JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1988, 50, 343-354 NUMBER 2 (SEPTEMBER)

DIFFERENTIATING THE BEHAVIOR OF ORGANISMS

GREGORY GALBICKA

WALTER REED ARMY INSTITUTE OF RESEARCH

The procedure used to differentiate responding is probably well understood by most behavior analysts. Skinner offers a description of response-intensity differentiation in *The Behavior of Organisms:*

> Originally the responses occur with their intensities distributed (say, normally) about a low value. Reinforcement of members in the upper part of the range shifts the mean upward and with it the whole curve. Responses in the upper part of the new range may then be reinforced, and so on. (1938, p. 338)

The analytic community also clearly understands the ability of this process to produce novel and/or extreme forms of behavior. Again, Skinner describes the process:

> Both topographical and quantitative differentiation follow the rule of original operant conditioning that responses of the required form must be available prior to reinforcement if differentiation or conditioning is to take place. Extreme forms or values are obtainable only through successive approximations. (1938, p. 338)

Skinner recognized that operant conditioning and response differentiation represented two facets of the same phenomenon. He noted that all operant conditioning involves some response differentiation. In the latter the experimenter explicitly specifies the characteristics a response must have in order to produce consequent event(s). This is done implicitly in the former by transducing responses only across a restricted range. Such distinctions are not functional to the subject, however. It makes no difference whether responses go unreinforced because they lie outside the criterional range or outside the range of the transduction device. As long as these unreinforced responses are consistently related, the result will be the same. The extinction of such responding, and the reinforcement of responses within the specified (criterional or transductional) range, selects responses more closely resembling those correlated with reinforcement. As such, response differentiation is an integral part of all operant conditioning. Response differentiation *is* operant conditioning, and vice versa.

Finally, for Skinner response differentiation was prima facie evidence for the necessity of a new type of learning.

> In a respondent the intensity of the response is determined by the intensity of the stimulus, and no differentiation of the response is possible. The intensity of the response to a constant stimulus is a direct measure of the strength of the reflex. . . . [However, t]he intensity of the response in an operant is significant only in relation to the differentiative history of the organism. . . . (1938, p. 338)

The relation between stimulus intensity and response magnitude was not necessarily a direct one for operants, as the law of stimulus intensity/magnitude required for respondents. Instead, the intensity of a response was a function of a correlation between response intensities and reinforcement.

The above has been offered as a preamble to a paradox. Response differentiation played a crucial role in the conception of the operant. It is extant to some degree in every operant study. And it probably played an important part in our personal histories with respect to the experimental analysis of behavior by providing a vivid demonstration in a laboratory class somewhere (long ago?) of just how powerful are the variables controlling behavior.

The degree to which my verbal behavior about differentiation has been shaped by my interactions with John Platt is, I hope, evident in the behavior itself. If not, I readily surrender all credit forthwith. In that this was not a collaborative effort, however, I asymmetrically reserve full responsibility for the views contained herein. The views of the author do not purport to reflect the position of the Department of the Army or the Department of Defense (para 4-3, AR 360-5).

Yet, although demonstrations abound, there is very little in the way of a quantitative analysis of response differentiation. The major impediment to such an endeavor appears to be the procedure most typically used to demonstrate the phenomenon. To understand why the procedure poses problems, it is necessary first to consider Skinner's analysis of the processes involved in differentiation.

RESPONSE DIFFERENTIATION: FUNDAMENTAL PROCESSES

Skinner's analysis of response differentiation in *The Behavior of Organisms* invoked three processes: reinforcement, extinction, and induction (generalization). Reinforcement increased the likelihood that a response, with its associated constellation of characteristics, would be repeated. Through induction, similar responses occurred even though never directly reinforced. Further induction from these widened the range of response values emitted. At some point, induction generated a response value outside the reinforced class. The effect of extinction was to decrease the rate of responses with these characteristics, and also to decrease (through induction) the rate of similar responses. The key to differentiation was that the direct effect of reinforcement or extinction was always considered greater than the indirect effect of induction. Hence, the two classes drew apart in frequency, as criterional responses continued directly to be reinforced and noncriterional ones continued to be extinguished. Induction, however, ensured that response values continued to vary around those actually reinforced.

This proposed mechanism suggests at first glance that reinforcement should increase the variability of responses observed (through induction) and that extinction should decrease it. The latter should be the case given the reduced strength of induced responses relative to directly reinforced ones, suggesting that induced responses should first be reduced to zero (i.e., partial extinction could more quickly eliminate some induced classes than any directly reinforced class).

Skinner countered the first notion by emphasizing the differential between direct and indirect strengthening. Response frequency within the criterional range would have a "self-strengthening" effect, in that the direct effect of reinforcement would always be greater than any inductive effect on nonreinforced members. Hence, "[t]he first form reinforced has an initial advantage and may persist as a 'fixation'" (Skinner, 1938, p. 309). In this way, repeated occurrences of a modal response value disproportionately strengthened responses with the same characteristics, encouraging concentration of responses around that value.

Skinner also noted that the latitude of the reinforced class might be decreased through concurrent punishment. That is, certain responses within the reinforced class "may supply [punishing] stimulation tending to reduce the net reinforcing effect. 'Difficulty' or 'awkwardness' may be expressed in terms of the [punishing] stimulation automatically produced by a response" (Skinner, 1938, p. 309). This "law of least effort" (Keller & Schoenfeld, 1950), in concert with the above differential in the strengthening capabilities of reinforcement and induction, tends to decrease in most situations the variability in response values when responses are reinforced across a wide range of values.

Skinner noted that with extended extinction, the variation in response values did indeed decrease, usually leading to a class of responses having minimal values (i.e., with the responses differentiated in *The Behavior of Organisms,* weak intensities or short durations). More important (with respect to differentiation), extinction initially increased the variability in responding. This effect is important because "even when there is little variation [in responding], differential reinforcement necessarily involves extinction, and some strong responses are made available" (Skinner, 1938, p. 314). This extinction-induced variability, then, increases the probability that restricting reinforcement to a subclass of response values induces a response with the necessary characteristics. Because extinction initially increases the range of values beyond that normally observed, the probability of a response with the appropriate characteristic(s) will be increased. It must be remembered that this is a local effect, in that responding will quickly return to its undifferentiated level, and may cease altogether, if reinforcement frequency remains reduced too long. Skinner expressly noted that "the tendency for the force to decline in extinction sets the limiting rate at which the critical value may be advanced" (Skinner, 1938, p. 318).

Response differentiation, then, involves balancing two opposing processes. On the one hand, reinforcement differentially assigned to a particular subclass of the response transduced (criterional responses) increases the frequency of that subclass and of the class in general. In opposition, extinction of a subclass of noncriterional responses decreases the frequency of the subclass as well as the frequency of the response itself. In order for differentiation to occur, reinforcement must be allocated to a sufficiently distinct subset of the response distribution. However, in creating this subset through differential reinforcement, care must be taken that the response class itself (and hence the to-be-differentiated subset) not be extinguished. The adaptive value of extinction-induced variability involves the laxity with which the criterional range may be set. Although different subjects respond only across a restricted (and possibly idiosyncratic) range of any dimension prior to differentiation, imposing a criterion and the concomitant extinction of all responses not meeting the criterion will increase the variability of responding. This, in turn, increases the probability that some responses will satisfy the reinforcement criterion. Put in the converse, extinction-induced variability increases the feasible discrepancy between current response values and those to be required under differential reinforcement, allowing behavior to be differentiated more rapidly and decreasing the likelihood that the response class will extinguish altogether prior to a response with the criterional property or properties.

Nothing has transpired in the 50 years since publication of *The Behavior of Organisms* to repudiate any of the above analysis. It is clear that, in the short term, extinction and/or intermittent reinforcement initially produces increases in response variability, along all dimensions studied to date, including force (Notterman & Mintz, 1965), duration (Millenson & Hurwitz, 1961), lever displacement (Herrick, 1963, 1965), interresponse times (IRTs) (Ferster & Skinner, 1957), latencies (Stebbins & Lanson, 1962), and response location (Antonitus, 1951; Eckerman & Lanson, 1969). It is also clear that reinforcing specific subclasses of responses usually increases the frequency of that subclass. This has been demonstrated for spatial location (Davis & Platt, 1983; Eckerman, Heinz, Stern, & Kowlowitz, 1980; Galbicka & Platt, in press; Scott & Platt, 1985), lever press durations (Kuch, 1974; Platt, Kuch, & Bitgood, 1973), latencies (Catania, 1970), numbers (Mechner, 1958; Webster, 1976), displacements (Herrick, 1964), IRTs (Richardson & Loughead, 1974), key peck IRTs (e.g., Alleman & Platt, 1973; Galbicka & Platt, 1986; Kuch & Platt, 1976), and vocalization durations (Lane, 1964; Lane, Kopp, Sheppard, Anderson, & Carlson, 1967). Punishment has also been used to differentiate lever press IRTs in monkeys (Galbicka & Branch, 1981) and rats (Sizemore & Maxwell, 1985) as well as differentiating pauses of the latter (Sizemore & Maxwell, 1983). Page and Neuringer (1985) and Morris (1987) have even differentially reinforced *variation* in response patterns on two keys, and Blough (1966) variation in IRTs on one, and obtained increased frequencies of variable sequences or IRTs, respectively. This list of response dimensions that have been differentiated is a partial one, and includes only studies in which more than one criterion was examined.

It should perhaps again be emphasized that all operant conditioning demonstrates response differentiation, in that response rate change requires a concomitant inverse change in the rate of all "other" behavior. (The methodological problems associated with measuring other behavior lead Herrnstein [1970] to propose looking at relative rates among multiple, explicitly specified responses, more clearly a differentiation.) Response differentiation is so pervasive a phenomenon that failures to demonstrate differentiation will likely cause greater controversy than successful ones (see the current debate over differentiation of least-frequent response patterns: Morris, 1987; Page & Neuringer, 1985; Schwartz, 1982). Any aspect of operant behavior that can be correlated differentially with reinforcement must by definition change in frequency (i.e., become differentiated). This is nothing more than the definition of the operant.

The success of these demonstrations and the soundness of Skinner's 1938 formulation have produced little along the lines of a quantitative formulation of response differentiation.

Quantitative modeling, characteristic of operant conditioning procedures in which criterional ranges are as wide as possible, has not been extended to situations in which particular subclasses of transduced responses are differentiated. It is not even clear how "differential" reinforcement must be for a differentiation to develop. Skinner noted, for example, that not every criterional response need be reinforced for differentiation to proceed. "When responses above a given value are *periodically* reinforced, many responses above that value also go unreinforced. Nevertheless, the mean intensity of the response quickly rises and is maintained at about the same value as when all responses above the value are reinforced" (1938, p. 325). This is true even with mean interreinforcement intervals of approximately 5 min. Apparently, frequency of reinforcement for criterional responses can be very low and still differentiate responding.

What proportion of responses must be reinforced, or at what rate must reinforcement be delivered, in order for differentiation to develop? What effect does reinforcing noncriterional responses have on differentiation? What interactions, if any, exist with overall reinforcement rate? With the dimension? With the subject species? The answers to these questions have been long in coming, and many have not even been posed experimentally.

The absence of such quantification almost certainly does not reflect disinterest in response differentiation as a process, given the primary importance Skinner ascribed to it in the development of the operant. Indeed, demonstrations of response differentiation run the gamut of response dimensions, from differentiation of a covert muscle-fiber contraction (e.g., Hefferline & Keenan, 1963) to patterns of spatially defined operants (e.g., Page & Neuringer, 1985). The effort involved in developing transducers for many dimensions (e.g., Hefferline & Keenan, 1963; Herrick & Karnow, 1962; Notterman & Mintz, 1965) attests to a strong interest in the area.

The most likely reason that response differentiation has not developed the same degree of quantitative rigor is that the procedure Skinner outlined in 1938 does not control the variables considered fundamental in differentiating responding, namely the degree of differential reinforcement. In shaping a response, it is necessary to wait until a criterional response occurs and then follow it with the appropriate consequence. In that the experimenter cannot specify when criterional responses will occur, he or she also cannot specify what proportion of responses will be reinforced.

One solution to this conundrum is to change that upon which reinforcement is contingent, from responses with particular characteristics, typically outside the experimenter's control, to responses in the presence of particular stimuli, which can be controlled. That is, the stimuli in the presence of which a particular response occurs are a differentiable characteristic of responding that can be controlled by the experimenter. Reinforcement in the presence of some but not other stimuli leads to differentiation in the rate of the subclasses of responding defined by the presence of different stimuli. Establishing control by stimulus dimensions through differential reinforcement has such a long history in operant conditioning that Skinner saw fit to establish separate vocabularies. "To avoid confusion I shall speak of sensory discrimination using the terms already presented. The discrimination of the form of a response, however, will from this point on be referred to as Differentiation" (1938, p. 309).

It is probably because stimulus characteristics can more easily be controlled that they have predominated in the analysis of differentiation and discrimination. The questions above about response differentiation remain, however, and answers with appeals to the stimulus-control literature will be valid only to the extent that the domains are indeed isomorphic. This requires response differentiation be analyzed with precision equaling that found in the area of stimulus control. And that leads back to the initial problem; namely, that the procedure Skinner used to differentiate responding cannot be used to quantitatively analyze it.

RESPONSE DIFFERENTIATION: THE FIXED-CRITERION PROCEDURE

The procedure Skinner outlined is very easy to program. A portion of the range of transducible response values defined by some fixed value(s) is designated as criterion, then consequent events are delivered following responses in the criterional (and less often also in the noncriterional) range. But often what is easy to program is a nightmare to analyze. The fixed-criterion procedure's interaction with behavior is precisely that.

The problems are many, but they all involve reinforcement frequency in some way. Differentiation, as has been noted, involves a balance between reinforcement and extinction. As such, the frequency of these mutually exclusive outcomes should be controllable. If reinforcement frequency is too low, the differentiation may be lost through extinction. The behavior modifiers' maxim, "take small steps," is an actuarial attempt to minimize this outcome. But taking small steps means that more of the distribution of responding should be considered criterional. Thus, reinforcement will be associated with a greater range of values along the differentiated dimension, and this ultimately must lead to a decrease in the differentiation observed (i.e., the limit to this policy is to consider every response criterional, which constitutes baseline). Hence, differentiation under a fixed-criterion procedure might be expected when some proportion of responses large enough to maintain responding but small enough to be exceptional are considered criterional.

What are the parameters of this range? The fixed-criterion procedure cannot provide an answer, because the frequency of reinforcement is dependent on the subject's behavior. By defining as criterional those responses that exceed a fixed value, the prime determinant of when reinforcement occurs is when a criterional response occurs. Hence, both sides of the equation relating behavior to the environment take on the status of dependent variables under fixed-criterion procedures. (The same is true under all operant reinforcement procedures, of course, because they all differentiate response rates across some range. But the result is less dramatic because the range of response values considered criterional is so large that the rate of "criterional" responses seldom has an appreciable affect on reinforcement frequency. Under differentiation procedures, the rate of criterional responding must be reduced, and the associated variance increased, because the criterional responses are a subset of transduced responses, and hence must represent a smaller sample of behavior. The magnitude of this effect will be correlated with the degree of differentiation, because the proportion of responses occupying the criterional range changes with changes in the criteria.)

Further, changes in the frequency of criterional responses, both within and across differentiations, will be confounded with changes in the frequency of reinforcement. Hence, a quantitative analysis of differentiation will forever be subject to analysis in terms of more molar relations between the transduced response and consequent events. In that some measure of reinforcement frequency (be it rate or probability, conditional or otherwise) has dominated the role of independent variable in the analysis of operant behavior, it is clear that analysis of a phenomenon that does not control this variable will be fraught with difficulty and always open to reinterpretation.

A vivid example of this problem is provided by the continuing controversy over the role IRT differentiation plays in determining schedule-controlled responding. That IRTs can be differentiated is argued by few. Since Anger's (1956) classic paper, however, numerous procedures have been developed to demonstrate conclusively that IRT differentiation can or cannot account for schedule-controlled performance (see Ferster & Skinner, 1957; Galbicka & Platt, 1984, 1986; McDowell & Wixted, 1986; Platt, 1979; Richardson, 1973). It is perhaps not too surprising that the two most recent citations in this list come down squarely on either side of the question. McDowell and Wixted argue that "[i]nterresponse-time theories of schedule performance . . . are not supported by the results of the present experiment" (1986, p. 326). In contrast, Galbicka and Platt (1986) note that response rate varied by an order of magnitude with changes in IRT-reinforcement contingency, even though reinforcement rate was changed little, and cautioned that "it seems unreasonable to preclude IRT reinforcement from the analysis of responding under [random-interval] and constant-probability [variable-interval] schedules" (1986, pp. 377–378). A similar debate over the contribution relative reinforcement rate and reinforcement for switching make in determining choice under concurrent schedules has appeared recently in this journal (cf. Silberberg & Ziriax, 1985; Vaughan, 1981, 1987).

Finally, fixed criteria are problematic in comparative analyses of differentiation across species, responses, and dimensions, because of differences in the normal range of variance characterizing these different populations. Differentiation involves the subject's genetic endowment and past history of reinforcement interacting with the current environmental circumstances. Each of these may affect the degree of differentiation obtained, through a variety of avenues, by providing different absolute values of baseline response values, different degrees of dispersion in response distributions, and differential sensitivities to reinforcement parameters. A fixed-criterion procedure must of necessity contact these various response–subject combinations in different manners, rendering the comparative analysis subject to confound along a number of dimensions (e.g., overall frequency of reinforcement, proportion of responses in the criterional subclass, etc.). Attempting a comparison without somehow first equating the classes is tantamount to comparing apples and oranges, and is probably just as fruitless an exercise. Platt (1984) provides a detailed discussion of these and other problems associated with comparing temporal differentiations.

RESPONSE DIFFERENTIATION: PERCENTILE PROCEDURES

The fixed-criterion procedure is sufficient for illustrative purposes, but for long-term, parametric examination of the variables controlling response differentiation, too many possibly important variables are determined by the subject under this procedure. The percentile schedule, developed and extended by John Platt and his colleagues (see Alleman & Platt, 1973; Davis & Platt, 1983; Galbicka & Platt, 1986; Kuch & Platt, 1976), provides a method of differentiating responding immune to the above problems. Percentile schedules gain direct control over one variable, the frequency of criterional responses. Once established, controlling criterional response frequency allows control of reinforcement frequency for both criterional and noncriterional responses.

Now it may seem that controlling the frequency of criterional responses as an independent variable and response differentiation are antithetic. In fact, this is not the case, because control is gained only relative to the distribution of responses, not absolutely. That is, criteria are set such that a fixed portion of the distribution defines the criterional range. This proportion remains fixed independent of the absolute values of responses comprising the distribution. As performance changes, percentile schedules continually update the absolute response value that must be exceeded to be considered criterional, such that a constant percentile of the current distribution is always estimated. Defining a constant proportion of the response distribution as criterional establishes criterional-response frequency, and any resulting reinforcement frequency, as *independent* variables. In so doing, they allow a quantitative analysis of response differentiation unencumbered by the problems associated with fixed-criterion procedures.

When first encountered, the machinations of percentile schedules can seem horribly complex. This is probably due in large part to the contrasting simplicity of the fixed-criterion procedure and to the fact that it helps to have a computer to program them. The process involved, however, is very simple. Percentile schedules estimate a specified percentile of a response distribution. That is, they attempt to define the point in a response distribution below which $x\%$ of the responses lies (and hence above which the complimentary percentage lies). Much the same thing is done when a pigeon's key peck is shaped, or when a child's verbal behavior is differentially reinforced, or any other response is shaped. Criteria usually are set such that criterional responses occur with the same frequency as behavior changes. When we get too good at what we do (i.e., when most of our responses meet criterion), "expectations" are raised with respect to our behavior such that criterional responses occur less frequently. When we suffer serious setbacks, behavior is judged less harshly, increasing the probability that a response will meet criterion. We continually change the absolute criterion for reinforcement such that criterional response frequency remains roughly constant.

Now, it should be clear that this is not going to be one unchanging absolute value, as in fixed-criterion procedures, but will depend instead on the distribution. Given that distributions will vary across subjects, responses, and even with a single subject–response combination at different times during a differentiation, it would be helpful if some way could be found to simplify all these distributions into the same form. That is easier than it seems, because a continuous distribution implies an ordinal dimension, which implies that values can be ranked and a "distribution" of rank orders created. The ranks will be distributed rectangularly, with each rank being represented once. Hence, given that we dispose (for the moment) of absolute values and deal only with ranks, all distributions comprised of the same number of observations will be identical.

But how does this help to estimate specific percentiles? Suppose we have one observation in our distribution. This observation should estimate the 50th percentile of responses sampled afterwards, provided sampling is (a) from the same population of values and (b) random and independent. This is just another way of saying the next value sampled has an equal chance of being ranked above or below the observation in our distribution. If there were two observations in the distribution, given the same sampling constraints, the next item sampled would have an equal probability of ranking between the two observations, below the lesser or above the greater, one third of the time falling into each interval. The generalized form given a distribution of m observations implies subsequent observations will fall into the $m + 1$ intervals defined by these observations with equal probability, or with a probability of $1/(m + 1)$ because the probabilities of falling into each of the $m + 1$ intervals must sum to 1.0. This is true independent of the absolute values comprising the distribution (Smith, 1953), because the intervals defined are between consecutive ranks, not specific absolute values.

The rank most closely estimating a particular percentile can, from this information, be obtained easily. In that the goal is controlling criterional-response probability, however, it is more convenient to think in terms of probabilities. The probability the next observation will fall into one of k intervals is the product of the probability it falls into each and the number of intervals involved, or $k/(m + 1)$. Because there are k intervals below the kth rank, this also represents the probability of a response less than this rank. The probability of one exceeding this rank is the complement, or

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

Rearranging and solving for k allows a priori specification of any probability w of a criterional response,

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

Equation 2 specifies, given a distribution of m observations (termed the *control memory*), the number of ranks (k) the current response must exceed to observe a criterional response with the probability specified (w).

Hence, criterional-response probability is transformed from a dependent to an independent variable, and with it probabilities of reinforcement for criterional and noncriterional responses as well. This is accomplished with a minimum of assumptions; current values must be sampled from the population comprising the control memory, and sampling must be random and independent (i.e., there should be no significant sequential dependencies in sampled values).

Obviously, a successful differentiation implies that values obtained after the differentiation will differ from baseline. To ensure that the distribution of values characterizes the current population, percentile schedules continually update the control memory with each response to include only the most recent m responses. In this manner, the changing distribution of response values always most closely approximates the population from which the current response is selected. This also means that, although the criterion may be a particular rank in the distribution, the absolute value of that rank will change as the absolute values comprising the control memory change.

The assumption concerning sequential dependencies is no doubt violated; sequential dependencies have been observed under schedules of reinforcement whether they were (e.g., Shimp, 1973) or were not (Weiss, Laties, Siegel, & Goldstein, 1966) reinforced differentially. The effects of such dependencies will be most acute when control memories are small. Hence, in practice control memories usually consist of the most recent 25 or 50 responses. In the experiments of Alleman and Platt (1973) and other unpublished experiments by Platt, control memories between 2 and 400 responses have been examined. Except at the two extremes, memory size had little effect. Control over criterional-response probability is lost for different reasons at the two extremes; short memories are influenced excessively by sequential dependencies, whereas extremely long memories include values no longer characterizing current performance, but determining the absolute value of the criterion nonetheless.

The above "basic," "simple," or "probability" percentile (because it served as the foundation for extensions, was [relatively] simple and controlled the probability of criterional responses, respectively) is only one of a class of percentile procedures. The rate of criterional responses can also be controlled under an "interval percentile" by adjusting the criterion as a function of the rate of responding. As more responses are emitted in a given period of time, a smaller proportion are considered criterional (i.e., the rank that must be exceeded is increased), leading to a relatively constant rate of criterional responses (see Kuch & Platt, 1976).

In earlier treatments, percentile schedules were seen as controlling reinforcement probability. This, however, is an adjunct of controlling criterional-response probability. Because criterional responses were always reinforced in early studies, and noncriterional ones never reinforced, criterional responses and reinforcement were isomorphic. Subsequently, the distinction has been exploited in a series of experiments examining how changes in conditional reinforcement probabilities for criterional and noncriterional responses, and the changes in the degree of contingency between criterional responses and reinforcement implied by these manipulations, affect response differentiation in discrete-trials (Galbicka & Platt, in press; Scott & Platt, 1985) and free-operant (Galbicka & Platt, 1986) procedures. In all these, criterional-response and reinforcement probabilities or rates, as appropriate, were specified a priori.

Finally, probability- and interval-percentile schedules are "open" in the sense that there is no endpoint specified in the differentiation, only a direction. In procedural contrast are the "targeted" percentiles (see Davis & