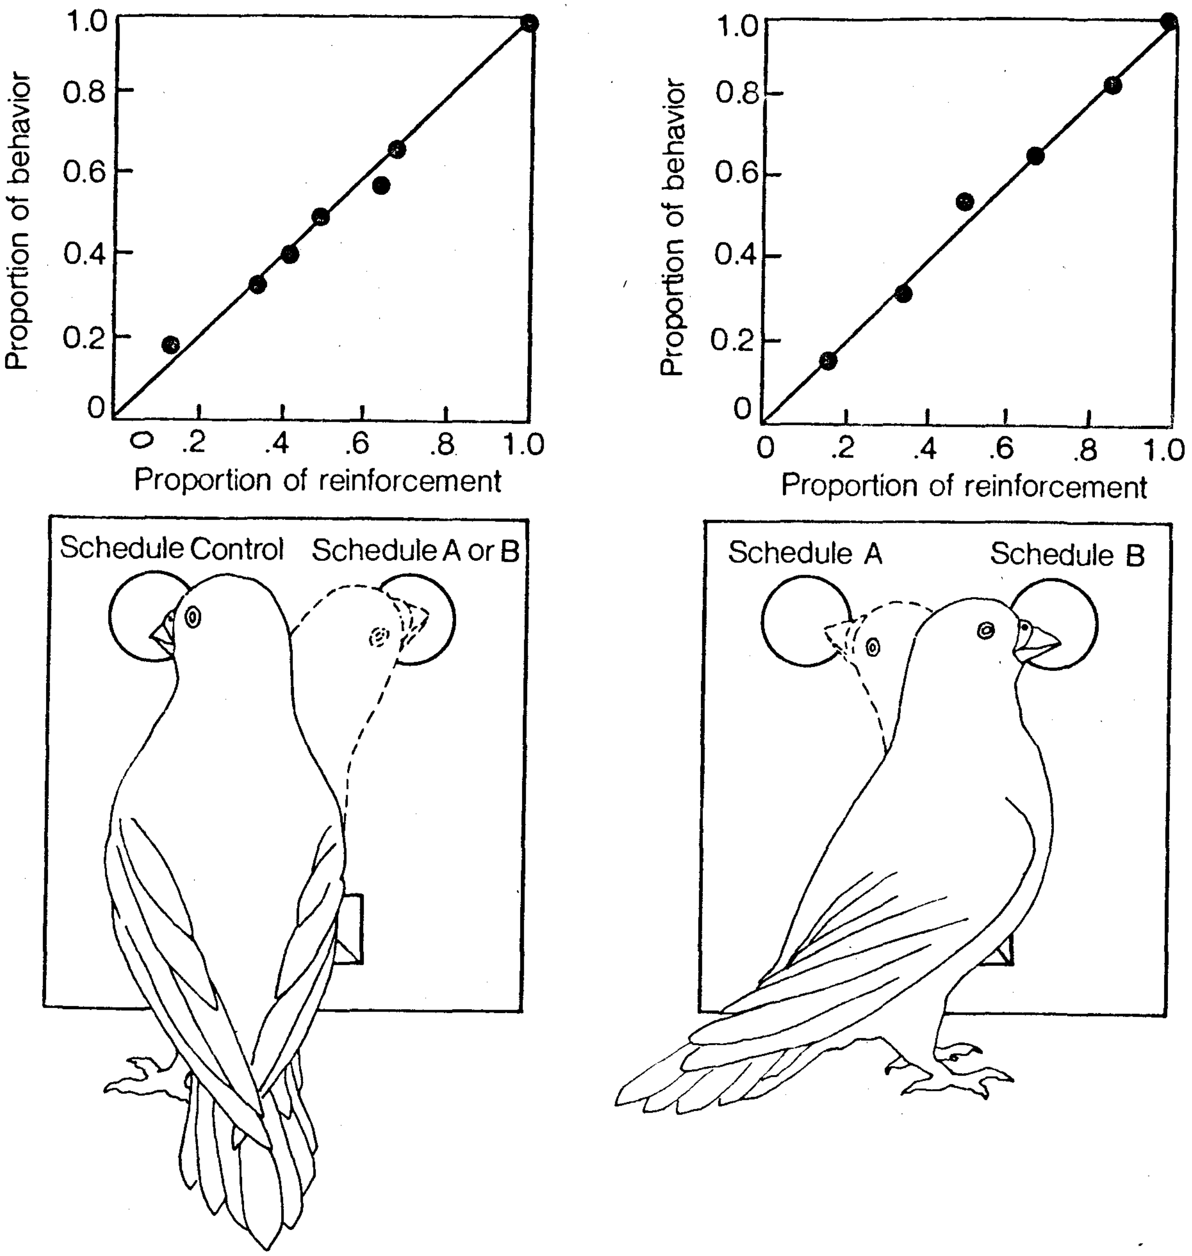

Fig. 1. Above: Proportion of behavior as a function of proportion of associated reinforcement (the diagonal is the matching line). Below: The CO-key (left) and the two-key (right) concurrent operants procedures. (Adapted from Herrnstein, 1970.)

The apparent necessity of a COD for obtaining matching has been viewed by some as possibly constituting a constraint on the generality of the matching relation. At least one researcher (e.g., Pliskoff, 1971) has suggested that matching is at least partly an artifact of the COD. In response, Herrnstein (1970) and Catania (1966) have argued that the COD is simply an experimental device for precluding the adventitous reinforcement of the switching response which, otherwise, might become the predominant response pattern. Further, de Villiers (1977, pp. 243-244) has pointed out that, if indeed the COD is the chief determining factor in matching, it follows that matching will occur only at certain CODs and not at others. In fact, it is generally the case that matching is obtained at all values of the COD beyond the required minimum (de Villiers, 1977, pp. 243-244). Of greater significance, some researchers have obtained matching without a COD in experimental subjects working at lower than standard drive levels. For example, Kulli (1977) obtained matching without a COD in pigeons run at close to their free-feeding weights. So too has Baum (1972).

According to Herrnstein (1970, p. 385), the COD most likely exercises its influence on concurrent responding through its effect on response definition. Paraphrasing Herrnstein (1970, p. 385), consider an experimental procedure which systematically reinforced alternation between two programmed alternatives such that responding first to the left alternative and then to the right was followed by food-reinforcement. Clearly, matching would not be expected. Responses would be distributed equally between the two alternatives in this kind of procedure but 100% of the reinforcement would arise from responses on the right alternative. Matching would not be expected here because, as Herrnstein (1970, p. 385) has pointed out, "Reinforcement for alternation is likely to give not two separate responses, but rather a single, albeit biphasic, response." The effect of such a procedure on behavior "... is thus not properly described by tallies of the individual response frequencies. Instead of so many left responses and so many right responses, there should be so many left-rights or the like."

This example is instructive insofar as response alternatives in standard operant experiments on choice are indeed situated near each other, thereby making it likely that some switching responses will be adventitiously reinforced. If such is the case, the practice of tallying responses to the individual alternatives in a choice procedure may inadvertently cut across the actual response classes being maintained by the reinforcement contingencies. In this respect, by making immediate reinforcement of the switching response impossible, the COD simply precludes the acquisition of unprogrammed and unwanted response classes.

However the COD exerts its influence over concurrent responding, it is clear that it is not a crucial determinant of matching. Indeed, insofar as matching has clearly been demonstrated without the benefit of a COD (e.g., Baum, 1972; Kulli, 1977), we may conclude that matching is fundamentally determined by the reinforcements associated with the respective target responses and, in particular, their value relative to each other.

The matching principle

When permitted to continuously choose between response alternatives of essentially equivalent topography, experimental subjects routinely distribute their responses between (or among) the alternatives in the same proportion as the reinforcements yielded therefrom. That is, choice comes to obey the matching equation, as follows:

$$\frac{B_1}{B_1 + B_2 + \ldots + B_n} = \frac{R_1}{R_1 + R_2 + \ldots + R_n} \qquad (1)$$

where B_1 through B_n enumerate the response alternatives and R_1 through R_n enumerate the reinforcements over some unit of time associated with each. According to equation (1), the relative frequencies of the alternative responses will tend to equal or "match" the relative rates of their associated reinforcements. That is, if 25% of a rat's within-apparatus reinforcement arises from pressing one of two available levers, then approximately 25% of the rat's total lever presses would be expected to occur on that lever (Herrnstein, 1979).[1] (See fig. 1.)

As noted above, Herrnstein (1961) was the first to demonstrate matching and a number of researchers have since that time successfully replicated that finding. Baum (1972), for example, conducted an experiment in which the subjects (pigeons) lived in the experimental apparatus and obtained all their food by pecking at either of two response keys continually available to them. Each response key was associated with a separate and independent variable-interval schedule of reinforcement and a 1.8 second COD was in effect throughout the experiment.

Baum systematically varied the number of reinforcers per session associated with the two keys. He found that the distribution of responses between the alternatives equalled or matched the proportion of reinforcement yielded by those alternatives for a broad range of relative reinforcement rates.

In another experiment, Baum (1974) autoshaped 20 wild and free-ranging pigeons to peck the keys of a makeshift operant conditioning apparatus installed in the attic of his house. Pecks were reinforced by occasional brief access to grain. The apparatus was constucted so as to permit only one pigeon at a time access to the response keys. No COD was programmed and the responses of the entire group of pigeons were treated as an aggregate. Baum found that the proportion of responses matched the proportion of reinforcers arising from them over a wide range of reinforcer distributions.

In a slight variation of the standard pigeon experiment, McSweeney (1975) trained pigeons to press either of two treadles for variable-interval reinforcement. A 2.0 second COD was in effect throughout the experiment. The rate of reinforcement yielded by one response alternative was held constant at 30 per hour while the rate of reinforcement yielded by the other was varied from a low of 15 to a high of 120 per hour. McSweeney found that relative frequency of responding in each alternative matched the associated relative rates of reinforcement over all experimental conditions.

Parameters of reinforcement other than rate have also been investigated by matching theorists. For example, Chung and Herrnstein (1967) examined the effects of various delays of reinforcement on choice in concurrent variable-interval schedules of reinforcement. Identical schedules were programmed to produce a "blackout" consequent upon the emission of a response after termination of the scheduled interreinforcement times. The blackout was followed by response-independent delivery of food. The duration of the blackouts was held constant at 8 seconds for one response alternative and systematically varied between 1 and 30 seconds for the other. Because the experimental subjects used in this study (pigeons) do not peck during a blackout, this procedure constituted a true delay of reinforcement. Chung and Herrnstein found that relative frequencies of response in this procedure were approximately proportional to the relative "immediacies" of reinforcement associated with the alternatives, with immediacy being defined as the reciprocal of delay of reinforcement.

Catania (1963) examined the effects of various magnitudes of reinforcement on responding in single and concurrent variable-interval schedules or reinforcement. He reported that magnitude of reinforcement had little effect on average response frequencies in single-schedule procedures. However, when two response alternatives and two magnitudes of reinforcement were made available, response distributions tended to match the relative magnitudes of the respective reinforcements. Using a time rather than a response allocation procedure, Brownstein (1971) has reported similar findings.[2]

Matching and molecular maximizing

As de Villiers (1977, p. 256) has recently noted, empirical matching appears to be a robust, general product of continuous choice procedures. But what kind of process is matching? As described by Herrnstein (1970, 1979), matching is a molar phenomenon having to do with the covariation of the relative rates of behavioral and environmental events. However, some writers - among them, Shimp (1966, 1969) and Mackintosh (1974) - have argued that matching is not a fundamental process but is, rather, a by-product of a molecular or momentary maximizing process. Mackintosh's (1974, p. 194) position below is fairly typical of those articulated by other molecular maximizing advocates.

> Subjects on concurrent schedules tend to respond at any moment to that alternative offering the highest probability of reinforcement ... The crucial point about concurrent (VI, VI) schedules is that both VI schedules continue to run while the subject is responding on one of the alternatives. Since, on any interval schedule, the probability of reinforcement increases with the time since the last reinforcement, it must necessarily be the case that the probability of reinforcement becoming available on the other alternative will eventually surpass the probability of reinforcement for continued responding to the same alternative. Hence, a shift to the other alternative is in accordance with a maximizing principle.

This line of argument is supported by research on the effects of concurrent interval schedules which do not run simultaneously (e.g., Killeen, 1970; Fantino and Duncan, 1972). Under this experimental condition, subjects tend to respond exclusively to the alternative associated with the shorter mean interreinforcement time, thereby maximizing rate of reinforcement. Additional support has been provided by an experiment conducted by Shimp (1966) in which reinforcement was scheduled probabilistically for choices in a discrete-trials two-alternatives procedure. An interval contingency was used to insure that a reinforcement set up for a particular alternative remained available to be taken eventually by the subject. Hence, the probability of reinforcement associated with one alternative increased while the subject worked the other. Shimp found that sequential changes in choice probability between scheduled reinforcements tracked differences in reinforcement probability for each programmed choice. Based on these findings, Shimp concluded that the overall matching relationship normally observed in concurrent interval schedules results from a more fundamental molecular maximizing process. In Shimp's view, matching is therefore best characterized as a derivative phenomenon.

The validity and viability of Shimp's analysis is, however, compromised somewhat by the fact that two subsequent studies (i.e., Herrnstein, 1970; Nevin, 1969), employing essentially the same experimental procedure, failed to replicate Shimp's initial finding. Although both Herrnstein (1970) and Nevin (1969) obtained matching, neither found any systematic relationship between sequential changes in choice probabilities and local probability of reinforcement. In a more recent study, Nevin (1979) reaffirmed his original analysis.

Heyman (1979) recently examined performances in concurrent variable-interval schedules for the purpose of detecting relationships between maximization strategies and molar matching. He found that matching was associated with haphazard and virtually random changeovers between the programmed alternatives. As Herrnstein (1979, p. 491) has recently noted, Heyman's results "seem to be conclusive evidence against the momentary maximizing hypothesis."

Other conceptions of reinforcement maximizing remain possible, however, and indeed, Herrnstein and Loveland (1975) have advanced a model of maximization based on the logic of the matching principle. They posit that, given a choice between two variable-interval schedules, subjects will strive to adjust their behavior so as to at least approximately equalize the ratio of responses to reinforcements across the two alternatives. This would result in matching between relative response frequencies and relative rates of reinforcement. According to Herrnstein and Loveland's maximization model, experimental subjects tend always to emit that response which is associated with the shortest response-reinforcement ratio. In doing so, however, the subject drives that alternative's response-reinforcement ratio up to a value exceeding that of the other alternative at which point the subject switches. It follows from this analysis that switching back and forth between alternatives is centrally dependent upon the subject's ability to affect response-reinforcement ratios. In situations where the experimenter fixes these ratios, as in the case of choice between ratio schedules of reinforcement, exclusive preference for the alternative with the shortest response-reinforcement ratio would be expected. This prediction has been empirically confirmed by Herrnstein and Loveland (1975).

It is not altogether clear that the matching principle, as distinct from theories of probability matching, is in any important sense at odds with the general logic of maximizing. Indeed, in the sense in which it is circumscribed by Herrnstein and Loveland (1975), "matching and maximization are just two words describing one outcome (p. 116)." As we saw above, the only kind of choice performance in concurrent ratio schedules that is consistent with the matching principle is reinforcement maximizing (Herrnstein and Loveland, 1975; Rachlin, 1976, pp. 565-567).

The matching principle is more at variance with the theory of probability matching (e.g., Estes, 1964) according to which subjects choosing between alternatives are expected to match reinforcement probabilities. Hence, in choosing between an alternative that reinforced every 5th response and one which reinforced every 10th response, subjects would be expected to choose the latter alternative half as many times as the former. In contrast, the matching principle predicts exclusive preference for the former alternative (e.g., Herrnstein and Loveland, 1975), a strategy which maximizes reinforcement.

The theory of probability matching is an outgrowth of statistical models of the conditioning process (e.g., Estes, 1959, 1964). According to that theory, the probability of some target response is determined by the proportion of stimulus elements previously conditioned to that response which are, in fact, prevailing on the subject at the time response probability is assessed. For probability matching theory, the occurrence of a reinforcing event simply conditions to the target response those stimulus elements sampled on a given trial. And, it is presumed that the more frequent the reinforcement associated with a target response, the greater the number of stimulus elements conditioned to it. In situations in which two responses are reinforced, some of the prevailing stimulus elements are conditioned to the more infrequently reinforced alternative and others are conditioned to both, thereby preventing the probability of the more frequently reinforced alternative from attaining unity.

The theory of probability matching and the matching principle differ in the way in which reinforcement is presumed to interact with behavior. The principle of reinforcement inherent in the matching principle is relativistic while its approximate counterpart in the probability matching model is, in a sense, implicitly absolutistic. Hence, when, in the presence of a set of stimulus elements a response yields a reinforcer, an increment in response probability _must_ result. Since both alternatives in a probabilistically reinforced choice situation are initially reinforced, the subject simply cannot, according to probability matching theory, develop exclusive preference for the richer of the two alternatives in that situation. In contrast, whether or not a response-produced reinforcer results in an increment in response probability is, according to the matching principle, focally dependent upon its value relative to the value of other reinforcers in the situation. Where the density of "other" reinforcers is great, it may be virtually without value and hence, of no consequence with respect to probability of response. The subject is therefore free, in such situations, to interact exclusively with the richer of the available alternatives thereby _maximizing_ reinforcement, a strategy which, as noted above, is the only outcome which is consistent with the matching principle.

The task of assessing the relative merits of these competing analyses is unfortunately complicated by the fact that the studies supporting them have characteristically employed different experimental procedures and different kinds of animal subjects. Most experiments supporting the matching principle have been on pigeons, rats, monkeys and, occasionally, human beings (de Villiers, 1977), while studies supporting probability matching have not infrequently been on "simpler" organisms such as fish or reptiles (Bitterman, 1965).

The study of probability matching may eventually assist in determining the phyletic generality of the matching principle. With respect to this possibility, Bitterman (1965) has suggested that maximization may hold only for relatively complex organisms, with probability matching being the rule for simpler creatures such as fish. Nonetheless, there are formidable nonempirical reasons for questioning the adaptive efficacy of the probability matching strategy for organisms at _any_ phyletic level. As Herrnstein and Loveland (1975, p. 107) have lucidly noted, probability matching

> ... is a psychological oddity, for it may be grossly at odds with the rational course of action to an extent rarely seen in nature. If the subject places value on the payoff itself, rather than on the sheer matter of guessing correctly, he is ill advised to match probabilities. Instead, he should always choose the higher probability alternative, a strategy which maximizes winnings per response.

Time matching

Several researchers (e.g., Brownstein and Pliskoff, 1968; Baum and Rachlin, 1969) have suggested that allocation of time, or time matching, may be more fundamental than response matching. Both of the foregoing research teams have demonstrated time matching in experimental situations in which no instrumental response was said to be required. In both situations the subjects had simply to choose which of two stimuli (and their corresponding rates of response-independent reinforcement) they preferred to be in the presence of.

Both studies reported matching of relative time allocation to the relative rates of reinforcement associated with the two alternative stimuli.

Using a standard CO-key procedure, Cantania (1963) and Fantino (1970) have obtained simultaneous response and time matching in pigeons working for food. Similar results have been reported by Shull and Pliskoff (1967) using rats as subjects and brain stimulation as the reinforcer.

