The Revealed Operant:
A Way to Study
the Characteristics
of Individual
Occurrences of
Operant Responses
Third Edition

Dr. Francis Mechner

Commentary:

Created, Revealed, and Imposed Operants

Dr. Donald M. Baer

The Shape of Things to Come:
A Commentary on Mechner's
"The Revealed Operant"

Dr. M. Jackson Marr

The Discriminated Operant
Revealed in Pigeon Multiple
Schedule Performance: Effects of
Reinforcement Rate and Extinction

Dr. John A. Nevin

Shaping in Terms of Parameter Shifts

Dr. Thom Verhave

Response to the Commentaries

Dr. Francis Mechner

PREFACE

About six years ago, upon accepting an appointment to the Advisory Board of the Cambridge Center for Behavioral Studies, I asked Robert Epstein to consider my acceptance an offer to contribute in some practical way towards the Center's goals. Not long after that, Dr. Epstein telephoned me to "take me up on my offer." He asked if I would serve as editor of a book comprised of chapters by authors whose work had been featured at the Center's annual meetings. My role would be to telephone the authors, round up the papers, solicit commentary, and do the substantive editing. Dr. Epstein wanted everything at the Center within 6-8 months. Having just assumed the editorship of *The Behavior Analyst*, I had some trepidation about taking on a concomitant editorial task. Reassured by Robert that this was a small task compared to serving as journal editor (and recalling that I had, after all, volunteered to help), I agreed to edit the proposed book.

During the course of the succeeding six years, the book became a monograph series; only one of the original chapters was ever delivered to me; and my own journal editorship was completed as was that of my successor, Sam Deitz. Tony Nevin, like his predecessor Robert Epstein, has completed countless initiatives undertaken as Executive Director of the Center. Similarly, we at the University of North Texas began a master's degree program in behavior analysis, developed a dozen courses, established a human operant laboratory and initiated research programs in several areas. The book cum monograph series appears to have provided the greatest challenge for us all.

The present monograph, the third in a series of three, began as a chapter on performance learning, the topic of Francis Mechner's presentation at a CCBS annual meeting. Despite a number of prior commitments, Francis Mechner (whom I knew only as a voice over the telephone) managed to work on the paper during the succeeding two years. After about the third revision, I convinced Francis to attend the annual meeting of the Association for Behavior Analysis; we agreed to meet there and discuss his chapter. During that conversation, Francis described to me his interest in re-opening a line of research he had begun nearly 30 years ago--research on what he later came to call "the revealed operant."

Although our faculty were involved in a number of other research projects, I encouraged them to consider collaborating with Francis to investigate the revealed operant. After a series of meetings, Dr. Cloyd Hyten and some of our graduate students began work in earnest on that topic. Whether or not research in this area remains on Dr. Hyten's agenda, I am pleased that he pursued the opportunity and gave our students the opportunity to work with Francis Mechner in the lab.

So it happened that the chapter on performance learning became a monograph on the revealed operant. The countless revisions which I read and commented on afforded me a rare opportunity to observe at very close range the development of a work that could well be seminal. As may be seen from Nevin's commentary, other researchers can use the revealed operant methodology in the creative exploration of topics of importance to them. The commentaries of Baer, Marr, and Verhave, as well as Nevin's, each bear on the theoretical relevance of the revealed operant to topics both new and old in behavior analysis.

Neither this monograph nor those featuring articles by Joe Brady and by Scott Geller would ever have come to fruition without the active support of Tony Nevin, Executive Director, and Adrienne Forstner-Barthell, Catherine I. Kaczowka, and Thomas Legg, Monograph Series Managers, of the Cambridge Center for Behavioral Studies. Their efforts have been unstinting, and I thank them for making it possible for me to realize my commitment to act in some way that supported the Center.

Sigrid S. Glenn
University of North Texas

TABLE OF CONTENTS

TITLE   **PAGE**

The Revealed Operant:
A Way to Study the Characteristics
of Individual Occurrences
# of Operant Responses

Francis Mechner

ABSTRACT

An operant is a "revealed operant" (rO) if it is specified as a sequence of more than two *recorded* sub-operants. The first and last sub-operants must be made on different manipulanda. An rO differs from traditional operants (those that are recorded as single instantaneous events) in that more than one of its sub-operants is recorded. This feature makes the rO a practical model for studying the structure of individual occurrences of operants. For each occurrence, several types of behavioral measures, including duration, sub-operant patterns, and the operant's internal structure, can be examined.

The monograph describes a practical method for implementing rOs, and discusses how the rO model can be used to address various issues in the field of behavior research. Among these are reinforcement, behavior shaping, variability, acquisition, extinction, automatization, punishment, choice behavior, equivalence, transition behavior, behavioral cyclicity, and drug effects.

CONTENTS

FOREWORD AND ACKNOWLEDGMENTS

This monograph is an important milestone for me because it is my first basic research publication in about 30 years. It picks up where I left off in 1961 when I left the laboratory for other endeavors, and builds in part on work I did prior to that time.

During the intervening years, I worked mostly in various areas of behavioral technology and technological entrepreneurship. My work in instructional technology and my life-long interest in pianism contributed to my growing preoccupation with performance learning. For over a decade now, I have been trying to understand how skilled performances are learned and how performance learning and practicing can be made more efficient. Those efforts led me back to some issues that seem to require the revealed operant as a research tool.

F.M.

SUMMARY

The "revealed operant" ("rO"), is a model of any operant. A single occurrence of an rO is comprised of a sequence of switch closures. The value of the rO as a research tool lies in the fact that it permits the internal structure and properties of individual occurrences of operant responses to be recorded and studied. Traditional operant behavior experiments, by contrast, record operant responses only as single instantaneous events ("iO"s), like a single closure of a switch. The justification for the rO model is that all operants, including iOs, are composed of sub-operants. An rO differs from an iO in that some of an rO's sub-operants can be recorded conveniently.

The rO makes it possible to address certain research questions that cannot easily be addressed by means of iOs, especially questions where molar phenomena require molecular data for their explanation. Examples of such questions are: What is the local effect of a single reinforcer presentation? What are the mechanisms of behavior shaping? Are there long-term drifts in the characteristics of operants, corresponding to automatization? What kinds of resurgence patterns occur during extinction? Is a change in response rate due to changed response durations or to changed IRTs? In what ways and how quickly is an operant impacted when a new contingency is introduced, as in acquisition?

Every occurrence of an rO consists of an initiating sub-operant, R_a, and a terminating sub-operant, R_c. Between the R_a and the R_c there can be a sequence of R_bs, a required wait, or any other specified behavior, according to how the experimenter chooses to specify the rO.

A variety of behavioral measures can conveniently be recorded for each individual occurrence of an rO. These fall into two categories, criterial and non-criterial measures, which have different properties. Criterial measures are taken in the dimension within which the operant is specified, like the number of R_bs required, and non-criterial measures are taken in unspecified dimensions. For rOs, examples of non-criterial measures are rhythmic and sequential patterns of R_bs within each rO, and execution speed (duration) of the rO. In a free operant situation, another measure is the time between the end of one rO and the start of the next. Thus, the rO technique provides a variety of measures that can be obtained for each rO occurrence. Some independent variables impact some of these measures in similar ways, and can therefore be grouped into families according to their impact profiles. The resulting groupings can then function as empirically-based constructs in quantitative theoretical formulations.

A technique for implementing rOs is described, and applications to various problems are discussed. Among these are: The experimental analysis of the shaping process by studying the effects of single reinforcer presentations on individual occurrences of operants; the independent variables on which these effects depend; the study of long-term changes in the properties of operants; the possibility of correcting for response duration in mathematical formulations like Herrnstein's matching law; and the speed with which abruptly introduced variables produce behavioral changes.

1.0 INTRODUCTION

1.1 Specified and Unspecified Dimensions of Operants

A particular operant is a class of behavior that produces a particular effect on the environment. In most traditional operant behavior experiments, the effect is one closure of a switch, and the operant is registered at the instant of that closure, as an all-or-none event (the effect).

Every operant is preceded by *sub*-operants. For example, the rat must place a paw on the response bar before it can press the bar and close the switch. Sub-operants are physically necessary antecedents of the final effect. Furthermore, the sub-operants can always consist of a wide range of unspecified movement patterns and response topographies (e.g., the rat can use the right paw, left paw, or teeth to press the bar).

Normally, experimenters do not specify the sub-operants deliberately, but nonetheless constrain them, usually inadvertently, by the physical construction of the response device and the surrounding environment. For example, the position and shape of the bar constrain the range of body and paw movements that can place a paw on top of the bar. And the distance the bar must then be moved downward, and the force required to move it, constrain the range of possible further movements that can result in switch closure. Those are examples of the physical constraints that determine the sub-operants in traditional operant behavior experiments.

In such experiments, the sub-operants are not easily recordable and are normally disregarded.[1] In general, only the last member of the sequence of sub-operants is recorded as the single instantaneous event that defines the operant.

1.2 Revealed Operants

The revealed operant permits the sub-operants to be recorded conveniently, thereby making the internal structure of the operant accessible for quantitative scientific study. The term "revealed operant" was chosen to stress the fact that the normally-unregistered sub-operants are specified in a way that reveals them, i.e., permits the experimenter to record them.

The notion that an operant can be specified and recorded either as a single instantaneous event or as a sequence of sub-operants is not new (Mechner, 1959a; Morse, 1966, p. 103). The purpose here is to describe some ways to use revealed

---

[1] In principle, it would be possible to take motion pictures of the movements comprising the sub-operants, and analyze them frame by frame, but such a process is too cumbersome and expensive to be practical.

operants as a practical research technique for addressing various problems that cannot be addressed by means of conventional instantaneously recorded operants.

An operant is a revealed operant if it is specified as a sequence of at least two *recorded* sub-operants. The one that initiates the operant must be made on a separate manipulandum, different from the one on which the other sub-operants are made, so that the start of the operant is marked by an unambiguous behavioral event. It is also generally useful to design the revealed operant so that a further distinctive behavioral event marks its termination. Examples of different types of revealed operants are presented and discussed in Chapters 3-5.

Henceforth, we will use the abbreviation **"rO"** for revealed operant, and **"iO"** for "instantaneously recorded operant."

1.3 Relationship Between rOs and iOs

The rO reveals phenomena that are present in iOs too, but not easily observable by ordinary methods. The sub-operants comprising an rO are analogous to the normally-unrecorded sub-operants of a conventional iO like bar pressing or key pecking. That is why rOs may be considered laboratory models of iOs or of any operants.

The fact that iOs are recorded as instantaneous events does not mean that they are necessarily simpler or shorter than rOs. An rO, like a typical iO, may take less than one second to execute. The difference between them lies only in whether the experimenter (a) specifies the sub-operants deliberately, or (b) constrains them inadvertently by way of the physical structure of the manipulandum and the experimental environment. Making the specification deliberate rather than inadvertent does not make the operant more complex, and should also not affect its generalizable properties in any way. (Donald Baer's commentary at the end of this monograph, and my response to it, take up the issue of whether an rO is a valid laboratory model of an iO, given that its sub-operants are deliberately installed.)

Some conventional reinforcement schedules in which iOs are used may appear to be superficially similar to rOs. But the behavior produced by a schedule is equivalent to an rO only if that behavior is initiated by a distinctive response, that being the minimal defining feature of an rO.

1.4 Behavioral Measures Provided by rOs

The rO makes available a variety of new behavioral measures that can be applied to individual occurrences of any rO (Refer to Figure 1 below).[2]

(a) D - "Duration"

This is the time consumed by the execution of the operant. When iOs are used, that time is unrecorded.

(b) L - "Latency"

This measure is like the IRT in the free operant situation for iOs, with duration subtracted out.[3]

(c) 1/L - "Tempo"

Since this measure excludes D, it requires a new name like "tempo." The term "rate" would imply a continuous time base. Tempo can replace iO response rate in some theoretical formulations (see Section 6.1). Straight iO response rate would be $1/(D + L)$.

(d) D/(D + L) - "Engagement"

This is the duration as a fraction of the total cycle time. It is the percentage, on a time basis, of all ongoing behavior comprised by the operant.[4]

(e) R_a-R_b Interval - "Initiation Time"

This time interval may prove to be selectively sensitive to such variables as complexity or aversiveness of the rest of the rO.

---

[2]Since these behavioral measures are applicable to individual occurrences of rOs, they can be used in either a free operant situation or a trial situation.

[3]The term IRT is generally applied to the end-to-end cycle time, the average IRT being the reciprocal of iO response rate. In two recent papers (Johnston and Hodge, 1989; Marr, 1990), IRT was defined as excluding response duration, like our present latency measure. However, we will continue to adhere to the conventional usage of the term IRT as referring to cycle time.

[4]The "kinetic output" measure proposed by Powell and Dickie (1990) is somewhat similar to this measure.

(f) Sub-Operant Patterns

A pattern is a specific recurring sub-operant sequence or rhythm within an rO. A series of occurrences of an rO contains a variety of patterns. Patterns are useful in studying variability and stereotypy, resurgence, and effects of individual presentations of reinforcers.

Figure 1. Generic diagram of a complete individual occurrence of an rO and the end of the preceding occurrence. The diagram shows the time relationships between the initiating sub-operant (called "**R_a**"), and the other possible sub-operants that can comprise the rO.

There are also other possible measures that rOs make available, and there are many ways to combine these measures into more elaborate ones.[5] For example, the measure $1/L^2$ may prove interesting. In physics, energy is $mv^2/2$, and since v is analogous to $1/L$, v^2 is analogous to $1/L^2$. If this expression were multiplied by a credible behavior analog of mass, we would have an energy analog. The usefulness of any such measure will depend on its selective sensitivity to some defined group of variables.

1.5 Criterial and Non-Criterial Measures of rOs

Every individual occurrence of an rO can be described in terms of the rO's *specified*, or criterial, measures, and in terms of its various *unspecified*, or non-criterial, measures (Herrnstein, 1966, p. 38). Most rOs have only one criterial dimension, the dimension in which the operant is specified.

---

[5]Johnston & Hodge (1989) discuss the possibility of using various behavioral measures based on the (unrecorded) iO response duration. Their paper takes on new significance in the context of rOs, where response duration can actually be recorded.

Examples of non-criterial measures are: Time the bar is held down when the operant is bar release (Margulies, 1961), where the time is unspecified; the sequence in which four keys are pressed, where the sequence is unspecified (Bruner & Revusky, 1961); position along a long horizontal slot where the rat could poke its nose through (Antonitis, 1951); and position along a long horizontal strip where a pigeon could peck (Herrnstein, 1961). In the last two studies, all positions were equivalent for reinforcement.

Familiar examples of criterial measures are: Force exerted on the bar (Skinner, 1938, p. 312; Notterman, 1959); length of time the bar is held down, where a minimum time was specified (Skinner, 1938, pp. 328-338); panel displacement, where a minimum displacement was required (Winnick, 1949); and thumb muscle potentials, where a minimum potential was required (Hefferline, Keenan, & Harford, 1959).[6]

Every operant has innumerable non-criterial dimensions. The rO technique makes it convenient to measure and record several of these for each individual occurrence of an rO. Among these are speed with which the required sub-operant sequence is executed, or the particular (optional) sequences and rhythms used by the subject for the R_bs, i.e., the sub-operant patterns. The measures described in Section 1.4 above are non-criterial except for L, which is not part of the rO itself. Note that a non-criterial measure does not become criterial just by virtue of "superstitious" effects resulting from unprogrammed impacts of reinforcer occurrences. Only the experimentally arranged contingencies determine an operant's criterial dimensions.

Criterial and non-criterial measures are differentially sensitive to various types of variables (for an example, see Section 3.4, last paragraph) and it is therefore important to distinguish between them. Guthrie was clearly referring to non-criterial measures when he wrote that for psychology to make real progress, psychologists will need to measure "movement series," "movement patterns," and "partial or subresponses," and not just the all-or-none occurrence of "acts" (e.g., Guthrie, 1959, pp. 184-185).[7]

---

[6]The dichotomy between criterial and non-criterial measures can be conceptualized as the two extremes of a continuum defined in terms of the range of allowable effects that define the operant. Thus, intermediate cases between criterial and non-criterial measures can be created by assigning different probabilities of meeting the criterion to different values of a dimension. For example, instead of *every* position of the slot or strip being 100% effective, some positions could made be less effective, thereby converting the position dimension into more of a criterial dimension.

[7]Catania (1973) used the term "descriptive operant" for the effect that defines the operant, i.e., the criterion, and the term "functional operant" for the actual occurrences, whose criterial measures may or may not meet the criterion.

Among the important non-criterial measures that rOs provide are sequential and rhythmic patterns of sub-operants. One can look for recurring patterns in successive individual occurrence of an rO. Such a recurrent pattern could consist of pressing the same 3 or 4 R_b keys in the same sequence, separated by time intervals that define a recognizable recurrent rhythm. Such patterns can be detected and analyzed by computer (Field, Mechner, Hyten, & Madden, 1991; Mechner, Hyten, Field, & Madden, 1992).

An important parameter of non-criterial measures is how far before the end of the operant they occur. It appears that the greater that distance, the greater the measure's sensitivity to certain variables (Mechner et al., 1992).

1.6 Comparison of rO Measures and iO Measures

Formulations that use rO measures are likely to have more theoretical generality than iO response rate or IRTs (the reciprocal of response rate). IRTs are the sum of the Ls and the Ds, which are not recorded separately for iOs. IRTs therefore confound L and D. So, if one records only IRTs, one cannot observe differences in the L's and D's differential sensitivity to different groups of independent variables. Also, one can expect the Ls of an rO to be more reliable and less variable from instance to instance than an iO's individual IRTs because:

(a) Variability of iO response topography can produce extremely long or extremely short IRTs. For example, key flicking by pigeons and bar gnawing by rats can result in bursts of rapid-fire switch closures. At the other extreme, various cumbersome types of iO topographies can result in very long IRTs.

(b) Any iO response topography can result in multiple consecutive abortive or incomplete responses that register only as very long IRTs. With rOs, abortive or incomplete occurrences are recorded as such. A succession of operants often shows a cyclic or wavy pattern, with the upper parts of the waves exceeding the criterion value, and the rest falling below it, (e.g., Mechner, 1958b; Mechner et al., 1992). With iOs, the response rate measure registers only the above-criterion instances. With rOs, on the other hand, both above- and below-criterion instances are registered.

Some variables that increase response rate (for example, certain drugs), also increase the frequency of abortive occurrences (e.g., Mechner, Snapper, & Ray, 1961; Mechner & Latranyi, 1963). An abortive occurrence is one where the criterial measure falls short of the criterion. With iOs, only those occurrences that meet or exceed the criterion are recorded. The recorded response rate will therefore decrease as the frequency of abortive instances increases. So, the iO response rate cannot be expected to reflect the effects of such variables in a simple way. With rOs, on the other hand, abortive or incomplete instances (as when the number of R_bs falls short of the criterion) *are* recorded, and can then be treated in any way desired (e.g., by being eliminated or included).

1.7 The Value of Molecular Information

The main value of the rO methodology is that it provides a practical way to obtain information about individual occurrences of operants. This type of information could be termed "molecular," in contradistinction to the molar type of information provided by response rate measures or other measures that depend on the pooling of multiple operant occurrences. Much of the present body of data comprising behavior analysis is of the molar type.

Molecular information often provides the key to understanding and explaining molar phenomena. Numerous examples are discussed in the chapters that follow. Examples can also be seen in other sciences: Our understanding of the properties of substances, such as hardness, viscosity, adhesion, lubricity, color, etc., is based on information about molecular and atomic structure. Similarly, our understanding of genetics is based heavily on information about genes and the structure of the DNA molecule. It is likely that molecular information about operants will similarly shed light on the molar behavioral phenomena that interest us.

2.0 FUNCTIONS OF THE INITIATING SUB-OPERANT R_a

2.1 Functions of R_a and its Essential Attributes

One of the keys to the rO's usefulness is its defining feature, that each instance must be initiated by a distinctive R_a. The purposes of the R_a are (1) to mark the end of the L and the initiation of the rO as cleanly as possible, (2) to make different types of rOs comparable, and (3) to make the "state of the organism" at the start of each rO as uniform as possible. In order to perform these functions, the R_a should have two attributes:

Attribute 1: Simplicity: It is simple and effortless compared to the rest of the rO.

Attribute 2: Standardization: It is made on a separate manipulandum, dedicated to R_a.

2.2 Functions of the Simplicity Attribute: Clean Separation of L and D

Keeping R_a simple and effortless minimizes the various contaminating effects that R_a will have on L and D. The recorded L should be minimally contaminated by behavior comprising the rO. This is important in order to protect the L's profile of sensitivity to independent variables.

The problem is that no matter how fast and effortless it is, the R_a still has its own sub-operants (which could thus be viewed as sub-sub-operants of the overall rO). The R_a's sub-operants are an unrecorded and unspecified part of the rO. Since they occur just prior to the recorded R_a, but are unrecorded and unspecified, their unknown duration is unavoidably included in the recorded L (see Figure 1).

That is why we want the duration of R_a's sub-operants to be as brief as possible in relation to the total L. We can assume that the simpler and less effortful the R_a, the shorter will be the duration of its sub-operants (for reasons discussed in Sections 7.2 – 7.4). And the less the recorded L is contaminated by R_a's sub-operants, the more selective will be its sensitivity to independent variables. The reason why R_a's sub-operants are considered a contaminant is that since they are part of the rO, their sensitivity profile would almost certainly not be identical to the sensitivity profile of the L.

In effect, the time taken up by the R_a's sub-operants adds an unknown constant to the recorded L. If we could measure that constant, we would be able to correct for it by subtracting it from the recorded L, and adding it to the D where it belongs. But since we can't measure it, we do the next best thing, which is to make it negligibly small by keeping the R_a as simple as possible.

2.3 Function of Standardization: Comparability of Ls for Different rOs

We want to partition the time interval between the termination of the preceding rO and the R_b that initiates the "work" portion of the next rO, into two sub-intervals: (1) the L, and (2) the time interval between the R_a and the very next recorded sub-operant, namely the first R_b (see Figure 1).

Certain types of research involve comparisons of different types of rOs in terms of the Ls that precede them, and in terms of their R_a–R_b intervals. Such comparisons are of interest because there is evidence that Ls and R_a–R_b intervals are differentially and selectively sensitive to certain family groupings of independent variables.

As explained above, the time taken up by R_a's sub-operants is unavoidably included in the recorded L. The Ls will be truly comparable only if that time is *the same* for different types of rOs, like a constant error. Using a standard R_a makes the error constant. If we tried to compare different rOs that used different R_as for different rOs, we would be adding the unknown durations of the *different* R_a sub-operants to the Ls of the rOs being compared, thereby making those Ls non-comparable to an unknown degree. When comparing different rOs, we want to be able to vary the specifications of the rest of the rO without affecting R_a and thereby contaminating the Ls. We can accomplish that by making sure the R_a is distinct from the rest of the rO, and by keeping it the same for different types of rOs.

Once we can meaningfully compare Ls and R_a–R_b intervals for different types of rOs, we have a way to answer quantitatively such questions as, "What are the possible equivalences and tradeoffs between such variables as probability, delay, density, and amount of reinforcement, various types of stressor variables, waiting time, difficulty of a discrimination, or amount of work required for the operant?" The methodology described in Sections 5.1 and 5.2 then becomes feasible.

We want the subject always to be in the same behavioral "state" at the start of the rest of the rO, as far as possible. In that sense, the R_a performs a function similar to the subject's turning on the equipment, or saying "I'm ready to do another one."

The R_a–R_b interval is of interest for many types of studies. It is likely to be sensitive to parameters that specify the rO's contingency structure. The Ls, on the other hand, are likely to be more sensitive to motivational variables like deprivation (Mechner, 1962), rate of reinforcement (Nevin, Mandell & Atak, 1983), amount of reinforcement, other types of establishing operations, or the schedule of reinforcement on which the rO is maintained (Hyten et al., 1991; Mechner et al., 1992).

2.4 How Are rOs Different From Reinforcement Schedules?

The initiating R_a is the basis for the distinction between rOs and certain reinforcement schedules. When a reinforcement schedule is in effect, the behavior that occurs between two consecutive reinforcements (sometimes called the "schedule unit") may or may not meet the definition of an rO. Such behavior qualifies as an rO only if it is initiated by an R_a.

In general, any reinforcement schedule can be converted into an rO by introducing distinctive sub-operants that initiate and terminate the schedule unit. There is a big difference between a schedule unit initiated by the experimenter (e.g., as when the experimenter presents a stimulus or manipulandum, or starts a timer) and one initiated by the subject's R_a. There is also a big difference between a schedule unit *terminated* by

the experimenter (e.g., by the presentation of a reinforcer or some other stimulus, as in a fixed ratio schedule), and one terminated by the subject.

When the schedule unit is initiated by the subject's behavior (an R_a), the unit qualifies as an rO. If the specification of such an rO also includes a terminating sub-operant (like our R_c) then each individual occurrence of the rO can be recognized and registered without regard to whether or not the occurrence was followed by a reinforcer presentation. When the unit is terminated by a distinctive behavioral event (the R_c), more interesting behavioral measures are made available for the rO. For example, abortive (incomplete) instances of the rO can be identified as such, and the duration measure becomes behaviorally meaningful.

2.5 Potential Value of Viewing Schedule Units as rOs

Viewing schedule units as rOs can be of interest because:

(a) Converting a reinforcement schedule to an rO procedure permits the obtained findings to be extended to all operants, thereby broadening their generality. The behavior generated by the schedule can then be compared with properties of other operants (see Mechner et al., 1992)

(b) It suggests experiments that would otherwise not be particularly interesting. For example, reinforcing the schedule unit intermittently rather than on the usual continuous reinforcement schedule, and then examining it under conditions of extinction, takes on new significance.

(c) It puts the spotlight on the easily overlooked implicit assumption that the behavior generated by "chained" schedules always remains a chain. To say that it is a chain implies that the first component of the schedule (or operant) functions as an S^D for the second component, and that the components remain chained in this manner, on a steady-state basis. But there is ample evidence that chaining, while it does operate initially, diminishes as a function of number of repetitions of the behavior, and eventually ceases. As the operant is repeated increasing numbers of times, as in long-term maintenance, there is progressively less cuing of successive components (from either exteroceptive or kinesthetic sources), and the behavior becomes fused into larger unchained units, although it often remains *open* to feedback from exteroceptive sources (Shiffrin, 1988, 740-767; Schmidt, 1988, p. 74; Restle & Burnside, 1972; Keele, Cohen, & Ivry, 1990, p. 96-98).

2.6 Operant Contingencies and Reinforcement Contingencies

Operant behavior experiments involve the use of two types of behavioral contingencies--the "*operant* contingency" and the "*reinforcement* contingency."

To date, the type of contingency on which virtually all of the attention in the field of operant behavior research has focussed is the *reinforcement* contingency (Skinner, 1938; 1969). Reinforcement contingencies specify the conditions that must be met for one or more operants to produce reinforcement or other specified environmental consequences (Weingarten & Mechner, 1960). The most familiar examples of reinforcement contingencies used in operant behavior research are reinforcement schedules.

The contingency that *defines the operant*, on the other hand, is an operant contingency, and not a reinforcement contingency. The operant contingency specifies the nature of the operant itself, i.e., the effect that must be produced for an instance of the operant to be considered as having occurred (Skinner, 1969, p. 13; Glenn, Ellis, & Greenspoon, 1992).

Thus, the contingency that defines a reinforcement schedule is a reinforcement contingency. But if the behavior generated by the schedule is viewed as an rO, then the same contingency is viewed as an operant contingency. That rO can then in turn be reinforced on any desired schedule, and that schedule is then the reinforcement contingency.

Different rOs permit different types of issues to be addressed and problems to be investigated. One can categorize rOs either according to their operant contingency structure, i.e., how the rO is defined, or according to the empirically observed sensitivities of their various measures to various independent variables. Such categorizations may, in some cases, have theoretical implications.

The variety of possible rOs that can be specified is clearly unlimited. The rOs described in this monograph have been implemented in human subjects using standard computer equipment, and can also readily be implemented in non-human species with appropriate adaptations of such equipment.

3.0 BASIC TECHNIQUES FOR IMPLEMENTING rOs

3.1 General Method for Implementing rO Experiments

Here is one practical way to implement rOs in human subjects (Field et al., 1991; Mechner et al., 1992): The subject sits in front of a personal computer. The screen displays stimulus S_O. Pressing the space bar on the keyboard (R_a) starts the rO and changes the screen display from S_o to S_1. The R_bs consist of pressing any key other than the space bar and the ENTER key, these being reserved for R_a and R_c respectively. If consecutive R_bs are made on same key, only the first press registers.

A single occurrence of the rO can be as short as a split second, depending on how the rO is specified. When the objective is to compare different rOs or to investigate the role of some rO parameter, the different rOs can be used in a multiple schedule. Alternation of the rOs can be signalled by easily discriminated auditory and/or visual cues (Nevin, 1992).

The R_as, R_bs, and R_cs shown in the diagrams are sub-operants of the overall rO. Some kind of feedback is always presented upon completion of each rO, but reinforcement may or may not be presented, according to the reinforcement schedule used. This point is stressed because the concept of specifying the operant contingencies (which define the operant itself, as opposed to its reinforcement contingencies) is an unaccustomed one.

The operant contingencies defining the various possible types of rOs can be described and conceptualized by means of a notation system for behavioral contingencies (Mechner, 1959b; Weingarten & Mechner, 1960).[8] Note that each diagram specifies the operant contingency for one type of rO. The Rs refer to the rO's sub-operants.

3.2 The Fixed-Ratio rO: rO$_{FR}$

This rudimentary rO is defined as a specified number N of R_bs made on any combination of the available keys. The particular keys chosen for the R_bs are one of the non-criterial measures of the rO. It has no R_c.

---

[8]The symbol $R_\rightarrow$ means "if R occurs, then. . .," and implies nothing about whether or not the response will actually occur. The vertical bracket indicates that all of the vertically aligned conditions inside the bracket go into effect at the same time.

S_E is the "effect" produced by a complete execution of the rO, signalled by a distinctive sound or visual stimulus presented on the screen (see diagram, p. 23). The only function of the S_E is to provide feedback for completion of the rO. S_E signals that the rO has been fully executed, but that no reinforcement is forthcoming this time. If a reinforcement *is* forthcoming, a different stimulus is presented, such as the words "25 cents" on the screen. When a single occurrence the rO_{FR} is compared to a single bar press, the S_E is analogous to the click of the microswitch when the bar

$$\begin{bmatrix} S_0 \\ R_a \end{bmatrix} \rightarrow \begin{bmatrix} S_1 \\ NR_b \end{bmatrix} \rightarrow S_E$$

is pressed and no reinforcer is delivered. Not providing S_E would be like eliminating the click of the microswitch, thereby depriving the rat of distinctive differential feedback for incomplete versus complete bar presses.[9]

In the rO_{FR}, the ratio run is comprised of the rO's sub-operants, while in conventional FR, the ratio run is comprised of iOs. The essential difference between the two is that in rO_{FR}, each ratio run must be preceded by an R_a. According to the definition of an rO in Sections 1.2 and 2.1, and Section 2.4 on schedules, a conventional FR can be converted into a type of rO by requiring an initiating R_a.

Since conventional FR schedules do not include an R_a, they provide no behavioral indication of when a ratio run begins or ends. Suppose, for example, that an iO fixed-ratio is set at 30, and the subject makes 20 R_bs, pauses, and then makes another 10, at which point S_E is presented. Without an R_a to initiate the run, there is no way of knowing whether the subject (a) made a run of 30 responses with a pause after 20, or (b) made a short run of 20, engaged in some other behavior, and then started a new run that was interrupted after the first 10.

The usefulness of the rO_{FR} is limited by the fact that it does not have an R_c to mark its termination. Its termination is marked only by the presentation of S_E, rather than by some active behavior. The rOs described in the next two sections are more useful because they are terminated by R_cs.

An example of a simple rO_{FR} (where N = 1) is the operant contingency where R_a is pressing the bar down and R_b is releasing it (Skinner, 1938, p. 328-338; Margulies, 1961;

---

[9]By our definition of a stimulus, only one can be present at one time. Thus, click plus reinforcement would be shown in the diagram as a different stimulus than click alone or reinforcement alone. The same applies to S_E plus reinforcement versus S_E alone.

Kutch, 1974). This R_b does perform the function of marking the end of the rO. However, since the R_as of those rOs involve more effort and probably also take longer to execute than the rest of the rO (i.e., releasing the bar), the R_a does not have the simplicity attribute. That rO's usefulness is therefore limited in the ways discussed in Section 2.2.

3.3 The Fixed Number rO: rO_{FN}

This rO requires completion of at least NR_bs followed by R_c to produce S_E. Premature R_cs have no effect. The R_c should use a specially designated key, like the ENTER key. No signal is provided when the $R_c {\rightarrow} S_E$ contingency goes into effect on completion of the N R_bs. Not signalling it is analogous to delivering food silently in fixed-ratio. If that were done in a rat experiment, the rat would learn to keep checking the food chute. In rO_{FN}, R_cs typically occur with increasing frequency as N is approached. This contingency is also analogous to the non-signalling of the end of the time interval in fixed-interval schedules, with the well-known result that responding often accelerates as the end of the interval approaches.

$$\left[\begin{array}{l} S_0 \\ R_a \end{array} \right. {\rightarrow} \left[\begin{array}{l} S_1 \\ NR_b \end{array} \right. {\rightarrow} R_c {\rightarrow} S_E$$

This contingency is similar to the FN (Fixed Number) schedule (Mechner, Snapper, & Ray, 1961; Mechner & Latranyi, 1963). It differs from it only in that the rO_{FN} includes a standard R_a to initiate the rest of the rO (R_a Attribute 2). Premature R_cs, i.e., R_cs made prior to completion of the N R_bs, have no effect.

3.4 The Fixed Consecutive Number rO: rO_{FCN}

In this rO, as in rO_{FN}, the subject must make R_c to obtain S_E. But here, a premature R_c, that is, an R_c made prior to the completion of N R_bs, *resets* the run and restores S_O. In other words, the N R_bs must be consecutive. It should be noted that whenever this type of reset contingency is used, an experimental option is to impose a further aversive consequence for premature R_cs, such as a "time out."[10]

---

[10]In the diagram, a vertical arrow cutting a horizontal arrow indicates that the effect of the horizontal arrow is prevented, i.e., a premature R_c terminates the NR_b-R_c-S_E contingency, and the completion of NR_bs prevents the next R_c from resetting the run and restoring S_o. Note, again, that only one stimulus can be present at one time.

This rO_{FCN} is similar to the FCN ("Fixed Consecutive Number") schedule (Mechner, 1958a). It is different only in that in rO_{FCN}, a standard and neutral R_a initiates the rO. The group of schedules called FN, FCN, FI, and FMI (Mechner, 1962; Mechner & Latranyi, 1963; Mechner, Guevrekian, & Mechner, 1963) are of particular interest in the present context because they are similar to rOs (Mechner, 1959a). They differ from rOs only in that their initiating R_as were not standard (Attribute 2). Their usefulness is therefore limited, for the reasons discussed in Section 2.3. Nonetheless, because their R_as do have the simplicity attribute (Attribute 1), these schedules are quasi-rOs, and their post-reinforcement pauses are analogous to rOs' Ls.

In rO_{FCN}, the S_E provides immediate feedback to distinguish between a premature R_c and an R_c made *after* N R_bs (let's call that a "correct" R_c). Since a premature R_c restores S_O, we want a correct R_c to produce some stimulus other than S_O. Otherwise, there would be no differential feedback for premature R_cs and correct ones.

In the rO_{FCN}, examples of non-criterial measures that can conveniently be recorded are (a) the R_a-R_b interval, (b) the elapsed time between the first R_b and the R_c, and (c) the sequential patterns of the actual R_bs chosen, and the rhythmic patterns of those R_bs. The number of R_bs made before R_c, on the other hand, is a criterial measure.

Criterial and non-criterial measures often have different sensitivity profiles. Evidence for that is shown in Figure 2 of Mechner (1962): FCN run length (the criterial measure) is not sensitive to deprivation level, while the FCN's run duration D, (a non-criterial measure), *is* sensitive to it. The data also show that the run duration D is only about one fourth as sensitive to deprivation level as the post-reinforcement pause (PRP). The PRP is neither a criterial nor a non-criterial measure. A similar result was reported by Gilbert (1958).

3.5 Designing rOs for Special Uses

The various permutations and combinations of the parameters available in different types of rOs can be used to investigate many kinds of questions, some of which were alluded to in the above examples, and some that are discussed in the chapters that follow.

Obviously, the complexity of an rO can be increased indefinitely by the process of nesting rOs within rOs. The rOs discussed above provide rudimentary examples of that. A related concept is discussed in the context of reinforcement schedules in Morse, 1966 (pp. 103-105). But, as was explained in Section 2.4, one of the crucial differences between rOs and most of the conventionally used reinforcement schedules, including those described by Morse, is that most schedules do not have an initiating R_a and therefore cannot be used in the ways rOs can be used (as was explained in Sections 2.2 and 2.3).

When using rOs as research tools, new rOs should be designed creatively to fit each new research problem. The next chapter provides examples of rOs designed for special purposes. The features of the rO used should always depend on the problem being addressed.

4.0 USING rOs TO STUDY VACILLATION AND MEDIATING BEHAVIOR

4.1 Recording Vacillation in Discrimination Tasks

Experiments in the areas of stimulus equivalence, psychophysics, or matching-to-sample have one thing in common: They all involve the presentation of a stimulus and the recording of a discriminative response. That response often involves classification or identification of the stimulus. The classification process begins at the moment of stimulus presentation and ends with the final overt classification response. The time interval between those two events is sometimes called the response latency.

In such experiments, it is often of interest to study the covert behavior that comprises the classification process, and the time course of that process. With traditional techniques, hesitancy or vacillation during the latency period occurs covertly and is recorded only in terms of the length of the latency. But the rO makes it possible to track vacillation and possibly mediating behavior as overt recorded behavior.

The time course of mediating behavior and vacillation should be far more sensitive and interesting as a measure of the classification response than just the latency. The latency measure reflects only the timing of the final recorded all-or-none classification response. But a continuous measure of the classification behavior may reveal that the subject begins to "lean" toward the choice corresponding to the all-or-none classification response before *making* the final choice, and perhaps continues to "have doubts," i.e., continues to vacillate, even *after* having made the choice.

A methodological problem in most psychophysics and stimulus discrimination experiments is the need to use statistical measures when the readings from a number of trials are pooled. The problem with such pooled measures is that the "state of the organism" may change over the course of the pooled trials due to the action of other variables (including the unavoidable one of number of prior stimulus presentations). This problem is particularly serious in equivalence relation studies, where the equivalence relations may actually *form* during, and as a result of, testing (Sidman, 1988; 1989). It is therefore important to obtain a reliable and complete reading for each individual stimulus presentation. This can be accomplished with rOs.

4.2 Using Rhythmic rOs to Record Vacillation in Discrimination Tasks

Vacillation can be tracked continuously only if these conditions are met:

(a) The subject maintains a continuous stream of *recorded overt* behavior (consisting of at least two operants per second) throughout that time interval.

(b) The subject is always able to make any one of two or more equally effortful classification responses, each response being the "name" of a stimulus category.

(c) All combinations of pairs of these "naming" responses lie at opposite ends of a continuum along which they are defined. Such a continuum allows intermediate response variants to occur and be recorded.

These three requirements can be met with a set of rOs whose specification are based on rhythms. Those rhythms are specified as part of the *internal structure* of the rO. They can be tapped out on an ordinary computer keyboard or on a three-button mouse, and are specified in terms of the ratios of the time intervals between consecutive taps. The purpose of using rhythms as the criterial dimension that define such rOs is that the execution of rhythms can take on intermediate values along a continuum. The resulting "hybrid" rhythms are the desired intermediate response variants referred to as condition (c) in the above paragraph.

For example, a rhythm within an rO could be specified as two consecutive time intervals defined by three consecutive taps on three different keys—R_a, R_b, and R_c. The ratio of the R_a-R_b interval to the R_b–R_c interval is required to be 2:1 (with a certain allowed margin of error). The subject learns to perform that rhythm through special training until the performance is well established. The result is a rhythmic operant (rO_{rh} for short) defined by the 2:1 ratio. That rO_{rh} can then be established by further training as the "name" of one of the stimulus classes. A second stimulus class can be named with an rO_{rh} that uses the same three keys but is defined by a 1:2 ratio for the two time intervals. Again, the subject would learn to perform that rO_{rh} through special training. These two rO_{rh}s could be represented like this:

If the time between R_a and R_c is specified as being the same for both $rO_{rh}s$, the only difference between the two $rO_{rh}s$ is the temporal position of the middle tap R_b. The subject can distort a rhythm by positioning R_b closer to the mid-point between R_a and R_c without changing the total time between R_a and R_c. Thus, shifts in the position of R_b represent the intermediate response variants that can provide the desired measure of vacillation. To ensure that the two $rO_{rh}s$ are comparable in terms of the amount of time and effort they consume for their execution, it may be necessary to include in their specifications the total time between R_a and R_c, which should be the same for the two $rO_{rh}s$.

4.3 Options in Implementing $rO_{rh}s$

In the rO_{rh} whose rhythm is specified as the 2:1 ratio, the R_b must occur two thirds of the way into the R_a–R_c interval. In the rO_{rh} whose rhythm is specified as the 1:2 ratio, the R_b must occur *one* third of the way into the R_a-R_c interval. One design issue is how large an error band to tolerate around the one third and two thirds points. Obviously, if the two tolerance bands are made wide enough, they will meet in the middle, and the specification of the two rhythms would be abolished. The result would be another type of rO--the one discussed in Section 6.5.

One dependent variable that can be recorded is where within the tolerance bands the R_bs are occurring. The intermediate rhythms we are calling vacillation may express themselves as R_bs occurring closer to the center edges of the bands. In extreme cases of vacillation, the R_b may also fall outside of those bands, somewhere in the no-man's land between the two tolerance bands, in which case the unit would not be a qualified occurrence of the rO_{rh}, but would still provide data of interest.

For some types of experiments, such as experiments on equivalence classes, more than two different rO_{rh}s are needed because more than two stimulus classes are being tested. To make the mediating behavior overt, at least one additional rO_{rh}, corresponding to the possible mediating behavior, would be needed. To meet such a need, the rO_{rh}s can be designed by specifying a larger number of taps and time intervals. For example, a set of rO_{rh}s could be defined in terms of four taps. Both the second and third taps would then be "moveable." Such an rO_{rh} could be used to define 5 distinct rhythms, any two of which have identifiable unique hybrid cases. The design challenge when more taps are used is to ensure that the rO_{rh} is still short enough to permit an emission rate of at least two rO_{rh}s per second to be maintained, for reasons explained in the next section.

4.4 Maintaining an Even and Continuous Stream of Rhythmic rOs

There are two reasons why it is important to maintain an even and continuous rapid stream of rOs.

We want hesitancy or vacillation to show up as intermediate rhythms rather than as longer latencies (Ls) between consecutive rO_{rh}s. If we put an upper limit on the Ls, any hesitancy will express itself through the intermediate rhythms. To permit continuous and uninterrupted tracking of the time course of the intermediate rhythms, the stream of rO_{rh}s must be rapid, continuous, and uninterrupted. The execution time of each rO_{rh} should be so short as to produce a rate of at least two rO_{rh}s per second. The maximum permissible L should be set at about 200 milliseconds. An avoidance contingency (e.g., avoidance of time out) could be used to maintain that pace.

When attempting to record mediating behavior in equivalence studies, we need a rapid and continuous stream of rOs for a different reason. The mediating behavior can be very fleeting, and occur at any moment during the response latency period. If the stream of rOs is too slow, the mediating behavior can occur between rO occurrences and not be recorded.

4.5 Establishing rO$_{rh}$s as Stimulus Names

All rO_{rh}s used in an experiment should require the same amount of time and effort so that there will be no basis for preference among them. As was mentioned earlier, the total time comprising the time intervals within a given rO_{rh} (i.e., the total time D between the first and last taps) should be independent of the rO_{rh}'s specified internal rhythm. Also, the Ds of different rO_{rh}s should remain the same even as the internal rhythms shift and drift beyond their specified limits.

Subjects should be given extensive pre-experimental training in the execution of each of the rO$_{rh}$s used, to make sure that the performance of the rO$_{rh}$s is highly stable and well established before the experiment is initiated. Once the rOs are well established as motor routines, the next phase of the training is to establish them as the "names" of stimulus classes. After such training, the rO$_{rh}$s constitute, in effect, a well-established vocabulary of stimulus names consisting of "words" that have intermediate values and that can be learned by non-human species as well as humans. The third phase of training is to establish a continuous and rapid stream of rO$_{rh}$s that always "name" whatever stimulus is being displayed.

4.6 Using rO$_{rh}$s To Study Mediation in Equivalence Research

In equivalence research, a focus of interest is the classification response that occurs the *very first time* a new stimulus is presented. Some of the most provocative research questions raised by equivalence phenomena concern the covert processes that mediate equivalence behavior and that occur during the response latency interval. The rO$_{rh}$s that appear during that interval may reveal covert mediation processes that are otherwise unobserved. Vacillation among the rO$_{rh}$s would reveal which of some previously-seen stimuli are serving as mediating links or are generalizing with the new stimulus being presented.

The stimuli named by the rO$_{rh}$s could be presented in a multiple schedule format where one of several possible stimuli is always shown in the stimulus display. The subject is asked to name the displayed stimulus. In such a format, *some* stimulus is always being displayed. The stimulus is changed from time to time, either by a timer or by the subject. The continuous stream of rO$_{rh}$s is continuously naming whatever stimulus is being displayed.

When a matching-to-sample procedure is used (e.g., a sample stimulus and two matching stimuli presented simultaneously in an array) the multiple schedule format could be used in the following way: First the sample stimulus is presented alone. The stimulus change consists of adding the two matching stimuli. The stream of rO$_{rh}$s never stops and remains continuous across the stimulus change. Before the change, it is the sample stimulus that is named by the rO$_{rh}$s. After the stimulus change, it is the matching stimulus that is named.

Starting at the moment when the display changes from one stimulus to another, the rO$_{rh}$ will begin to mutate from one rhythmic pattern to another with a certain time course or lag. The more familiar the stimuli and the easier the classification response, the shorter that lag will be. If the rO$_{rh}$s are occurring at a rate of, say, two per second, and the classification time for the new stimulus is half a second, then the lag would be absorbed within a single rO$_{rh}$ occurrence. On the other hand, if the subject takes as long

as three seconds to classify the new stimulus, the lag would encompass about six rO_{rh} occurrences, and these would presumably contain rO_{rh}s that name mediating stimuli, or hybrid rhythms that reflect response vacillation. Recording the classification process by using rO_{rh}s does not interfere with the concurrent use of a traditional all-or-none classification response, such as the utterance of a word corresponding to the choice or the pressing of a special choice key. It should be interesting to examine the phasing of the rO_{rh}-based vacillation and/or mediating behavior, and the all-or-none choice response.

For example, the internal rhythm of an rO_{rh} may already begin to shift *before* the subject identifies the stimulus verbally or reports perception of it at the level of verbal awareness. It may even become fully transmuted into the other rhythm corresponding to the name of the presented stimulus prior to any verbal report. Also, rO_{rh} vacillation, reflecting uncertainty, can continue *after* the occurrence of a definitive verbal choice or classification response. Patterns in the mediating behavior prior to the choice response may predict whether the choice will be correct or incorrect.

4.7 Using rO_{rh}s in Psychophysics Research

In most psychophysics experiments, the objective is to test absolute or differential thresholds. In commonly used presentation formats, a test stimulus is presented from time to time, usually as a brief stimulus presentation. The rest of the time there is no stimulus present. In such experiments, the subject is usually required to identify either the presence of absence of a stimulus (or a stimulus difference), or to classify the stimulus with a "name" response.

For the subject to maintain a continuous and uninterrupted stream of rO_{rh}s in such experiments, there is a need for one rhythm that means "no stimulus is present" and one or more additional rhythms that correspond to the stimulus classes being investigated. The rO_{rh} technique makes it possible to track the perceptual process in terms of a gradual shift, or a vacillating shift, from the rO_{rh} that means "I don't see it" to the one that means "I see it," and vice versa. On the other hand, the English words "I don't see it" and "I see it" do not have intermediate cases along a continuum. When difference thresholds are investigated, one of the rhythms would mean "the stimuli are the same" and the other rhythm would mean "the stimuli are different."

As in other stimulus discrimination experiments, there is the possibility that the internal rhythms of the rO_{rh}s will begin to shift before there is perception at the level of verbal awareness.
Equivalence research and psychophysics are merely the most obvious areas where rO_{rh}s can be used. rO_{rh}s can also be used in other areas of research where vacillation can occur, such as approach-avoidance conflict and certain types of choice situations where

each of the rO$_{rh}$s can be associated with a different reinforcement condition.

5.0 USING rOs FOR GROUPING INDEPENDENT VARIABLES

5.1 Grouping Behavioral Measures and Independent Variables

The rO provides a methodology for grouping independent variables into "families" empirically, according to the behavioral measures they impact and the ways in which they impact them. Such family groupings would be defined by their impact profile.

Examples of independent variables that are candidates for such groupings are "establishing operations" which subsume motivational variables (Michael, 1982); probability, frequency, amount, or delay of reinforcement, sometimes called "reinforcement value" (Williams, 1988, p. 184; Rachlin & Krasnoff, 1983); the group that may include the operant's work requirement, complexity, or aversiveness; the group that may include such acquisition variables as number of times the operant has previously been executed, number of times it has been reinforced, and the behavioral contingencies used during acquisition. The rO methodology is particularly useful in this type of research because some of the behavioral measures described in Section 1.3 may prove to be selectively sensitive to these groups of independent variables, as well as to certain yet-to-be-discovered ones (see Nevin's commentary on the present monograph).

As this type of research proceeds, one must expect the groupings of independent variables to be continuously revised. The spiral of interplay between the formulations used for the independent and dependent variables should result, over time, in an iterative convergence and gradual stabilization of the groupings. The "chicken-and-egg" bootstrapping character of this methodology is usual in the development of theoretical formulations.

5.2 Relationship of These Groupings to Constructs

Empirically-based groupings of independent variables can lead to the discovery of simple relationships (that are invariant across different classes of rOs) between behavioral measures and such groupings. Such relationships would provide more promising constructs for use as variables in quantitative theoretical formulations. The usefulness of a construct depends in large part on its degree of correspondence with measurable variables, or its promise of such correspondence.

For example, there are measures that predict behavioral "disruptability" or "resistance to change" when the behavior is impacted by various types of independent variables. Such measures could be viewed as being related to Nevin's behavioral momentum construct (Nevin, 1979; Nevin, Mandell, & Atak, 1983; Nevin, 1988). Also, new empirically-based

constructs may help improve the generality of existing mathematical formulations of behavioral principles. An example is the possible reformulation of Herrnstein's matching law by including the D variable (Section 6.1, 6.2).

The research strategy described above, which basically involves the search for invariances by accounting for observed variances, has been amply discussed in the methodology literature (e.g., Stevens, 1951; Sidman, 1960; Nevin, 1984). The measures provided by the rO, because they can be applied to individual occurrences of rOs, lend themselves particularly well to this type of strategy.

5.3 The rO$_{FCN}$ with a Time Limit

For rO$_{FR}$, rO$_{FN}$, and rO$_{FCN}$, it is possible to require all N of the R$_b$s to be made within a certain limited time T once the "run" has been initiated. If they are not completed in that time, the count resets. Here is what the diagram would look like if there is a limited time to make the R$_b$s in the rO$_{FCN}$ procedure:

The time limit contingency is of interest because it may belong to the family of variables that increase an operant's stressfulness or aversiveness. With the rO technique, we can determine whether the imposition of a time limit has different effects on the L and the R$_a$–R$_b$ interval. The question is whether that difference is similar to the effects produced by other variables that have differential impacts on the L and the R$_a$–R$_b$ interval.

It would also be interesting to determine whether such a family of variables, which could be called "stressors," impacts behavior differently than the "response cost" or "work requirement" family. A related question is whether there are trade-off relationships between these families and between such "reinforcement value" variables as amount, frequency, probability, or delay of reinforcement, as well as between the various variables in the establishing operations family.

In performing such studies, consideration must be given to the possibility that if R$_a$ is permitted to start a limited-hold timer (thereby putting the subject under immediate time pressure), the subject may well dally a bit longer before making the initiating R$_a$, because an R$_a$ that starts such a timer is probably more aversive than one that doesn't, especially if there are aversive consequences for premature responses. We would want such dallying to fall into the R$_a$–R$_b$ interval, not into the L. As suggested above, a time limit

contingency may be similar in its behavioral effects to such variables as an increase in the operant's work requirement. Thus, if a time limit contingency is used in the specification of the rO, the timer should be started not by the R_a, but by an additional sub-operant response, which could be called R_{a2} or R_{b1}.

5.4 rOs with a "Window of Opportunity": rO_w

Another "limit" parameter that can be used is one that sets an *upper* limit of $N_1 + N_2$ (in addition to the minimum requirement of N_1) on the *permissible* number of R_bs. Thus, if $(N_1 + N_2)$ R_bs are made, the next R_c has the same effect as a premature R_c would have had, i.e., it does not produce S_E and resets the rO. In effect, this contingency creates a "window of opportunity" between N_1 and N_2. Here is the diagram for this type of rO, which we can name rO_w:

$$\boxed{\begin{array}{l} S_0 \\ R_a \end{array}} \rightarrow \boxed{\begin{array}{l} S_1 \\ N_1R_b \uparrow \\ R_c \end{array}} \rightarrow \boxed{\begin{array}{l} R_c \longrightarrow S_E \\ N_2R_b \rightarrow R_c \end{array}}$$

The time limit contingency discussed earlier can also be used in conjunction with the "window of opportunity" contingency. When those contingencies are used together, the expiration of the time limit T can reset the rO_w in either of two ways: (a) only if it expires prior to the completion of N_1R_bs, or (b) if it expires at any time prior to $R_c{\rightarrow}S_E$. In case (a), reaching the "window" results in "safe harbor," i.e. the time pressure is off, and R_c will produce S_E whenever it is made after that point. In case (b), the time pressure remains until R_c has been made.

In both cases, the subject has no exteroceptive cue that safe harbor has or has not been reached. R_c then becomes a probe that fetches either good or bad news according to whether it results in S_E or in reset/time out. Because of that uncertainty, making the R_c should entail some approach-avoidance conflict. In case (a), there might be some hesitancy in making the R_c. In case (b) there is less time for hesitancy, as the clock continues to run. These two cases, when used together, therefore make it possible to isolate the amount of hesitancy due to that uncertainty, and also to study the many obvious potential parameters of the hesitancy.

In addition to producing the different degrees of hesitancy in making R_c, the cases (a) and (b) above may also have differential effects on the Ls and on the R_a-R_b interval, effects that would suggest into which group of variables each of these operant contingencies should be placed.

5.5 Using rOs to Study the Functions of Stimuli

In the above example, the "risk" in making the R_c can be eliminated completely by the presentation of an exteroceptive stimulus that signals the completion of N_1R_bs. Such a stimulus (S_1 changing to S_2) eliminates the uncertainty inherent in relying exclusively on response-produced cues, and constitutes a kind of "safe harbor" signal (Nevin, 1992). In the Mechner, 1958a study on FCN, such a stimulus (actually, the reinforcer) was presented after exactly NR_bs in some fraction of the instances, that fraction ranging from zero to .75. Experiments that use such an S_2 could be used to assess the effect of the risk-related stress in the time-limited rO_w by comparing the Ls and the R_a-R_b intervals with and without the S_2.

The presentation of S_2 need not coincide with the completion of exactly N_1R_bs. An alternative would be to present S_2 after fewer or more than N_1R_bs (Mechner & Guevrekian, 1960). The "informational value" of such an S_2 depends on where in the run it is presented, on the costliness or aversive consequence of an R_c made too early or too late, and on other variables.

These and others of the various parameters available for the rO_w provide ways to determine whether the window contingency puts an rO into the "stressor" family or the "work cost" family of variables.

In general, rOs provide some new ways to classify the functions of stimuli, in terms of their effects on the rO's various measures. The initiating R_a in effect "turns on the equipment," and an additional response is required to initiate the operant contingency that defines the rest of the rO. There is no limit to the variety of stimulus functions that can be incorporated into the design of the rest of the rO, while still keeping the R_a-R_b intervals of different rOs commensurable.

For example, the rest of the rO could consist of the contingency in which pigeons are required to sequence the R_bs based on color rather than spatial position (Terrace, 1987). Or, it could consist of the inside/outside discrimination task like that described in Herrnstein, Vaughan, Mumford, & Kosslyn (1989); of constancy tests for pattern perception and discrimination of perspective (Cerella, 1990a; 1990b); or of the tasks used in the numerous types of stimulus equivalence studies that have been done in recent years (see reviews in Sidman, 1989; and Fields, Verhave, & Fath, 1984).

The types of studies mentioned above involve a rather diverse set of tasks. But diverse as these tasks are, when used as part of rOs, they can all be compared in terms of their effects on the Ls and R_a-R_b intervals.

5.6 rOs with Time Contingencies: The Fixed-Interval Analog, rO_{FI}

Another possible rO parameter is the time that must elapse after the first R_b before the $R_c \rightarrow S_E$ contingency goes into effect (See diagram on p. 36). These rOs are interesting in part because the R_a separates the L from the D so that behavior involving time discrimination can be studied.

The T parameter can be used either in conjunction with an NR_b requirement, or without any N (which is like setting N equal to zero). The time contingency would then completely replace the number contingency. In the following diagrams, T is used without N.

$$\begin{array}{cc} \boxed{\begin{array}{l} S_0 \\ R_a \end{array}} & \boxed{\begin{array}{l} S_1 \\ R_b \end{array}} \rightarrow T \rightarrow R_c \rightarrow S_E \end{array}$$

This rO_{FI} is an FI schedule in which R_a sets up a condition where an R_b will initiate the fixed interval T (Mechner, Guevrekian, & Mechner, 1963). As explained previously, there are two reasons for having R_b rather than R_a initiate the interval: (a) to keep R_a "clean" and its sub-operants unaffected by the temporal contingency it sets up, and (b) to provide an R_a-R_b interval that can meaningfully be compared with the R_a-R_b intervals of other types of rOs.

5.7 Fixed Minimum Interval: rO_{FMI}

When a time interval T replaces the NR_b requirement in the rO_{FCN}, we have an rO counterpart of the fixed minimum interval (FMI) schedule (Mechner, Snapper, & Ray, 1961; Mechner, 1962). That rO could be called rO_{FMI}.

$$S_0 \quad S_1$$
$$R_a \rightarrow R_b \rightarrow T \uparrow \rightarrow R_c \rightarrow S_E$$
$$R_c \downarrow$$

Both rO_{FI} and rO_{FMI} may prove to be useful tools for studying the question of whether time discrimination uses an internal clock or is based on the execution of a particular behavior routine, a question that has already received attention (e.g., Laties & Weiss, 1969; Laties, 1972). A potentially fruitful new way to approach this question is to compare the criterial and non-criterial measures in rO_{FCN} and rO_{FMI}.

Tantalizing in this regard is finding that the duration D (a non-criterial measure) in FCN is affected by deprivation level, while the criterial measures in both FCN and FMI are not (Mechner, 1962). This study could be repeated with rOs with attention to the non-criterial measures. One technique would be to alternate rO_{FMI} and rO_{FCN} in a multiple schedule format and compare the effects of motivational or drug variables on the measures relevant to the time discrimination issue. Some of these variables are defined in Mechner & Latranyi (1963).

Note that the "Window of Opportunity" contingency described above for response number can also be used for minimum time intervals.

6.0 CORRECTING FOR D IN MATHEMATICAL FORMULATIONS

6.1 The Undermatching Phenomenon in Matching Law Experiments

Any mathematical formulation that includes iO response rate as a variable can be corrected for response duration by subtracting D from the iO IRT, leaving L. One can then replace 1/IRT (which is the iO response rate) with 1/L. Correcting for D in this way may increase the generality of certain mathematical formulations.

A possible example is Herrnstein's matching law for choice situations (Herrnstein, 1961; 1970), which states that response rates tend to be proportional to the reinforcement values that prevail in each of the conditions being tested. Reinforcement value can refer to probability of reinforcement, amount of reinforcement, or other variables related to reinforcement.

But most experiments that have attempted to fit the matching equation to the data have discovered a deviation from theory that has come to be called "undermatching" (Baum, 1974; Williams, 1988). J.J. McDowell states that the phenomenon

> . . .has no obvious explanation. It is produced by the tendency of responding to deviate from perfect matching in the direction of indifference. This tendency seems to occur to at least some degree in most choice situations (Baum, 1979; Myers & Myers, 1977; Wearden & Burgess, 1982)...the causes of undermatching are unclear. (McDowell, 1989, p. 161)

Several explanations for undermatching have been advanced and some reformulations of Herrnstein's law proposed (see Williams, 1988 and McDowell, 1989 for reviews), but none has proven quite satisfactory.

6.2 Adjusting Herrnstein's Matching Law for D

One explanation that does not yet seem to have been put forward is the following: It may be possible to eliminate undermatching by correcting the two iO response rates by subtracting the Ds from the IRTs. Changes in response rate may be due more to changes in the Ls than to changes in the Ds. There is reason to believe that the Ds depend primarily on the specification of the operant (i.e., the operant contingency), while the Ls depend primarily on the reinforcement values (Mechner, 1962), or the prevailing schedules of reinforcement (Nevin, 1992; Hyten et al., 1991; Mechner et al., 1992).

The rO technique provides a way to test that explanation directly if we accept the premise that an rO's D is an experimental model of an iO's response duration, with the rO's L + D being the analog of the iO's IRT.

The adjusted matching law for rOs is derived by substituting L_2/L_1 for Herrnstein's response rate ratio B_1/B_2:

$$\frac{L_2}{L_1} = \frac{r_1}{r_2} \quad \text{Adjusted Matching Law}$$
$$\text{applied to rOs}$$

where L_1 and L_2 are the measured Ls in the two reinforcement conditions or values, r_1 and r_2, that are used.

6.3 Testing the Adjusted Matching Law Using iOs

The L of an iO is its IRT minus its D. Therefore, if the quantity (IRT-D) is substituted in the above equation for each L, and the two IRTs are then replaced with the reciprocals of the response rates, $1/B_1$ and $1/B_2$, the resulting restatement of Herrnstein's matching law (solved for D) is:

$$D = \frac{r_2}{B_2(r_2-r_1)} + \frac{r_1}{B_1(r_1-r_2)} \quad \begin{array}{l}\text{Adjusted}\\ \text{Matching}\\ \text{Law for iOs}\end{array}$$

where D is the (unrecorded) duration of any iO.

This formula provides an indirect way to test the theory by means of a test that uses iOs. A number of matching experiments could be conducted using different ratios of r_1 to r_2, but the same iO for both rs. The formula would then be used to calculate D for each of these ratios.[11] If D turned out to be invariant, as it would if the same iO retained its characteristic D regardless of r, then this theory would be supported. In addition to supporting the theory, such a result would also mean that the adjusted matching law provides a way to calculate, from experimental measurements, the unrecorded behaviorally operative duration of any iO.

A weakness of this type of indirect test is that if the calculated Ds turn out *not* to be the same for the different conditions, the theory is not refuted. There could be various possible reasons for their not being the same: (a) The different reinforcement conditions changed the Ds in the same way they changed the Ls, at least to some degree; (b) Different iO response rates produce different response topographies and hence

---

[11]A mathematical analysis of the equation for D reveals that to obtain statistically reliable measurements of D, the values of r_1 and r_2 used in the experiments should not be too close to each other, and the response rates B_1 and B_2 should not be too low.

different Ds; (c) The required adjustment requires the subtraction of a quantity other than D (see Chapter 7 for an analysis of the possibilities) and (d) The theory is wrong.

Reason (b) is a particularly serious possibility, as there is substantial evidence that different reinforcement schedules do indeed produce different response topographies (Herrnstein & Heyman, 1979; Platt, 1979). In ratio schedules, for example, the iOs tend to become fused into larger functional response units of indeterminable size (Skinner, 1938; Keller & Schoenfeld, 1950; Boren, 1961). As a result, higher fixed ratios and higher (recorded) rates of iOs would tend to produce the appearance of *over*matching, as has in fact been observed (e.g., Pliskoff & Fetterman, 1981; Dunn, 1982). On the other hand, if a higher iO response rate produced more efficient topographies and hence shorter Ds, the observed result would be more *under*matching. Thus, variables that affect iO response topography in unknown ways would also distort the indirect test in unknown ways.

Therefore, the adjusted matching law could be tested more directly by matching experiments that use rOs rather than iOs, so that Ds and Ls can be recorded directly. One would then see how well equation (1) fits the data for various values of r_1 and r_2, and for various different types of rOs. However, this test too entails some of the problems of measurement discussed in the next chapter.

6.4 Correcting for Bias in the Adjusted Matching Law

In the same McDowell (1989) paper referenced earlier, the asymmetry case is described as "a distortion of matching" or "bias" resulting from a possible inequality of the qualitative character (operant contingencies) of the two operants (such as their relative effortfulness) or an inequality of the reinforcement contingencies prevailing in the two conditions.

It is likely that these variables fall into groupings that have different kinds of effects on L and D, based on the nature of the various functional relationships describing those effects. Therefore, the mathematical adjustments that could correct for bias in a generalized matching law cannot be fewer in number, or simpler, than the separate effects of each of the possible biasing variables on L and D.

A systematic approach to the identification of predictively useful bias adjustments would require the performance of parametric studies (in non-matching situations) in which D and L for various kinds of rOs are used as the dependent variables. The independent variables (the presumptive biasing variables) in such studies would be, for example, effortfulness, aversiveness, or complexity of the operant, difficulty of required discriminations, all kinds of stress factors, and various reinforcement contingencies. As was mentioned above, such variables may fall into groupings according to the similarities and differences of their effects on various behavioral measures. The corresponding functional relationships, when inserted into the adjusted matching law, should then

predict the amount of bias that various combinations of those biasing variables will produce, thereby making the matching law correspondingly more general.

6.5 Testing the Matching Law with an rO that Has No D

Several formats have been used for matching law experiments. A common one involves two responses on two separate manipulanda that are simultaneously available. Another format makes available only one response on one manipulandum, and uses stimuli to signal the prevailing reinforcement value. Switching from one reinforcement value to the other can then be under the subject's control (as in a concurrent schedule), or it can be programmed (as in a multiple schedule).

When there are two manipulanda, switching to the other manipulandum involves some reorientation of a body part, and hence some time and effort. With concurrent VI schedules, switching requires a special response. The time and effort required by that response generates a certain amount of inertia that may inhibit switching and produce perseveration. On the other hand, switching is normally reinforced, as a reinforcement is usually waiting to be picked up. Switching may therefore be controlled more by the balance between the inertia and the lure of the waiting reinforcement, than by the matching contingency. Such procedures therefore provide tests of the matching law that are imprecise at best.

The matching law can be tested cleanly with a set of two equivalent rOs that have no duration and no rate. They have the additional feature that when the subject switches from one rO to the other, no reorientation of any body part and no special switching response is involved.

Such a pair of rOs would be comprised of the same three sub-operants R_a, R_b, and R_c described in Sections 4.2-4.5 for the rO_{rh}, but here there is no specified rhythm. The R_a-R_c interval is simply partitioned into two parts, and the two rOs are defined in terms of whether the R_b occurs *before* or *after* the midpoint of the R_a-R_c interval. One of the two reinforcement values applies to rOs whose R_bs occur before the midpoint, and the other reinforcement value applies to rOs whose R_bs occur after the midpoint. We will call this rO a "partitioned rO," or rO_p.

As in the case of $rO_{rh}s$, the response rate is constant because of the enforced rapid pacing via the maximum time allowed for the L. Thus, the response rates under the two reinforcement values are equalized.

When the rO_p is used, dependent variables that can be examined, and plugged into the matching law, are (a) the proportion of rOs for which the R_b falls to each side of the midpoint; (b) the distribution of positions where the R_bs occur; and (c) possible cyclic effects in the positions of the R_bs. The independent variables are always the pairs of reinforcement values used.

Differential visual or auditory feedback could be associated with each of the two rOs. To guard against possible preferences for one of the two rOs, the reinforcement value associated with each of the rOs could periodically be switched and retrained, as in discrimination reversal procedures. Another refinement of the procedure would be to require that R_b occur within a narrow window located in the middle of the R_a-R_c interval.

If undermatching can be eliminated by correcting for D, then it should also be possible to eliminate it by eliminating D, as in the rO_p.

6.6 Explanation of Nevin's Law by Reference to Ls and Ds

Nevin made the following discovery: In the type of experiment to which it applies, there are two concurrent or alternating conditions that differ both in terms of the frequency with which reinforcements are delivered and the iO response rates generated. At some point in time, a variable that affects the response rates is introduced. The resulting percentage change in the two response rates is much more highly correlated with the pre-introduction frequencies of reinforcement than with the pre-introduction response rates (Nevin, 1979; 1988; Nevin et al., 1983).

The rO provides an approach to explaining this provocative finding. One way would be to conduct a parametric experiment that uses rOs that have a range of different Ds. One parameter would be the attribute of the rO that determines the Ds. The other parameters would be several types of variables that affect L but not D. Frequency of reinforcement would clearly be one of these, but it would not be the only one.

One significance of Nevin's finding is its suggestion that frequency of reinforcement may belong to the same family of variables as deprivation or amount of reinforcement. Variables in that family have their main impact on L. See Sections 5.1 and 5.2 on the methodology involved.

7.0 DEFINING THE QUANTITY TO BE SUBTRACTED

7.1 Subtracting D May Not Provide the Desired Correction

In the preceding chapter, D was proposed as the quantity to be subtracted from the IRTs as a correction that might increase the generality of certain theoretical formulations like the matching law. D has the advantage of being clearly defined as the R_a-R_c interval when rOs are used.

But the problem, alluded to in the last paragraph of Section 6.3, is that D, defined as the R_a-R_c interval, may not be the right entity to subtract if the desired correction is to be achieved. Intuitively, it would seem that we need to define and subtract a quantity that could be called "fraction of all ongoing behavior comprising the operant." D, as defined, may not fill the bill for three types of reasons which are discussed in the next sections.

The definition of a quantity that provides the desired correction when subtracted should find use in any field in which the matching law is applicable, such as in the signal-detection area (Nevin, 1981; Davison & Tustin, 1978; McCarthy & Davison, 1980; Williams, 1988, pp. 231-234), foraging theory, and the analysis of behavior in natural environments (McDowell, 1988).

7.2 The Operant's Pre-Overt Phase

The Ra marks only the *overt* beginning of an rO. Recent work in neuropsychology has shown that the execution of individual occurrences of operants begins at the neuronal level in the CNS, before there is any muscular engagement or movement (Rolls, 1981; Conrad, Benneke, & Goodman, 1983; Wurtz, 1985; Requin, 1985; Mortimer, Eisenberg, & Palmer, 1987; Schmidt, 1988 (pp. 179-181); Georgopoulos, 1990; 1992; Semjen & Gottsdanker, 1990; Wiesendanger, 1990; Decety, 1992). This "pre-overt" neuronal phase of any rO must certainly be considered part of the rO, though an unrecorded part.

In the diagram below, our uncertainty regarding the time taken up by the pre-overt phase of any given type of rO is indicated by the alternative brackets drawn with dotted lines. As the diagram shows, we don't know whether the pre-overt phase of the rest of the rO is fully included in the Ra-Rb interval, or starts before it.

Figure 2. Schematic diagram of the time relationships of the pre-overt phases of the rO

The figure shows that the R_a has its own pre-overt phase. That pre-overt phase precedes the R_a's sub-operants, which are also unrecorded. The rest of the rO too has a pre-overt phase of its own, different from R_a's pre-overt phase, and presumably of longer duration (since the R_a is by definition simpler than the rest of the rO). Therefore, the pre-overt phase of the rest of the rO may begin either before or after the recorded R_a. If it begins before R_a, then it begins during the recorded L, as does the R_a's own pre-overt phase. Those are the reasons why the recorded R_a does not mark the true boundary between the behaviorally operative D and the L.

7.3 Limitations of R_a as a Demarcation Event

As was explained in Section 2.2, one of the main functions of R_a is to provide a boundary between the L and the D by marking the beginning of the rO. However, as was explained above, the recorded R_a marks the beginning of only the overt phase--not of the pre-overt phase--of an rO.

We do not yet have enough knowledge to predict the length of the pre-overt phase of any given operant from a knowledge of the operant's characteristics. Neurophysiology cannot yet even provide valid and reliable ways to *record* the length of a given rO's pre-overt phase, leave alone predict it. Given that lack of knowledge, we have to consider the R_a a "fuzzy" marker, the fuzziness being in part a reflection of our ignorance. But the fuzziness may also be inherent in the behavioral processes involved: In addition to the obvious fact that different types of overt behavior can proceed concurrently, it is

44

also well established (see the above-cited references) that pre-overt neuronal activity can begin while other behavior is still in progress.

7.4 Making an Operant's Pre-Overt Phase Negligibly Short

Notwithstanding the uncertainties and problems discussed in Section 7.3, there is a body of experimental data that permits us to make certain qualitative statements: The nature and duration of the pre-overt activity depends on such characteristics of the operant as the complexity and type of discrimination being made, the coordinations involved in executing the operant, the complexity and length of the operant, the number of times the operant has previously been executed, the consequences of executing the operant, its "automaticity" and "attention requirements," etc. (Klapp and Erwin, 1976; Sternberg, Monsell, Knoll, & Wright, 1978; Schmidt, 1988, pp. 474-475, in addition to some of the other references cited in Section 7.2). Additional support for those statements is provided by numerous reaction time studies. In the neuropsychology and motor behavior literatures, that pre-overt activity has variously been referred to as "preparatory," "feedforward," "planning," "programming," "pre-motor," "priming," and "loading." We shall use the more descriptive and agnostic term "pre-overt phase" to avoid imputing any particular function to pre-overt activity.

The above references establish only the qualitative fact that there is a positive relationship between the complexity or effortfulness of the operant and the length of its pre-overt phase. Even though that information is only qualitative, it suggests that we can make the rO's pre-overt phase negligibly short by making the R_a as simple and effortless as possible compared to the rest of the rO.

That is part of the justification for R_a's Attribute 1——simplicity. The purpose of the simplicity requirement is to minimize both R_a's sub-operants (as discussed in Section 2.2), and the pre-overt phases of R_a and of R_a's sub-operants, which have their own pre-overt phases. All of these pre-overt phases as well as the R_a's sub-operants are unrecorded, and precede the recorded R_a.

7.5 Target and Non-Target Behavior Are Usually Concurrent

In recent years, behavior theorists have referred to "target behavior" and "nontarget behavior" (e.g., Herrnstein, 1970; Williams, 1988; McDowell, 1989). Target behavior is the operant that is specified and recorded in an experiment, in our case the rO; nontarget behavior is all other behavior. As stated in Section 7.1, we need a measure that reflects the portion of the total flow of behavior attributable to the rO. The rO is obviously never the only ongoing behavior. The subject is always engaged in collateral and concurrent extraneous behavior such as breathing and moving other parts of the body. But there is little we can say about the *relative amounts* of behavior that are target

45

and non-target, as we have no units in terms of which amounts of different types of behavior can be compared.

In contrast to the effect discussed in Section 7.2, this effect may be far from negligible, as the rO may account for only a small fraction of all ongoing behavior. Suppose the rO accounts for, say, one third of all ongoing behavior during the time that the rO is in progress. The quantity used in the correction would then be one third of D (assuming that the amount of behavior attributable to the pre-overt phase has been rendered negligibly small). If D and L were equal in length, the rO would then comprise about one sixth of the total amount of behavior.

Here is a diagrammatic way of visualizing this effect:

7.6 Fraction of Total Behavior Varies Over an rO's Time Course

The later portions of rOs are less susceptible to the effects of variables than the earlier portions (Hyten et al., 1991; Mechner et al., 1992; Nevin, 1992). There are also data that show similar phenomena for chains.

The following question now arises: Which part of the rO——the sensitive beginning or the insensitive end——comprises the larger percentage of all simultaneously ongoing behavior?   The following line of reasoning can be invoked to answer that question: When behavior A is insensitive and insusceptible to the effects of variables (we could say "automatized"), other behavior B can go on at the same time, without interfering with A. Conversely, behavior A tends not to interfere with other simultaneously ongoing behavior. On the other hand, when behavior A is relatively unstable, and sensitive and susceptible, simultaneously ongoing behavior B would be more likely to interfere with it and disrupt it. Therefore, to avoid such interference, simultaneously ongoing behavior would need to be minimized.

According to that line of reasoning, it would appear that the greater the behavior's susceptibility to the effects of variables, the greater the fraction of the total amount of the ongoing behavior flow it accounts for. We can conclude from our data that the final portions of the rO are more automatized. Accordingly, we would expect them to involve less of the total ongoing flow of behavior than the earlier portions would. We could therefore conclude that the earlier portions of operants involve larger cross-sectional fractions of the total ongoing flow of behavior than the later portions.

Here is a way to visualize such an effect diagrammatically (The pre-overt phase, though it may be negligibly short, is included in the target behavior):

7.7 Research Implications for Neuropsychology

In Section 7.2, reference was made to the relationships between pre-overt neuronal activity on the one hand, and the properties of the overt behavior produced by the operant contingency on the other hand. Decety (1992) and Georgopoulos (1992) are currently conducting such studies.

The rO technique provides a practical way to study those relationships. For example, relationships between the rO's pre-overt phase and the overt behavior comprising the rO could be explored directly by means of laboratory preparations that permit concurrent recording of the rO's pre-overt neuronal activity and the overt behavior. Skinner discussed the methodological status and importance of this type of research (Skinner, 1938, p. 422; 1974, p. 236).

Furthermore, when rOs are used for the concurrently recorded operant behavior, the brain areas from which the neural recordings are obtained can be related to behavioral categories and measures that are more meaningful, and have more generality across species, than the categories and measures that have been used heretofore. For example, it would be interesting to examine relationships between such neural measures as (a) the length (and other attributes) of the neurally recorded pre-overt phase of an rO, (b) the relative amount of neural activity involved (by PET or MRI scans), and (c) the brain areas involved, and such behavioral measures as (i) how the rO is specified, (ii) the rO's various criterial and non-criterial measures, (iii) the values used for the various types of

47

independent variables discussed earlier (reinforcement value, work cost, stressfulness, etc., and (iv) the rO's degree of automatization. Behavioral and neural measurement would always proceed concurrently. Such studies could provide a strong impetus, and point to new directions, both for the study of brain function and the neural mechanisms of operant behavior. The rO technique is a uniquely practical and convenient tool for such research because it permits all relevant properties and dimensions of operants to be manipulated and recorded.

8.0 BEHAVIORAL SUSCEPTIBILITY AND MOMENTUM

8.1 Speed With Which Variables Impact Behavior ("Susceptibility")

When an experimental variable is changed abruptly during a steady state performance, the usual result is a gradual performance change and restabilization. Most experiments on the behavioral effects of independent variables examine only the effects on steady state patterns. They pay little attention to the *speed* with which behavioral measures change from one steady state to another. They rarely focus on the transition process itself. The reason for this has mainly been a practical one: Analysis of a transition process requires the examination of individual occurrences of the operant that is undergoing the transition, especially when the transition occurs rapidly. But when response rates of iOs are used as the dependent variable, reliable measurement of behavioral effects requires the observation of behavior over a number of successive responses, and cannot be measured for individual occurrences of responses.

With rOs on the other hand, where each occurrence of the operant yields a complete set of behavioral measures, one can address the question of *how fast* any given independent variable produces its particular behavioral effects. The speed with which the effects manifest themselves, and the shape of the transition function, can be examined by plotting the impacted behavioral measures as a function of successive occurrences of the rO.

This type of study could be called a transition study, as it examines the transition of the behavioral measures from one steady state to another when some relevant independent variable is changed from one value to another (Mechner, 1959a; Mechner et al., 1992).

8.2 Examples of Transition and Restabilization Studies

In some recent pilot studies on susceptibility conducted with the use of rOs, the number of R_bs required was abruptly changed. Following initial instability in the various behavioral measures of the rO, the performance restabilized (Mechner et al., 1992). The number of R_bs per rO (run length) generally restabilized within 50 to 100 rOs when the required number of R_bs was changed from 10 to 20. But when the criterion was changed back from 20 to 10, the run length required several sessions and over 1,000 rOs to return to its previous level of about 11 R_bs per rO.

Transition processes can usually be described mathematically by an exponential function that is either fully damped and monotonic, or that exhibits damped oscillations as it approaches the new steady state. The coefficient of the exponent of e (the base of natural logarithms) would describe the overall rate at which the new asymptote is approached. That coefficient would correspond to the transition speed, and could be

subsumed under Nevin's behavioral momentum construct, which he defines as resistance to change (Nevin, 1984; 1988; Nevin et al., 1983). Momentum could be viewed as the opposite of susceptibility. The term susceptibility has the advantage of possessing the two-faced janusian feature of "susceptibility *to*" and "susceptibility *of*," as in the statement "susceptibility of x to y."

8.3 Use of Adjusting Techniques to Measure Susceptibility

With rOs, the transient characteristics of behavior can also be studied in a steady state form, though that sounds at first like a contradiction of terms. The key is to use the "adjusting schedule" technique (Ferster & Skinner, 1957, pp. 72-721), where the value of a known impactful variable (like amount of reinforcement or work requirement) is increased or decreased by one notch at each successive occurrence of the rO. The direction of the adjustment is made to depend on the registered value of the impacted behavioral measure at the immediately-preceding occurrence. For example, the size of a fixed ratio could be used as the impactful variable. The fixed ratio would be increased a notch if the immediately preceding post-reinforcement pause (PRP) was less than 10 seconds, and decreased a notch if the PRP was more than 10 seconds. This procedure causes the impactful variable to fluctuate, in effect reversing the usual roles of the independent and dependent variable.

The PRP, or L, is a good choice for the impacted variable when adjusting procedures are used. The L of an rO is far more sensitive to the effects of independent variables than other rO measures (Mechner et al., 1992).

In research conducted in 1960, Mechner and Snapper used an adjusting technique in three separate studies (8 rats per study). In all three studies, the impacted measure was the PRP, which was maintained at a particular fixed critical value for 21 experimental hours (7 days, 3 hours per day). Five critical PRP values were used as the independent variable. The values used were 10, 20, 30, 40, and 50 seconds. These five values were tested first in ascending and then in descending order, being shifted every 7 days. Accordingly, each animal's series took about two months to be completed.

In the first study, the impactful variable (the one that fluctuated as the dependent variable in accordance with recorded PRPs) was the number of responses required for reinforcement (fixed-ratio); in the second study, it was the length of the interval used in the rO_{FI} schedule described in Mechner et al. (1963); in the third, it was the amount of liquid delivered per reinforcement, the rO being rO_{FN} with an N of 1. The setting of the impactful variable was automatically readjusted after each PRP (i.e., after each occurrence of the operant), in small logarithmic steps, upward or downward, according to whether the immediately-preceding recorded PRP was longer or shorter than the critical PRP value.

In all three studies, smooth cyclic fluctuations with a stable period of oscillation were observed and recorded in the latter portion of each 7-day stabilization period, for all three impactful (adjusted) variables as well as for the actual recorded PRPs. The PRPs fluctuated around their critical values. When the average stabilized values of each of the three impactful variables were plotted against the 5 critical PRP values, reliable orderly positive relationships were found. These studies were never published, though they demonstrated the feasibility of the technique.

8.4 Steady State Measures of Behavioral Susceptibility

When the adjusting technique is used for studying susceptibility to the impact of variables, the key measures to examine are:

(a) The cycle length (or period of oscillation),[12] of the cyclic fluctuations of the impactful and impacted variables (since they are yoked, their cycle lengths would have to be the same),

(b) The phase lag of their respective cycles, and

(c) The amplitude of the oscillations.

In such studies, a behavioral measure's susceptibility to the effects of any given impactful variable should be reflected in the period of oscillation. It seems likely that the step size used for the instance-to-instance adjustments would determine the amplitude of the oscillations. Step size may also have some effect on phase lag, but (when varied within limits) should not affect the period of oscillation.

The measures of cycle length, phase lag, and amplitude of oscillations recorded by the experimental techniques described above make it possible to describe behavioral susceptibility in quantitative terms.

---

[12]Here the reference dimension for period of oscillation, or cycle length, is the train of successive rOs, rather than time. However, it may also be worth examining time cycles, using either cumulative latencies or real time as the time base.

All physical and biological systems exhibit a period of transient behavior when a relevant variable is changed. Then they may restabilize at a new steady state. There is a body of mathematical techniques, including exponential equations, differential equations, and the Laplace transform, that are generally useful in describing transient and steady states of any system. In chemistry, mechanics, and electronics, transients are often analyzed by means of second order differential equations, which describe the amplitudes, periods, and decay rates of damped oscillation functions. It is likely that behavioral transients too can fruitfully be described and analyzed by these techniques. Marr (1989, p. 147) made the point that behavior analysis is ripe for the application of differential equations. An excellent starting point is the analysis of transient patterns and periodicities, for which differential equations provide a ready-made analytic tool.[13] Periodicities are best analyzed by means of the Fourier analysis or autocorrelation.

8.5 Comparison of Transition Data and Steady State Data

Steady-state oscillation data obtained with the adjusting technique are more reliable than transition study data, because steady state data can be based on a larger number of readings. But the most important advantage of the adjusting technique is that the period-of-oscillation measure which it provides is probably not affected by the step size used (within a reasonable range of step sizes). The adjusting technique also provides the potentially important phase lag measure. Phase lags of more than one rO would presumably reflect higher-order sequential effects of the impactful variable on the impacted one.

Data obtained with the transition study procedure, which involve a scheduled one-step change in the value of an independent variable, can be and should be compared with steady state data obtained by comparable adjusting technique studies. It is always important to check any measurement by more than one method.

---

[13]Nevin has looked to physics as a model for ways to describe behavior's resistance to change, using the provocative mv definition of momentum (Nevin et al., 1983; Nevin, 1988). With the availability of the periodicity and wavelength measure provided by the rO, the physics model suggests the possibility of looking at some additional variables from physics, such as density of the medium through which the wave is propagated, the elasticity or deformability of the medium, and wave velocity. When searching for lawful relationships in relatively uncharted terrain, using a well-understood model from another domain (like physics) can have heuristic value by suggesting potentially fruitful new constructs, relationships, and formulations (Reichenbach, 1938; Smith, 1990, p.13).

8.6 Higher-Order Sequential Effects in Adjusting Procedures

An impactful variable may have an impact not only on the immediately following rO, but also on rOs that are two or more occurrences removed. When rOs are used, the impactful measure would normally be the latency L. A given L may be affected by the value of the impacted variable that prevailed at the two prior or the three prior occurrences of the rO, and may reflect the composite effect. The value of the impactful variable prevailing at any given rO occurrence may affect not only the immediately following L but also, to a lesser degree, the L after that one, and to a still lesser degree the ones after that.

The rO provides ways to study these possible higher-order sequential effects on the Ls or on whatever behavioral measure is used as the impactful measure. One could study such effects by means of parametric studies in which the parameter is the degree of removal of the rO whose measure was used as the impactful variable. In such studies, there are several sources for the adjustment instruction: (a) the immediately preceding rO, (b) the one before that, or (c) the one before that one. In effect, more than one rO would go by before the adjustment was made.

Higher order sequential effects may affect the phase lag primarily, and the periodicities only secondarily or not at all. If it turned out that they do affect the periodicities, then these higher order sequential effects would have to be taken into account as parameters.

8.7 Filtering Out Background Periodicity

When the adjusting procedure is used in the types of experiments described above, the periodicities that are observed are not necessarily due only to the adjusting procedure. A background periodicity may be present even in the absence of any adjusting procedure or experimentally arranged feedback.

The data in Mechner (1958b, Figures 2 and 3) show that a sequence of successive occurrences of the quasi-rO$_{FCN}$ under regular reinforcement exhibits definite periodicity in the absence of any adjusting procedure. Also, Mechner et al. (1992) reported data on background periodicities in various behavioral measures of the rO. (A possible explanation of such background periodicity is presented in Section 9.5). In the present context of studying periodicities produced by adjusting procedures, such periodicities are viewed as background noise that needs to be filtered out, since the goal is to obtain clean measures of behavioral susceptibility.

In studies that use the adjusting procedure, it seems highly unlikely that the periods of the spontaneous background oscillation would ever be the same as, or even harmonics of, the periods produced by the adjusting procedure. The background oscillation can therefore be filtered out by either Fourier analysis or autocorrelation. Once these are filtered out, the periodic oscillations due to the adjusting procedure would be left behind.

8.8 The Broader Context of Susceptibility Analysis

The general methodology for measuring susceptibility described here provides a way to group variables according to their effect on the susceptibility of various behavioral measures.

In a broader context, susceptibility analysis provides a way to study the dynamic interplay that reflects an organism's behavioral adaptation to its environment. The adjusting technique simulates an important aspect of that interplay. It provides a laboratory method for studying that interplay in the spirit of the "ecological" orientation of J.J. Gibson (1979), and others (e.g., Kelso, 1990; Kugler & Turvey, 1987; Schmidt & Turvey, 1989; Turvey, 1990).

9.0 WHAT ARE THE EFFECTS OF REINFORCER PRESENTATIONS?

9.1 Questions Relating to Reinforcement

The term "reinforcer" refers to a type of stimulus event that produces certain behavioral effects in some situations and circumstances, some of the time. The need to use the word "some" twice in the above sentence, bears testimony to the fact that the concept of "reinforcer," though often used as if it were well understood, actually conceals many as-yet unanswered questions.

One of the important categories of variables on which the effects of stimuli termed "reinforcers" depend (besides the preceding behavior) are "establishing operations" (Michael, 1982). They depend also on the recency and recent density of similar stimulus events, on the organism's history relating to similar events in similar circumstances, on the history, type, and frequency of behavior on which the event impacts (Morse, 1966, pp. 54-55), on prevailing reinforcement contingencies and stimuli, and on the organism's level of arousal (Killeen, Hanson, & Osborne, 1978).

The term "reinforcer" itself is semi-descriptive, as it implies that the presentation of such a stimulus produces a type of behavioral effect that can be described as "reinforcement." But one of the frontiers of operant behavior research is the description and analysis of the various types of behavioral effects *actually* produced by the presentation of presumptive "reinforcers," and the independent variables on which those effects depend (Morse, 1966, p. 55). The main categories of such independent variables are set forth in the above paragraph. The likely reason why this area of research has remained unexplored is that it cannot be addressed by means of iOs. The sections that follow attempt to show how it can be addressed by means of rOs.

9.2 Do Reinforcer Presentations Affect Individual Occurrences of Operants?

An example of an open question regarding reinforcement is which, if any, behavioral effects of the presentation of a presumptive reinforcer (let's call that a "PPRf" for short) can be observed at the level of individual occurrences of operants.[14] A closely related question is what (if any) behavioral effects are produced by a single PPRf, and how these effects summate to produce the known effects of multiple PPRfs. Multiple PPRfs can result from a reinforcement contingency maintained for an extended time.

Since the traditional iO techniques do not provide a way to address these questions, some researchers have, from time to time, attempted to approach them by visually observing and comparing the topography of the behavior immediately preceding and following a PPRf (e.g., Muenzinger, 1928; Skinner & Morse, 1958; Iversen, 1982). Such visual comparisons have never established that the PPRf produced a repetition of the immediately preceding behavior.

This failure may have been due to the fact that every PPRf has a dual function: that of reinforcement and that of a discriminative stimulus (S^D) for some behavior that was previously reinforced in a similar situation. It is plausible that the effect due to the S^D function normally overwhelms and obliterates the possible reinforcement effect, leaving unanswered the question of what (if any) "reinforcement" effect is *ever* present at the individual occurrence level.

A word of explanation may be in order regarding the important concept of the "S^D function of a PPRf." Every PPRf is preceded by and is concurrent with certain other events. Those events include the subject's own pre-PPRf behavior and the recent density of PPRfs. For example, a long stretch of behavior without PPRfs can comprise part of the recent events. Thus, the PPRf is a compound stimulus with several identifiable

---

[14]It would be ungrateful of me not to mention here one of the unforgettably provocative comments that William N. Schoenfeld made in a graduate psychology seminar at Columbia University in 1952. "We don't even know the effects of a single reinforcement presentation on an individual response", he said, to drive home the point that learning theory is still in its infancy. The comment would be as valid today as it was then.

components: (a) the PPRf itself, (b) the subject's immediately preceding behavior, (c) the schedule on which reinforcers have recently been received, and (d) the exteroceptive stimuli comprising the physical environment in the presence of which the PPRf is presented. That compound stimulus will always generalize, to some degree, with sets of circumstances that occurred earlier in the subject's history, perhaps minutes, hours, days, or months earlier. The degree of generalization will depend on the similarity of those sets of circumstances. One must therefore expect the compound stimulus inherent in any PPRf to act like an S^D, setting the occasion for whatever behavior was shaped by the contingencies that prevailed right after a previous occurrence of a similar compound stimulus and PPRf. To the degree that this happens, every PPRf functions as an S^D.

The rO technique provides a way to separate the S^D effects from the other effects of PPRfs, and for observing and measuring the effects of PPRfs at the micro or molecular level of individual occurrences, rather than merely at the macro or molar statistical mass-action level. The rO also provides a way to investigate how these effects depend on the independent variables listed in Section 9.1.

9.3 Mechanisms of Shaping

There is also a question that transcends the molar versus molecular one: *How* do PPRfs shape operant behavior (Morse, 1966, p. 56)? That question is valid regardless of whether shaping operates at the individual occurrence level or only at the molar level. The "how" question calls for an explication of mechanisms.

The shaping process is explained in the literature by reference to the molar process of "successive approximations" and "response differentiation." According to this explanation, shaping occurs when reinforcers selectively impinge on the response variants that fall to the chosen side of the variability distribution for a chosen criterion, thereby progressively shifting that distribution in the desired direction by a cumulative statistical action (e.g., Skinner, 1938, pp. 312-338; Keller & Schoenfeld, 1950, pp. 164-190; Wilson & Keller, 1953; Herrick, 1964; Morse, 1966, p. 55; Pear & Legris, 1987).

Although that explanation is qualitatively consistent with much experimental data, it does not account persuasively for the remarkable speed and efficiency with which the shaping process often proceeds. Skilled animal trainers can shape behavior so fast that it sometimes seems almost as if they were telling the animal what to do. A skilled trainer clearly does not rely on the inherently slow progressive statistical shifting of variability distributions (Pierrel & Sherman, 1963). The explanation also leaves open the question of whether selection of response variants is indeed the mechanism that underlies shaping.[15]

---

[15]Skinner's observation that selection operates both in the shaping of operant behavior and in the shaping of species did not require an explication of the mechanisms by which the selected behavioral variants are generated, any more

To be satisfying, an explanation of the shaping process would have to describe (a) the mechanisms that generate the variants from which selection can take place, and (b) the proximal (molecular) behavioral effects of a selection event. The remainder of this chapter attempts to show how the rO permits such questions to be addressed. Section 9.6 below outlines a possible alternative mechanism for the shaping process, and Section 9.9 addresses issue (a) above. The difficulty of addressing these issues with iOs may be the reason why even an approach or strategy for addressing them has so far remained elusive.

9.4 Reinforcement as a Parameter Shifter

Research in the field of motor behavior, most of it done since 1975, teaches us that well-established operant behavior routines become linked and coordinated in ways that allow them to be flexibly specified (at the CNS level) by attribute parameters. Examples of important attribute parameters are response force (which generally corresponds to muscle potential or degree of muscle engagement, which in turn corresponds to placement along the dimension of overtness-covertness); the timing and phasing of the muscular contractions: and the particular system of effectors that produces the operant's defining effect (Stelmach, Mullins, & Teulings, 1984; Kelso, Tuller, Vatikiotis-Bateson, & Fowler, 1984; Summers, Sargent, & Hawkins, 1984; Rosenbaum, 1985; Ivry, 1986; Schmidt, 1988, pp. 187-298; Pew & Rosenbaum, 1988; Keele, Cohen, & Ivry, 1990; Semjen & Gottsdanker, 1990; Wiesendanger, 1990).[16] Thinking of operants as behavior routines modified by parameter settings suggests some possible mechanisms for the action of reinforcement.

One such mechanism is that a PPRf results in the repetition not of the most recent *behavior*, but rather of its most recent *direction of change*. In other words, the PPRf operates on the operants' parameter settings more like a vector than like a duplicator. It perpetuates its most recent *shift*, not its most recent setting. By way of an oversimplified illustration, if a certain operant has recently occurred twice, and if, in those two occurrences, the setting of one of its parameters shifted from 7 to 8, for whatever reason, a reinforcer presented after the second of those two occurrences would shift the setting again, this time from 8 to 9. (These numbers are only illustrative, of course). Presenting the reinforcer says, in effect, "Keep going in that direction," as one says in the children's parlor game "warmer/colder" in guiding the player toward a chosen object. This analogy also points up the ecological adaptation value of a parameter shift mechanism of reinforcement, and may explain the speed with which the shaping process

than Darwin needed to explain the origins of biological variants.

[16]These references, several of which are reviews of the literature, constitute a sampling of the recent literature in the fields of motor behavior and neurophysiology on which the summary statement in this paragraph is based. It is far from complete.

often proceeds.

9.5 Explaining Cyclic Behavior Patterns

The FCN operant contingency used in the Mechner, 1958b study (see Sections 3.4 and 8.7) is a quasi-rO_{FCN} under continuous reinforcement, its defining criterion being a certain run length. Figure 2 of that study shows that if the length of a run (the criterial measure) deviates from the mean run length in either direction, then the length of the next run deviates from the mean still more, in the same direction. To the degree that the FCN procedure is an rO, this finding could be regarded as preliminary suggestive evidence for the operation of the parameter shift mechanism.

The parameter shift mechanism also predicts cyclic fluctuations of run lengths, and these are certainly evident in Figure 3 of that study. The lengths of successive reinforced runs should keep shifting away from the mean until a reversal occurs. If we assume that the perpetuation of a parameter shift has a certain probability p which is less than 1.0, then a reversal will soon occur. The average number of consecutive shifts depends on the effective p. At the point of reversal there is a single initial parameter shift in the downward direction, and that is then the direction in which further parameter shifts are perpetuated, until a further reversal occurs. That next reversal, this time from the downward to the upward direction, usually occurs somewhat below the criterion. The average number of consecutive shifts in a given direction depends on p, and would be the same in the upward and downward directions. The average rO_{FCN} run length, which is the mean criterial measure across the cycles, normally tends to fall about 5-10% above the criterion.

The process just described is responsible for the oscillations we see in the lengths of consecutive response runs, and in most other "steady-state" behavior.

9.6 The Parameter Shift Mechanism and Behavior Shaping

As was implied above, the parameter shift mechanism may be important during shaping. Shaping procedures usually involve setting progressively more stringent criteria for reinforcement. The dimension along which the criterion is set corresponds to a parameter of the operant. Reinforcers are presented each time the operant's criterial measure has just shifted in the desired direction. When the parameter shift mechanism is operating, the result is a further shift of that parameter in the same direction. This sometimes creates the impression that the subject "understands what is desired," even though the underlying mechanism is quite mechanical. Since the parameter shift mechanism can produce very rapid behavioral changes, it may well be responsible for the speed with which shaping often proceeds.

Reinforcement does not necessarily shift all parameters equally. For example, the force or "overtness" parameter (i.e., degree of muscle engagement)[17] may be impacted more strongly or more frequently than other parameters (Morse, 1966, p. 54). The overtness parameter spans the range from the covert level, where there is no movement at all, to the overt level. When a covert response becomes more forceful, the degree of muscle engagement can reach a level where there is movement, at which point the response is overt. Hefferline & Keenan (1963) showed that when the criterion is a certain thumb muscle potential, operant contingencies can shift the overtness parameter from below to above the threshold for movement.

This would explain how a skilled animal trainer can evoke an operant that may be occurring at a covert level but has not yet occurred overtly. For example, the trainer knows from experience that when the animal fixates an object without yet moving its body, the behavior of moving toward the object may already be occurring at a covert level. A PPRf at that instant tends to impact the overtness parameter of that movement, with the result that an overt movement toward the object may follow.

9.7 Questions Regarding the Effects of PPRfs

These are some questions and plausible conjectures that can be investigated with rOs:

(a)   What determines *which* operants will be impacted most strongly by a particular PPRf? In Section 9.4 I suggested that it can be the operant whose parameters have just shifted, and that the PPRf's impact is to produce a further shift in the same direction.

---

[17]We must distinguish between two type of measures: One, which requires multiple instances, is the probability, frequency, or rate of the operant, regardless of its level of force or level of overtness, and the other, which is applicable to single instances, is the operant's level of force or overtness if and when it occurs. It may prove useful to subsume both of these types of measures under the construct of "response strength," but only if it is found that variables that increase one also increase the other.

But a PPRf may shift the parameters of other operants too. For example, it may selectively increase the overtness level of operants that have recently been at high levels of strength, or that have been followed by PPRfs in the past. Thus, a possible variation of the parameter shift mechanism is that a PPRf *produces* a parameter shift in certain operants regardless of whether or not those operants have just had parameter shifts.

(b)   How recent is the behavior (or parameter shift) that is repeated when a reinforcer is presented? The experimenter or trainer may have a certain operant and criterion in mind, and presents the reinforcer right after an instance of that operant has occurred. But the reinforcer's impact is not necessarily confined to that operant and that criterial measure (Catania, 1971; 1988). Recency of the targeted behavior may not be the only factor that determines which behavior is impacted. It is plausible that a PPRf can call forth behavior that occurred some time before, including behavior in non-criterial dimensions, because outside the laboratory, operants often produce delayed effects, even when shaping is occurring.

(c)   The recency of the behavior (or parameter shift) that is repeated may depend on the recent *density* of PPRfs. Thus, when the PPRf density has recently been high, as in an active shaping session, the behavior shifts that are called forth tend to be relatively recent ones, while in situations where PPRfs are sparse, or where there has been no PPRf for a long time, the behavior called forth may tend to be of older vintage.

(d)   The effect of a PPRf is not necessarily confined to a single instance of a parameter shift. There may be circumstances in which more than one parameter shift is impacted and perpetuated, i.e., where an entire block of preceding behavior (or parameter shifts) is impacted.

(e)   It is possible that the parameter shift mechanism operates only during shaping sessions and not at other times. For example, it has previously been observed that PPRfs, when they occur after a long period without a PPRf, tend to have an arousing or excitatory effect: The subject starts moving faster and more vigorously (Killeen et al., 1978). Once arousal has occurred, the parameter shift mechanism may swing into action, but not until then.

(f)   PPRfs have a stronger parameter shift effect on behavior (or parameter shifts) that have received PPRf at least once before, than on parameter shifts that are receiving a PPRf for the first time. In fact, the sensitivity of parameter shifts to PPRfs may increase as a function of the number of times they have previously been followed by PPRfs.

(g)   Do PPRfs have different types of effects at different stages of the shaping or automatization process, as the malleability and susceptibility of the behavior undergoes changes or diminishes?

(h)   When the parameter shift mechanism is *not* operative, PPRfs may function as S^Ds only, by selectively evoking behavior that was at high strength in similar situations in the past. As was explained in Section 9.2, every PPRf also functions as part of a compound S^D, with the effect of that S^D depending on the subject's earlier history.

(i)   PPRfs have a greater impact on parameter shifts that involve overtness increases than overtness decreases. This conjecture is plausible because outside the laboratory, an operant is effective more often when it increases in force, or when a previously covert operant becomes overt, than when it decreases in force or becomes covert.

9.8 Effects of Punishment

Punishment may work by *reversing* parameter shifts that were in the direction of greater overtness. It is possible that just as positive reinforcement can *increase* the degree of overtness, punishment can *decrease* it.

Punishment, like positive reinforcement, may affect not only the immediately preceding behavior but an extended block of preceding behavior. That would also explain the well-documented phenomenon of regressive resurgence produced by punishment or stress (Epstein, 1985; Mechner et al., 1992). If punishment has the effect of decreasing the overtness parameter of an extended block of preceding behavior, then the overtness of many individual operants in that block would drop below threshold, and cease to occur overtly or at all. Older behavior would then resurge, because its overtness level would become higher relative to the recently depressed behavior. The automatic result is regressive resurgence.

The punishment literature is replete with statements to the effect that punishment does not alter the *strength* of the punished operant, and suppresses it only temporarily (Keller & Schoenfeld, 1950; Azrin & Holz, 1966). The conjecture that punishment shifts the overtness parameter in the direction of increasing covertness would explain how punishment depresses or suppresses behavior without eliminating it. The conjecture is also plausible from the ecological adaptation standpoint: The same behavior that is punished in overt form is not punished in covert form ("Think it but don't say it"). Punished operants can occur in covert form and be retained in the behavior repertory for use at a future time when conditions for that behavior may be more favorable.

9.9 Implications for the Origins of Response Variability

The normally-observed variability of operant behavior may be due to the continuous action of reinforcements perpetuating the behavior's recent direction of change, thereby producing cyclic fluctuations and sequential effects like those seen in Figures 2 and 3 of Mechner (1958b). We know that during an organism's continuous normal interaction with its environment, large and small reinforcements, in various motivational modalities, constantly impinge on all of its behavior. When the parameter shift mechanism is operative, these reinforcements would generate response variability by continuously shifting the parameters of all ongoing behavior routines in their recent directions of change.[18] The operation of such a mechanism may also explain why the topography of superstitiously conditioned behavior tends to become cyclic, rather than fluctuate randomly (e.g., Skinner, 1948).

Thus, the parameter shift mechanism can explain how reinforcement generates response variants and how it can shape behavior.

9.10 A Research Program to Study the Effects of PPRfs

This section outlines the dependent and independent variables of an rO-based research program for studying the types of issues discussed above. The experiments should be done parametrically, because the interactions among the independent variables are likely to be important for the interpretation of the observed effects.

Dependent variables that should be examined are:

(a)   Comparisons of pre- and post-presentation occurrences of criterial and non-criterial measures. As stated above, the criterial and non-criterial measures represent the parameter settings for each occurrence of the rO. Compare the parameter *shifts* and the parameter *settings* before and after the PPRfs, and determine which (if either) of the two is more strongly perpetuated by the presentation (or reversed if punishment is used). The parameter shift mechanism would produce repetitions of parameter *shifts* rather than parameter *settings*.

---

[18]This view of variability follows Sidman's admonition (Sidman, 1960) that variability should not be viewed fatalistically as a manifestation of nature's indeterminacy, and then savored as a comfortingly reliable dependent variable, but should instead be viewed as a scientific challenge, with success measured by the degree to which the observed variations become predictable.

(b)   Relationships between the dozen-or-so criterial and non-criterial measures that immediately follow the PPRf, and those that occurred earlier in the subject's history, particularly just before and after previous PPRfs. See especially if there is a preponderance, or disproportionate representation, of the non-criterial measures that were most heavily represented in previous blocks of rOs with attention to the possible S^D effects of the stimulus compound that includes the PPRf.

(c)   Proximity of a non-criterial measure's position to the end of the rO in which it occurs. A non-criterial measure's sensitivity to the effects of PPRfs may depend on its proximity to the end of the rO (Mechner, et al. 1992).

(d)   To investigate the mechanisms of reinforcement at the neurological level, all of the behavioral measures obtained can and should be correlated with concurrently obtained neurological measures.

Independent variables that should be investigated are:

In the Reinforcer Category

(a)   Reinforcer value (e.g., amount of the reinforcer, or motivational level) and valence of the reinforcement. An aversive consequence, like time out or loss of money, can be used instead of positive reinforcement.

(b)   Density of PPRfs (a) just prior to the presentation being studied and (b) in previous sessions with which the post-presentation criterial and non-criterial measures are being compared.

(c)   Number of times the reinforcer has previously been presented (a) in the subject's history, and (b) in the experiment, under the reinforcement contingency being used.

In the Pre-Presentation Behavior Category: Type of behavior on which the presentation impinges

(a)   Total number of times the subject has previously emitted that rO.

(b)   Behavioral susceptibility, independently measured by one of the methods described in Chapter 8.

(c)   Type of rO used (i.e., its classification in terms of the operant contingency).

(d)   Time, or number of elapsed rOs, since the last PPRf.

(e)   Level of activity (e.g. keystrokes per minute, 1/L, etc.) since the last PPRf, or in the preceding block of rOs.

(f)   Reinforcement schedule on which the rO is being, and has previously been, maintained.

In the Reinforcement Contingency Category

(a)   Reinforcers are presented only when a certain specified shift has just occurred in the selected non-criterial measure. (Note that instituting such a contingency converts the non-criterial measure into a criterial one).

(b)   Reinforcers are presented only when a certain selected non-criterial measure has just occurred. (Again, note that instituting such a contingency converts the non-criterial measure into a criterial one).

(c)   Repeat *a* and *b* above for criterial measures.

(d)   The PPRf is contingent on *two consecutive* shifts in the same direction in two criterial or non-criterial measures. While two successive shifts will be rarer than single shifts, they may be more sensitive to the presentation. It would make sense from the ecological adaptation standpoint for them to be more sensitive, as presentations after double shifts would confirm more selectively and with a firmer basis that the shift is "on the right track."

   Note: The criterion for when to present the reinforcer is always based on the shift from a base reading to a comparison reading of criterial or non-criterial measures. The computer makes the determinations by monitoring the shifts on an on-line basis. The base reading can be the average obtained in an immediately-preceding block.

Procedure used to generate and maintain the baseline behavior

(a)   It is possible to generate a relatively stable stream of rOs by the use of an

intermittent reinforcement schedule, like VI or RI, that sustains long stretches of unreinforced rOs. Or a stable stream can be maintained by the presentation of PPRfs in another (weaker?) motivational modality. With human subjects, stable long streams can be sustained by verbal instruction.

(b)   Another way is to use continuous reinforcement, with PPRfs that are in the same motivational modality, but where each PPRf consists of an amount of reinforcement that is very small compared to the occasional PPRfs of larger amounts, the latter being the PPRfs that are being studied.

10.0 ACQUISITION, EXTINCTION, AUTOMATICITY, AND INTERFERENCE

10.1 How These Four Terms Are Related

The rO technique provides a new way to study the group of related phenomena various aspects of which have been referred to in the psychology literature as acquisition, extinction, automaticity, and interference. Though each of these terms has a different focus, the effects they refer to have in common their dependence on the number of prior repetitions of the behavior,[19] the number of times it has been reinforced, and certain other behavioral history variables.

The term "acquisition" focuses on the new behavior being learned, and on the process of its emergence. The dependent variables commonly used in acquisition studies have been the amount of time, number of trials, or number of reinforcements required to reach some criterion. Acquisition studies have usually disregarded the attributes and history of the nontarget behavior that is being replaced or displaced in the process.

The terms "automaticity" and "interference," on the other hand, focus on some target behavior's tendency to disrupt or interfere with other behavior. A commonly used dependent variable is the target behavior's own susceptibility to disruption by certain stimuli or by other ongoing behavior.

The term "extinction" focusses on changes in the behavior that is no longer being reinforced. Extinction studies have generally tended to disregard the nontarget behavior that is replacing it. This tendency is discussed in the paper "Extinction-Induced Resurgence" (Epstein, 1985) which makes the point that under conditions of extinction,

---

[19]The colloquial term "practice," though much used in the motor behavior literature, is avoided here because it implies that any type of repetition will produce performance "improvement," which is not the case. The term tends to draw attention away from the behavior that is being replaced or modified, and from the dependence of the progressive behavioral changes on the shaping process.

previously established behavior patterns resurge in ways that should be investigated systematically. In the past, the resurgence phenomenon has sometimes been observed informally. The Field et al. (1991) and Mechner et al. (1992) studies demonstrate resurgence under conditions of extinction using rOs. These results and Epstein's paper highlight the importance of taking into account the history (both experimental and pre-experimental) of the behavior being impacted by an experimentally introduced variable when the impact of that variable is being assessed. A closely related point was made by Mowrer (1940).

The Mechner (1959a) paper conceptualizes acquisition and extinction as special cases of transition behavior, as both involve transition from a starting performance to a new one.

10.2 A Research Framework That Encompasses These

The rO provides a research framework that encompasses acquisition, extinction, automaticity, and interference. These are some of the main dependent and independent variables that define that framework:

Dependent variables:

(a)   Criterial and non-criterial measures obtainable with rOs.

(b)   Susceptibility (as defined in Chapter 8) of the rO's various behavioral measures to various independent variables;

(c)   Concurrently recorded neurological data;

(d)   Degree to which emission of the rO disrupts or interferes with other concurrent behavior, or with the learning of new rOs;

(e)   Resurgence characteristics of the rO when learning histories have been experimentally installed and controlled.

Independent variables:

In the "behavioral history" category:

(a)   Number of prior repetitions of the rO;

(b)   Number of times the rO has previously been reinforced;

(c)   Amount of *differential* reinforcement the rO has received (for example: past shaping procedures, number of previous discrimination reversals, similarity and variety of other rOs from which the rO has previously been differentiated.

In the rO attribute category:

(a)   Type of rO it is (i.e., its classification in terms of the operant contingency);

(b)   Parameter settings of the rO (i.e., criteria that specify the operant contingency).

In the establishing operation category:

(a)   Motivational variables;

(b)   "Value" of the reinforcements used (amount, delay, probability, etc.);

(c)   Chemical/physiological variables.

In the interference category:

(a)   Introduction of various types of stimuli that have known ability to trigger, or set the occasion for, certain other operant behavior or sensory responses;

(b)   Standardized mechanical or chemical interventions.

10.3 Implications for Neuropsychology and Species Comparisons

There is now a rapidly growing literature, both theoretical and experimental, on the neurological correlates of the automatization process and its relationships to interference phenomena and attention (e.g., Carpenter and Grossberg, 1990; Kesner & Olton, 1990; Jeannerod, 1990; Colley & Beech, 1988; MacKay, 1987). Some earlier literature is reviewed by Shiffrin (1988). In particular, Decety (1992) showed by means of PET scans that there is a dramatic drop in the amount of brain activity as a function of number of repetitions of the behavior. The rO methodology may permit neuropsychology research generally to make even better contact with operant behavior research, and to deal with a broader spectrum of behavior.

The rO methodology also provides a way to extend neuropsychological research more broadly to non-human species, thereby making species comparisons feasible. For example, many of the important studies in the areas of automaticity, attention, perception, and interference, and some of the studies often characterized as "cognitive," use verbal instructions as independent variables. The methodological problem with verbal instructions is that their effect is often unpredictable because they rely on unstated assumptions regarding the verbal, social, and cultural histories of the subjects, assumptions that are normally so complex that they would defy any attempts to describe them.

The rO methodology makes it feasible to replace verbal instructions with experimentally installed operant contingencies that do not depend on verbal behavior for their effectiveness. The advantage of replacing verbal instructions with known operant contingencies is that doing so improves the specification of the experimental procedure and the degree of control over an important independent variable. Also, doing so offers a way to perform with non-verbal species many of the same experiments that have previously been performed only with human subjects, thereby broadening the generality of the empirical underpinnings of the emerging constructs.

REFERENCES

Antonitis, J. J. (1951). Variability of response in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, *42*, 273-281.

Azrin, N. H. & Holz, W. C. (1966). Punishment. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 790-826). New York: Appleton-Century-Crofts.

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

Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, *32*, 269-281.

Boren, J. J. (1961). Resistance to extinction as a function of the fixed ratio. *Journal of Experimental Psychology*, *61*, 304-308.

Bruner, A., & Revusky, S. H. (1961). Collateral behavior in humans. *Journal of the Experimental Analysis of Behavior*, *4*, 349-350.

Carpenter, G. A., & Grossberg, S. (1990). Neural dynamics of category learning & recognition: Structural invariants, reinforcement, and evoked potentials. In M. L. Commons, R. J. Herrnstein, S. M. Kosslyn, and D. B. Mumford (Eds.), *Quantitative analyses of Behavior, Vol. IX* (pp. 23-49). Hillsdale, NJ: Lawrence Erlbaum Associates.

Catania, A. C. (1971). Reinforcement schedules: The role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior*, *15*, 271-287.

Catania, A. C. (1973). The concept of the operant in the analysis of behavior. *Behaviorism*, *1*, 103-116.

Catania, A. C. (1988). Reinforcement schedules: Retroactive and proactive effects of reinforcers inserted into fixed-interval performances. *Journal of the Experimental Analysis of Behavior*, *49*, 49-73.

Cerella, J. (1990a). Shape constancy in the pigeon: The perspective transformations decomposed. In M. Commons, R. J. Herrnstein, S. M. Kosslyn, & D. B. Mumford (Eds.), *Quantitative analyses of behavior: Behavioral approaches to pattern recognition and concept formation.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Cerella, J. (1990b). Pigeon pattern perception: Limits on perspective invariance. Perception, 19, 141-159.

Colley, A. M. & Beech, J. R. (1988). *Cognition and action in skilled behaviour.* New York: Elsevier Science Publishing.

Conrad, B., Benneke, R., & Goodman, M. (1983). Pre-movement silent period in fast movement initiation. *Journal of Experimental Brain Research*, *51*, 310-313.

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

Decety, J. (1992). *Cortical and subcortical structures participating in the preparation phase in a visuomotor task: positron emission tomographic measurements of regional cerebral blood flow. International journal of psychology*, *27*, Issues 3 & 4, 16. Hillsdale, NJ: Lawrence Erlbaum Associates.

Dunn, R. M. (1982). Choice, relative reinforcer duration, and changeover ratio. *Journal of the Experimental Analysis of Behavior*, *38*, 313-319.

Epstein, R. (1985). Extinction-induced resurgence preliminary investigations and possible applications. *The Psychological Record*, *35*, 143-153.

Ferster, C. B. & Skinner, B. F. (1957). *Schedules of reinforcement.* New York: Appleton-Century-Crofts.

Field, D., Mechner, F., Hyten, D., & Madden, G. (1991). Paper presented at the Annual Convention of the International Association for Behavior Analysis, Atlanta GA.

Fields, L., Verhave, T., & Fath, S. (1984). Stimulus equivalence and transitive associations: A methodological analysis. *Journal of the Experimental Analysis of Behavior*, *42*, 143-157.

Georgopoulos, A. P. (1990). Neurophysiology of reaching. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control* (pp. 227-263). Hillsdale, NJ: Lawrence Erlbaum Associates.

Georgopoulos, A. P. (1992). Neural correlates of motor cognitive function. *International journal of psychology*, *405*, Issues 3 & 4, 16. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Gilbert, T. F. (1958). Fundamental dimensional properties of the operant. *Psychological Review*, *65*, 272-278.

Glenn, S. S., Ellis, J., & Greenspoon J. (1992). On the evolutionary nature of the operant as a unit of behavioral selection. *American Psychologist*. In press.

Guthrie, E. R. (1959). Association by Contiguity. In S. Koch (Ed.), *Psychology: A study of a science, Vol 2.* New York: McGraw-Hill Book Company.

Hefferline, R. F., Keenan, B., & Harford, R. A. (1959). Escape and avoidance conditioning in human subjects without their observation of the response. *Science*, *130*, 3385, 1338-1339.

Hefferline, R. F., & Keenan, B. (1963). Amplitude-induction gradient of a small-scale (covert) operant. *Journal of the Experimental Analysis of Behavior*, *6*, 307-315.

Herrick, R. M. (1964). The successive differentiation of a lever pressing response. *Journal of the Experimental Analysis of Behavior*, *7*, 211-215.

Herrnstein, R. J. (1961). Stereotypy and intermittent reinforcement. *Science*, *133*, 2067-2069.

Herrnstein, R. J. (1966). Superstition: A corollary of the principles of operant conditioning. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 33-51). New York:

Appleton-Century-Crofts.

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

Herrnstein, R. J., & Heyman, G. M. (1979). Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio? *Journal of the Experimental Analysis of Behavior*, *31*, 209-223.

Herrnstein, R. J., Vaughan, W., Mumford, D. B., & Kosslyn, S. M. (1989). Teaching pigeons an abstract relational rule: Insideness. *Perception and Psychophysics*, *46*, 56-64.

Hyten, C., Mechner, F., Field, D., & Madden, G. (1991). Paper presented at the Annual Convention of the International Association for Behavior Analysis, Atlanta, GA.

Iversen, I. H. (1982). Contiguity and response generalization. *Behavior Analysis Letters*, *2*, 185.

Ivry, R. B. (1986). Force and timing components of the motor program. *Journal of Motor Behavior*, *18*, 449-474.

Jeannerod, M. (1990). *Attention and performance XIII: Motor representation and control* (pp. 227-263). Hillsdale, NJ: Lawrence Erlbaum Associates.

Johnston, J. M., & Hodge, C. W. (1989). Describing behavior with ratios of count and time. *The Behavior Analyst*, *12*, 177-185.

Keele, S. W., Cohen, A., & Ivry, R. (1990). Motor programs: Concepts and issues. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control* (pp. 77-110). Hillsdale, NJ: Lawrence Erlbaum Associates.

Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology*. New York: Appleton-Century-Crofts.

Kelso, J. A. S. (1990). Cooperative phenomena in biological motion. In H. Haken (Ed.), *Synergistics of complex systems in physics, chemistry, and biology*. New York: Springer-Verlag.

Kelso, J. A. S., Tuller, B., Vatikiotis-Bateson, E., & Fowler, C. A. (1984). Functionally specific articulatory cooperation following jaw perturbations during speech: Evidence for coordinative structures. *Journal of Experimental Psychology*, *10*, 812-832.

Kesner, R. P., & Olton, D. S. (1990). *Neurobiology of comparative cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Killeen, P., Hanson, S. L., & Osborne, S. R. (1978). Arousal: its genesis and manifestation as response rate. *Psychological Review*, *85*, 571-581.

Klapp, S. T. & Erwin, C. I. (1976). Relation between programming time and duration of the response being programmed. *Journal of Experimental Psychology*, *2*, 591-8.

Kugler, P. N. & Turvey, M. T. (1987). *Information, natural law, and the assembly of rhythmic movement*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kutch, D. O. (1974). Differentiation of press durations with upper and lower limits on reinforced values. *Journal of the Experimental Analysis of Behavior*, *22*, 275-283.

Laties, V. G. (1972). The modification of drug effects on behavior by external discriminative stimuli. *Journal of Pharmacology and Experimental Therapeutics*, *183*, 1-13.

Laties, V. G., & Weiss, B. (1969). Further observation on overt "mediating" behavior and the discrimination of time. *Journal of the Experimental Analysis of Behavior*, *55*, 85-91.

MacKay, D. G. (1987). *The organization of perception and action: A theory for language and other cognitive skills*. New York: Springer-Verlag.

Margulies, S. (1961). Response duration in operant level, regular reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, 4, 317-321.

Marr, J. S. (1989). Some remarks on the quantitative analysis of behavior. *The Behavior Analyst*, *12*, 143-152.

Marr, J. S. (1990). Dimension in action: A commentary on Powell & Dickie. *The Behavior Analyst*, *13*, 11-14.

McCarthy, D. & Davison, M. (1980). Independence of sensitivity to relative reinforcement rate and discriminability in signal detection. *Journal of the Experimental Analysis of Behavior*, *34*, 273-284.

McDowell, J. J. (1988). Matching theory in natural human environments. *The Behavior Analyst*, *2*, 95-109.

McDowell, J. J. (1989). Two modern developments in matching theory. *The Behavior Analyst*, *12*, 153-166.

Mechner, F. (1958a). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior*, *1*, 109-121.

Mechner, F. (1958b). Sequential dependencies of the lengths of consecutive response runs. *Journal of the Experimental Analysis of Behavior*, *1*, 229-233.

Mechner, F. (1959a). *Some properties of a multi-response operant*. Unpublished manuscript. New York: Behavioral Science Applications.

Mechner, F. (1959b). A notation system for the description of behavioral procedures. *Journal of the Experimental Analysis of Behavior*, *2*, 133-150.

Mechner, F. (1962). Effects of deprivation upon counting and timing in rats. *Journal of the Experimental Analysis of Behavior*, *5*, 463-466.

Mechner, F., & Guevrekian, L. (1960). *Effects of exteroceptive feedback on "timing" and "counting" behavior*. Unpublished manuscript. New York: Behavioral Science Applications.

Mechner, F., Snapper, A. G., & Ray, R. (1961). Behavioral effects of methamphetamine and methylphenidate in rat and man. In E. Rothlin (Ed.), *Neuro-Psychopharmacology, Vol. 2*. (pp. 167-171). Amsterdam: Elsevier Publishing Company.

Mechner, F., Guevrekian, L., & Mechner, V. (1963). A fixed interval schedule in which the interval is initiated by a response. *Journal of the Experimental Analysis of Behavior*, *6*, 323-330.

Mechner, F., & Latranyi, M. (1963). Behavioral effects of caffeine, methamphetamine, and methylphenidate in the rat. *Journal of the Experimental Analysis of Behavior*, *6*, 331-342.

Mechner, F., Hyten, C., Field, D., & Madden, G. (1992). Using revealed operants to study the structure and properties of individual occurrences of operant responses. Paper submitted to *JEAB* in May, 1992.

Michael, J. (1982). Distinguishing between discriminative and motivational functions of stimuli. *Journal of the Experimental Analysis of Behavior*, *37*, 149-155.

Morse, W. H. (1966). Intermittent reinforcement. In W.K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52-108). New York: Appleton-Century-Crofts.

Mortimer, J. A., Eisenberg, P., & Palmer, S. S. (1987). Pre-movement silence in agonist muscles preceding maximum efforts. *Journal of Experimental Neurology*, *98*, 542-554.

Mowrer, O. H. (1940). An Experimental analogue of "regression" with incidental observations on Reaction-formation." *Journal of Abnormal and Social Psychology*, *35*, 56-87.

Muenzinger, K. F. (1928). Plasticity and mechanization of the problem box habit in guinea pigs. *Journal of Comparative Psychology*, *8*, 45-70.

Myers, D., & Myers, L. (1977). Undermatching: A reappraisal of performance on concurrent variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *27*, 203-214.

Nevin, J. A. (1979). Reinforcement schedules and response strength. In M. D. Zeiler & P. Harzem (Eds.), *Advances in analysis of behaviour, Vol. 1: Reinforcement and the organization of behaviour*, (pp. 117-158). Chichester, England: Wiley.

Nevin, J. A. (1981). Psychophysics and reinforcement schedules: An integration. In M. L. Commons & J. A. Nevin (Eds.). *Quantitative analyses of behavior, Vol. 1: Discriminative properties of reinforcement* (pp. 3-27). Cambridge, MA: Ballinger.

Nevin, J. A. (1984). Quantitative analysis. *Journal of the Experimental Analysis of Behavior*, *42*, 421-434.

Nevin, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin*,

103, 44-56.

Nevin, J. A. (1992). The discriminated Operant Revealed in Pigeon Multiple Schedule Performance: Effects of Reinforcements Rate and Extinction. In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Nevin, J. A., Mandell, C., & Atak, J. R. (1983). The analysis of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, *39*, 49-59.

Notterman, J. M. (1959). Force emission during par pressing. *Journal of Experimental Psychology*, *58*, 341-347.

Pear, J. J. & Legris, J. A. (1987). Shaping by automated tracking of an arbitrary operant response. *Journal of the Experimental Analysis of Behavior*, *47*, 241-247.

Pew, R. W. & Rosenbaum, D. A. (1988). Human movement control: Computation, representation, and implementation. In R.C. Atkinson, R. J. Herrnstein, G. Lindzey, R.D. Luce (Eds.), *Stevens' handbook of experimental psychology, Second edition, Volume 2: Learning and cognition* (pp. 473-509). New York: John Wiley & Sons.

Pierrel, R., & Sherman, J. G. (1963). Train your pet the Barnabus way. *Brown Alumni Monthly*, Feb. 1963, pp. 8-14.

Platt, J. R. (1979). Interresponse time shaping by variable-interval-like interresponse-time reinforcement contingencies. *Journal of the Experimental Analysis of Behavior*, *31*, 3-14.

Pliskoff, S. S., & Fetterman, J. G. (1981). Undermatching and overmatching: The fixed-ratio changeover requirement. *Journal of the Experimental Analysis of Behavior*, *36*, 21-27.

Powell, J., & Dickie, R. F. (1990). Kinetic output: A conceptual, dimensional, and empirical analysis. *The Behavior Analyst*, *13*, 3-10.

Rachlin, H., & Krasnoff, J. (1983). Eating and drinking: An economic analysis. *Journal of the Experimental Analysis of Behavior*, *39*, 389-404.

Reichenbach, H. (1938). *Experience and prediction.* Phoenix Books, 1961. Chicago: University of Chicago Press.

Requin, J. (1985). Looking forward to moving soon: Ante factum selective processes in motor control. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 147-167). Hillsdale, NJ: Lawrence Erlbaum Associates.

Restle, F. & Burnside, B. L. (1972). Tracking of serial patterns. *Journal of Experimental Psychology*, *95*, 299-307.

Rolls, E. T. (1981). Processing beyond the inferior temporal visual cortex relating to feeding, memory, and striatal function. In Y. Katsuki, R. Norgren, & M. Sata (Eds.), *Brain mechanisms of sensation.* New York: John Wiley.

Rosenbaum, D. A. (1985). Motor programming: A review and scheduling theory. In H. Heuer, U. Kleinbeck, & K. M. Schmidt (Eds.), *Motor behavior: Programming, control, and acquisition* (pp. 1-33). Berlin: Springer-Verlag.

Schmidt, R. A. (1988). *Motor control and learning: A behavioral emphasis*. Champaign, IL: Human Kinetics Publishers.

Schmidt, R. C., & Turvey, M. T. (1989). Absolute coordination: An ecological perspective. In S.A. Wallace (Ed.), *Perspectives on the coordination of movement* (pp. 123-156). New York: Elsevier Science Publishers.

Semjen, A. & Gottsdanker, R. (1990). Rapid serial movements: Relations between the planning of sequential structure and effector selection. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control* (pp. 409-427). Hillsdale, NJ: Lawrence Erlbaum Associates.

Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, R.D. Luce (Eds.), *Stevens' handbook of experimental psychology, Second edition, Volume 2: Learning and Cognition* (pp. 739-811). New York: John Wiley & Sons.

Sidman, M. (1960). *Tactics of scientific research*. New York: Basic Books.

Sidman, M. (1989). Equivalence relations: Some basic considerations. (VB Institute, Aguas de Lindoia. In S. Hayes & L. Hayes (Eds.), *Proceedings of Verbal Behavior Institute* (pp. 15-27). 1992. Reno, Nevada: Context Press.

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

Skinner, B. F. (1948). "Superstition" in the pigeon. *Journal of Experimental Psychology*, *38*, 168-172.

Skinner, B. F. (1969). *Contingencies of Reinforcement*. New York: Appleton-Century-Crofts.

Skinner, B. F. (1974). *About Behaviorism*. New York: Vintage Books.

Skinner, B. F., & Morse, W. H. (1958). Fixed interval reinforcement of running in a wheel. *Journal of the Experimental Analysis of Behavior*, *1*, 371-379.

Smith, L. D. (1990). Models, mechanisms, and explanation in behavior theory: The case of Hull versus Spence. *Behavior and Philosophy*, *18*, 1-18.

Stelmach, G. E., Mullins, P. A., & Teulings, H. L. (1984). Motor programming and temporal patterns in handwriting. New York Academy of Sciences conference on timing and time perception: Timing of motor programs and temporal patterns. *Annals of the New York Academy of Sciences*, *423*, 144-157.

Summers, J. J., Sargent, G. I., & Hawkins, S. R. (1984). Rhythm and the timing of movement sequences. *Psychological Research*, *46*, 107-119.

Sternberg, S., Monsell, S., Knoll, R. L., & Wright, C. E. (1978). The latency and duration of rapid movement sequences: Comparisons of speech and typewriting. In G.E. Stelmach (Ed.), *Information control in motor processing and learning*. New York: Academy Press.

Stevens, S. S. (1951). *Handbook of experimental psychology.* New York: John Wiley & Sons, Inc.

Terrace, H. (1987). Chunking behavior by a pigeon in a serial learning task. *Nature*, *325*, 149-151.

Vaughan, W., Jr. (1988). Formation of equivalence sets in pigeons. *Journal of Experimental Psychology*: Animal Behavior Processes, *14*, 36-42.

Wearden, J. H. & Burgess, I. S. (1982). Matching since Baum (1979). *Journal of the Experimental Analysis of Behavior*, *38*, 339-348.

Weingarten, K., & Mechner, F. (1960). The contingency as an independent variable of social interaction. In T. Verhave (Ed.), *The experimental analysis of behavior, 1966*. New York: Appleton-Century-Crofts.

Wiesendanger, M. (1990). The motor cortical areas and the problem of hierarchies. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control* (pp. 59-75). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, R. D. Luce (Eds.), *Stevens' handbook of experimental psychology, Second edition, Volume 2: Learning and cognition* (pp. 167-244). New York: John Wiley & Sons.

Wilson, M. P., & Keller, F. S. (1953). On the selective reinforcement of spaced responses. *Journal of Comparative and Physiological Psychology*, *46*, 190-193.

Winnick, W. (1949). Response vacillation in conflict situations. Unpublished doctoral dissertation. New York: Columbia University.

Wurtz, R. H. (1985). Stimulus selection and conditioned response mechanisms in the basal ganglia of the monkey. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 441-456). Hillsdale, NJ: Lawrence Erlbaum Associates.

AUTHOR INDEX

Skinner, 1969 (21)
Skinner, 1974 (46)
Smith, 1990 (50)
Stelmach et al., 1984 (55)
Sternberg et al., 1978 (44)
Stevens, 1951 (32)
Summers et al., 1984 (55)
Terrace, 1987 (35)
Wearden & Burgess, 1982 (37)
Weingarten & Mechner, 1961 (21), (22)
Wiesendanger, 1990 (42), (55)
Williams, 1988), (37), (42), (44)
Wilson & Keller, 1953 (54)
Wurtz, 1985 (42)

Created, Revealed, and Imposed Operants

Donald M. Baer

The University of Kansas

IN DEVELOPING THE CONCEPT of the revealed operant, Mechner offers a methodology for pursuing a long-neglected line of molecular behavior analysis, one that he specifies in properly systematic terms. In essence, he repeats a technique used earlier by Schwartz (1980, 1982) and Neuringer (1984; Page & Neuringer, 1985) for the pursuit of different questions, a technique of creating a class of alternative chains in order to study (Schwartz) or control (Neuringer) choice among them. Since Schwartz, Neuringer, and Mechner have already done that, I may rely on their precision, and here only loosely characterize their argument as explaining how a behavior gets from here to there--from stimulus control to functional consequence.

Mechner offers his method and then pursues the path it opens; he traverses an admirable number and choice of dimensions. Many of those dimensions are unspecified by the methodology but all of them are suddenly broadened by that application, sometimes in the sense that suddenly they are open to further analysis, and sometimes in the sense that suddenly they are clearer than they were.

How a behavior gets from here to there is not a totally neglected question in behavior analysis or epistemology, of course. Behavior analysis has its concept and technique of chaining, and academic epistemology has its technique of molecular clarification of arguments through footnotes. One of the most interesting aspects of Mechner's contribution is itself revealed in the traditional revealed operant of molecular academic behavior, the footnote. His introductory text (Section 1.1) reminds us that most of the research on which behavior analysis is based does not reveal the chain of behavior--the sub-operants--that precede the final operant. Instead, the final operant is the only one measured and scheduled for a functional consequence. Mechner also implies that there must always be such a chain occasioned by the stimulus control for that final operant.

He then writes, "In such experiments, the sub-operants are not easily recordable and are normally disregarded" (Section 1.1). The footnote that follows reveals a sub-operant chain that does indeed clarify how Mechner's argument will get from here to there: "In principle, it would be possible to take motion pictures of the movements comprising the sub-operants, and analyze them frame by frame, but such a process is too cumbersome and expensive to be practical" (Footnote 1, Section 1.1). That footnote posits the real existence of "the" sub-operants of an existing chain. It first *creates* a specially constrained class of possible sub-operant chains between a stimulus control and its operant, then *requires* that one of them (typically, but not necessarily, any one of them) occur before the operant occurs, and thus *imposes* a class of easily recorded sub-operants. That allows the study of how selections from this class occur, and Mechner has used it for just that purpose. Thus the argument has moved from positing that an organism acquires often difficult-to-observe sub-operants that we would like to study but cannot, to instead imposing on the organism a choice of easy-to-observe sub-operants so that we can study the choices, what happens to them, and what we can make happen to them.

Depending on context, that kind of movement is sometimes criticized, sometimes admired, by scientists. Our standard criticism often is embodied in a metaphor. At night, a passer-by notices a drunk searching the ground under a lamp post. The passer-by offers to help, and the drunk explains that he is looking for his dropped car keys. They look together for a while but do not find the keys. Finally the passer-by asks, "Are you sure you lost them here?" and the drunk replies, "No, I dropped them further up the street." "Then why," the passer-by asks in anger, "are we looking for them here?" and the drunk, pointing to the lamp post, replies, "Because the light is better here."

We invoke the metaphor as criticism when we, sober, know that the keys are not where they are most easily sought. We admire the tactic of looking only where the light is or can be made to shine, rather than being defeated by darkness elsewhere, when we are ignorant that the keys are not there--when sober analysis suggests that they might be there. Clearly, this is the second context: Although these sub-operants, when they are studied experimentally, are created and imposed rather than revealed, they do exist, given this methodology, and they do yield interesting answers. They may not be the keys we are seeking, but they are the keys to something.

At least, they are the keys to something we have missed so far, despite the earlier studies by Schwartz and Neuringer. Two classes of potential demonstrations are especially arresting: (1) certain experimental variables create changes in the sub-operants not obvious in the final operant, and (2) certain experimental variables create changes in the final operant yet leave its sub-operants intact for almost immediate recovery of that final operant, should the environment change. We may soon better understand how an operant changes getting from here to there, while still arriving at the same there, and how stopping an operant from getting there does not diminish the imminent possibility of its getting there again. As Mechner begins to show, looking into these classes of effects could resolve a number of contradictions and failed replications of the behavior-analytic literature (and others).

One might argue that this work is simply a resumption of study within the realm of chaining, and that it introduces a needlessly elaborate new terminology to discuss and extend an old concept--the chain schedule--which Ferster and Skinner (1957) already endowed with the necessary descriptive terms. Mechner's comments about automaticity--the eventual loss of cuing functions within chains--do not in themselves establish the need for new terms. On the other hand, we do not necessarily know a lot about chains, just because we know how to create and use chains. True, that kind of knowledge has allowed compelling explanations of acquired reinforcement and otherwise mysterious conceptualizations of self-control, and it has enabled numerous token-reinforcement systems in myriad successful applications. But we know little about creating classes of functionally equivalent chains, or about the contextual variables that can temporarily make all members of the class functionally equivalent or can temporarily make some functionally different from others. So, we may reasonably entertain some supplementary terminology to begin that exploration, especially if we treat it as supplementary by asking steadily, as the research it prompts unfolds, if it is still useful.

In his essay about the theoretical implications of revealed operants, Mechner is exceptionally comprehensive; he unfolds a world of conceptual analysis. The obvious reactive question is how that conceptual analysis will be made experimental, thereby allowing evaluation of its generality and applicability. Only his Chapter 3 deals with this question directly, and it is a short chapter. Furthermore, as argued above, Chapter 3 displays a methodology for imposing sub-operants rather than revealing them. That is not a criticism; imposing an analogue of a phenomenon often leads to useful analysis, and is a time-honored research tactic. However, it is only a tactic, in that it leaves the generality question still largely unanswered.

True, some of the precedents of analogue tactics are good. For example, a great deal was learned about the operant (the final operant, in Mechner's terms) by initial experimental tactics that restricted the operant to rats' bar presses and pigeons' key pecks, the consequences mainly to food and water access and electric shock, and the establishing operations mainly to food and water deprivation. Subsequently, those findings were tested for generality across a very wide variety of responses, settings, organisms, consequences, and establishing operations (a test that continues today). Largely because those operants were final operants, in Mechner's terms, the generality of earlier findings was obviously available for experimental evaluation.

In the present case, the conceptual status of the revealed operant is seen to be broadly analytic, but its experimental analysis is restricted to operants better characterized as imposed than revealed. Inevitably, if the experiments Mechner proposes for imposed sub-operants accomplish the analytic potential he has sketched for them, the generality question will become imperative, unless the role of revealed-operant analysis is to remain purely conceptual. Do the sub-operants we do not impose follow the same principles as those we do impose? Is Nature's programming just a broader case of our own experiments? For Mechner's case, the experimental evaluation of generality looks inherently problematic, as indicated in his first footnote. Whereas the analysis of final operants could be pursued intensively for decades, secure in the obvious fact that its generality could be tested later, here there may be no such security. Then it may follow that early research into revealed operants should target that difficult question of generality almost concurrently with its very tempting entry into easily imposed and studied chains of keyboard operants.

Finally, we may ask for a criterion of when to ignore an invitation to infinite regress. The revealed operant may prove dramatically analytic of what we know about the final operant, but its logic is extendable to itself. Does not every revealed operant have a not-yet-revealed operant preceding it? Would revealing that currently hidden operant clarify any variability of contradictions found in the analysis of the previously-revealed operant? And surely the currently hidden operant has an even-more-hidden operant preceding it. This line of reasoning can go on forever without ever reaching a conclusive answer.

Perhaps the answer will emerge as the demonstration that at some point, no further analytic power is gained by further revelation. Perhaps not. Clearly, demonstration is the necessary tactic of restraint. However, even the first demonstrations implicit in Mechner's essay look exceptionally interesting; their realization should have urgent priority.

REFERENCES

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement.* New York: Appleton-Century-Crofts.

Neuringer, A. (1984). Melioration and self-experimentation. *Journal of the Experimental Analysis of Behavior, 42,* 397-406.

Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 429-452.

Schwartz, B. (1980). Development of complex stereotyped behavior in pigeons. *Journal of the Experimental Analysis of Behavior, 33,* 153-166.

Schwartz, B. (1982). Failure to produce response variability with reinforcement. *Journal of the Experimental Analysis of Behavior, 37,* 171-181.

The Shape of Things to Come:
A Commentary on Mechner's
"The Revealed Operant"

M. Jackson Marr

Georgia Institute of Technology

FRANCIS MECHNER HAS from his earliest published work displayed a considerable talent in the *analysis* of behavior both in the sense of detailing characteristics of behavior-contingency interactions and in the quantitative approach and display of the outcomes of those interactions. He was, for example, one the first investigators to measure and emphasize the sequential properties of operant behavior. His method of sequential dependency display, namely plotting properties of the *n+1*th response against the *n*th response has recent application in the search for chaotic dynamical properties of behavior (Marr, 1992; Palya, 1992).

His article, "The Revealed Operant: A Way to Study the Characteristics of Individual Occurrences of Operant Responses," is clearly an outgrowth of his earlier work and is an impressive and stimulating treatment of an astonishing variety of fundamental issues and problems in behavior analysis. The paper invites comment at every turn, but I shall limit my remarks to three general areas. The first is towards the basic scheme of distinguishing behavioral units. The second is directed at the discussion of behavioral susceptibility to change. The third pertains to the treatment of shaping.

The Lively Unit

In attempting to expand and deepen our perspective on the operant, Mechner is clearly reacting to the long history of treating an operant as a functional class whose behavioral details are of little concern, so long as the behaviors converge on a final common effect--for example, a switch closure. Moreover, traditional emphasis has been almost entirely on the ongoing rate of those repeated switch closures as the primary, if not exclusive, feature to assess in relation to contingencies of consequences. Skinner never wavered from the primacy of rate as embodying the essential dimension of a behavioral analysis. This program of research has, without question, been immensely successful in the exploration of contingencies, and in engendering an enormous domain of principles, issues, and controversies reaching from the foundations of behavior to the most vital of applied concerns. Yet, I would argue, to some extent in parallel with Mechner, that from the beginning and at an accelerating pace, a significant number of investigators (Mechner, of course, was a pioneer) have explored a variety of conditions addressed in the monograph as "revealed operants." By this remark, I refer to the establishment and analysis of conditions that require detailed specification of behaviors beyond the simple rates of what Mechner is calling the iO. This work ranges from Findley's (1962) "multi-operant behavioral repertoires" to Palya's (1992) analysis of sequential dependencies and harmonic structure of literally tens of millions of key pecks under various contingencies.

Also, in at least the spirit of Mechner's approach, is the considerable work done with second-order schedules (following from Kelleher and Gollub, 1961 and Findley, 1962) exploring the extremely challenging problem of behavioral units (see, for example, Marr, 1979 and Shimp, 1979). Here interest was not focused primarily on rates, per se, but on patterns and even patterns of patterns of responding, going far beyond the individual key peck to see, for example, how one might scale up properties of ordinary schedules of single key pecks to complex sequences of those key pecks. To some extent, second-order schedules have demonstrated certain invariances under change of scale in characteristic

patterns of responses engendered by reinforcement schedules. Thus, a schedule contingency engenders a pattern of responding that may act as a unit, in that its temporal sequence under a second-order contingency will share properties with the pattern of individual responses under the same contingency class. For example, reinforcing a ratio performance under an interval schedule produces a typical fixed-interval pattern of ratio units (Kelleher, 1966). This statement is, of course, one way to *define* a unit. Such a definition arises from an already established functional unit, for example, the key peck. The essential requirement is that some specifiable behavior can enter into a functional relation with some specifiable consequence. Presumably, any definable dimension or combination of dimensions of behavior might satisfy that requirement and thus qualify as a behavioral unit.

Exemplifying this approach to unit generation, several studies of Zeiler dealt with attempts to control, through differential reinforcement of ongoing performance, features of schedule behavior that cannot be reduced to a single rate measure. Zeiler (1970; 1971) and DeCasper & Zeiler (1974) conducted a series of studies with fixed-ratio schedules showing that pause times and run times could be differentially (and independently) reinforced. Zeiler's work also revealed the limitations of differential reinforcement of schedule performance features. Attempts to control response number in fixed-interval schedules or pattern indices, such as quarter-life, resulted in behaviors moving in directions opposite to those required. What is more, other experiments demonstrated that even though discriminative control could be demonstrated for differences in performance features, those features themselves were unmodified by differential reinforcement (Zeiler, 1979).

A number of other examples of analysis and synthesis of behavioral units could be cited (e.g., Catania, 1983; Epstein, 1991; Thompson & Lubinski, 1986; Thompson & Zeiler, 1986), but my principal point is that, as ingenious as Mechner's methods and suggested studies are, they fit neatly into a general research program with a long history of exploring relatively complex contingencies. Given that assessment, some of Mechner's nomenclature seems misleading. For example, "sub-operant" implies a kind of subordinate status to classes of behavior that, in fact, could be the primary focus of functional control. Perhaps the rO should be called a "super-operant." The deeper issue here is raised by Mechner's distinction between "operant contingencies" and "reinforcement contingencies." The term "contingency" should be reserved to describe the *functional* relation between behavior and consequence (with appropriate inclusion of discriminative stimulus arrangements). A behavior's *description* has no clear consequences for that behavior, and thus is not a contingency. At best, it is a characterization of nominal structure. The relationship between a behavior class and a consequence in operant conditioning is similar to the relation between mass and force in Newton's second law of motion. The terms have no definition independent of the relationship between them. Like a rope, there are two ends; one cannot have a rope with only one end. Moreover, we may *arrange* a contingency, but fail to modify the specified behavior appropriately, as Zeiler's work and that of others have taught us.

Mechner's revealed operant bears closest relation to second-order schedule units. As mentioned earlier, such units may act in a similar way to single responses (e.g., key pecks). However, and Mechner may already be aware of this fact, the scheduling of such putative units can produce interactive effects (I hesitate to say "emergents"), that is to say, patterns that reflect combinations of contingencies (Marr, 1979). A similar kind of effect was seen in DeCasper & Zeiler's (1974) experiment with selecting run

times in fixed-ratio schedules. As the run time requirement was increased, the typical fixed-ratio pattern was replaced by a slow, but relatively steady, dribble of responses. What emerged was no longer a fixed-ratio performance. Whether it can be described as a unit is a matter for debate. Behavioral units are not static or consistent entities; they are lively, dynamic, capricious, and even ghostly. The principal concern I have about Mechner's procedures is interpreting the outcomes from the perspective of such slippery units, as well as their generality for any class of response characteristics beyond that expected from arranging consequences.

Behavior Dynamics

I am delighted to see a treatment of issues of behavior change and transition effects. Happily, there is a growing trend in behavior analysis towards theory and experiment in the direction of behavioral dynamics as opposed to the traditional behavioral statics. Behavior analysis began with a strong focus on dynamical processes, as Skinner's *The Behavior of Organisms* (1938) demonstrates. However, these kinds of investigations receded along with the ascension of interest in the steady-state properties of schedules of reinforcement. Perhaps this trend contributed to a general and largely intuitive notion that the construct "response strength" was embodied in rate of responding. For example, high-rate performances produced by random-ratio schedules seem somehow "stronger" than a moderate rate performance such as random-interval even if reinforcement frequency is the same in both schedules. We know now, of course, that this is absurd. But it took the imposition of changing conditions to begin to teach us that point. The correspondence with classical mechanics is close: to understand motion, you must study forces, either by imposing them, or by analyzing their effects as you find them in nature. Likewise, to understand behavior you must change it. Nevin has made a truly great contribution to behavior analysis by drawing our attention to this elegantly simple principle (e.g., 1979; 1992).

Concepts like "behavioral momentum" (and its apparent inverse, "susceptibility") require by their very nature methods for the continuous assessment of behavioral change under a variety of conditions. Mechner emphasizes adjusting procedures to study the variations in behavior under their control. These kinds of procedures have a history reaching back to Ferster & Skinner (1957) who used fixed-ratio schedules, as Mechner points out. There are examples as well using extended operants under second-order schedules (e.g., Marr, 1971; 1979). Unfortunately, no attempt was made in studying second-order schedules to examine quantitative features of cyclicity engendered by adjusting contingencies. What is clear is that the effects of those contingencies on maintaining otherwise "weak" behavior can be enormous. For example, chained schedules can produce very extended periods of non-responding, but appropriate adjusting arrangements can virtually abolish those periods (Marr, 1971).

As interesting and potentially significant as adjusting procedures are (whether applied to iOs or rOs), they contain complexities that can mislead a research program devoted to the experimental analysis and quantitative modeling of transition effects. First of all, as Mechner reminds us, there are "background periodicities" in ongoing operant behavior. He would prefer to look upon these as "noise" in relation to his adjusting contingencies. But to the contrary, these background behavioral periodicities are not only an inherent outcome of the imposed contingencies, they *drive* the adjusting procedures themselves. To take a simple example, under an adjusting fixed-ratio schedule, the ratio requirement is adjusted on the basis of the pre-ratio pause (PRP) duration. But the ratio contingency controls variations in the PRP. Thus, as the ratio requirement is increased, the PRP does not simply increase in length, it increases in variability. It is primarily the variation in PRP with which the adjusting procedure makes contact. This process is the essence of shaping. The variations in PRP under an ordinary fixed-ratio schedules have received little quantitative attention, yet such variations lie at the heart of the schedule dynamics. The same holds true about pause time and response number in fixed-interval schedules, although more data are available on this topic (e.g., Dews, 1970; Gentry & Marr, 1982; Lowe & Wearden, 1981; Shull, 1971; Wearden, 1979). At present, no dynamical theory can account for the variations in performance features under schedules (see, e.g., Zeiler, 1979, for a fascinating experimental treatment of this problem).

The second difficulty with adjusting procedures in studying transition effects arises from the feedback dynamics, which can be very complex because the behavior of the system determines the transition effects. Because we are only beginning to understand feedback functions as they describe contingency effects, the exploration of very complex systems might provide extraordinary outcomes, but not necessarily interpretable ones. To take an example from classical mechanics, consider the simple pendulum. Assume for the moment that you had no theory of the pendulum, but only set out to explore its properties by careful measurement. An effective strategy might be to fix the length and mass of the pendulum and measure the period as a function of the displacement. Next, one might change the length and repeat the period measurements. Then one might change the mass, etc. What you would *not* do is arrange a pendulum system where, for example, the pivot moved with the pendulum, or instead of using say, a thin rigid rod to support the mass, you used a spring! The results would be interesting and orderly, but you wouldn't understand them, or worse, you would misunderstand them. Only after you had developed a theory (e.g., classical mechanics) dealing effectively with the "simple" case (be ready to deal with elliptic integrals for large displacements), would you be likely to tackle the more complex cases.

In studying transition effects, we might take some lessons from control theory, and treat the organism as a black box which we explore by seeing how it transforms inputs into outputs (see, e.g., McDowell & Kessel (1979), for a pioneering effort in this direction, and McDowell, Bass, & Kessel (1992), for a recent application to transition effects). For relatively simple cases, it is possible to formulate models that provide at least first approximations to available data. A primary example is the change in rate of responding under a random-interval schedule when the interval parameter changes (i.e., the frequency of reinforcement changes) (Hunter & Davison, 1985; Marr, 1992; McDowell, et al., 1992).

Shaping

I found this section of the monograph to be wonderfully stimulating. It deals with the most vexing and fundamental problem in behavior analysis: How does reinforcement act to change behavior? Attempts to answer this question are the bases for such issues as molar versus molecular accounts, melioration versus maximization, reinforcers as strengtheners versus reinforcers as selectors, reinforcers as events versus reinforcers as behaviors of differential probability, etc., etc. In the sense of the blind men confronting the elephant, all these views are correct, but the phenomenon of reinforcement must encompass much more than we now are able to discern.

Mechner's "vector analysis" approach is creative and exciting, if only because of the images it engenders. In one vision I see the total behavior of the organism as a gigantically complex, multi-dimensional vector field; each point in this field has components (the magnitudes of Mechner's parameters) and a direction given by the resultant of the components. The field is not stationary, but whirls and swirls and ebbs and flows. At times there are regions of torrents and turbulence; at other times and places quiet, and even stagnant pools. Reinforcers disturb this field selectively by acting on those points or regions with the greatest rates of change. In another image, I see an abstract phase space like those used to picture dynamical systems. Here positions (i.e., parameter magnitudes) are plotted against velocity (i.e., rate of change in parameter magnitudes). "Attractors," that is, regions of local stability, continuously shift in this behavior space. Again, forces preferentially acting on points with the highest velocities drive the system.

There are many difficulties with these pictures. One is that they are not complex enough! A principal reason is that the set of behavior parameters interact with each other differentially in space and time, thus they cannot form an orthogonal set to comprise a comprehensible coordinate system. Even if we ignore this fact for sake of simplification, Mechner's model runs into problems because, given that reinforcement operates preferentially on those parameters that are changing, part of the system would grind to a halt and the rest would run away in a behavioral frenzy. Such a system lacks dynamic stability. Mechner, of course, recognizes this problem and has to add other possibilities to his model. Some of these are plausible; for example, all classes of parameters not being equally sensitive to reinforcement, or that reinforcement might act on some currently unchanging parameters. The problem is: how do we predict when these kinds of things will happen, as opposed to other kinds of things?

Other suggested constraints seem less plausible. The stability of a parameter-shift system depends on mechanisms that ultimately bring about a dynamic equilibrium. Mechner suggests that cyclic fluctuations (a *limit cycle* in dynamical terms) may occur because the perpetuation of a parameter shift may have a probability less than one, so that reversals are possible. Perhaps I am misunderstanding something, but except under very special circumstances, this arrangement would also be unstable. The closest analogy I can think of is a random-walk problem. Imagine a particle initially sitting at the origin. Every second, say, with a certain probability, p, it jumps one place to the right or, alternatively, with probability $q = (1 - p)$, it jumps one place to the left. If, for example $p > q$, then the particle will ultimately be as far to the right as you can imagine from its starting place. Only if $p = q$ will the particle wander aimlessly back and forth

across the starting place forever. Needless to say, this circumstance is unlikely in any realistic system.

One way to place limits on changes is to invoke a dynamic where the effect of a reinforcer presentation is proportional to the difference between present conditions and an asymptotic state (an example is the Rescorla-Wagner model of conditioning). This suggestion seems contrary to Mechner's suggestion that the sensitivity to parameter shifts might *increase* with exposure to reinforcer presentations. But that, in turn, seems contrary to acquisition data which always has an asymptotic level.

An asymptotic dynamic would not, by itself, account for cyclicity. Perhaps another possibility is that a reinforced shift in one parameter may interact with another parameter such that the first will be counteracted by the second, and so on. Aside from cyclic effects, any account of shaping has to deal with selectivity in that features of behavior that initially dominate may disappear, while other, less dominant behaviors come forth (see, for example, Zeiler, (1977) on acquisition of fixed-ratio responding). Another problem refers to behavioral units. In the course of acquisition, units change, coalesce, fragment, disappear, and otherwise transform in perplexing ways. The behavioral unit structure of a skill, for example, helps define its state of acquisition, and its "automatic," as well as "controlled" features. We seem to be a long way from understanding the dynamics of acquisition and maintenance of even the simplest skill, despite the long and continuing history of investigation by most able researchers. Mechner will have made a major contribution if either his theoretical or his empirical investigations can lead us closer to the many fundamental problems he has addressed. Whatever the ultimate correctness of his approach, we will certainly learn much in the doing.

REFERENCES

Catania, A. C. (1983). Behavior analysis and behavior synthesis in the extrapolation from animal to human behavior. In G. C. L. Davey (Ed.), *Animal models of human behavior* (pp. 51-69). Chichester, England: John Wiley & Sons.

DeCasper, A. J. & Zeiler, M. D. (1974). Time limits for completing fixed ratios, III: Stimulus variables. *Journal of the Experimental Analysis of Behavior*, *22*, 285-300.

Dews, P. D. (1970). The theory of fixed-interval responding. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 43-61). New York: Appleton-Century-Crofts.

Epstein, R. (1991). Skinner, creativity, and the problem of spontaneous behavior. *Psychological Science*, *2*, 362-370.

Ferster, C. B. & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

Findley, J. D. (1962). An experimental outline for building and exploring multi-operant repertoires. *Journal of the Experimental Analysis of Behavior*, *5*, 113-166.

Gentry, G. D. & Marr, M. J. (1982). Intractable properties of responding under a fixed-interval schedule. *Journal of the Experimental Analysis of Behavior*, *37*, 233-241.

Hunter, I., & Davison, M. (1985). Determination of a behavioral transfer function: White-noise analysis of session-to-session response-ratio dynamics on concurrent VI VI schedules. *Journal of the Experimental Analysis of Behavior*, *43*, 43-59.

Kelleher, R. T. (1966). Chaining and conditioned reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 160-212). New York: Appleton-Century-Crofts.

Kelleher, R. T. & Gollub, L. R. (1962). A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, *5*, 543-597.

Lowe, C. F. & Wearden, J. H. (1981). A quantitative model of temporal control on fixed-interval schedules. In C. M. Bradshaw, E. Szabaldi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behavior* (pp. 177-188). Amsterdam: Elsevier.

Marr, M. J. (1971). Sequence schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *15*, 41-48.

Marr, M. J. (1979). Second-order schedules and the generation of unitary response sequences. In M. D. Zeiler & P. Harzem (Eds.), *Reinforcement and the organization of behaviour* (pp. 223-260). Chichester, England: John Wiley & Sons.

Marr, M. J. (1992). Behavior dynamics--One perspective. *Journal of the Experimental Analysis of Behavior*, *57*, (in press).

McDowell, J. J. & Kessel, R. (1979). A multi-variate rate equation for variable-interval performance. *Journal of the Experimental Analysis of Behavior*, *31*, 267-283.

McDowell, J. J., Bass, R. & Kessel, R. (1992). Applying linear system analysis to dynamic behavior. *Journal of the Experimental Analysis of Behavior*, *57*, (in press).

Nevin, J. A. (1979). Reinforcement schedules and response strength. In M. D. Zeiler & P. Harzem (Eds.), *Reinforcement and the organization of behaviour* (pp. 117-158). Chichester, England: John Wiley & Sons.

Nevin, J. A. (1992). An integrative model for the study of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, *57*, (in press).

Palya, W. L. (1992). Dynamics in the fine structure of schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior*, *57*, (in press).

Shull, R. L. (1971). Sequential patterns in post-reinforcement pauses in fixed-interval schedules of food. *Journal of the Experimental Analysis of Behavior*, *15*, 221-231.

Shimp, C. P. (1979). The local organization of behaviour: Method and theory. In M. D. Zeiler & P. Harzem (Eds.), *Reinforcement and the organization of behaviour* (pp. 261-298). Chichester, England: John Wiley & Sons.

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

Thompson, T. & Lubinski, D. (1986). Units of analysis and the kinetic structure of behavioral repertoires. *Journal of the Experimental Analysis of Behavior*, *46*, 219-242.

Thompson, T. & Zeiler, M. D. (Eds.). (1986). *Analysis and integration of behavioral units*. Hillsdale, NJ: Erlbaum.

Wearden, J. H. (1979). Periodicities within a fixed-interval session. *Journal of the Experimental Analysis of Behavior*, *31*, 345-350.

Zeiler, M. D. (1970). Time limits for completing fixed ratios. *Journal of the Experimental Analysis of Behavior*, 14, 275-286.

Zeiler, M. D. (1972). Time limits for completing fixed ratios, II. Stimulus specificity. *Journal of the Experimental Analysis of Behavior*, *18*, 243-251.

Zeiler, M. D. (1977). Schedules of reinforcement: The controlling variables. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 201-232). Englewood Cliffs, NJ: Prentice-Hall.

Zieler, M. D. (1979). Output dynamics. In M. D. Zeiler & P. Harzem (Eds.), *Reinforcement and the organization of behaviour* (pp. 79-115). Chichester, England: John Wiley & Sons.

The Discriminated Operant Revealed in Pigeon Multiple Schedule Performance: Effects of Reinforcement Rate and Extinction

John A. Nevin

University of New Hampshire

MECHNER'S MONOGRAPH SETS FORTH a research program for the future, but it also reminds me of the past--specifically, the Columbia laboratories in the late 1950s and early 1960s. Mechner had completed his Ph.D. a year or two before I started graduate study, but his fixed-consecutive-number and fixed-minimum-interval procedures were alive and well in the Berryman-Cumming-Keller lab where I worked. Two of the articles that resulted from our analyses--one on mediating behavior in the fixed-minimum-interval procedure (Nevin & Berryman, 1963), and the other on the relation between observing-response latency and choice accuracy during fixed-ratio reinforcement of matching-to-sample (Nevin, Cumming, & Berryman, 1963)--explicitly analyzed the properties of individual occurrences of complex operant units. Reading Mechner's monograph revives that concern for me.

I was converted to molarism in the years after I left Columbia, and indeed I am persuaded that a molar analysis in which individual instances of operant behavior are aggregated over time is essential for the development of a cumulative, quantitative science of behavior that can be generalized beyond the laboratory (for a thorough discussion, see Baum, 1989). However, it can be informative and exciting to look at the local details of behavior and note the occasional striking instance that perfectly illustrates--or even better, utterly destroys--one's conception of what is going on within the molar aggregate. Here, I will describe an experiment inspired by an early draft of Mechner's monograph and referred to in the present version. It combines a molar approach to aggregated responding with an analysis of its components in relation to Mechner's "revealed operant" paradigm.

The subjects of the experiment were three white Carneaux pigeons, all thoroughly experienced with various multiple schedules, that were maintained at 85% of their free-feeding body weights. They worked for wheat reinforcers in a three-key chamber where the left key served as R_a in Mechner's "revealed operant" paradigm, while the center key served as R_b and the right key as R_c. To be specific: When the left key was lighted, a single peck turned it off and simultaneously lighted the center key. When the center key was lighted, a fixed ratio of ten pecks turned it off and lighted the right key. When the right key was lighted, a single peck turned it off and produced either a 3-second period of access to wheat or a 3-second blackout, after which the left key was relighted so that another operant unit could be initiated. Pecks at unlighted keys had no scheduled consequences. In Mechner's terminology, this is a fixed-number rO with explicit discriminative stimuli signaling each component of the unit.

Each daily session was segmented into alternating 3-minute periods during which the key lights differed. If the keys were lighted green, a variable-interval (VI) 40-second schedule was in effect such that the first R_c (signifying a completed occurrence of the operant) after an average of 40 seconds was followed by 3 seconds' access to wheat, whereas all other completed occurrences were followed by blackout. If the keys were lighted red, the schedule varied between two experimental conditions. In Condition 1 the schedule was VI 200 seconds, and in Condition 2 it was CRF (that is, every completed occurrence of the operant was reinforced). Condition 1 was in effect for 49 36-minute sessions, after which wheat reinforcers were discontinued for five sessions of extinction. Condition 2 was in effect for 29 18-minute sessions (where session length was reduced to keep the birds from gaining weight), after which there were another five sessions of extinction (each lasting 36 minutes so as to be comparable to extinction after Condition 1). The latency of left-key pecks (T_a), the time required for completion of the center-key ratio (T_b), and the latency of right-key pecks (T_c) were cumulated for each schedule component and then divided by the number of completed operants to give average values. In Mechner's terminology, T_a is

equivalent to L, the latency to initiate an operant, and the sum of T_b and T_c is equivalent to D, the duration of the operant.

Data for the final three sessions before extinction in Condition 1, and the final six sessions before extinction in Condition 2, were pooled to provide estimates of baseline performance in those conditions. The results are summarized in Table 1. Inspection of the table shows that T_a, the time to initiate an operant (or L, its latency) depended on the reinforcer rate. Within each condition, T_a was shorter in the component with the higher reinforcer rate, and between conditions, T_a in the varied, red-key component decreased markedly when the schedule was changed from VI 200 s to CRF. There was also evidence of behavioral contrast: T_a in the constant, green-key component was shorter when the varied schedule was lean than when it was rich. However, neither T_b nor T_c depended consistently on reinforcer rate either within or between schedules.

_____

Table 1. Mean time (in seconds) to initiate an operant (T_a), to emit a fixed ratio of ten pecks (T_b), and to complete the operant (T_c) for three pigeons in two multiple schedule conditions.

Condition 1:		VI 40 s (green)			VI 200 s (red)		
	Bird	T_a	T_b	T_c	T_a	T_b	T_c
	W33	2.16	1.91	0.57	4.86	2.82	0.55
	W34	1.61	2.33	0.59	6.31	2.84	0.58
	W35	2.18	2.05	0.56	8.52	4.39	0.79

Condition 2:		VI 40 s (green)			CRF (red)		
	Bird	T_a	T_b	T_c	T_a	T_b	T_c
	W33	2.81	3.62	0.54	1.86	2.88	0.57
	W34	2.55	2.67	0.57	1.28	2.98	0.56
	W35	5.10	1.84	0.59	1.60	1.84	0.62

_____

The data in Table 1 were reexpressed as relative speeds (the relative reciprocals of times, analogous to relative response rates), and plotted in Figure 1 in relation to relative obtained reinforcer rates. For example, $(1/T_a \text{ green})/(1/T_a \text{ green} + 1/T_a \text{ red})$ was plotted against (rft/session green)/(rft/session green + rft/session red). Figure 1 shows that relative speed to initiate the operant with R_a was highly sensitive to relative reinforcer rate, approximating matching--an unusual result for multiple schedules with long components (see Davison & McCarthy, 1988, for review)--whereas relative speed to emit the FR 10 (R_b) was less sensitive and relative speed to complete the operant (R_c) was virtually unaffected by relative reinforcer rate.

Average results for extinction are presented in Figure 2, where the data are reexpressed as speeds for each extinction session relative to the preceding baseline. The figure shows that the major effect of extinction was to decrease the speed with which operants were initiated $(1/T_a)$, although there is evidence of slowing in the speed of R_b during extinction after multiple VI 40 s, CRF (Condition 2). The major differences in changes of speeds of the components of the operant during extinction were consistent across subjects. However, there were no consistent differences within or between conditions in the rate of decrease of the speed with which operants were initiated.

The results for baseline performance and extinction illustrate Mechner's suggestion (Section 6.2) that the latency to initiate an operant depends on reinforcer rate and that the internal components of the operant are relatively insensitive to reinforcement schedules. Thus, reinforcer rate joins the family of variables such as deprivation and drug dosage that have their primary effects on the latency rather than the structure of an operant.

This conclusion is consistent with a good deal of previous research on schedules of reinforcement. In effect, the revealed operant is a chain (R_a --> R_b --> R_c), so the data reported here may be related to the literature on chained schedule performance. For example, Findley (1954) reported that the rate of responding in the constant initial link of a chain VI VI schedule was more sensitive to the reinforcer rate in the terminal link than was the rate of responding in the varied terminal link itself (see Kelleher & Gollub, 1962, for a review of the early literature).


Figure 1. Relative speeds of three successive components of a revealed operant--R_a, R_b, and R_c--as functions of relative obtained reinforcer rate in multiple schedules. The diagonal line in each panel indicates matching.
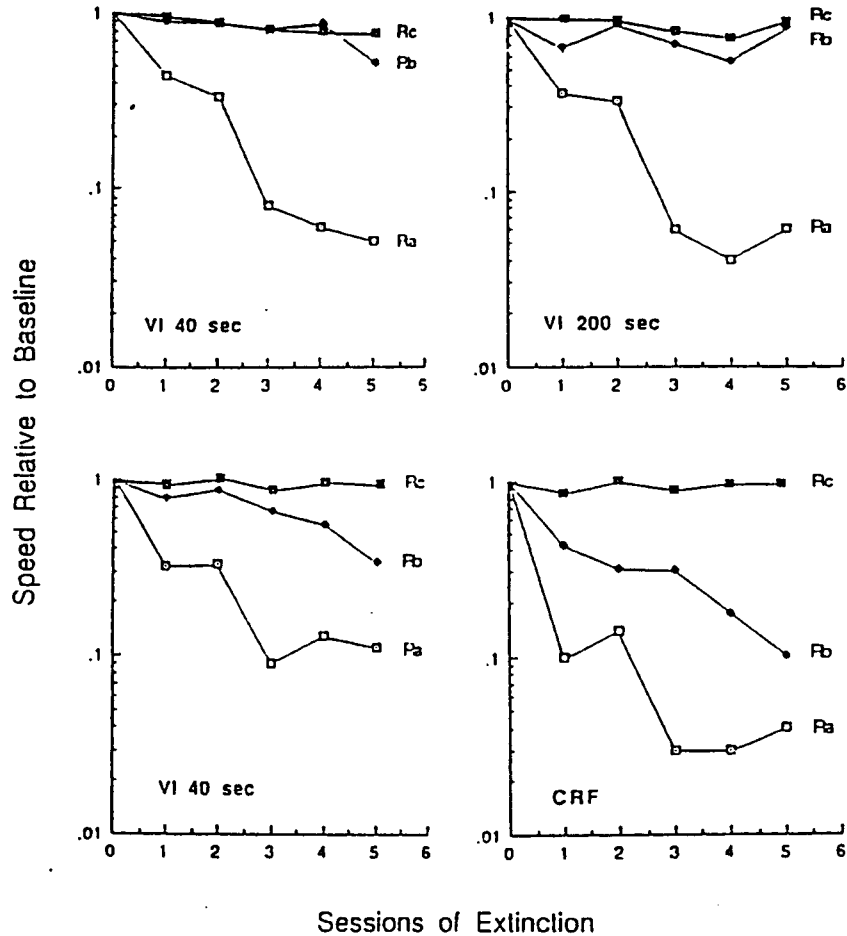
Figure 2. Speeds of three successive components of a revealed operant--Ra, Rb, and Rc--during five successive sessions of extinction, expressed relative to baseline speeds for each component. The upper sets of functions describe the course of extinction after training on multiple VI 40 seconds, VI 200 seconds, and the lower sets describe the course of extinction after training on multiple VI 40 seconds, CRF.

Likewise, Nevin (1964) found that the rate of responding in a constant FI or VI initial link was more sensitive to the probability of reinforcement for completing the chain than was the terminal-link response probability. More recently, Baum (1974) arranged a pair of chained schedules concurrently, and found that sensitivity to relative reinforcement was greater in the initial than in the terminal link, as reported above for multiple schedules. Finally, Nevin, Mandell, and Yarensky (1981) found evidence of initial-link contrast in two-link multiple chained schedules, and observed that initial-link performance was more easily disrupted by a change of conditions (analogous to extinction) than was terminal-link performance (see also Ferster & Skinner, 1957, pp. 678-680). If the initial and terminal links of a two-link chain are construed as R_a and R_b, and (unrecorded) departure from the operandum as R_c, and if response rate is taken to be analogous to speed, the present data are entirely consistent with previous findings.

A revealed-operant analysis also opens up some exciting new possibilities. Mechner describes some ways

in which revealed operants can be employed to address questions in areas as diverse as the nature of the reinforcement process, the matching law for choice behavior, and stimulus equivalence in conditional discriminations, and I will suggest another: The research practices that divide "instrumental" and "operant" behavior analyses.

The area traditionally termed "instrumental conditioning" relies heavily on discrete-trial procedures and within-trial speed data that can be characterized in revealed-operant terms. For example, for a rat in a runway, each trial can be construed as a revealed operant where leaving the start box is R_a, running down the alley in a series of little steps is R_b, and entering the goal box is R_c. The principal datum is running speed, $1/T_b$. The discrete-trial character of the method precludes measurement of T_a, the time to initiate a run after completion of the preceding run. By contrast, the area termed "operant conditioning" arranges conditions in which each response can be initiated at any time after the preceding response. The response itself is instantaneous and unstructured (as usually measured), and the principal datum is response rate (or its reciprocal, the interresponse time, T_a).

Because these two methods--identified with different research and theoretical traditions in the study of behavior--measure different components of behavior, it is hardly surprising that their results sometimes disagree. For example, research on extinction after discrete-trial training almost universally finds that responding persists longer after intermittent than after continuous reinforcement (the ubiquitous partial-reinforcement extinction effect). By contrast, free-operant responding falls off more slowly during extinction after continuous than after intermittent reinforcement, at least after extensive baseline training (Nevin, 1988). Revealed-operant methods permit the study of both initiation rate and within-operant speed, and may therefore be able to bridge the methodological chasm that has separated the study of instrumental and operant behavior. They may also serve to integrate molar and molecular approaches to behavior, as illustrated in the analyses reported here. As such, revealed-operant methods can help to unite some of the disparate research styles and their associated findings into a more coherent science of behavior.

REFERENCES

Baum, W. M. (1974). Chained concurrent schedules: Reinforcement as situation transition. *Journal of the Experimental Analysis of Behavior*, **22**, 91-101.

Baum, W. M. (1989). Quantitative prediction and molar description of the environment. *The Behavior Analyst*, **12**, 167-176.

Davison, M., & McCarthy, D. C. (1988). *The matching law*. Hillsdale, NJ: Erlbaum.

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

Findley, J. (1954). Unpublished doctoral dissertation, Columbia University. Data reprinted in Findley, J. (1962), An experimental outline for building and exploring multi-operant behavior repertoires. *Journal of the Experimental Analysis of Behavior*, **5**, 113-166.

Kelleher, R. T., & Gollub, L. R. (1962). A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, **5**, 543-597.

Nevin, J. A. (1964). Two parameters of conditioned reinforcement in a chaining situation. *Journal of Comparative and Physiological Psychology*, **58**, 367-373.

Nevin, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin*, **103**, 44-56.

Nevin, J. A., & Berryman, R. (1963). A note of chaining and temporal discrimination. *Journal of the Experimental Analysis of Behavior*, **6**, 109-113.

Nevin, J. A., Cumming, W. W., & Berryman, R. (1963). Ratio reinforcement of matching behavior. *Journal of the Experimental Analysis of Behavior*, **6**, 149-154.

Nevin, J. A., Mandell, C., & Yarensky, P. (1981). Response rate and resistance to change in chained schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 278-294.

Shaping in Terms of Parameter Shifts

Thom Verhave

Queens College

City University of New York

ONE OF THE MORE INTRIGUING proposals of Mechner's article is the theoretical reinterpretation of the shaping process. Since the earliest account I am aware of, by Alfred Smee (1850, quoted in Verhave, 1966), shaping has been described in terms of successive approximations to a final target behavior as specified by the trainer. According to Skinner's analysis, the trainer differentially reinforces successive approximations, and the property of behavior that makes shaping possible is the fact that behavior is variable. As Mechner points out, this explanation does not "account persuasively for the remarkable speed and efficiency with which the shaping process often proceeds." I have long shared Mechner's dissatisfaction with the statistical interpretation of shaping. He has offered, for the first time, a novel interpretation, which, given modern technology, seems to be eminently testable.

I wish I could provide hard quantitative evidence to show that there is indeed a parameter shift mechanism at work. Unfortunately, I cannot. I can, however, testify on the basis of personal experience as a trainer, that shaping can proceed with remarkable speed indeed. In fact it can proceed so smoothly that the traditional account, in terms of a vague unspecified spectrum of response variants from which the trainer selectively reinforces successive approximations, makes little sense at all. The trouble has been that there was no alternative interpretation until now.

My attempts at shaping long predate the introduction of the video recorder. Thus there is no objective record of what happened when I got a rat called Bozo to haul in a small metal chain with the fingers of his paws. The chain was attached at one end to a wire mesh basket and at the other end to a small platform on which the rat was standing. Once the chain was hauled in sufficiently far, the basket was in reach. By holding the chain with one paw, the rat could grasp the basket with the other paw and bring it into position where he could jump into the basket and swing to another platform. This component was but one of a much larger chain of tricks. The Bozo Box was inspired by Skinner's rat Pliny and an article in *Life Magazine* around 1953-54 about the work of Lo Tseng Chai, then at Tulane University. He had trained rats to perform a number of tasks which more or less mimicked those for which Sultan, Kohler's ape, had become famous.

The nature of response chains was a subject under active discussion among the graduate students at Columbia University. One question dealt with the maximum possible length of a chain, considering that all of its serial components were maintained by secondary reinforcers. Why not construct a chain consisting of all or some of the separate tasks that Lo Tseng Chai had trained his rats to perform? That idea led to the creation of the Bozo Box which was submitted as a contribution to the commemoration of Columbia University's 200 years of existence. The Bozo Box was later imitated by a number of colleagues such as Pierrel and Sherman (1963) who left out the most intricate part, the chain-hauling sequence described above. When you shape such a perceptual-motor sequence and do it in a half-hour session, it might well appear to an observer watching over your shoulder that you "are telling the rat what to do." What happened was that the desired behavior unfolded rapidly and smoothly with each sharp metallic click produced by a dime-store metal frog which was used as the secondary reinforcer. There is no waiting for a desirable "next successive approximation" while undesirable variants occur; they don't, as long as the trainer maintains concentration and keeps reinforcing the correct progressive shifts. Note, there are incorrect progressive shifts; the trainee after all does not know where the trainer wants to go. The point is that the spectrum of progressive shifts is far more narrow and far different from what a Gaussian distribution would lead one to expect.

At one time or another, I have shaped equally intricate acts with parakeets, mice, hamsters, and monkeys. My sense is that Mechner's new interpretation of shaping is well worth exploring with the objective and quantitative methodology now available. I look forward to reading relevant articles in the journal literature in the near future!

REFERENCES

Pierrel, R., & Sherman, J. G. (1963). Train your pet the Barnabus way. *Brown Alumni Monthly*, February, 8-14.

Verhave, T. (1966). *The experimental analysis of behavior*. New York: Appleton-Century-Crofts.

Response to the Commentaries

Francis Mechner

Verhave On the Parameter Shift Mechanism

I greatly value Thom Verhave's comments regarding my "parameter shift" hypothesis. I was there, in 1956, when Verhave created his spectacular "Bozo" demonstration. That demonstration, along with some of Verhave's other work, e.g., training pigeons to inspect pharmaceutical ampules, established Verhave in my eyes and those of our colleagues as "Mr. Shaping" incarnate. His comments therefore carry particular weight for me.

Nevin's Experiment on rO Component Sensitivity

I am flattered that Nevin not only commented on the rO method, but actually *used* it to study an important problem. As he states, he was at the Columbia University Department of Psychology in 1959 when I was circulating my "Multi-Response Operant" paper among the department's graduate students. That paper was the predecessor of the present monograph by some 33 years. Nevin reports a 1991 experiment in which he used the rO method to examine the relationship between (a) a behavioral measure's proximity to the operant's termination point and (b) the measure's sensitivity to an independent variable. Particularly exciting to me is the fact that the results he obtained parallel those we obtained with human subjects at the University of North Texas (Mechner et al., 1992).

Marr's Comment Regarding Contingencies

Marr's comments suggest that the distinction between operant contingencies and reinforcement contingencies needs clarification. I certainly agree with Marr that the term "contingency" should be reserved for the functional relationship between behavior and its consequences, and should not be used to refer to any description of behavior (Mechner, 1959b; Weingarten & Mechner, 1960). According to my definition, neither an operant contingency nor a reinforcement contingency refers to a description of behavior.

The concepts of operant contingency and reinforcement contingency, as I attempted to explain them in Section 2.6, apply to both rOs and iOs. A familiar example of an operant contingency for iOs is the requirement to press the bar down a certain distance until the switch is tripped. Also part of the operant contingency is the visual and tactile feedback generated as the bar moves downward and at the moment the switch is tripped.

Schedule-generated behavior qualifies as an rO only if it is initiated by an R_a, as I stated in Section 2.4. The reinforcement contingency that defines the schedule can then be viewed as an operant contingency as well, with the "schedule unit" being the rO. Such an rO could then in turn be reinforced on any desired schedule, and that schedule is then a reinforcement contingency. The Mechner et al. (1992) studies, and the Nevin (1992) study illustrate the reinforcement of rOs on fixed ratio and variable interval schedules. Those rOs are defined by their operant contingencies, and those schedules are

therefore reinforcement contingencies.

The Classification of Operants

The point I did not make sufficiently clear is that all operants can be classified according to their operant contingencies, whether they are studied as rOs or iOs.

Here are some examples of parameters that define different classes of iOs, just for bar pressing (or button pushing): The criterion can be met either by the down-press or by the release of the bar; the force may increase, remain the same, or decrease as the bar is pressed down; the bar may have to be held down a minimum time, or released within a given time, for the criterion to be met; feedback may or may not be provided when the criterion has been met; the criterion may be defined by pressure on a bar that remains immobile (e.g., a strain gauge); continuous visual or auditory feedback may or may not be provided as the bar moves through an excursion; in the case of key pecking, factors that define possible classes may include the distance the pigeon must move its head to reach the key, and what, if anything, it sees as it moves its head toward the key.

Baer On the Generality of rO Findings

The Generalizability of rO Data to iOs

Baer identified one of the key issues regarding the revealed operant: the issue of generality. As Baer points out, this issue must be raised for any model used in science, not just for the rO.

I see the generality issue as having two distinct parts: (1) whether what one learns about one class of rOs is applicable to other classes, and (2) the degree to which knowledge gained with *any* class of rOs is applicable to iOs. Since Baer focussed his attention on part (2), I will discuss it first.

In analyzing the generalizability of rO findings to iOs, the first task is to examine possible differences between deliberately installed (or "imposed" as Baer describes them) sub-operants, and "natural" ones. Does the process of deliberately installing sub-operants confer new properties on the operant? In other words, do rOs have properties that iOs don't have, properties that may interact with the phenomena we are studying?

Both rOs and iOs Constrain the Range of Sub-Operants

It is true that in designing an rO and its operant contingencies, the experimenter constrains the range of possible sub-operants. But the same is true for iOs. The only difference is that in the case of iOs, the constraints are imposed inadvertently.

When experimenters design an iO manipulandum and the other physical parameters (heights, distances, illumination levels, spring tensions, etc.) of the experimental situation, they determine aspects of the iO's operant contingencies. Some of these parameters are enumerated above in the section entitled *The Classification of Operants*. These parameters, though they are normally determined inadvertently and by happenstance, still constrain the iO's sub-operants. For example, the physical parameters of the manipulandum limit the range of movements and body positions that determine the domain of possible sub-operants. But even in spite of such constraints, and within that limited range, there are still an infinite number of useable movements and body positions.

The general issues seem to be: (a) Do rOs constrain their sub-operants *more*, or in different ways, than iOs do? And if the answer to that is yes, then (b) does the degree of constraint have implications for generality? I will try to show that rOs do not constrain their sub-operants any more or any differently than iOs do, and that we therefore do not reach question (b).

Comparing the Ways rOs and iOs Are Constrained

It seems to me that the sub-operants of iOs and those of rOs are constrained in comparable ways. I will compare those constraints for the standard bar-pressing type of iO and for the type of rO whose sub-operants are defined by a requirement to press any keys from a specified set. In this rO, any one of eight simultaneously available keys can be selected at the subject's option at each press.

One way the iO is constrained is that only a limited number of body parts can be used to press the bar (the four paws, the elbows, teeth, chin, etc.). This dimension of constraint is analogous to the various fingers that can be used on any of the eight available keys. Furthermore, regardless of which key is chosen at any particular

press, there is an infinite number of ways the key can be pressed (involving different movement topographies, sectors of the key surface pressed, finger positions, body positions, etc.).

Another seeming dimension of constraint in the case of the rO is that there are only eight rather than an infinite number of keys among which the subject can choose at each press. This constraint too has a counterpart in iOs. The rat bar (or pigeon key) can be thought of as being subdivided into eight sectors, corresponding to the eight keys. If it is then argued that the rO keyboard keys are discrete, whereas the sectors of the rat bar are continuous, that distinction can easily be eliminated by placing a heavy cloth over the eight available rO keys and eliminating the distinctive click of the switch that is produced when one of the keys is pressed. The subject would then be pressing down on a bar-like stretch of cloth, and

could apply pressure at any of an infinite number of points. The experimenter would still be recording depressions of the eight keys in the same way, but the subject would not see or feel any keys and would have no distinctive feedback when a key was pressed. In principle, analogous recordings could be obtained for iOs by means of the videotape method, which would make it possible to identify and record bar sectors pressed (or pigeon key sectors pecked) or movement paths traversed to reach the bar or key and to close the switch. Thus, the issue of discreteness versus continuity of the sub-operants appears to be simply one of how feedback is arranged.

Note that I am not advocating the actual use of any of the procedures described above. I am using them only as props in a conceptual analysis of the generality issue.

The Constraints of rOs' and iOs' Sub-Operants Are Equivalent

The preceding analysis leads me to conclude that:

(a)   There are no fundamental differences between rOs and traditional iOs. We install sub-operants in all cases, whether by design as in the case of rOs, or inadvertently and by happenstance as in the case of most iOs.

(b)   There is an infinite (though constrained) number of ways to execute each sub-operant, both for rOs and for iOs.

(c)   The degrees and types of constraints are comparable in the two cases.

(d)   There is no difference between "natural" versus experimentally installed (or "imposed") sub-operants. The two are equally natural or unnatural. The subject has no way to distinguish the experimentally intended operant contingencies from the rest of nature.

(e)   rOs are not necessarily more complex than iOs.

The above analysis of the rO-to-iO transferability issue is basically a conceptual one. Is there a way to address this issue empirically? The inherent methodological obstacle to doing so is that iOs are defined by the fact that their sub-operants are *not* recorded or analyzed. To settle the matter, we would have to record and analyze them. But the moment we record and analyze them so as to reveal the sub-operants, for instance by the cumbersome videotape method, the iO is turned into an rO. This consideration itself suggests that the distinction between rOs and iOs may be more one of perspective than of substance.

Why the Conclusions (a)-(e) Above Are Counter-Intuitive

I am aware of the fact that many operant behavior researchers, even if they find this analysis logically convincing, will still be left with an uncomfortable, intuitive feeling that the rO is *nonetheless* in some way more complex than an iO, and *nonetheless* in some way less "natural," as Baer implies. I empathize with these intuitive feelings, and, in fact, share them. But I think I can identify their origins and explain them away. They have their roots in years of working with iOs, chains of iOs, iO reinforcement schedules and iO response rates, and of viewing the iO's single recorded event as exemplifying the unit of behavior. The rO model proposes a different perspective, one that extends the operant unit back a bit to encompass some normally unrecorded pre-final movements, the "sub-operants." The sub-operants are "revealed" by being installed as markers. What bothers us is that the markers themselves exemplify iOs, our traditional response units. This fact, and this fact alone, makes it seem to one who is accustomed to working with iOs that the rO is more complex and more artificial than the iO. I believe that if the markers were something other than our traditional iOs (for example, successions of body positions during transitional movements revealed by a videotape recording), the illusion would vanish.

I hope that the above analysis of rO-to-iO generalizability will convince Baer and others, as it has convinced me, that the issue is an illusory one that stems from deeply ingrained perspectives and thought habits, rather than from any underlying reality.

Generality Across Classes of Operants

This is the second part of the generality issue to which I referred earlier. In the article, I attempted to show that with rOs we can address types of

questions and collect types of data that we cannot address or collect with iOs. One such question is that of inter-class generalizability.

As I stated above, classes of operants are defined by their operant contingencies. Chapters 3-5 describe some classes of rOs. Is what we learn about any one of these classes applicable to the other classes? This is an important question because the "real-life" situations to which we want to apply our findings encompass a wide range and variety of operant contingencies. Thus, the issue of inter-class transferability must always be an active concern when operant research findings are being extrapolated. Perhaps the reason this concern has rarely been expressed is that iOs do not provide a practical way to address it.

The only way to examine the question of inter-class generality empirically is first to identify and describe major classes of operants, and then to determine whether various otherwise equivalent experiments produce different results for those various classes.

Sub-Operants of iOs and Inter-Class Generality

For some types of iO experiments, different operant contingency classes may yield the same data, and for others, especially iO experiments that use sophisticated dependent variables, different classes may produce different data. There are certainly many types of iO experiments that would produce the same data for all classes, especially if the data being recorded were sufficiently crude and "molar." But the same experiments may produce different data for different classes when we look at the behavior's fine-grain structure and at subtler (but not necessarily unimportant) phenomena that can only be seen if pre-final (normally-unrecorded) sub-operants are recorded too.

In general, when iOs are used, the installation of the sub-operants is not deliberate, and the operant contingencies that define the iOs are not designed deliberately. Both are usually unknown and rarely discussed. The widespread unstated assumption is that they are unimportant.

In order to address empirically any question of generality regarding the equivalence of iO classes, it would be necessary (a) to specify the iO's operant contingency so that the iO can be classified, and (b) to record its pre-final sub-operants so that its distinguishing properties can be observed. The latter would require a frame-by-frame analysis of videotaped iO performance--a fairly impractical undertaking.

If the rO can be considered a sufficiently valid model of iOs to obviate the need for such a cumbersome and expensive procedure, then it provides a practical way to define classes of operants according to their operant contingencies. Conventional iOs do not provide a practical way to do so. Hence our need for the rO model: It provides a practical way to study the transferability of data between different classes (as defined by different operant contingencies).

Other Points Made by Marr

What Is a "Behavioral Unit"?

Marr, by implication, raises the important question of how behavioral units have been, can be, and should be defined. It seems to me that the term "behavioral unit" has traditionally been used in three senses:

(a)   Self-formed units consisting of behavior that has coalesced and jelled into a recurring sequence or topography, usually during the course of long-term repetition. Such self-formed units do not conform to the definition of an operant, since they do not produce a specified or necessarily consistent environmental effect.

(b)   Units produced by the experimenters' interrupting ongoing behavior. For example, ongoing behavior can be cut off at the end by the delivery of a reinforcement. So-called "schedule units" are often truncated in this manner. Similarly, the beginning of a unit can be defined arbitrarily by

the end point of the preceding "unit" (as when IRTs are regarded as units) or by the delivery of a reinforcement, or can be experimentally initiated by a programmed event such as the presentation of a stimulus (which creates an approximation to a trial situation). The issue of when such units do and do not conform to the definition of an operant was discussed in Section 2.4.

(c)   Units of behavior initiated by the subject and defined only by their effect on the environment, i.e., by an operant contingency. Such units conform to our definition of an operant. If their beginning and end points are experimentally defined by sub-operants that initiate and terminate the unit, and that are readily recordable, the resulting operant is useable as an rO.

I am drawing these distinctions between types of behavioral units in the spirit of elaborating Marr's observation that the rO concept is relevant to research on behavioral units. My point here is that not all types of behavioral units are operants. Operants are a special type of behavioral unit. It follows that results obtained with the rO model are not necessarily applicable to other types of "units," and vice versa.

How Cyclic Behavior Patterns Stay Within Boundaries

Based on Marr's commentary, *Explaining Cyclic Behavior Patterns* (Section 9.5) evidently needs clarification.

I believe that cyclic fluctuations are kept within their boundaries, and the system is kept stable, by the following mechanism: The average number of consecutive shifts in a given direction depends on the probability of a shift. For a given operant, a criterial parameter initially drifts upward with probability p until there is a reversal. In the case of rO_{FCN}, for example, the lengths of successive reinforced runs — the criterial parameter—keeps increasing until a reversal occurs. If we assume that the perpetuation of the parameter shift has a certain probability p less than 1.0, a reversal will eventually occur. At the point of reversal there is (by definition) a single initial parameter shift in the downward direction. This is the direction in which further parameter shifts are then perpetuated, this time probably with a probability greater than p (due to a "least effort" principle?). If and when the parameter value crosses the criterion in the downward direction, the operant ceases being reinforced, with a resulting jump in variability that should increase the likelihood of a new reversal from the downward to the upward direction. Such reversals usually occur somewhat below the criterion, perhaps explaining why the average rO_{FCN} run length, which is the mean value of the criterial measure across the cycles, usually tends to fall about 5-10% above the criterion.

The types of mechanisms outlined in the above paragraph may keep cyclic oscillations within bounds in many instances of "steady-state" behavior. It should be noted that observed cyclic fluctuations are usually not clean and monotonic, presumably because other types of variability-causing effects (which we could view as "noise") are superimposed on the cyclic fluctuations.

Cyclic Oscillations and Adjusting Schedules

On a different point, Marr makes reference to "Mechner's suggestion that the sensitivity to parameter shifts might *increase* with exposure to reinforcer presentations." What I had in mind was the well-documented "arousal" effect of a reinforcer presentation after a long stretch of non-reinforcement (e.g., Killeen, Hanson, & Osborne, 1978). I do not yet have reason to believe that there is a *continuous* relationship between an operant's susceptibility to reinforcement and the number of reinforcers recently received, but we are currently investigating the nature of this relationship at the University of North Texas.

Finally, I found Marr's discussion of adjusting schedules quite illuminating. It points out some complexities of the adjusting technique that I had not been aware of. And I strongly resonated with his point that acquisition is one of the frontiers of our field.

REFERENCES

Baer, D. M. (1992). Created, Revealed, and Imposed Operants. In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Killeen, P., Hanson, S. L., & Osborne, S. R. (1978). Arousal: its genesis and manifestation as response rate. *Psychological Review*, *85*, 571-581.

Marr, M. J. (1992). The Shape of Things to Come: A Commentary on Mechner's "The Revealed Operant". In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Mechner, F. (1959b). A notation system for the description of behavioral procedures. *Journal of the Experimental Analysis of Behavior*, *2*, 133-150.

Mechner, F. (1992). The Revealed Operant: A Way to Study the Characteristics of Individual Occurrences of Operant Responses. In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Mechner, F., Hyten, C., Field, D., & Madden, G. (1992). Using revealed operants to study the structure and properties of individual occurrences of operant responses. Paper submitted to *JEAB* in May, 1992.

Nevin, J. A. (1992). The discriminated Operant Revealed in Pigeon Multiple Schedule Performance: Effects of Reinforcements Rate and Extinction. In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Verhave, T. (1992). Shaping in Terms of Parameter Shifts. In Sigrid S. Glenn (Ed.). *Cambridge Center for Behavioral Studies Monograph Series, 6, Second Edition.* Cambridge, MA.: Cambridge Center for Behavioral Studies.

Weingarten, K., & Mechner, F. (1960). The contingency as an independent variable of social interaction. In T. Verhave (Ed.), *The experimental analysis of behavior, 1966*. New York: Appleton-Century-Crofts.