In making their case for the primacy of time matching, Baum and Rachlin (1969) pointed out that, when working on response-dependent schedules of reinforcement, pigeons tend to peck at a fairly constant frequency, with the majority of interresponse times falling within the range of .3 to .5 seconds. It follows from this putative fact that overall sessional response frequencies result, not from graded fluctuations in local response frequencies but rather from periods of nonresponse between bursts of responses occurring at a fairly constant frequency (Blough,1963). Presumably then, time engaged in responding is a more fundamental measure of preference, with frequency of responding being regarded as a derivative of time matching (cf. Baum and Rachlin, 1969).

There are both advantages and disadvantages to viewing choice in the usual concurrent schedule procedure as time allocation. One conspicuous disadvantage is that measures of time spent in the presence of one of several discriminative stimuli invariably includes time spent in activities other than the target response, thus complicating the task of assessing the time actually spent engaging in that response.

As a result, relative time allocation would be expected to match relative rates of reinforcement only in those situations where time spent in other activities was essentially fixed across all alternatives. Ironically, this consideration has prompted Baum and Rachlin (1969) to suggest that relative frequency of response indeed may be the best measure of the relative amount of time spent in an activity. This, of course, assumes that the time required to emit the instrumental response is fairly constant across all instances of that response. While this is surely the case with such responses as the pigeon's key-peck or the rat's lever press, it may not be the case with more complex response forms. Moreover, empirically equating response and time allocation accounts of choice makes it very difficult indeed to fully distinguish between them.

Another body of data bearing on the question of the primacy of time matching has been gathered by researchers interested in choice between different types of reinforcement schedules. Since most of the initial research on the matching principle involved choice in concurrent variable-interval schedules, many researchers have been prompted to assess its relevance for experimental arrangements of a slightly different type. For example, Herrnstein (1970) examined choice in concurrent variable-interval, variable-ratio schedules. He obtained response matching to relative reinforcement rates but did not obtain time matching. In a similar study, La Bounty and Reynolds (1973) examined choice in concurrent fixed-interval, fixed-ratio schedules. They obtained approximate response matching but not time matching. Both of these findings indicate that time matching is clearly not more fundamental than response matching since the latter can hold without the former.

The manner in which the time versus response matching issue is characteristically discussed almost always presupposes that these two perspectives are fundamentally incompatible. However, one might quite reasonably argue that time and response allocation are both important, partially inseparable properties of choice. While this is most obviously the case in procedures in which subjects repeatedly engage in a target activity for a period of time, it is no less true for time allocation data gathered in experimental situations which are said to require no instrumental response. In fact, there are fairly obvious reasons for questioning the claim that no instrumental response is required in time allocation procedures. To be sure, concurrent response-independent schedules of reinforcement neither require nor encourage a significant frequency of a pre-established target response. Neither is any one reinforcer dependent upon the occurrence of any one response. As a result, choice responses in these types schedules may not be easily related to measures of reinforcement in ways which are regarded as useful. But this is clearly beside the point. What is important is the fact that subjects performing on either response-dependent or response-independent concurrent schedules must behave in ways that are ordinarily taken to be evidence of choice, with the particular form of that behavior being determined by the explicit and implicit reinforcement contingencies bearing on the subject. Without such differential behavior, time matching could not be obtained.

Although of some methodological interest, it is not clear that the time versus response matching issue is one of real and substantial theoretical significance. Arguments advanced in support of one or the other position almost invariably imply, probably erroneously, that one of the two positions is the fundamental basis of choice. One might just as easily argue, however, that, in its purest distillation, choice is inherently neither. Nevertheless, it is surely true that choice frequently and indeed, quite naturally, manifests itself in behavior over time. As such, response and time allocation are best seen as two fairly useful measures of choice applicable to two different types of experimental procedure. To the extent that behavior is not instantaneous, both procedures clearly involve instrumental responding and allocation of time, although not always in forms which are convenient or useful.

Single-schedule performances

As written in equation (1) (above), the matching principle predicts relative response frequencies in concurrent schedules of reinforcement. However, equation (1) can be brought to bear on absolute response frequencies in such schedules by setting the sum of the various Bs in equation (1) to K, a constant. The matching principle is then written

$$\frac{B_1}{K} = \frac{R_1}{R_1 + R_2 + \ldots + R_n}$$

Multiplying both sides of the foregoing by K yields

$$B_1 = \frac{KR_1}{R_1 + R_2 + \ldots + R_n} \qquad (2)$$

where B_1 is the frequency of some pre-established target response, R_1 through R_n are the reinforcements yielded by all of the possible alternative responses and K is the sum of all behavior. According to equation (2), the absolute frequency of the target response, B_1, is proportional to the relative rate of its associated reinforcement. Absolute response frequencies for alternatives B_2 through B_n are described by analogous equations. Equation (2) is widely supported by the existing concurrent performances data (Herrnstein, 1970; Catania, 1966).

While successful in accounting for absolute response frequencies in concurrent schedules, equation (2) seems to run into difficulty in accounting for response frequencies in single-response or single-schedule situations, for when R_2 through R_n equal zero, that equation reduces to

$$B_1 = K \qquad (3)$$

according to which the target response, B_1, goes to asymptotic frequency irrespective of the specific rate of reinforcement yielded by it. Equation (3) is empirically incorrect (see, e.g., Herrnstein, 1970, p. 388; Fantino and Logan, 1979, p. 223).

However, equation (3) is based on the dubious assumption that the absence of a programmed alternative to B_1 implies that the subject has no viable alternatives to choose among. But clearly, the subject can, at the very least, choose between emitting B_1 and not emitting B_1, thus sparing the effort associated with that response. And in most experimental settings, alternatives such as preening, scratching, resting, defecating and so on are clearly available to the subject. Hence, it may be reasonably assumed that responding always takes place in a context of extraneous, spontaneous, or, simply, alternative reinforcement (Herrnstein, 1970, pp. 388-389). The matching equation for single-schedule performances may therefore be written as follows:

$$B_1 = \frac{KR_1}{R_1 + R_e} \qquad (4)$$

where B_1, K and R_1 are the same as in equation (2) above and R_e summarizes all sources of reinforcement <u>not</u> associated with B_1 (i.e., "background" reinforcement, "all the other reinforcers that a subject brings with itself or finds in the experimental setting" (de Villiers, 1977, p. 257).

Herrnstein (1970) assessed equation (4) using data from two experiments which examined the relationship of rate of reinforcement to the frequency of key-pecking in pigeons working on single schedules. The first of these was conducted by Catania and Reynolds (1968). They exposed pigeons to variable-interval schedules ranging in value from eight to 300 reinforcements per hour. Response frequencies were monitored throughout. Depending on the pigeon, equation (4) accounted for 77 to 98.8% of the data variance. With K and R_e set, respectively, at 66.3 responses per minute and 7.3 reinforcements per hour, equation

(4) also accounted for 91.3% of the data variance of the group response frequencies.

In the other study, Chung (1966) exposed pigeons to a tandem fixed-ratio 1, fixed-interval X schedule of reinforcement, where "X" represents a duration timed from the first key-peck following delivery of a scheduled reinforcer. This response starts a fixed-interval clock and the first response occurring after the clock has timed out is reinforced. Chung examined the effects of various FI durations (and their associated rates of reinforcement) on frequency of key-pecking. Herrnstein (1970) has shown that equation (4) accounts for 94.7% of the data variance in the group response frequencies of Chung's pigeons.

Equation (4) has also been tested against data collected for the purpose of determining the behavioral effects of parameters of reinforcement other than rate. For example, Davenport, Goodrich and Hagguist (1966) examined the effects of various magnitudes of reinforcement on response frequency in macaque monkeys working on a variable-interval 1-minute schedule of reinforcement. The number of food pellets delivered to subjects per reinforcement was systematically varied and response frequencies were monitored. Equation (4) accounts for between 90.1 to 99.9% of the variance in the individual response frequencies (de Villiers, 1977).

Guttman (1954) investigated the effects of various concentrations of response-contingent sucrose and glucose on frequency of lever-pressing in rats. Concentrations were varied between 2 and 32% and delivered on a variable-interval 1-minute schedule. Response frequencies were monitored throughout. Although the rats responded more rapidly for sucrose than for glucose when concentrations were comparable, equation (4) nonetheless provides a remarkably good fit to the data, accounting for 93.7% (sucrose) and 98.7% (glucose) of the data variance (de Villiers and Herrnstein, 1976).

Pierce, Hanford and Zimmerman (1972) studied the effects of delay of reinforcement on frequency of response in a single-alternative procedure. Rats were trained to work for food on a variable-interval 1-minute schedule of reinforcement. Delivery of food was delayed from .5 to 100 seconds over experimental conditions and response frequencies were monitored throughout. A signal light in the apparatus was lit for the duration of the delay. Using immediacy of reinforcement as the measure of reinforcement associated with the target response, equation (4) accounts for 96.1% of the variance in the average lever-pressing frequencies in Pierce et al.'s subjects (de Villiers, 1977).

De Villiers and Herrnstein (1976) examined some forty instrumental conditioning experiments on pigeons, rats, monkeys and human beings. With rare exception, they found that the obtained functional relationships between some measure of response level (response frequency, running time, etc.) and some parameter of reinforcement tended to closely approximate the function specified by equation (4). That is, strength of responding was approximately proportional to associated relative reinforcement.

Conclusion

As Herrnstein (1979, p. 493) has recently pointed out, the matching principle "is a simple, quantitative statement of the law of effect, within the limits set by the available data." As noted earlier, few researchers now question its approximate accuracy (Herrnstein, 1979, p. 486). Hence, no further effort will be made to demonstrate that fact here (see de Villiers, 1977, for an exhaustive review of the matching literature).

The territory circumscribed by the matching principle is, as Herrnstein (1970, p. 399) has succinctly noted, "sizeable, expandable, and susceptible to precise measurement." As a highly successful framework within which the data of operant conditioning can be elegantly and usefully ordered, it would be of some value, on the grounds of parsimony alone, to demonstrate that the matching principle applies, at least partially, to the data of operant extinction. It is toward this task that we now turn.

FOOTNOTES

[1] Herrnstein's is not, of course, the only current quantification of the law of effect. Several other writers (e.g., Baum, 1973; Catania, 1973) have advanced similar models. However, Hernnstein's model is clearly the most extensively studied and, for our purposes, the most general.

[2] Not all concurrent operants research is precisely consistent with the matching principle. Indeed, there is a body of research literature on deviation from matching (cf. de Villiers, 1977) which is concerned in large part with the exact shape of the function relating response frequencies to rates of reinforcement. However, since we are concerned here with matching as an *approximate* principle, no further discussion of this literature is required here.

CHAPTER III

MATCHING AND EXTINCTION

Since Pavlov (1927), psychologists have routinely distinguished between conditioning, or the acquisition of conditioned responses, and extinction, the elimination or attenuation of conditioned responses. Theoretical accounts of extinction have not infrequently invoked processes unlike those posited to explain the acquisition and maintenance of conditioned reponses. Thus, Pavlov (1927) and Hull (1943) invoked inhibitory processes to account for the response decrements observed in extinction. While commitment to this practice has not been unanimous (see, e.g., Guthrie, 1952; Estes, 1955), contemporary accounts of extinction and extinction-related phenomena tend to emphasize factors which at most play only secondary or marginal roles in the explanation of acquisitions and maintenance, e.g., frustration and other emotional responses (Amsel, 1967; Skinner, 1938), long-term memory processes (Capaldi, 1971), proactive interference (Gleitman, 1971; Spear, 1971) and cognitive dissonance (Lawrence and Festinger, 1962).

One straightforward justification for the conditioning-extinction distinction is rooted in the contrasting operational bases of these two phenomena. By definition, conditioning and extinction refer to the behavioral effects of, respectively, the presentation and omission of reinforcing events.

In contrast, the putative basis of the need for separate theories of extinction is multifaceted, sometimes convoluted, and anchored in part in the various systematic orientations of learning theorists.

However, a common thread running through many accounts of extinction has to do with the presumed theoretical inadequacy of steady-state and acquisition processes in explaining the facts of extinction. For example, Mackintosh (1975, p. 405) has argued that there is nothing inherent in the principle of reinforcement (or the law of effect) which can explain why omission of a reinforcing event should result in a decrement in responding (see the introductory quotation, chapter one).

Mackintosh's reservation about the law of effect in this regard is either cogent or not depending on the particular assumptions made about the nature of experimental extinction and the process of reinforcement. From the logical perspective of modern relativistic formulations of the law of effect - and in particular, the matching principle (Herrnstein, 1970; de Villiers, 1977) - Mackintosh's reservation is partly or wholly without substance. Indeed, from such a perspective, the various phenomena of experimental extinction necessarily occur within the general framework circumscribed by the matching principle and are, in large part, manifestations of the ongoing tendency to allocate and re-allocate responses in proportion to the relative value of their consequences.

Before going into greater detail about the implications of the matching principle for free-operant extinction, it may be instructive to first briefly discuss the formal or logical properties of that principle. It is to this task that we now turn.

Formal properties of the matching principle

A formal analysis of what is logically entailed by the matching principle begins with a delineation of the reinforcement principle inherent in it and with the specific interpretations placed on the R_e and K parameters in equation (4), which is repeated here for the convenience of the reader.

$$B_1 = \frac{KR_1}{R_1 + R_e} \tag{4}$$

As noted earlier, K is simply the sum of the various Bs (B_1 through B_n) in equation (1). That is, it is a measure of total behavior. Also, however, K is the asymptote of whatever target response is under observation. Or rather, it is total behavior scaled in units commensurate with the target response (Herrnstein, 1974, p. 162). Hence, the K parameter is simply "the modulus for measuring behavior" (Herrnstein, 1974, p. 163), the sole influence on the numerical size of which is the choice of the target response. Herrnstein (1974) has shown that a formal implication of the matching relation is that K must remain invariant across all qualities and quantities of reinforcement or drive as long as the topography of the target response does not change.

On the surface R_e, like K, is first and foremost a free parameter that must be inferred from data, but as noted previously, it has a definite empirical reference. R_e denotes the sum of the "background" reinforcement yielded by responses other than the target response (i.e., the various B_es), or reinforcement occuring response-independently or spontaneously (Herrnstein, 1970, p. 389; de Villiers, 1977, p. 257).

The fundamental efficacy of R_e is based on the putative fact that there is no such thing as a single, isolated schedule of reinforcement. As Herrnstein (1970, p. 388) reminds us below, organisms always have alternatives available to them.

> Even in a simple environment like a single-response operant-conditioning chamber, the occurance of the response is interwoven with other, albeit unknown, responses, the relative frequencies of which must conform to the same general laws that are at work whenever there are multiple alternatives. In fact, it seems safe to assume that all environments continually demand choices in this sense, even though in many cases the problem of identifying and measuring the alternatives may be insoluble. The problem is, however, the experimenter's, not the subject's. No matter how impoverished the environment, the subject will always have distractions available, other things to engage its activity and attention, even if these are no more than its own body, with its itches, irritations, and other calls for service.

In addition to its intuitive basis, experimental work bearing on R_e, though not extensive, suggests that R_e "behaves in an orderly and empirically predictable fashion when varied" (Fantino and Logan, 1979, p. 225). For example, Rachlin and Baum (1972) maintained pigeons on a variable-interval 3-minute schedule of reinforcement. R_e, reinforcement not yielded by key-pecks, was manipulated by varying the amount of "free" or response-independent food given over and above delivery of response-dependent food. As predicted by the matching principle, Baum and Rachlin found that frequency of key-pecks was inversely related to the amount of free food given.

From inspection, it is clear that R_e influences response allocation through the role it plays in the principle of relative reinforcement inherent in equation (4), i.e.

$$\frac{R_1}{R_1 + R_e}$$

As R_e grows in value, progressively less of K (total behavior) is allocated to the target response and progressively more is allocated to the various B_es. Reinforcement, then, simply governs the distribution of responses over available alternatives (Herrnstein, 1979) and it achieves this end whether, by any absolute standard, it is large or small in quantity. Thus (and bearing in mind the invariance of K), when an increment in reinforcement increases the frequency of a target response other, alternative responses must decrease in frequency. Likewise, when a decrement in reinforcement results in a decrease in the frequency of a target response other, alternative responses must increase in frequency. In either case, the net change in behavior is zero because, as Herrnstein and Loveland (1974) have pointed out, "each increment in behavior caused by reinforcement is precisely counterbalanced by decrements in the behavior's competitors, ... and *vice versa*." Which is to say, total behavior, K, remains invariant.[1]

As noted above, the frequency of a target response is, within the logical framework of the matching principle, controlled by its relative reinforcement; that is, the fraction of total reinforcement yielded by it. The effect on behavior of a change in its reinforcement depends on the level of reinforcement already associated with it, with the effect of such a change being inversely related to the initial level of reinforcement. Thus, responses commanding a large fraction of total reinforcement are less sensitive to an increment in reinforcement than those commanding a less substantial fraction (Herrnstein, 1979). The effect of a reinforcer on behavior is therefore not invariant: it depends, rather, on the context of reinforcement in which it occurs (Donahoe and Wessels, 1980, pp. 147-148).

In way of a general summary, the matching principle acknowledges that organisms in virtually any setting engage in a variety of behaviors in order to produce valued consequences. And it may be assumed that instrumental behavior occurs as a result of an explicit or implicit choice between or among alternatives (cf. Hamblin and Miller, 1977) based on the relative value of their consequences. That is to say, the matching principle suggests that organisms assign value to behavioral consequences based, not on some absolute scale of reinforcement but, rather, relativistically or with reference to the prevailing context of reinforcement. Hence, a consequence which in one setting might be of considerable value may, in another setting, be virtually without value. Likewise, a consequence which is ordinarily meager in value may take on considerable value in extraordinary settings. Finally, the matching principle implies that total behavior (B_1 through B_n or, simply, K) is invariant. Changes in the frequency of some target response must therefore be accompanied by changes in the frequencies of alternative responses such that the net change in behavior is zero.

Implications for extinction

Let us begin this section by examining, from the vantage point of the matching principle (equation (4) above), precisely what happens when an instrumental response is first conditioned and then, subsequently, extinguished. During acquisition, B_1, a target response, is raised in frequency by virtue of its association with R_1, a programmed quantity of reinforcement. Concomitantly, the various nontarget responses (or B_es) associated with the sources which make up R_e, compensatingly assume a relatively low frequency as a consequence of R_e's small value relative to R_1. At asymptote, the subject of the foregoing conditioning procedure may be seen to be choosing between B_1 and the various B_es in a manner consistent with the matching principle. If, for whatever reason, R_1 was abruptly halved, we would expect behavior to re-distribute itself among the alternatives so as to re-establish the relationship called for by equation (4). The same would be expected should R_1 be halved again or indeed, repeatedly re-halved. Now, should we expect anything strikingly different if R_1 is omitted altogether? From the perspective of the matching principle, the answer is no: because the various B_es would then command virtually all available reinforcement, we would expect a decrement in B_1 and an increment in the B_es such that the adjusted frequencies again obey the matching principle.[2] (Because of the differing topographies of B_1 and the B_es, empirical matching would, of course, not be expected.)

In contrast to inhibitory accounts of extinction, of noteworthy interest here is this: the foregoing account invokes no special response-decrementing process to account for the decline undergone by

B_1 in the above example. Nor need it. Given a discriminated change in reinforcement, "behavior simply redistributes itself so as to re-establish matching" (Herrnstein, 1979, p. 493). We may therefore surmise (at least in way of a hypothesis) that the changes in response frequencies observed in extinction reflect substantially nothing more or less than the fundamental predilection to allocate and re-allocate behavior in proportion to its associated relative reinforcement.

It may be tempting to conclude that the matching principle treats extinction as simply another case of acquisition. Such is not entirely the case, however, for from the perspective of that principle, whether an experimental procedure is to be identified as acquisition or extinction is determined <u>entirely</u> by the experimenter's choice of a target response. In the case above (where reinforcement was omitted for B_1) if, rather than B_1, one of the B_es had been selected as the target response, the procedure would have been identified as acquisition, not extinction. Although the procedure in question resulted in a decrement in the frequency of B_1, the various B_es <u>increased</u> in frequency. The upshot of all of this is that acquisition and extinction are not, in theory, fundamental and mutually exclusive behavioral processes. They are, rather, two experimenter-dependent facets of essentially the same experimental scenario, namely, response allocation as guided by the matching principle.[3]

Theoretical status

According to the foregoing analysis, acquisition and extinction are little more than experimental procedures for shifting behavior back and forth between experimenter-defined and undefined response categories. In either case, observed changes in the allocation of responses are yielded solely by changes in the relative reinforcement values associated with the various responses.

More generally, the foregoing analysis establishes a theoretical superstructure within which organisms may be presumed to be working when interacting with the particulars of acquisition and extinction procedures. In casual language, it tells us what the experimental subject is doing, or trying to do, when coping with the exigencies of such procedures. However, the matching principle is not, in any orthodox sense, a theory of extinction for by implication it questions the very existence of extinction as a separate and distinct process.

Having articulated this noteworthy qualification, we may nevertheless discuss the relationship of the foregoing matching-based interpretation of extinction (hereafter referred to as response reallocation theory) to so-called interference theories of extinction. According to Mackintosh (1974, p. 410), interference theories are those which posit that "extinction establishes some set of responses whose occurrence competes with, and eventually prevents the appearance of, the originally reinforced response." As Tarpy and Mayer (1978, p. 127) have recently pointed out, theories of this sort have a great deal of appeal because it is clear that subjects indeed <u>do</u> engage in other behaviors during extinction (e.g., Bindra, 1961; Jones and Bridges, 1966) and because, when subjects explicitly do not perform alternative behaviors during extinction, the rate of decline of the target response is diminished (e.g., Davenport, 1964).

Of course, the mere appearance of new responses during the extinction of a target response is not proof that such responses compete with and eventually displace that response. Clearly, the appearance of such responses could be a result of the extinction of the target response rather than its cause (Mackintosh, 1974, p. 411). Proponents of interference theories must demonstrate that there indeed is such a causal relationship. And, moreover, they must explain where the competing responses come from and why they become strong enough during extinction to compete with and displace the target response (Mackintosh, 1974, p. 412). As Mackintosh (1974, p. 412) points out below, classical interference theorists were notably unsuccessful in accomplishing these aims.

> It is not at all clear that the theories advanced by Guthrie (1935) and Wendt (1936) succeed in doing this, and Guthrie (1935, pp. 69-70) was willing to accept that the original response might have to be first inhibited before the competing response could be securely established. Even then, it is not clear what process is responsible for the later establishment of the competing response, unless it is simply a dominant response in the experimental situation that was suppressed by the appearance of the CR or instrumental response during initial acquisition.

Perhaps because of these difficulties, frustration theory (Amsel, 1958, 1967) is now the only seriously accepted interference theory of extinction (Mackintosh, 1974, p. 413; Tarpy and Mayer, 1978, p. 127). According to frustration theory, subjects come to expect reinforcement during acquisition. The subsequent omission of reinforcement in the presence of that expectation results in frustration (and frustrative behavior) and a host of escape and avoidance behaviors which are reinforced by frustration reduction (Tarpy and Mayer, 1978, p. 127).

These behaviors are said to compete with and displace the originally reinforced response (Mackintosh, 1974, p. 413).

Although not fully successful as a theory of extinction (see, e.g., Tarpy and Mayer, 1978, pp. 135-136), frustration theory succeeds in formally accounting for the original or initial strength of potentially competing responses. However, because frustration is only a transitory phenomenon (cf. Mackintosh, 1974, p. 413), frustration theory cannot readily account for the more permanent cessation of responding routinely produced by extinction procedures.

The present account of extinction, response re-allocation theory, also accounts for the origin and strength of potentially competing responses; they are the various B_es (nontarget responses) which are associated with R_e (background reinforcement) in equation (4). Although the value R_e is apt to be relatively small during acquisition, it must assume a relatively large value when reinforcement is omitted for B_1, the target response. Hence, the various B_es must increase in frequency up to the point where the increments undergone by them counterbalance the decrement undergone by the target response. Moreover, unlike classical interference theory (e.g., Guthrie, 1935), the process responsible for the establishment of the putative competing responses is clearly articulated in the present account; it is response allocation or matching.

Although response re-allocation theory readily satisfies the theoretical requirements outlined above, it is nevertheless mistaken to conceive of it as an interference theory. To be sure, the present account clearly predicts an increment in the frequency of nontarget responses during extinction. However, whether or not these responses physically compete with the target response is of little or no consequence. From the perspective of response re-allocation theory, the frequency changes observed in target and nontarget responses over the course of extinction result from the change in relative reinforcement values arising from the extinction operation, and nothing more.

How, then, is the present account to be characterized? As noted above, response re-allocation theory is not a theory of extinction in any ordinary sense and indeed, whether or not it constitutes a complete alternative to such theories remains to be seen. In succinct language, the present account, if true, specifies when and to what extent an organism in a particular environment is interested in interacting with an alternative and why, when that environment changes, it may be no longer interested. As such, it is simply a motivational account of response allocation and re-allocation brought to bear on the data of experimental extinction.

Matching and resistance to extinction

It was argued in the previous section that a theoretical account of experimental extinction is formally entailed in the matching principle. Like theories that have gone before it (e.g., Pavlov, 1927; Guthrie, 1935; Hull, 1943), that theory attempts to concisely articulate the essential or fundamental nature of extinction. Such theories aim to detail the processes or mechanisms that are presumably at work when a response undergoes a decline in extinction. However, such theories have not routinely shed light on the fact that subjects take more or less time to achieve extinction criteria and that, most particularly, subjects under some experimental conditions evidence a tendency to resist the response-decrementing effect of nonreinforcement.

The plan of attack of this section is to show first how the matching-based account of extinction sketched in the previous section may be conceptually brought to bear on resistance to extinction effects and then demonstrate how it might plausibly account for resistance to extinction data. The present account will then be briefly contrasted with two prominent theories of extinction, Amsel's (e.g., 1962) frustration theory and Capaldi's (e.g., 1971) sequential aftereffects theory.

R_e and resistance to extinction

According to the above matching-based account of extinction, extinction may be viewed as an instance of response re-allocation guided by the matching principle. The extinction operation may therefore be regarded as one of a number of procedures which change the relative reinforcement values associated with the target and nontarget responses occurring in the experimental setting. As a result of this change in value, the various responses are presumed to redistribute themselves so as to re-establish the matching relation. Hence, the changes in response frequencies observed in extinction occur because of subjects' tendencies to maintain a relationship between response allocation across available alternatives and the relative reinforcement values associated with each.

A distinctive feature of the matching principle is the R_e parameter which, as noted earlier, enumerates all sources of reinforcement not yielded by the target response (i.e., background reinforcement). It is this source of reinforcement which chiefly motivates response re-allocation during extinction. Indeed, in the absence of background reinforcement, extinction could not occur; the target response would resist extinction indefinitely.[4]

The various sources which make up R_e may be positive or negative. For example, if it could be shown that a rapid diminution in rate of reinforcement produces an aversive state (e.g., frustration) which could subsequently be avoided by engaging in this or that alternative behavior, its effect on behavior would be regarded as negative and it would be quantified as R_e.

Given that extinction may be viewed as an instance of response re-allocation, what determines the rate at which the process takes place? According to the account of extinction advanced herein, it is the size of R_e. Or, more specifically, it is the size of R_e relative to the deteriorating value of the target response. According to the present account, it is only in the presence of some nonzero value of R_e that a decline in the absolute value of the target response can result in a decline in the frequency of that response, and it is only with reference to the relative value of R_e that response re-allocation more or less rapidly proceeds. For any given response and training procedure, resistance to extinction would therefore be generally expected to diminish as R_e assumed larger values and increase as it approached zero.

As outlined above, the present account imagines that subjects in an extinction procedure are continually striving to readjust response frequencies to match the declining relative value of the target response and the escalating relative values of the R_e-yielding nontarget responses. Factors affecting the absolute value or strength of the target response may therefore be expected to affect resistance to extinction. Hence, high drive levels and large magnitudes of reinforcement would ordinarily be expected to enhance resistance to extinction and, with notable exception, they do (see Tarpy, 1975, p. 144; Tarpy and Mayer, 1978, p. 131).[5]

An assumption implicit in the foregoing account is that the changes coincidental with an extinction operation are fairly quickly discriminated by the subject. However, this is a safe assumption only in the case of subjects previously trained on schedules of continuous reinforcement. Subjects previously trained on schedules of <u>intermittent</u> reinforcement may fail to discriminate the extinction operation for some time following its implementation (Millenson, 1967, p. 153). This fact is of significance here because matching occurs with respect to <u>discriminated</u> reinforcement rates (Herrnstein, 1974, p. 160, 1979, p. 493). That is, it is presupposed that the subject "knows" that the target response yields this or that quantity of reinforcement and this clearly may not be the case shortly following an abrupt, unsignalled change in reinforcement, especially after training on a schedule of intermittent reinforcement. After having been trained on, say, a fixed-ratio 50 schedule of reinforcement, a subject cannot <u>begin</u> to discriminate a change to a leaner schedule, or to extinction, until the unreinforced occurrence of the 50th response. It must therefore be assumed that the schedule previously in effect continues to guide behavior until such a change is at least partly discerned by the subject.

The upshot of all of this is simply that, even though an intermittently reinforced response may be of relatively less inherent value than one reinforced continuously, it may under some conditions persist longer in extinction. Indeed, there is ample evidence that they routinely do (see, e.g., Mackintosh, 1974, pp. 434-435). One can nevertheless imagine R_e values so large that an intermittently reinforced response would extinguish very rapidly indeed due, not so much to discrimination of the extinction operation but to lack of interest in the target response. This, of course, assumes that the foregoing R_e value is not so large that initial conditioning of the target response is impaired or precluded. And this, surely, is a consideration not without theoretical significance, for in a milieu rich in viable alternatives, an infrequently reinforced alternative would be of relatively little value according to the matching principle[6] (c.f. Herrnstein, 1970, p. 392; de Villiers and Herrnstein, 1976, p. 1133).

However, due in large part to the efficiency of learning laboratory procedures, R_e is usually rather small in value relative to the levels of programmed reinforcement normally yielded by target responses in instrumental conditioning experiments. Extinction data in particular is virtually always gathered in a context of exceedingly small R_e values. And this, as we shall see below, has implications for the generality of the so-called partial-reinforcement effect.

The partial-reinforcement effect

As suggested above, intermittently reinforced training results in greater subsequent resistance to extinction than that engendered by continuously reinforced training, a finding widely known as the partial-reinforcement effect (PRE) (Tarpy and Mayer, 1978, p. 130). As Mackintosh (1974, p. 435) has noted, the effect of intermittent reinforcement on extinction "is well established, usually substantial and occurs over a relatively wide range of conditions." According to Marx (1969, p. 64), the PRE "is one of the most robust results in instrumental conditioning."

Skinner (1938) and Humphreys (1939) were among the first to experimentally observe the PRE which, as Kimble (1961, p. 287) has pointed out, was then regarded as paradoxical because it was inconsistent with commonly held assumptions concerning the relationship of reinforcement to response strength. Intermittent reinforcement should, it was assumed, result in weaker conditioning and therefore, lower resistance to extinction.

A number of factors contribute to the PRE. For example, the PRE is an increasing function of reinforcement magnitude, training level, and level of drive (Tarpy, 1975, pp. 152-153). However, the most extensively studied factor affecting the PRE is the schedule of reinforcement on which the target response is initially trained (Mackintosh, 1974, p. 434). According to Tarpy (1975, p. 152), the foremost factor influencing the PRE is the percentage of reinforced responses in prior acquisition training, with resistance to extinction generally being inversely related to percentage of reinforcement. This is clearly not universally the case, however, for some researchers (e.g., Bacon, 1962; Coughlin, 1970) have found that resistance to extinction was an inverted U-shaped function of percentage of reinforcement, with the leaner percentages failing to support a robust PRE.

Many theories have been advanced to explain the PRE. Of those, most have died a natural death due to one or another critical deficiency or flaw (Mackintosh, 1974, p. 436). Only two theories of the PRE now receive major attention (Tarpy and Mayer, 1978, p. 133). These are Amsel's (1962, 1967) frustration theory and Capaldi's (1966, 1971) sequential aftereffects theory. Capaldi's is clearly the more widely accepted of the two (Mackintosh, 1974, p. 466; Tarpy and Mayer, 1978, p. 136).

According to Tarpy and Mayer (1978, p. 136) the basic concepts invoked by these two theories are nearly identical. Both place considerable emphasis on the discriminative role given nonreinforcement by intermittently reinforced training. In Amsel's theory, the occurrence of unreinforced responses in training produce frustration, the discriminative properties of which form part of the stimulus complex in which a subsequent response occurs and is reinforced. Hence, frustration becomes simply one of the events which evokes the target response. The repeated occurrence of unreinforced responses in extinction simply introduces an event (frustration) which, due to intermittently reinforced training, <u>increases</u> the probability of the target response. Thus, resistance to extinction is enhanced.

In Capaldi's theory, responses are presumed to have one of two distinctive aftereffects depending on whether they are reinforced

("R" aftereffects) or not ("N" aftereffects). These aftereffects, which, in contrast to Amsel's theory are assumed to be nonemotional in nature, persist through time via memory and subsequently form part of the stimulus complex in which a response is eventually reinforced. With respect to the effects of intermittently reinforced training on extinction, Capaldi notes that such training repeatedly exposes subjects to "N-R transitions", transitions from one or more unreinforced response to a reinforced response. Such transitions guarantee that the aftereffects coincidental with nonreinforcement will acquire discriminative control over the target response. The occurrence of unreinforced responses in extinction therefore introduces events (N aftereffects) which increase the probability of the target response, thus enhancing resistance to extinction.

Capaldi's is a generalization decrement theory of extinction. It posits that subjects will continue to respond in extinction to the extent that the conditions experienced in extinction are similar to those encountered in acquisition. And, according to Capaldi's theory, responses eventually undergo a decline in extinction because of the generalization decrement arising out of the omission of a focally important aspect of acquisition training, the reinforcing event. Hence, for Capaldi, testing for resistance to extinction is formally similar to testing for stimulus generalization (Hilgard and Bower, 1975, p. 567).

Relationship to response re-allocation theory

Although Capaldi's sequential aftereffects theory is the most widely accepted of the two foregoing theories (Tarpy and Mayer, 1978, p. 136), there can be little question that both frustration and generalization decrement play some role in determining the rate at which an intermittently reinforced response subsequently extinguishes. However, the question is, under what conditions do they play determining roles and to what extent are those roles central in accounting for the PRE? In and of itself, frustration, as a transitory emotional phenomenon, can play no substantial role in explaining the permanent effects of an extinction procedure (see Mackintosh, 1974, pp. 413-414). In contrast, it is difficult to imagine the generalization decrement arising from the omission of reinforcement not playing a significant role. However, generalization decrement clearly cannot, by itself, fully explain the response-decrementing effects of an extinction procedure, for when generalization decrement is precluded by scheduling response-independent reinforcers in extinction, responses nevertheless undergo a decline (Mackintosh, 1974, pp. 408-409).

Generalization decrement theories such as Capaldi's may be cogent explanations of why subjects fail to clearly discriminate the transition from maintained responding to extinction, especially following intermittently reinforced training. They do not, however, explain why and to what extent a subject is interested in interacting with an alternative and why and to what extent that interest diminishes in extinction. In contrast, it is precisely these issues that the theory advanced herein, response re-allocation theory, focally addresses.

Although the foregoing theories seem to coexist nicely - and indeed, may compliment each other to some extent - only response re-allocation theory has significant implications for the generality of the PRE. As noted previously, in an environment rich in viable alternatives, an infrequently reinforced response would not be expected to attain a significant frequency to begin with. Moreover, having for whatever reason attained such a frequency, that response would be expected to extinguish relatively rapidly if the level of background reinforcement was substantial.

A similar position has been taken by Burstein (1976). He pointed out that, in every case where the PRE has been demonstrated, the experimental subject's behavioral repertoire was in some sense restricted. In Burstein's view, the PRE is centrally dependent upon the absence of viable alternatives. A consequence of this position is that, to the extent that it is true that organisms in natural settings virtually always have alternatives available to them, the PRE is confined to unrepresentative situations and may indeed be little more than an artifact of learning laboratory procedures, not a "natural, robust and pervasive effect of intermittent reinforcement" (Burstein, 1976, p. 501).

The similarity between Burstein's position on the PRE and response re-allocation theory is straightforward. Both views acknowledge the importance of alternative behaviors - or rather, the lack thereof - in the determination of the PRE. Response re-allocation theory goes on to emphasize the crucial importance of the relative reinforcement associated with those alternatives, a factor Burstein does not stress.

The previous section was concluded with the reminder that response re-allocation theory is not, in any conventional sense, a theory of extinction. Neither may it be a complete theoretical account of extinction. According to that theory, organisms strive to distribute their responses among available alternatives in proportion to the reinforcements arising from them within the limits of their ability to discriminate reinforcement rates and changes in reinforcement rates. If organisms always promptly and accurately discriminated changes in reinforcement rates response re-allocation theory would, in principle, completely account for extinction data. Clearly, however, prior training on schedules of intermittent reinforcement usually makes quick discrimination of a subsequent change (e.g., extinction) very difficult indeed. Hence, response re-allocation theory complements analyses which explain why, under some conditions, organisms do not swiftly and effectively discriminate the changes in reinforcement associated with the extinction operation. It cannot, as presently constituted, replace them. By the same token, such theories are not themselves complete, for they are in large part silent about the motivational basis of response allocation.

Summary and conclusion

According to the matching-based account of extinction advanced herein, the changes in response frequencies observed in extinction are products of a process of response re-allocation. According to this account, responses are divided between target and nontarget response categories in order to achieve or maintain the matching relationship and, in the face of a change in the distribution of reinforcements, are *re*-allocated in order to re-establish or preserve that relationship.

Furthermore, as specific instances of response re-allocation, acquisition and extinction are, in their purest distillations, merely two sides of the same coin, not distinct and unique processes requiring separate explanation.

With respect to resistance to extinction, the present account implies that responses will tenaciously carry on in the face of non-reinforcement if the relative value of R_e, "background" reinforcement, is meager. If the relative value of R_e is substantial, as it probably is in most natural settings, resistance to extinction would be expected to be modest. However, because response re-allocation takes place with respect to discriminated changes in reinforcement, training procedures which obfuscate such changes would be expected to prolong extinction. Thus, training on schedules of intermittent reinforcement increases resistance to extinction in standard laboratory settings.

Schedules of intermittent reinforcement may therefore be presumed to foster two factors affecting the tenacity with which a subject resists extinction. First, as intermittent reinforcement becomes progressively leaner, the subject may be expected to be progressively less interested in emitting the target response in both acquisition and subsequent extinction. Indeed, with some modicum of background reinforcement and an exceedingly lean schedule of intermittent reinforcement, it may be impossible to interest the subject in the target response to begin with. Secondly, however, given an interested subject, training on intermittent schedules of reinforcement may be expected to make discrimination of a subsequent extinction operation difficult, thus impeding extinction. Response re-allocation theory clearly addresses those processes affecting the first of the foregoing factors while Capaldi's (e.g., 1971) sequential aftereffects theory addresses those factors affecting the second.

As an interpretation of the matching principle, response reallocation theory establishes a cogent, empirically rooted motivational framework within which the data of instrumental acquisition and extinction may be ordered and tentatively understood. More particularly, it is a coherent and elegant, albeit possibly incomplete, account of extinction and resistance to extinction. But how does it bear up under experimental test? It is toward the beginnings of such a test that we now turn.

FOOTNOTES

[1] As Deluty (1977) has recently pointed out, the matching principle is conceptually similar to several other models of conditioning. With respect to the formal properties of the K parameter, the matching principle is similar to the Rescorla-Wagner model of classical conditioning. Implicit in both is a "pie" assumption according to which the quantity of behavior available for conditioning is fixed. Hence, as the quantity of behavior controlled by one stimulus increases, the quantity of behavior controlled by alternative stimuli must commensurately decrease. Thus, the size of the behavioral pie remains invariant. Because of this feature, the matching principle is similar to the field theory of instrumental behavior advanced by Schoenfeld and his colleagues, as articulated below.

> We take it as axiomatic that behavior is a continuous stream ... The continuousness of behavior means that the organisms can be thought of as "always doing something," so that at any instant the probability of occurrence of the (response) R under observation, or P(R), is 1.00 minus the probability of occurrence of any response-other-than R... (Schoenfeld and Farmer, 1970, p. 222)

[2] Extinction may therefore be regarded as a limiting case of reinforcement intermittency and not something distinctively different. Indeed, it could quite reasonably be argued that, no matter how rigorous the extinction procedure, it is impossible to eliminate all of the reinforcement associated with the target response. In fact, efforts to eliminate all sources of conditioned reinforcement during extinction are rare. Hence, extinction may be better characterized as a change from one reinforcement schedule to another rather than to no reinforcement whatsoever.

[3] One could, of course, argue that, while extinction is unquestionably an instance of response re-allocation, it is nevertheless distinguishable from acquisition on the basis of the emotions coincidental with the extinction operation. Apart from the questionable advisability of discriminating between categories of behavior on the basis of presumed correlated experiences, it is not clear that the emotions said to characterize these operational categories in fact carefully do so. The emotions commonly identified with extinction also occur in procedures in which reinforcement is reduced but not eliminated or when responses are both reinforced and punished (cf. Terrace, 1972). However, this is not to dismiss the possibility that an abrupt shift in the distribution of reinforcements such as the one produced by the extinction operation creates correlated changes in the emotions felt. However, from the present perspective, it must be borne in mind that, while the target response undergoes a decrement in the relative

3 continued ...

reinforcement yielded by it in extinction, the various nontarget responses undergo <u>increments</u> in associated relative reinforcement. Indeed, both acquisition and extinction inevitably involve increments and decrements in relative reinforcement if both target and nontarget responses are considered. Hence, whatever emotions are posited to characterize acquisition or extinction may be presumed to occur in both procedures and exclusively typify neither.

4 Needless to say, it is difficult to imagine a setting totally devoid of background reinforcement, but if such a setting existed - or were somehow synthetically produced - target responses established therein would be expected to subsequently resist extinction without limit according to the model of extinction proposed herein. For although the absolute value of the target response would unquestionably deteriorate over the course of extinction in such a setting, its relative value would not. For clearly, despite the declining absolute value of the target response, <u>no</u> point could be reached at which that response controlled anything less than all of the available (albeit meager and mostly conditioned) reinforcement. Hence, the subject would be expected to continue to allocate its behavior to the target response despite the absence of programmed reinforcement.

5 The present account of extinction is not conspicuously consistent with the finding (Tarpy, 1975, pp. 141-142) that resistance to extinction is inversely related to magnitude of reinforcement following continuously reinforced training. However, this finding is neither demonstrably general nor pervasive. As Mackintosh (1974, p. 427) has pointed out, "in situations other than the alley, with reinforcement other than food and in species other than the rat , ... there is virtually no evidence of an inverse relationship between magnitude of reinforcement and resistance to extinction."

6 If essentially correct, the matching principle has significant implications for the viability and long-term workability of behavior modification in natural settings and in particular, the use therein of intermittent reinforcement contingencies. Clearly, from the perspective of the matching principle, whatever reinforcers a behavior modifier arranges for a person or persons in a natural setting will take on motivational significance or not depending on the context of background reinforcement in which they are presented. If the level of background reinforcement is substantial, the reinforcement contingencies arranged by behavior modifiers may have only marginal effects on behavior, or no effects at all. Or, they may have the desired effects in controlled institutional or clinical settings (where their relative value may be great) but have no enduring effects in subsequent uncontrolled natural settings (where they may be of relatively little value).

CHAPTER IV

THE EXPERIMENTS

EXPERIMENT I

INTRODUCTION

The response enhancing effect of intermittently reinforced training on subsequent resistance to extinction is, on the surface, one of the great paradoxes of modern psychology, for it seems to imply that behavioral persistence or tenacity is inversely related to rate of payoff (cf. Tarpy and Mayer, 1978, p. 130). However - and as maladaptive as it may seem to be - the empirical reality of this effect is exceedingly well corroborated (Mackintosh, 1974, p. 435).

Within the framework established by response re-allocation theory, the robust resistance to extinction reportedly resulting from prior intermittently reinforced training is neither paradoxical nor particularly puzzling. According to that theory, resistance to extinction is inversely related to the level of prevailing background reinforcement (i.e., R_e, all reinforcement <u>not</u> yielded by the target response). And indeed, it is only in the context of some nonzero value of R_e that a decline in the absolute value of a target response can result in a decline in the frequency of that response. Hence, in settings virtually lacking in background reinforcement, a target response would be expected to evidence substantial resistance to extinction <u>whether or not</u> it was initially trained on a schedule of intermittent reinforcement. By the same token, however, resistance to extinction would be expected to be <u>impaired</u> in the context of some modicum of background reinforcement, a feature, presumably, of most natural environments. If substantiated, this latter implication of response re-allocation theory would inpose serious constraints on the generality of the reported inverse relationship of percentage of reinforcement in training to subsequent resistance to extinction.

The purpose of the experiment below is to examine the effects of R_e magnitude on extinction following intermittently reinforced training with an eye to assessing the hypothesis that resistance to extinction is inversely related to the value of R_e, as predicted by response re-allocation theory. Since response re-allocation takes place with respect to discriminated changes in reinforcement rates, we ought also to expect that factors impairing the discriminability of the extinction operation will, given a constant level of background reinforcement, enhance resistance to extinction. Since the most prominent of such factors is the percentage of reinforcement on which the target response is trained, the following experiment is specifically aimed at investigating the effects of three levels of background reinforcement on resistance to extinction following training on schedules of three percentage of reinforcement levels of the target response.

A problem arising in parametric studies of the effects of background reinforcement has to do with its accessability and manipulability. Much of what counts as background reinforcement in standard experimental settings - e.g., such things as preening, scratching, wing-flapping, exploring and the like - do not really lend themselves to precise experimental manipulation (see, e.g., Rachlin, 1976, p. 595). Other background reinforcers are presumably somewhat idiosyncratic and therefore not reprsentative of all subjects and, in any case, clearly difficult to identify and control (Herrnstein, 1970, pp. 388-389).

Fortunately, these problems can be largely circumvented by employing a two-alternatives concurrent operants procedure (see Catania, 1966) wherein one alternative is designated the "target response" and the other is designated an "alternative response" (i.e., one of the B_es). This procedure readily allows for the precise manipulation of a major source of background reinforcement (i.e., the reinforcement yielded by the alternative response) while also permitting the measurement of a major alternative response during the extinction of the target response. For reasons such as these, this basic procedure was adopted for the research undertaken below.

METHOD

Subjects:

The subjects were 45 male and female white king pigeons reduced to and maintained at 80% of their free-feeding weights. Body weights were maintained at this level throughout the experiment by sessional and post-sessional feedings.

Apparatus:

The experiment was conducted in a Grason-Stadler two-key operant conditioning station (model E1100PE). White noise was piped into the station from a Lafayette white noise generator to mask extraneous laboratory sounds. The left key, which served as the target response key, was lit white throughout the experiment and the right key, which served as the alternative response key, was lit red. Target response reinforcement, which consisted of 3 seconds access to food (Noyes pellets), was delivered on a fixed-ratio schedule arranged by a standard electro-mechanical pre-determining counter. Alternative response reinforcement, which consisted of 3.5 seconds access to food (also Noyes pellets), was delivered on a variable-interval schedule arranged by a Grason-Stadler celluloid tape reader (model 1079). Standard electro-mechanical equipment controlled all timing and switching sequences and responses, reinforcements and session durations were tallied on electro-mechanical event counters.

Procedure:

Subjects were shaped to peck a response key lit, first white and then red. They were then assigned to one of nine groups differing in the ratio requirement associated with the target response and/or the rate of reinforcement yielded by the alternative response. In session "0", the rates of reinforcement yielded by the two responses were gradually diminished so as to correspond to the intermittencies characteristic of the groups to which the subjects had been assigned. The reinforcement schedule associated with the target response was either a fixed-ratio 4, 8 or 16 while the schedule associated with the alternative response was either no reinforcement at all, a variable-interval 22.5 second, or a variable-interval 45 seconds. The conditions of target and alternative response reinforcement were experimentally juxtaposed in a 3 x 3 factorial format. In the groups in which no alternative reinforcement was available, pecks on the alternative response key had no programmed consequences. A changeover delay (COD) imposed a minimum duration (.5 seconds) between a switch from one to the other alternative and reinforcement forthcoming from that latter alternative. Subjects were run for 8 consecutive sessions in their respective conditions and sessions automatically terminated following delivery of the 60th target response reinforcer. On the day following completion of the 8th session all reinforcement associated with the target response was omitted and, upon return of the subject to the experimental apparatus, that response was allowed to undergo extinction. Extinction criteria was set at no more than 3 responses over a six minute block of time. The conditions of reinforcement associated with the alternative response were maintained without interruption throughout the extinction session and alternative and target response counts were tallied over the course of that session.

RESULTS AND DISCUSSION

The main results, shown in fig. 2, generally support the hypothesis advanced in the introduction. Resistance to extinction declined as background reinforcement rose. At each level of background reinforcement, a partial-reinforcement effect was observed. A two-way analysis of variance revealed that the effect of level of background reinforcement was statistically significant ($F(2/36) = 10.99$, $p. < .001$). So, too, was the effect of percentage of reinforcement ($F(2/36) = 7.345$, $p < .05$). The background reinforcement x percentage of reinforcement interaction was not significant ($F(4/36) = .189$, $p > .05$). According to a multiple classification analysis, the main effects accounted for 30% (background reinforcement) and 20% (percentage of reinforcement) of the data variance, with the two jointly accounting for half of that variance.

With respect to the groups in which the target response was reinforced on a FR-4 schedule, the presence of background reinforcement resulted in a 35% (VI-45) and a 55% (VI-22.5) reduction in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement (which was essentially a control group). With respect to the groups in which the target response was reinforced on a FR-8 schedule, the presence of background reinforcement resulted in a 30.5% (VI-45) and a 51% (VI-22.5) reduction in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement. And, with respect to the groups in which the target response was reinforced on a FR-16 schedule, the presence of background reinforcement results in a 27% (VI-45) and a 49% (VI-22.5) reduction in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement (see table 1). On the whole, resistance to extinction was reduced an average of 30.8% in the VI-45 groups and 51.7% in the VI-22.5 groups relative to the groups receiving no programmed background reinforcement. The effects of the respective levels of background reinforcement on resistance to extinction were therefore approximately constant across the three target response training schedules.

Given the clear trend in the present experiment, greater values of background reinforcement than those used here could, up to a point, quite reasonably be expected to produce even greater decrements in resistance to extinction. But the apparent limit of this trend, i.e., virtually instantaneous extinction, would not be expected if response re-allocation theory (and indeed, the matching principle itself) are essentially correct. The reason has to do with the environmental conditions necessary for the acquisition and maintenance of an instrumental response antecedent to an extinction procedure. As pointed out previously, an environment rich in viable alternatives would not, according to the matching principle, be expected to engender and robustly maintain an infrequently reinforced response. In such a setting, the relative value of such a response would ordinarily be trivial. Thus, a formal upshot of the matching principle (and, of course, response re-allocation theory) is that the so-called partial-reinforcement effect is confined to settings which engender and sustain intermittently reinforced behavior in the first place due to the paucity of viable alternatives. Although such settings surely exist, they may not be typical of the ecologies in which animals (including human beings) normally or routinely find

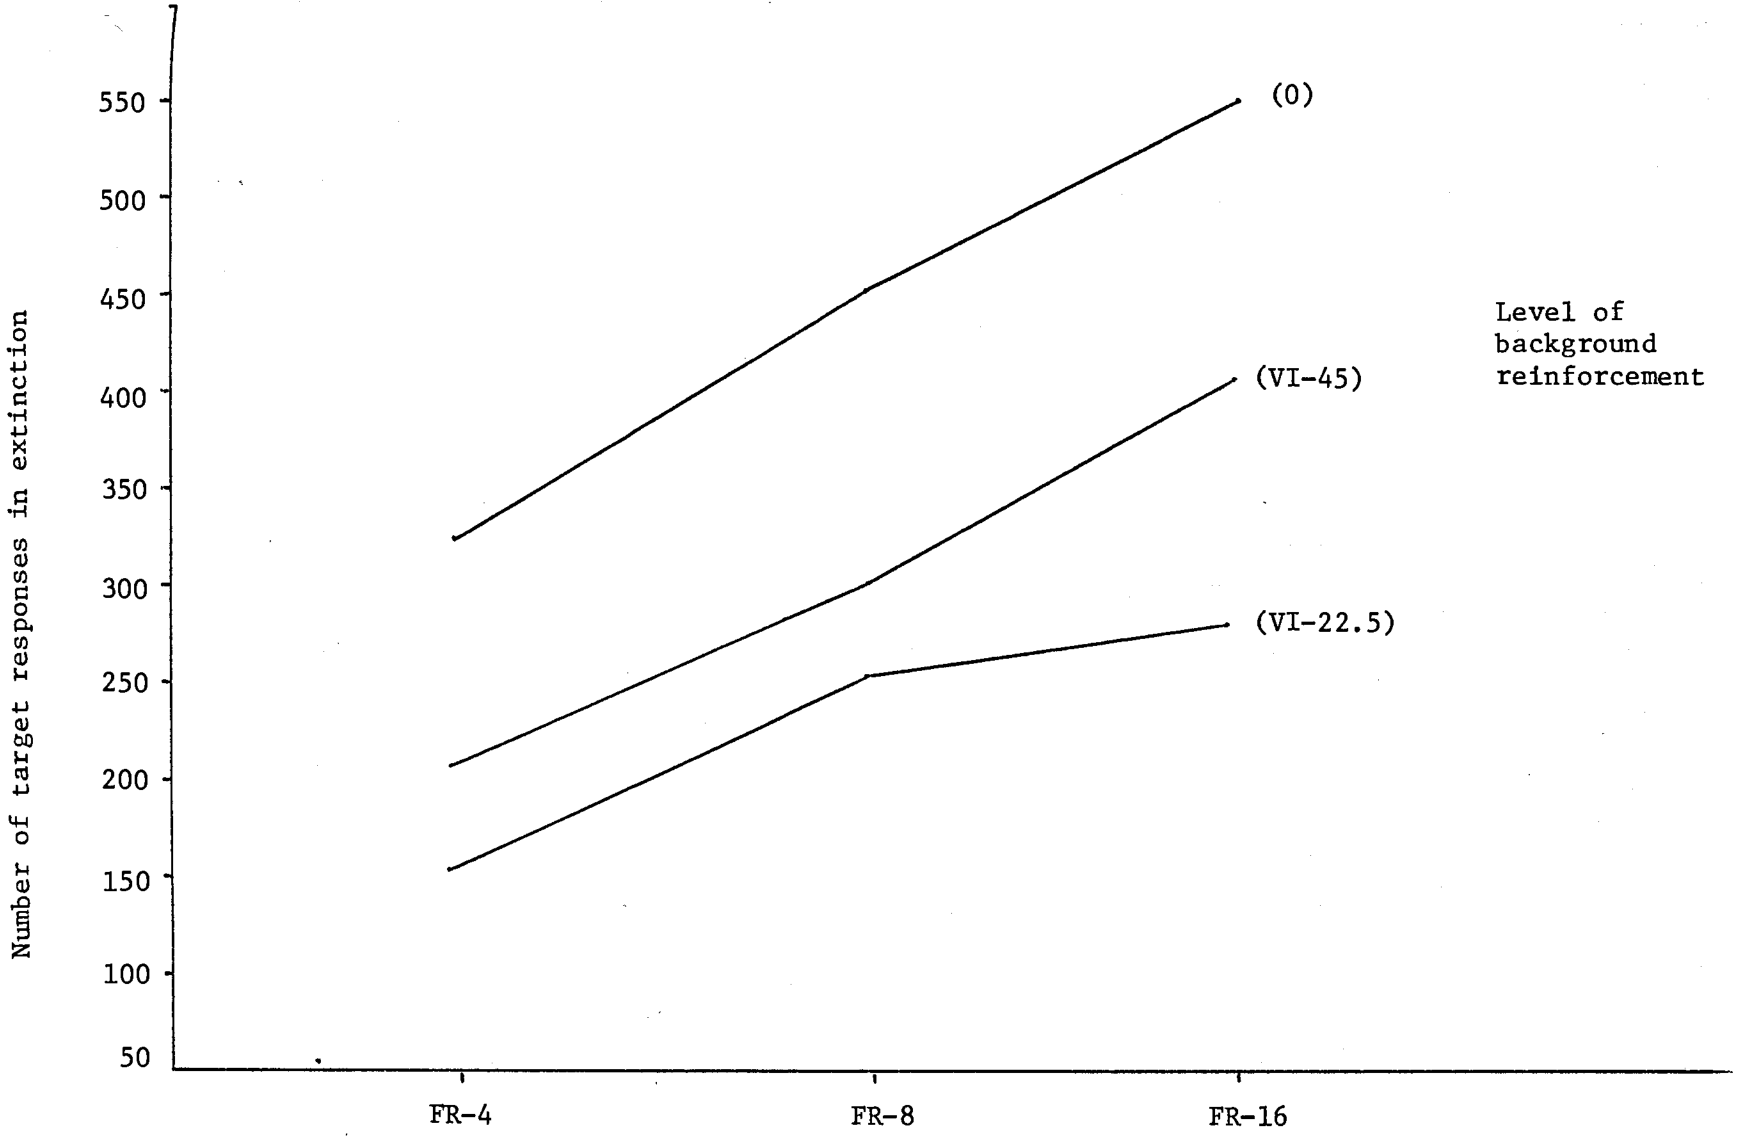

Fig. 2. Mean resistance to extinction as a function of percentage of target response reinforcement in training (FR-4, 8 or 16) and level of background reinforcement.

Table 1: Individual scores (responses to extinction) by group, and group means, experiment I.

BACKGROUND REINFORCEMENT		INDIVIDUAL SCORES					GP MEANS
Zero	FR-4	379	299	350	331	297	331.2
	FR-8	398	288	398	819	273	435.2
	FR-16	328	816	449	655	421	533.8
VI-45 sec.	FR-4	195	255	218	198	211	215.4
	FR-8	300	314	268	361	270	302.6
	FR-16	387	190	576	437	365	391
VI-22.5	FR-4	110	153	183	176	142	152.8
	FR-8	348	154	293	310	173	255.6
	FR-16	383	321	205	133	352	278.8

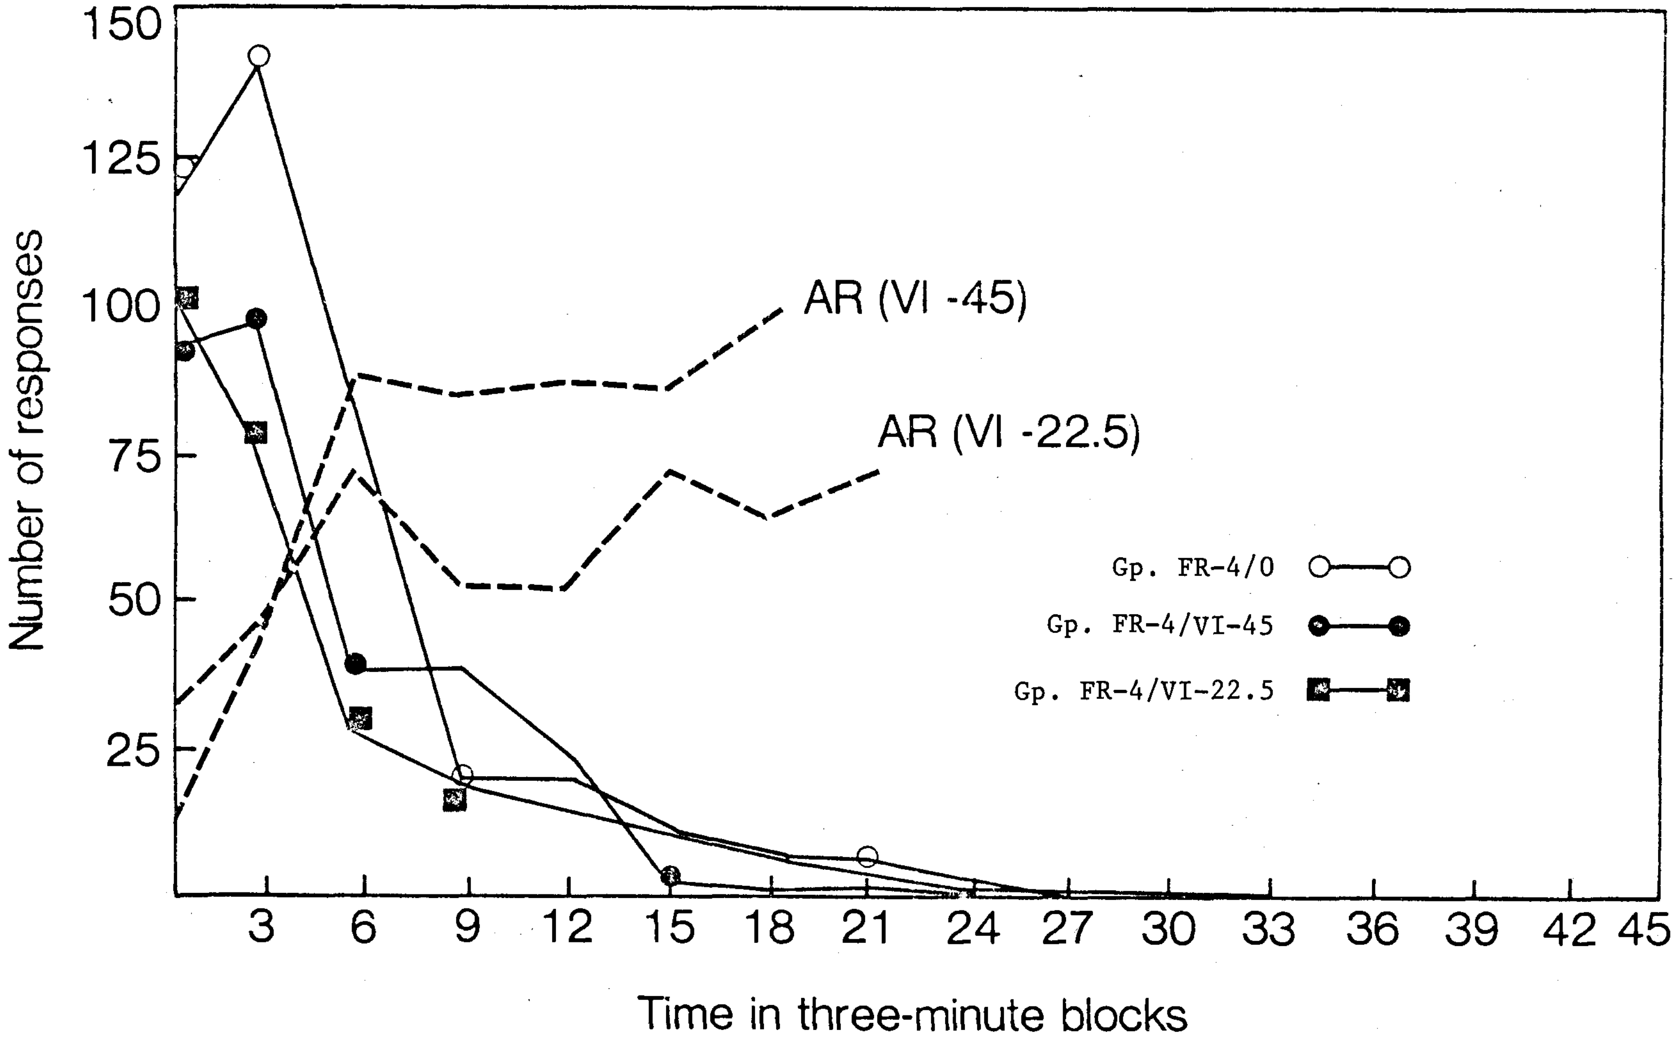

Fig. 3. Level of responding over consecutive 3-minute blocks in the target and alternative responses (AR) during extinction, groups FR-4/0, FR-4/VI-45 and FR-4/VI-22.5. (AR levels calculated to first target response achieving extinction criterion.)

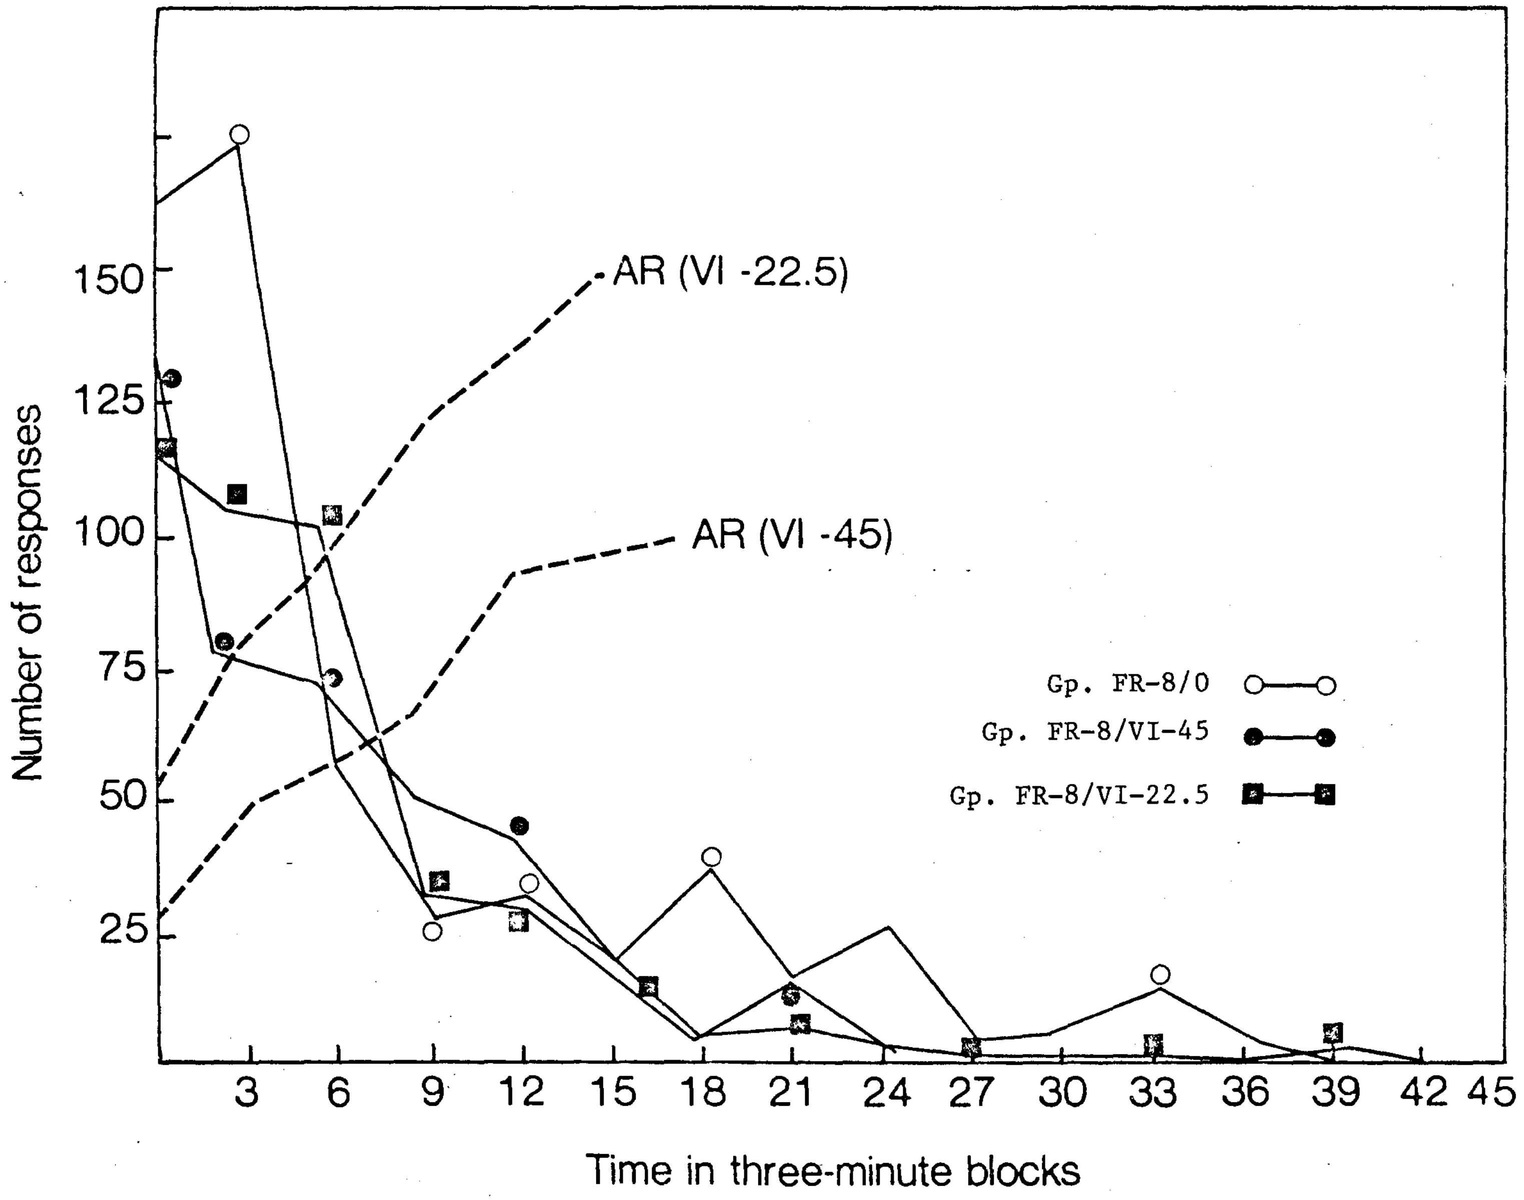

Fig. 4. Level of responding over consecutive 3-minute blocks in the target and alternative responses (AR) during extinction, groups FR-8/0, FR-8/VI-45 and FR-8/VI-22.5. (AR levels calculated to first target response achieving extinction criterion.)

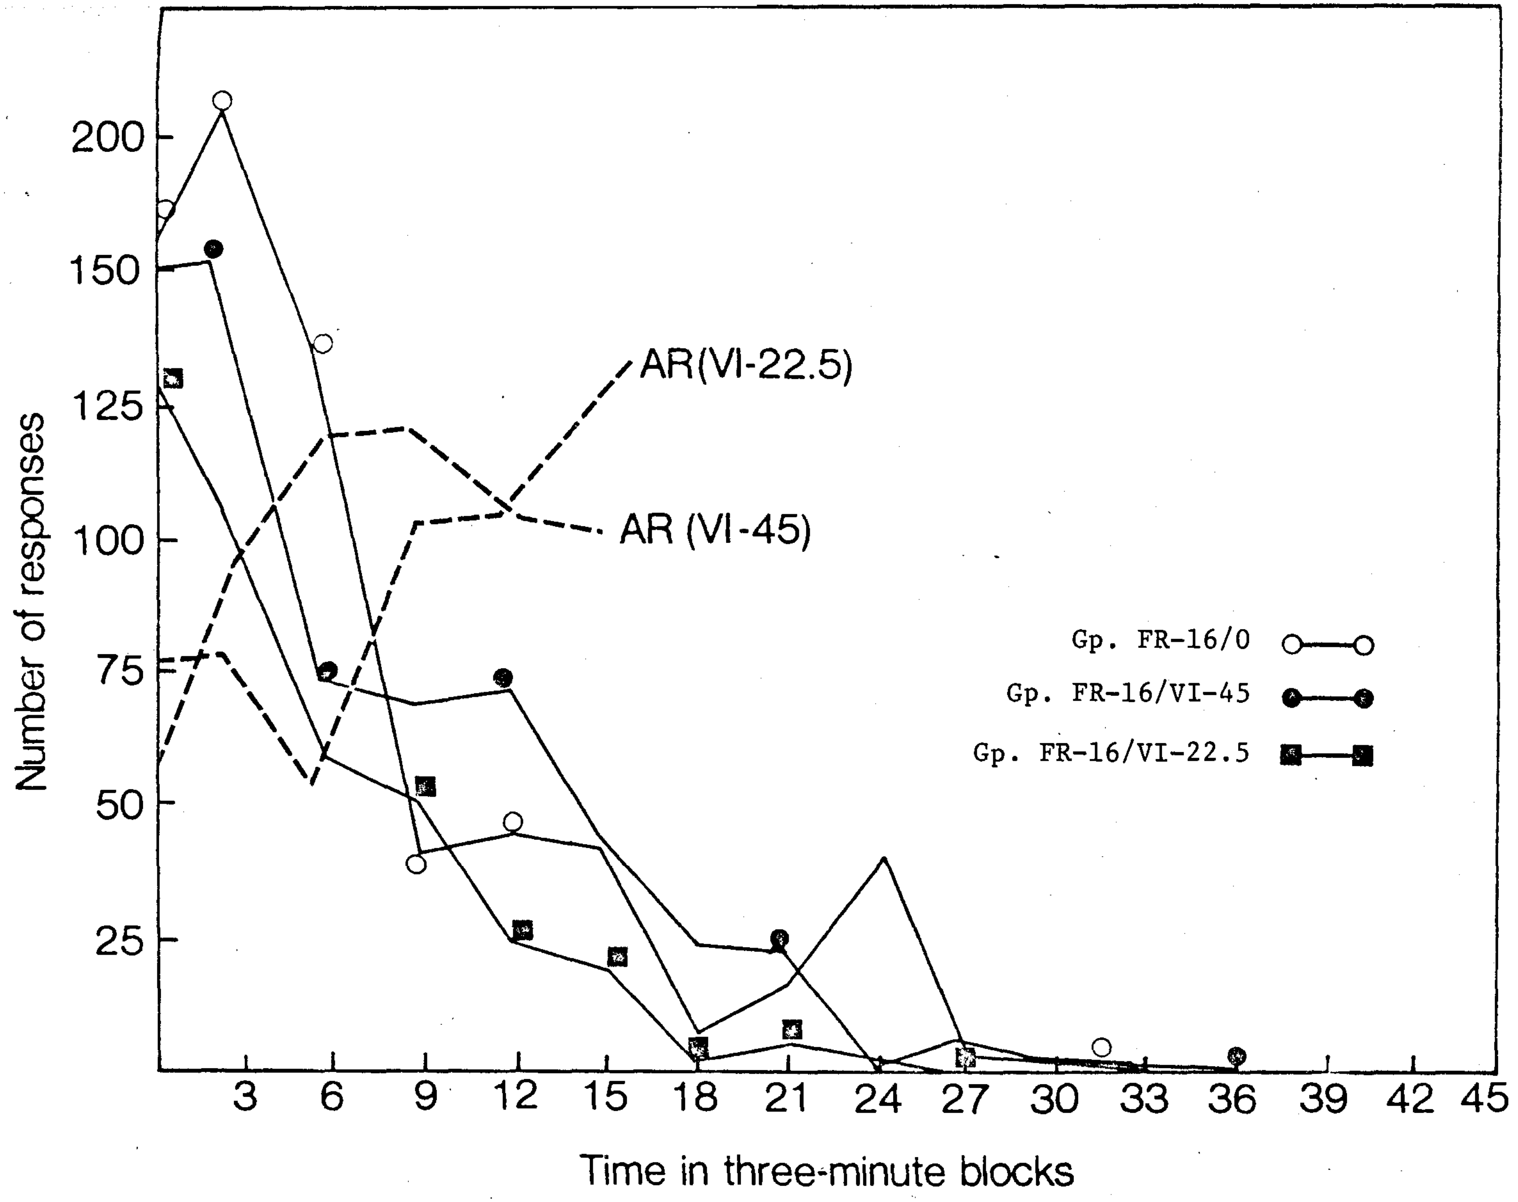

Fig. 5. Level of responding over consecutive 3-minute blocks in the target and alternative responses (AR) during extinction, groups FR-16/0, FR-16/VI-45 and FR-16/VI-22.5. (AR levels calculated to first target response achieving extinction criterion.)

themselves (zoos, prisons and mental institutions notwithstanding).

When background reinforcement is greater than zero, response re-allocation theory predicts a decrement in the frequency of the target response over the course of its extinction and a corresponding increment in the frequencies of the various nontarget responses. With the exception of the so-called alternative response, no nontarget responses were monitored in the present experiment. It is therefore not clear that the changes in the frequencies of the nontarget responses commensurately counterbalance those observed in the target response during its extinction. However, as shown clearly in figs. 3-5, the frequency of the alternative response generally increased as the target response underwent decline in extinction.

The results of the present study are generally consistent with the results of an experiment reported by Leitenberg, Rawson and Bath (1970). They trained rats to press a lever for variable-interval reinforcement during alternating periods of "house light on" and "house light off." After five days training, a second level was added to the apparatus and reinforcement was no longer provided for the original response when the house light was off. Reinforcement continued to be provided for the original response when the house light was on. For half the subjects, presses on the second lever during light off periods produced fixed-ratio reinforcement while for the other half, presses on this lever had no programmed consequences. Rate of extinction during light off periods was monitored in the original response for all subjects. Leitenberg et al. found that extinction criteria was achieved more rapidly in subjects receiving second-lever reinforcement than in those receiving no such reinforcement. However, when second-lever reinforcement was subsequently discontinued in the former group, emission of first-lever responses resumed and, overall, there were no savings in total responses to extinction.

While Leitenberg et al. made no effort to provide a theoretical explanation of the foregoing effects, both are consistent with the logic of response re-allocation theory. According to that theory (as the reader will recall), responses undergo a decline in extinction because the extinction operation affects a change in the relative reinforcement values associated with the various target and nontarget responses. As the absolute value of the target response deteriorates in extinction, the relative values of the nontarget responses rise due to their association with the elements which make up R_e, the only source of ongoing reinforcement. Clearly, the larger the value of R_e, the sooner the deteriorating value of the target response will sink to the level at which almost all available reinforcement (conditioned or unconditioned) is yielded by the various nontarget responses and extinction criteria is, as a result, achieved. However, if at this point, the value of R_e is abruptly and substantially reduced, the target response may be expected to again manifest itself, for its deteriorating absolute value, <u>no matter how small</u>, would then constitute a larger proportion of total available reinforcement. Hence, the effects of Leitenberg et al.'s second-lever reinforcement, and the effects of its subsequent omission, are neither surprising nor puzzling from the perspective of the view of extinction advanced herein.

In its purest distillation, response re-allocation theory is a motivational account of the response decrements normally observed in extinction. As such, it accounts for those decrements by showing why the subject becomes progressively less interested in the target response as a result of the extinction operation. Another interpretation of the results of experiment I is possible, however. One could argue that competition from the alternative response interfered with and eventually prevented the emission of, the target response. According to this argument, the reinforcement yielded by the alternative response is important only insofar as it fuels the response which competes with and displaces the target response. This, of course, is the kind of account interference theorists would likely advance (cf. Mackintosh, 1974, pp. 410-413). To be sure, both of the foregoing accounts predict an increase in the frequency of alternative responses, but it is only in the latter account that the appearance of alternative responses (for whatever reason) causes the decrements in the target response routinely observed in extinction. According to response re-allocation theory, the decrements observed in the target response and the corresponding increments observed in the various alternative responses are consequences of the shift in the relative value of those responses coincidental with the extinction operation. Hence, according to that theory, whether or not alternative responses compete with the target response is immaterial.

In experiment I, the reinforcement schedule associated with the alternative response was continuously correlated with a "red" discriminative stimulus. The subjects of that experiment could not therefore know when reinforcement had been arranged by that schedule and when it had not. A fairly robust frequency of response is normally required (and usually evidenced) in such procedures in order to collect each programmed reinforcer as it becomes available for the taking. As regards a target response undergoing extinction, such a frequency of response is clearly a plausible source of competing of interfering responses. It follows that any experimental procedure which reduced the frequency of the alternative response to an insubstantial level while, at the same time, maintaining the levels of reinforcement allotted to it in experiment I would constitute an exceedingly rigorous test of response re-allocation theory, for the decrements undergone by the target response in extinction could not be attributed to the interfering effects of the alternative response.

The separation of response and reinforcement rates envisaged above is readily accomplished by a so-called _signalled reinforcement procedure_ (Catania, 1963). In such a procedure, one discriminative stimulus is in effect when _no_ reinforcement has been set up by the schedule programming, and a second is in effect when reinforcement has been set up by that programming to be taken by the subject's next response. In a signalled reinforcement procedure, subjects "know" when reinforcement is and is not available and consequently confine their responses to periods when reinforcement is available (i.e., signalled). As one might expect, the result is an exceedingly low frequency of response.

The plan of the experiment below was to replicate experiment I in a context in which the reinforcements associated with the alternative response were signalled both during training and extinction of the target response. Since the rates of reinforcement yielded by the alternative response in this procedure would be largely the same as those in experiment I, response re-allocation theory would, for reasons given earlier, predict essentially the same result. However, if the occurrence of the alternative response played a substantial role in suppressing or displacing the target response in the extinction phase of experiment I, the sharp reduction in its frequency expected in the present experiment should, in contrast to the results of experiment I, result in a marked enhancement of resistance to extinction in the target response.

EXPERIMENT II

METHOD

Subjects:

The subjects were the same male and female White King pigeons used in experiment I assigned to groups corresponding to those to which they had been assigned in that experiment. The pigeons' body weights were maintained at 80% of their free-feeding weights throughout the experiment by sessional and post-sessional feedings.

Apparatus:

The apparatus was the same as that used in experiment I.

Procedure:

Following the extinction session in Experiment I, subjects were returned to their home cages and fed. The following day, target response reinforcement was restored and subjects were returned to the experimental apparatus for session "0". During this session, scheduled target response reinforcement was supplemented, when required, with experimenter-delivered reinforcers in order to expedite re-acquisition of the target response. As in experiment I, the target response key was lit "white" throughout the experiment. However, unlike experiment I, the alternative response key was lit "red" only when the schedule programmer had set up a reinforcer for the next response. It was otherwise lit "green." The schedule of reinforcement associated with the alternative response was either no reinforcement whatsoever, a variable-interval 45 seconds, or a variable-interval 22.5 seconds. The schedule associated with the target response was either a fixed-ratio 4, 8 or 16. The reinforcement magnitudes associated with the two responses were the same as in experiment I. The conditions of target and alternative response reinforcement were experimentally juxtaposed in a 3 x 3 factorial format. In the groups in which no alternative reinforcement was available, pecks on the alternative response key had no programmed consequences. A changeover delay (COD) imposed a minimum duration (.5 seconds) between a switch from one to the other alternative and reinforcement forthcoming from that latter alternative. Subjects were run for six consecutive sessions in their respective conditions and sessions automatically terminated following delivery of the 60th target response reinforcer. On the day following completion of the sixth session all reinforcement associated with the target response was omitted and, upon return of the subject to the apparatus, that response was allowed to undergo extinction. Extinction criteria was set at no more than three responses over a six minute block of time. The conditions of reinforcement associated with the alternative response were maintained without interruption throughout the extinction session and alternative and target response counts were tallied over the course of that session.

RESULTS AND DISCUSSION

The results are shown in fig. 6 and table 2. A two-way analysis of variance revealed that the effect of level of background reinforcement was statistically significant (F(2/36) = 10.92, p < .001). So too was the effect of percentage of target response reinforcement (F(2/36) = 20.81, p < .001). The background reinforcement x percentage of reinforcement interaction was also statistically significant (F(4/36) = 4.492, p < .005). According to a multiple classification analysis, the main effects accounted for 18% (background reinforcement) and 36% (percentage of target response reinforcement) of the data variance, with the two jointly accounting for 54% of that variance.

The results of experiment II are inconsistent with response reallocation theory as developed in chapter three of this text and generally consistent with an interference or competing response interpretation of extinction. Unlike the results of experiment I, the presence of a source of alternative reinforcement, when explicitly signalled, markedly enhanced resistance to extinction in the present study.

With respect to the groups in which the target response was reinforced on an FR-4 schedule, the presence of signalled background reinforcement resulted in a 3.5% decrement (VI-45) and a 28% increment (VI-22.5) in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement (which was essentially a control group). With respect to the groups in which the target response was reinforced on a FR-8 schedule, the presence of signalled background reinforcement resulted in a 430% (VI-45) and a

649% (VI-22.5) increment in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement. And, with respect to the groups in which the target response was reinforced on a FR-16 schedule, the presence of signalled background reinforcement resulted in a 320% (VI-45) and a 214% (VI-22.5) increment in resistance to extinction relative to the corresponding group receiving no programmed background reinforcement (see figs. 7-9).

As shown by the robust background reinforcement x percentage of reinforcement interaction, the response enhancing effects of signalled background reinforcement was not uniform across all experimental conditions. The presence of signalled background reinforcement resulted in a decrement in resistance to extinction in one of the FR-4 groups (VI-45) and produced only a modest increment (28%) in the other (VI-22.5) relative to the corresponding group receiving no programmed background reinforcement. The most robust effects were observed in the FR-8 and FR-16 groups wherein resistance to extinction was elevated an average of 539.9% and 267% respectively relative to corresponding groups receiving no programmed background reinforcement (see fig. 6). The presence of an interaction in experiment II in the context of no such interaction in experiment I may prove to be of some theoretical significance. However, the account of extinction advanced in these pages makes no formal prediction regarding the foregoing interaction (or its absence in experiment I) and little in way of a tentative account of it will be offered here.

In a sense, the present results replicate those of Rescorla and Skucy (1969) and Boakes (1973). These researchers showed that response-independent delivery of reinforcement during extinction retarded the

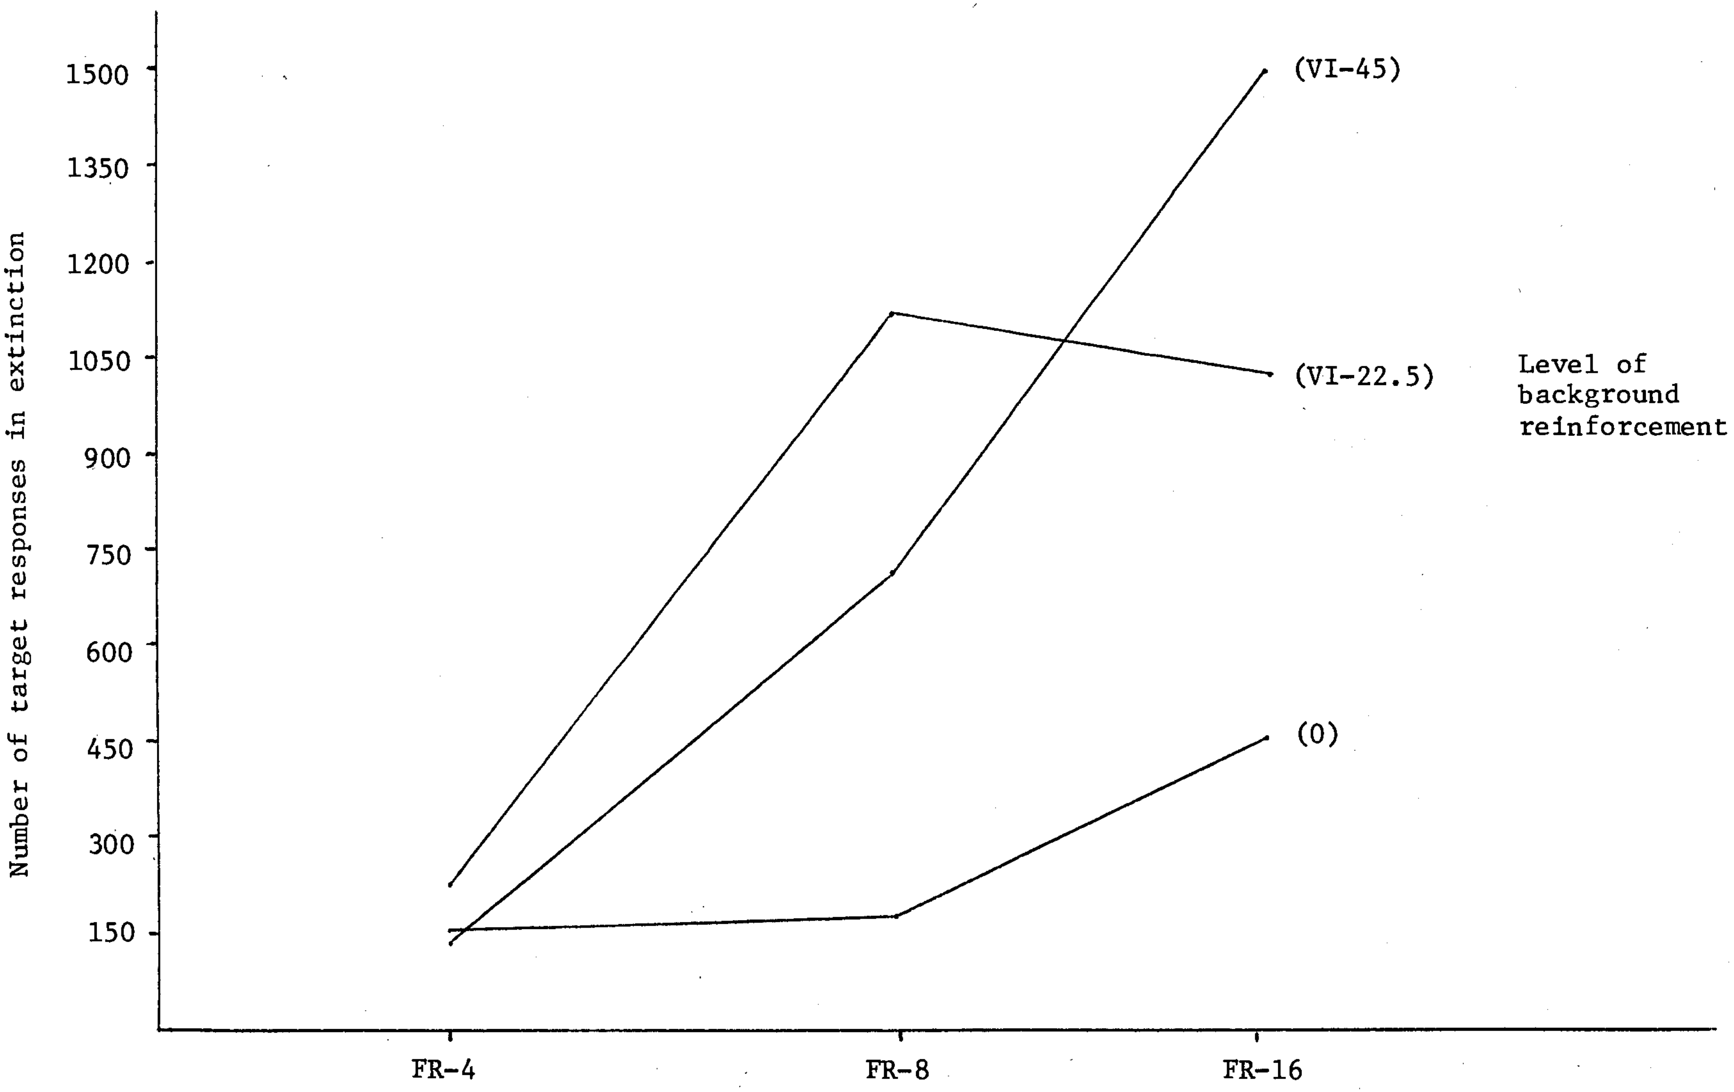

Fig. 6. Mean resistance to extinction as a function of percentage of target response reinforcement in training (FR-4, 8 or 16) and level of signalled background reinforcement.

Table 2: Individual scores (responses to extinction) by group, and group means, experiment II.

BACKGROUND REINFORCEMENT		INDIVIDUAL SCORES					GP MEANS
	FR-4	216	172	124	173	136	164.2
Zero	FR-8	119	178	213	217	124	170.2
	FR-16	401	656	554	341	422	474.8
	FR-4	229	205	94	92	173	158.6
VI-45 sec.	FR-8	921	532	722	684	800	731.8
	FR-16	658	2231	1809	597	2293	1517.6
	FR-4	235	238	288	138	154	210.6
VI-22.5	FR-8	1174	1223	626	1393	1107	1104.6
	FR-16	286	703	1698	1207	1188	1016.4

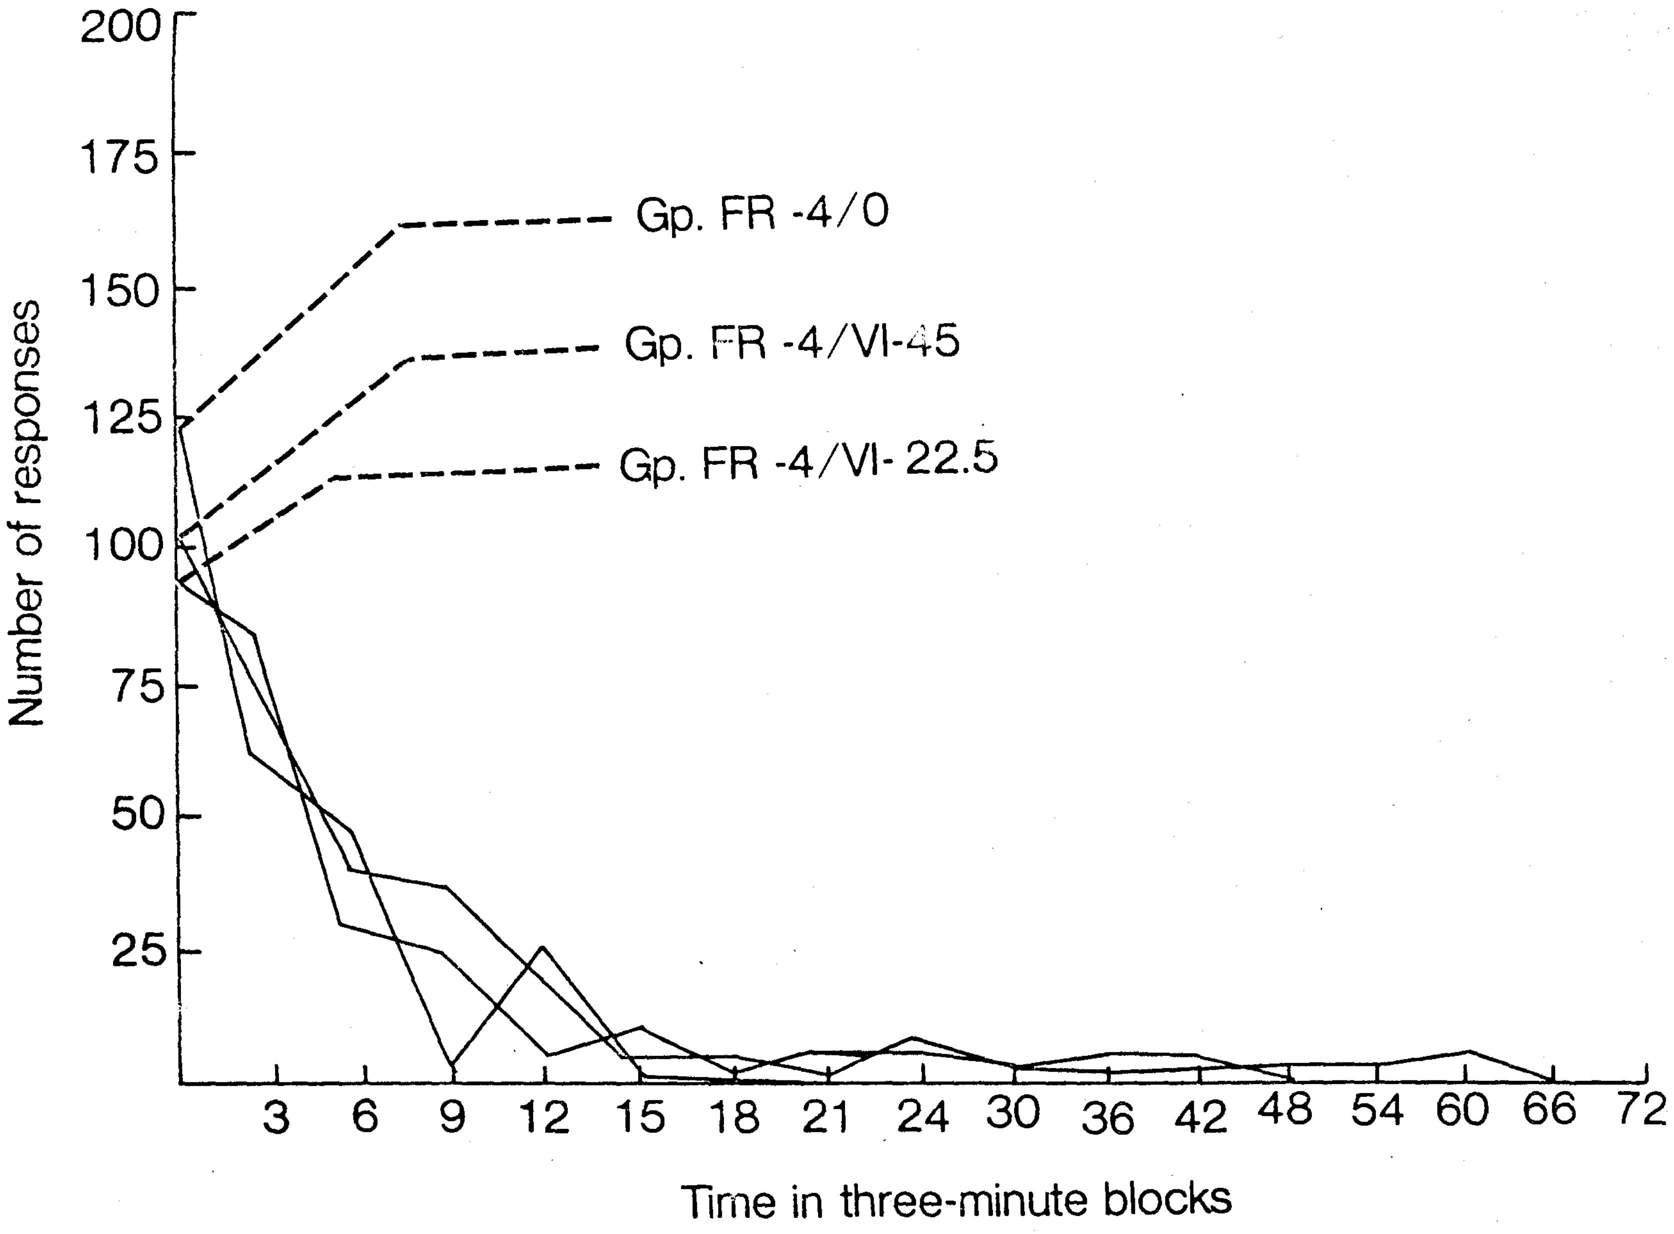

Fig. 7. Level of responding over consecutive 3-minute blocks in the target response during extinction, groups FR-4/0, FR-4/VI-45 and FR-4/VI-22.5.

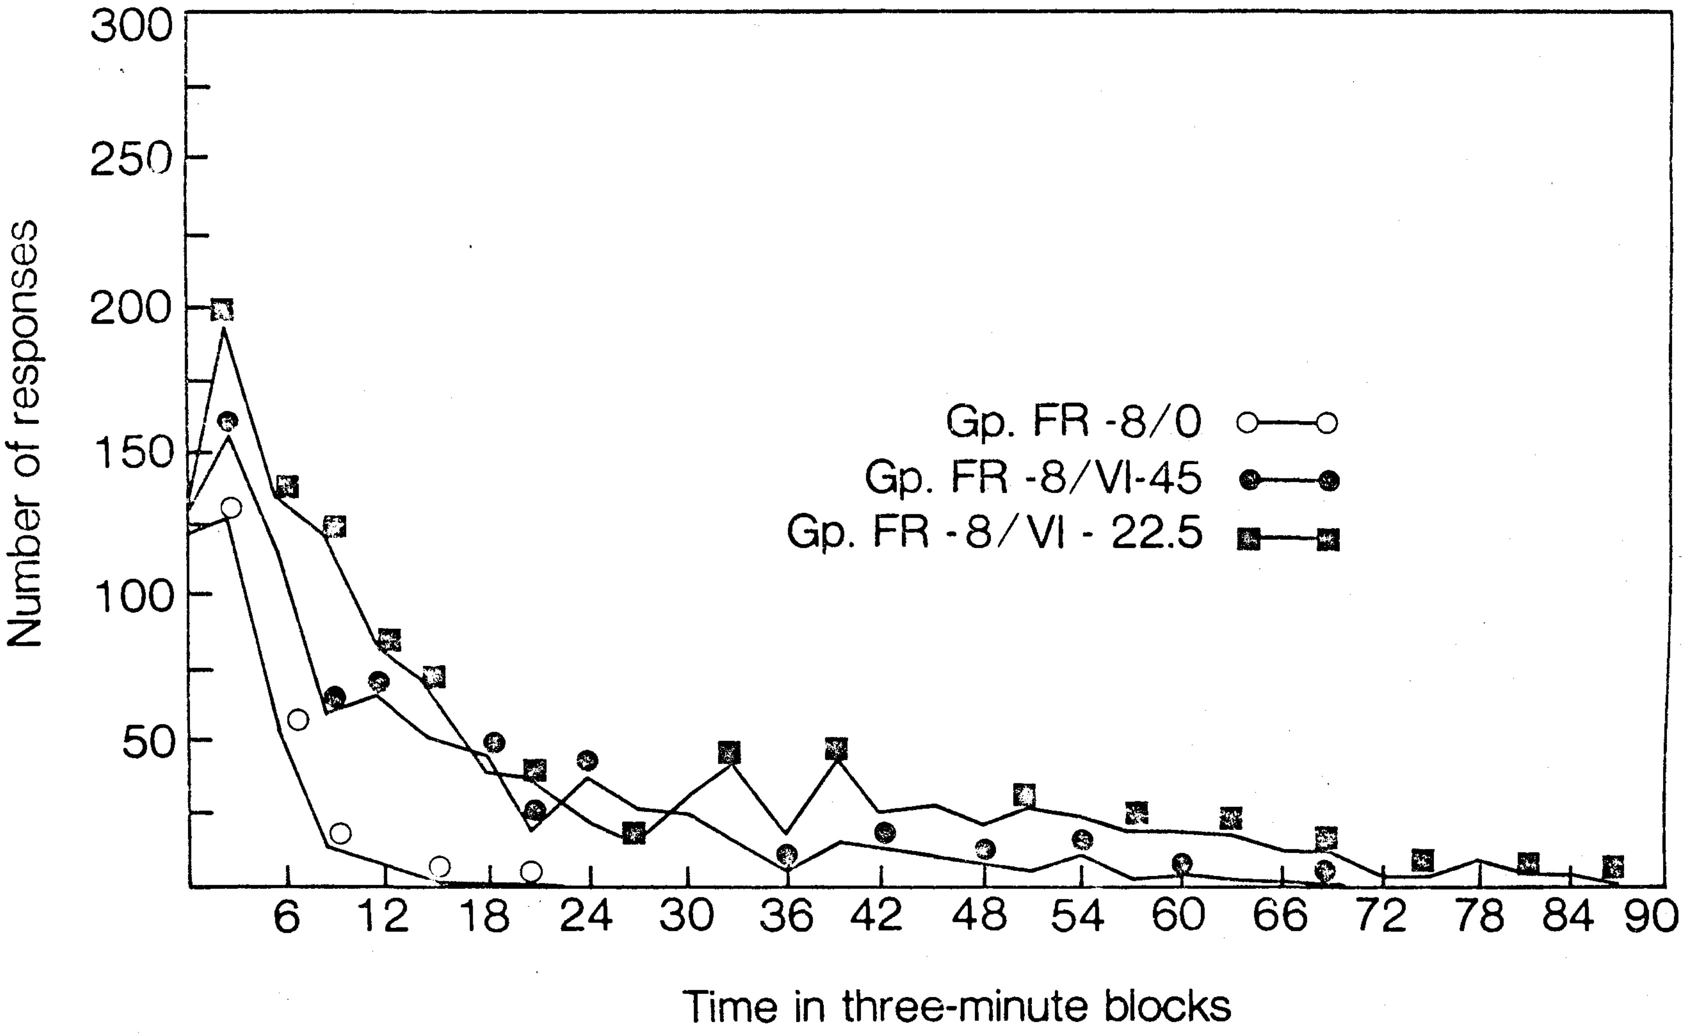

Fig. 8. Level of responding over consecutive 3-minute blocks in the target response during extinction, groups FR-8/0, FR-8/VI-45 and FR-8/VI-22.5.

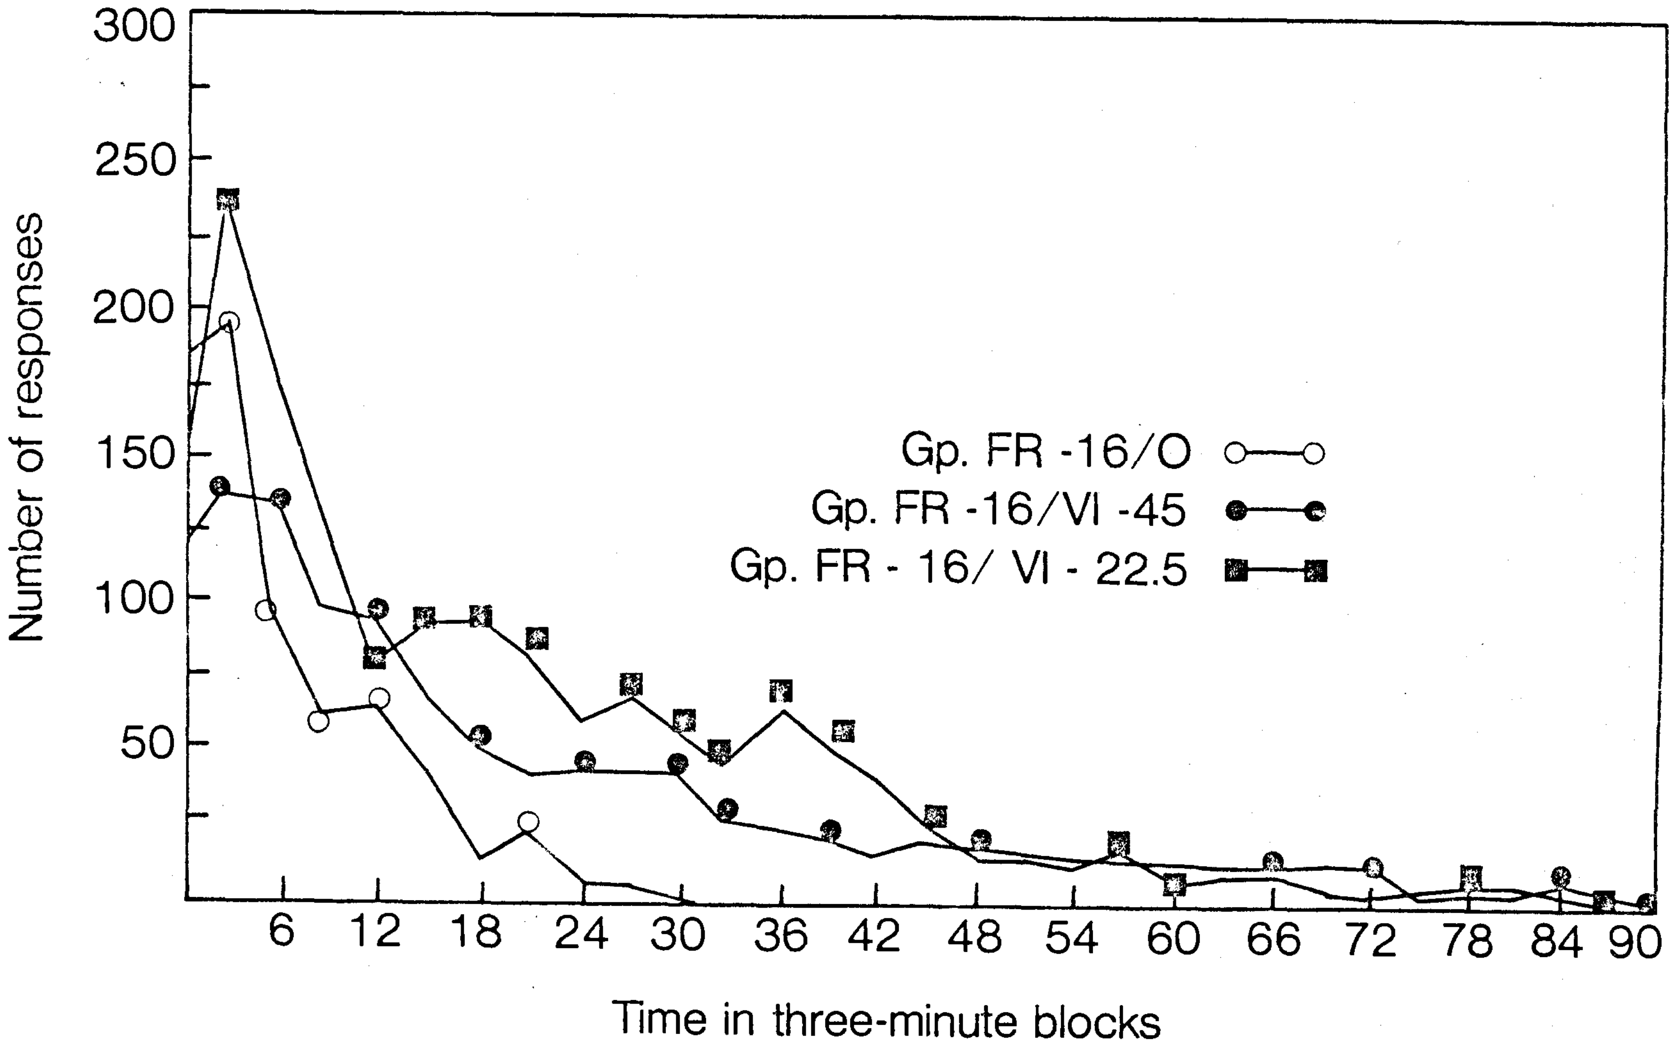

Fig. 9. Level of responding over consecutive 3-minute blocks in the target response during extinction, groups FR-16/0, FR-16/VI-45 and FR-16/VI-22.5.

rate at which the target response underwent decline compared to an extinction group that received no such reinforcements. Uhl and Garcia have obtained similar results even, paradoxically, when the presentation of reinforcement was explicitly dependent upon <u>not</u> emitting the target response for a period of time. They went on to show, according to Mackintosh (1974, p. 409), "that the basis for this effect was the maintenance of certain aspects of the condition prevailing during initial reinforcement of lever pressing." Hence, response-independent delivery of reinforcement precluded the generalization decrement that would ordinarily be expected when an important feature of an acquisition procedure is eliminated. Uhl and Garcia went on to demonstrate that prevention of the target response for 30 seconds following delivery of each response-independent reinforcer eliminated the effect of such deliveries on resistance to extinction.

Since both the foregoing procedures and the one employed in experiment II involved presentation of reinforcers during the extinction of some pre-established response, it is clearly possible that occurrence of those reinforcers played some role in enhancing resistance to extinction in the target response in the study undertaken herein. However, it is not abundantly clear why the presentation of such reinforcers had no such effect in experiment I, where they were not signalled.

A related explanation of the results of experiment II is suggested by a study conducted by Catania and Cutts (1963). They trained a pigeon to peck at two concurrently available response-keys for variable-interval reinforcement. Reinforcement was then omitted for pecks on one of the two keys and responses to that key were allowed to undergo extinction.

Catania and Cutts found that, despite the absence of reinforcement, the declines undergone by that response were remarkably superficial. On the assumption that these responses were somehow being maintained by the reinforcers yielded by the other alternative, Catania and Cutts imposed a changeover delay (COD, see chapter two) contingency of 1.5 seconds. The results were striking: the unreinforced response underwent rapid decline and ultimately extinguished.

The possible implications of the Catania and Cutts study for the results of experiment II are straightforward: the COD used in experiment II was .5 rather than 1.5 seconds and a .5 second COD may not have sufficiently promoted a clear and stable discrimination of the two programmed alternatives. If true, the occurrence of alternative reinforcement during extinction of the target response may have partially maintained the target response, thus enhancing its resistance to extinction. Again, however, it is not clear why a .5 second COD was effective in experiment I and not uniformly effective in experiment II.

A final consideration has to do with the contrasted conditions of responding inherent in experiment II. The reader will recall that the subjects of experiment II were confronted with two sets of choices during extinction of the target response. During most of the extinction session, they were permitted to choose between a "white" response-key, which was previously associated with a fixed-ratio schedule of reinforcement, and a "green" key, which was <u>negatively</u> correlated with alternative reinforcement. However, they were sometimes confronted with a choice between the "white" key and a "red" key, which was <u>positively</u> correlated with alternative reinforcement. While casual observation indicated that the "red" key was highly preferred by the subjects (i.e., responded to whenever present), the "green" key clearly was not. If conjecture is not out of place here, it may well be that the routine presence of a stimulus negatively correlated with reinforcement during the extinction of the target response enhanced, by contrast, the relative value of the "white"-signalled alternative, thereby augmenting resistance to extinction. This may make adaptive good sense: when confronted with a choice between an alternative which has been correlated positively with reinforcement and a second alternative which is correlated negatively with reinforcement, the best strategy may be to interact with the former alternative.

However explained, the results of experiment II may be limited to what might be called "extinction-naive" subjects. That is to say, subjects which are unfamiliar with the experimental particulars of the extinction operation. Since much of what has been proposed above in way of possible explanations of the results of experiment II stem in part from a "confusion" on the part of the subjects with respect to the causal details of the procedure used in experiment II, a more intimate knowledge of that procedure may produce a result more consistent with response re-allocation theory.

A preliminary investigation of this possibility was recently conducted by the author. Two groups of three pigeons each were trained on the same concurrent operants procedure used in experiments I and II. The target responses were reinforced on a fixed-ratio 8 schedule of reinforcement while the alternative response was reinforced on a VI-22.5 schedule. Alternative reinforcement was signalled in group "A" and unsignalled in group "B". Subjects were run through five cycles of conditioning and extinction sessions consisting of four sessions of conditioning followed by an extinction session and then, again, four sessions of conditioning and so on. The reinforcement magnitudes used were 3.0 seconds (target response) and 3.5 seconds (alternative response) access to food. Deprivation levels were the same as those used in experiments I and II. Training sessions ended following delivery of the 60th target response reinforcer and extinction sessions ended following emission of no more than three target responses during a six minute block of time.

The results of the foregoing experiment were provocative: although the group in which alternative reinforcement was unsignalled initially yielded fewer responses in extinction, the number of target responses emitted in extinction in the "signalled" group plummeted over repeated conditioning-extinction sessions to a point slightly beneath that of the unsignalled group in the fifth extinction session (the differences were not significant, however). Should this prove to be a robust, replicable finding, it seems to suggest that, while there seems to be a response competition component in extinction in extinction-naive subjects, the motivational interpretation of extinction inherent in response re-allocation theory may be sufficient to account for extinction in extinction-experienced subjects.

CHAPTER V

GENERAL DISCUSSION

Chapter one of the present work opened with an introductory quotation by N. J. Mackintosh (1974) in which that writer boldly called into question the explanatory effacacy of the principle of reinforcement in formally accounting for the facts of experimental extinction. The point of the present exercise was to show that there are logical and empirical reasons for questioning the cogency of Mackintosh's position. The point of departure of the present undertaking was the so-called matching principle for asymptotic operant behavior according to which organisms may be expected to distribute their responses across available alternatives in proportion to the fraction of available reinforcement yielded therefrom. Moreover, when that fraction of reinforcement is discriminably changed, organisms may be expected to redistribute their responses so as to preserve or reestablish the matching relation (cf. Herrnstein, 1979). We have referred to this process in the present text as "response re-allocation." Based on the logic of the matching principle, we have advanced a response re-allocation account of free-operant extinction according to which the changes in response frequencies observed in extinction are posited to be manifestations of the fundamental tendency to preserve the relationship between response output and reinforcement specified by the matching principle. Although the present focus is experimental extinction, the logical purview of the present account extends without modification to acquisition. Indeed, as with extinction, the changes in target and nontarget response frequencies observed in acquisition are, according to the present account, attributable to the fundamental tendency to maintain the matching relation. Hence, in a sense, acquisition and extinction do not exist except insofar as they are useful components of the behavioral taxonomies psychologists bring to bear on their laboratory experiences. As noted in chapter three, it is only after an experimenter defines a target response (explicitly or implicitly) that a behavioral procedure may, for the benefit of the experimenter, be identified as either "acquisition" or extinction." While such taxonomies are surely useful - and in some sense, appropriate in the learning laboratory - their utility and significance in natural settings may virtually pall in comparison.

It was noted earlier that an environment rich in viable alternatives would not, according to the present view, be expected to engender and robustly maintain an infrequently reinforced response. The fraction of total reinforcement yielded by such a response would simply be too small. By the same token, however, no single alternative in such a setting would be expected to occur at the single-minded, sometimes frenetic frequencies so often observed in the conditioning laboratory. To the extent that such animated performances are regarded as virtually synonomous with operant conditioning or, more generally, as the sine qua non of what is entailed by "conditioning", one might be tempted to conclude that operant conditioning is confined solely to the learning laboratory, which is surely not the case. However, it is clearly possible that many behavioral phenomena under scrutiny in the laboratory (and the partial-reinforcement effect in particular) have few specific analogues in the settings in which human and infrahuman organisms routinely conduct their affairs.

Returning for a moment to the position articulated by Mackintosh in the quotation opening chapter one, I submit that the principle of relative reinforcement inherent in the matching principle has clear and conspicuous formal implications for the nature of extinction and the robustness of resistance to extinction effects, as shown in detail herein. Apart from formal implications, the results of experiment I lend considerable empirical credence to this position. The results of experiment II do not, however, and much of what must be done in way of future research involves determining the generality of the effects observed in this experiment. As noted previously, there are cogent reasons for guardedly presuming that those effects are not particularly general. Indeed, they may be partly or wholly abolished in a procedure using longer CODs than those used herein, which were, it should be noted, shorter than those ordinarily used in concurrent operants research in order to preserve the ratio of responses to reinforcements established by the various fixed-ratio schedules for their respective target responses.

The response-enhancing effects of signalled alternative reinforcement observed in experiment II may also be partly or wholly abolished by associating discriminably different reinforcers with the target and nontarget responses. Rescorla and Skucy (1969) have shown that the delivery of response-independent reinforcers over the course of extinction enhances resistance to extinction by maintaining an aspect of the conditions prevailing during prior acquisition training. Using an alternative reinforcer which is discriminably different from that yielded by the target response may, in the experimental setting employed herein, preclude this effect by promoting a clearer discrimination of the alternatives.

The point of procedures such as the one above is to reasonably assure that the subjects of behavioral experiments "perceive" the relevant features of an experimental procedure in the way in which the experimenter presumes they do. Thus, for example, tones and colors are not used in experiments on subjects that cannot detect such stimuli. A presumption made in concurrent operants research, and in particular, the research undertaken herein, is that the programmed manipulanda constitute separate and separately reinforced alternatives. Changeover-delays are sometimes used in such research to underline this fact. Clearly, if it cannot be presumed that the experimenter and subjects at least roughly "agree" on what the particulars of the experimental procedure and setting are, the conclusions drawn from the data subsequently yielded may be puzzling and misleading. It is appropriate to make this point at this time simply because it may be premature to conclude with any finality that the results of experiment II falsify response reallocation theory. The subjects of that experiment may or may not have perceived the relevant features of that experimental procedure in the way in which the experimenter presumed it did. Clearly more research is needed here before this issue can be decisively resolved.

In any event, the author's preliminary investigation of the effects of signalled and unsignalled alternative reinforcement on resistance to extinction in extinction-experienced subjects would seem to indicate that the response-enhancing effects of signalled alternative reinforcement is transitory. As the reader will recall, after repeated cycles of conditioning and extinction, there were no statistically significant differences in resistance to extinction between "signalled" and "unsignalled" groups in that study. An upshot of this finding is that, insofar as extinction in natural settings ordinarily take place in organisms which are, in some sense, extinction-experienced, response re-allocation theory is probably a considerably more general account of extinction than those theories stressing response competition or interference.

This, however, is not to say that interference can play no role in the present account of experimental extinction. Future research may show that the present account must degenerate, at least in part, into a reinforcement-based interference theory of extinction. As noted earlier on response re-allocation theory bears some interesting similarities to interference theory and indeed, in some cases, improves upon that theory. For example, as Mackintosh (1974, p. 412) has noted, "interference theory must not only show that competing responses are the cause of the suppression of the originally reinforced response, it must also show why any new response should ever become sufficiently strengthened during extinction to compete with the original response." As Mackintosh (1974, p. 412) subsequently points out, it is not clear that interference theory succeeds in doing this. The present model of extinction, however, does. As pointed out previously, response re-allocation theory explains why the target response must decline in extinction and why the various nontarget responses must rise in frequency over the course of extinction. In its present form, response re-allocation theory does not invoke competing responses, but it does presume that responses compete for available reinforcement. However, should it become necessary to invoke competing responses, the present model would have no difficulty explaining why they "... become sufficiently strengthened during extinction to compete with the original response."

A final comment is in order regarding the relationship of response re-allocation theory to generalization decrement theory (in its generic sense) and, in particular, to Capaldi's (1971) sequential-aftereffects theory of extinction. According to generalization decrement theory, responses persist in extinction because the stimulus conditions prevailing therein are similar to those which prevailed during acquisition. The more similar the former conditions are to the latter, the greater the resistance to extinction expected. However, the repeated emission of unreinforced responses in extinction guarantees that stimulus conditions will eventually change enough to become discriminably different from acquisition, at which point responses undergo progressive decline. From the perspective of generalization decrement theory, resistance to extinction may be enhanced by either making extinction more similar to acquisition (as Rescorla and Skucy (1969) did by presenting response-independent reinforcers during extinction) or by making acquisition more similar to extinction (by, e.g., introducing a large number of unreinforced trials in acquisition). Capaldi's sequential-aftereffects theory, which is now the most widely accepted theory of extinction (Tarpy and Mayer, 1978), is aimed at explaining why and how different patterns of unreinforced trials during acquisition subsequently enhance resistance to extinction.

There can be little doubt that generalization decrement, in one or another form, plays some role in extinction. Certainly, the welldocumented empirical viability of Capaldi's sequential-aftereffects theory is not questionable. However, from the perspective on response re-allocation theory, its generality (as noted previously) surely is. As noted in several places above, highly intermittently reinforced responses are apt not to become very interesting to subjects in settings rich in viable alternatives according to the account of extinction advanced herein. Hence, as pointed out previously, the robust resistance to extinction effects associated with prior training on schedules of intermittent reinforcement are, according to the present account, limited to those settings wherein, for an identifiable target response, the level of background reinforcement is relatively meager. Environments rich in viable alternatives would ordinarily not be expected to robustly give rise to intermittently reinforced responses in the first place. It follows that the variables most emphasized by Capaldi, dealing, as they do, with the effects of patterns of intermittent reinforcement on subsequent resistance to extinction, are for the most part limited in their more marked effects to settings in which background reinforcement is, and remains, meager. The generality of Capaldi's theory is therefore dependent on the typicality of such environments. According to some writers (e.g., Burstein, 1976), such environments are notably *atypical*.

Conclusion

The substance of the present work is, in comparison to the considerable empirical work that remains to be done, largely theoretical and methodological. The present work has, however, succeeded in laying the foundation for that remaining work and has taken the initial empirical steps in an experimental program that will either substantiate or ultimately falsify the matching-based account of extinction articulated within these pages. It is upon the results of that program that the theoretical future of that account in large part rests.

REFERENCES

Amsel, A. The role of frustrative nonreward in noncontinuous reward situations. Psychological Bulletin, 1958, 55, 102-119.

Amsel, A. Frustrative nonreward in partial reinforcement and discrimination learning. Psychological Review, 1962, 69, 306-328.

Amsel, A. Partial reinforcement effects on vigor and persistence. In K. W. Spence and J. T. Spence (Eds.), The psychology of learning and motivation (Vol. 1). New York: Academic Press, 1967.

Bacon, W. Partial-reinforcement extinction effect following different amounts of training. Journal of Comparative and Physiological Psychology, 1962, 55, 998-1003.

Baum, W. Choice in a continuous procedure. Psychonomic Science, 1972, 28, 263-265.

Baum, W. The correlation-based law of effect. Journal of the Experimental Analysis of Behavior, 1973, 20, 137-153.

Baum, W. Choice in free-ranging wild pigeons. Science, 1974, 185, 78-79.

Baum, W., & Rachlin, H. Choice as time allocation. Journal of the Experimental Analysis of Behavior, 1969, 12, 86-874.

Bijou, S., & Baer, D. Child Development I: A systematic and empirical theory. New York: Appleton-Century-Crofts, 1961.

Bindra, D. Components of general activity and the analysis of behavior. Psychological Review, 1961, 68, 205-215.

Bitterman, M. Phyletic differences in learning. American Psychologist, 1965, 20, 396-410.

Boakes, R. Response decrements produced by extinction and by response-independent reinforcement. Journal of the Experimental Analysis of Behavior, 1973, 19, 293-302.

Brownstein, A., & Pliskoff, S. Some effects of relative reinforcement rate and changeover delay in response-independent concurrent schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 1968, 11, 683-688.

Burstein, K. On the generality of the partial reinforcement effect. Bulletin of the Psychonomic Society, 1976, 7, 501-502.

Catania, A. C. Concurrent performances: A baseline for the study of reinforcement magnitude. Journal of the Experimental Analysis of Behavior, 1963, 6, 299-300.

Catania, A. C. Concurrent operants. In W. K. Honig (Ed.), Operant behavior: Areas of research and application. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1966.

Catania, A. C. Self-inhibiting effects of reinforcement, Journal of the Experimental Analysis of Behavior, 1972, 18, 231-241.

Catania, A. C. Learning. Englewood Cliffs, N.J.: Prentice-Hall, 1979.

Catania, A. C., & Cutts, D. Experimental control of superstitious responding in humans. Journal of the Experimental Analysis of Behavior, 1963, 6, 203-208.

Catania, A. C., & Reynolds, G. A quantitative analysis of the responding maintained by interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 1968, 11, 327-383.

Capaldi, E. J. Partial reinforcement: a hypothesis of sequential effects. Psychological Review, 1966, 73, 459-477.

Capaldi, E. J. Memory and learning: a sequential viewpoint. In W. K. Honig and P. H. R. James (Eds.), Animal memory. New York: Academic Press, 1971.

Chung, S.-H. Some quantitative laws of operant behavior. Unpublished doctoral dissertation, Harvard University, 1966.

Chung, S.-H., & Herrnstein, R. J. Choice and delay of reinforcement. Journal of the Experimental Analysis of Behavior, 1967, 10, 67-74.

Coughlin, R. The frustration effect and resistance to extinction as a function of percentage of reinforcement. Journal of Experimental Psychology, 1970, 84, 113-117.

Davenport, J. Instrumental runway performance as a function of temporal distribution of reinforcement. Journal of Comparative and Physiological Psychology, 1964, 57, 205-210.

Davenport, J., Goodrich, K., & Hagguist, W. Effects of magnitude of reinforcement in Macaca speciosa. Psychonomic Science, 1966, 4, 187-188.

de Villiers, P. Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig and J. E. R. Staddon (Eds.), Handbook of operant behavior. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1977.

de Villiers, P., & Herrnstein, R. J. Toward a law of response strength. Psychological Bulletin, 1976, 83, 1131-1153.

Dohahoe, J., & Wessells, M. Learning, language and memory. New York: Harper and Row, 1980.

Deluty, M. Similarities of the matching law to other models of conditioning. Psychological Record, 1977, 3, 599-602.

Estes, W. K. Statistical theory of spontaneous recovery and regression. Psychological Review, 1955, 62, 145-154.

Estes, W. K. The statistical approach to learning theory. In S. Koch (Ed.), Psychology: A study of a science (Vol. 2). New York: McGraw-Hill, 1959.

Estes, W. K. Probability learning. In A. W. Melton (Ed.), Categories of human learning. New York: Academic Press, 1964.

Fantino, E., & Duncan, B. Some effects of interreinforcement time upon choice. Journal of the Experimental Analysis of Behavior, 1972, 17, 3-14.

Fantino, E., & Logan, C. The experimental analysis of behavior: A biological perspective. San Francisco: Scott-Foresman, 1979.

Gilbert, R. Variation and selection of behavior. In R. M. Gilbert and J. R. Millenson (Eds.), Reinforcement: Behavioral analyses. New York: Academic Press, 1972.

Gleitman, H. Forgetting of long-term memories in animals. In W. K. Honig and P. H. R. James (Eds.), Animal memory, New York: Academic Press, 1971.

Guthrie, E. The psychology of learning. New York: Harper, 1935 (revised 1952).

Guttman, N. Equal-reinforcement values for sucrose and glucose solutions compared with equal-sweetness values. Journal of Comparative and Physiological Psychology, 1954, 47, 358-361.

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

Herrnstein, R. J. Method and theory in the study of avoidance. Psychological Review, 1969, 76, 49-69.

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

Herrnstein, R. J. Quantitative hedonism. Journal of Psychiatric Research, 1971, 8, 399-412.

Herrnstein, R. J. Formal properties of the matching law. Journal of the Experimental Analysis of Behavior, 1974, 21, 159-164.

Herrnstein, R. J. Derivatives of matching. Psychological Review, 1979, 86, 486-495.

Herrnstein, R. J., & Loveland, D. Maximizing and matching on concurrent ratio schedules. Journal of the Experimental Analysis of Behavior, 1975, 24, 107-116.

Heyman, G. Markov model description of changeover probabilities on concurrent variable-interval schedules. Journal of the Experimental Analysis of Behavior, 1979, 31, 41-51.

Hilgard, E., & Bower, G. Theories of learning (4th ed.). Englewood Cliffs, N.J.: Prentice-Hall, 1974.

Homans, G. Social behavior: Its elementary forms. New York: Harcourt Brace Jovanovich, Inc., 1974.

Hull, C. Principles of behavior. Englewood Cliffs, N.J.: Prentice-Hall, 1943.

Humphreys, L. The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. Journal of Experimental Psychology, 1939, 25, 141-158.

Jones, E., & Bridges, C. Competing responses and the partial reinforcement effect. Psychonomic Science, 1966, 6, 483-484.

Killeen, P. Preference for fixed-interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 1970, 14, 127-31.

Kimble, G. Hilgard and Marquis' Conditioning and learning (2nd ed.). New York: Appleton-Century-Crofts, 1961.

Kuhn, . The structure of scientific revolutions. Chicago: University of Chicago Press, 1970.

Kulli, T. Unpublished experiment reported in de Villiers, P., Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig and J. E. R. Staddon (Eds.), Handbook of operant behavior. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1977.

LaBounty, C., & Reynolds, G. An analysis of response and time matching to reinforcement in concurrent ratio-interval schedules. Journal of the Experimental Analysis of Behavior, 1973, 19, 155-166.

Lawrence, D., & Festinger, L. Deterrents and reinforcement. Stanford: Stanford University Press, 1962.

Leitenberg, H., Rawson, R., & Bath, K. Reinforcement of competing behavior during extinction. Science, 17 July, 1970.

Lundin, R. Personality: A behavioral analysis. New York: Macmillan, 1974.

Luthans, F., & Kreitner, R. Organizational behavior modification. Glenview, Ill.: Scott, Foresman and Co., 1975.

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

McSweeney, F. Matching and contrast on several concurrent treadle-press schedules. Journal of the Experimental Analysis of Behavior, 1975, 23, 193-198.

Millenson, J. R. Principles of behavioral analysis. New York: McMillan, 1967.

Nevin, J. Interval reinforcement of choice behavior in discrete trials. Journal of the Experimental Analysis of Behavior, 1969, 12, 875-885.

Nevin, J. Overall matching versus momentary maximizing: Nevin (1969) revisited. Journal of Experimental Psychology: Animal Behavior Processes, 1979, 5, 300-306.

Pavlov, I. Conditioned reflexes. Oxford: Oxford University Press, 1927.

Rachlin, H. Behavior and learning. San Francisco: W. H. Freeman and Co., 1976.

Rachlin, H., & Baum, W. Effects of alternative reinforcement: Does the source matter? Journal of the Experimental Analysis of Behavior, 1972, 18, 231-241.

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

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

Revusky, S., & Garcia, J. Learned associations over long delays. In G. Bower (Ed.), The psychology of learning and motivation (Vol. 4), New York: Academic Press, 1970.

Reynolds, G. On some determinants of choice in pigeons. Journal of the Experimental Analysis of Behavior, 1963, 6, 53-59.

Schoenfeld, W., & Farmer, J. Reinforcement schedules and the "behavior stream." In W. Schoenfeld (Ed.), The theory of reinforcement schedules. New York: Appleton-Century-Crofts, 1970.

Schroeder, S., & Holland, J. Reinforcement of eye movement with concurrent schedules. Journal of the Experimental Analysis of Behavior, 1969, 12, 897-903.

Shimp, C. Probabilistically reinforced choice behavior in pigeons. Journal of the Experimental Analysis of Behavior, 1966, 9, 433-455.

Shimp, C. Optimum behavior in free-operant experiments. Psychological Review, 1969, 76, 97-112.

Silberberg, A., & Fantino, E. Choice, rate of reinforcement, and the changeover delay. Journal of the Experimental Analysis of Behavior, 1970, 13, 187-197.

Skinner, B. F. The behavior of organisms. New York: Appleton-Century-Crofts, 1938.

Skinner, B. F. Science and human behavior. New York: Macmillan, 1953.

Skinner, B. F. Verbal behavior. New York: Appleton-Century-Crofts, 1957.

Skinner, B. F. About behaviorism. New York: Knopf, 1974.

Spear, N. Forgetting as retrieval failure. In W. K. Honig and P. H. R. James (Eds.), Animal memory. New York: Academic Press, 1971.

Spence, K. The role of secondary reinforcement in delayed reward learning. Psychological Review, 1947, 54, 1-8.

Staats, A. Learning, language and cognition, New York: Holt, Rinehart and Winston, 1968.

Staddon, J. E. R., & Simmelhag, V. The "superstition" experiment: A re-examination of its implications for the principles of adaptive behavior. Psychological Review, 1971, 78, 3-43.

Stubbs, D., & Pliskoff, S. Concurrent responding with fixed relative rate of reinforcement. Journal of the Experimental Analysis of Behavior, 1969, 12, 887-895.

Tarpy, R. Basic principles of learning. Palo Alto: Scott, Foresman, 1975.

Tarpy, R., & Mayer, R. Foundations of learning and memory. Palo Alto: Scott, Foresman, 1978.

Terrace, H. By-products of discrimination learning. In G. Bower and J. Spence (Eds.), The psychology of learning and motivation. New York: Academic Press, 1972.

Uhl, C., & Garcia, E. Comparison of omission with extinction in response elimination in rats. Journal of Comparative and Physiological Psychology, 1969, 69, 554-562.

Ullman, L., & Krasner, L. A psychological approach to abnormal behavior. Englewood Cliffs, N.J.: Prentice-Hall, 1969.

Wendt, G. An interpretation of inhibition of conditioned reflexes as competition between reaction systems. Psychological Review, 1936, 43, 258-281.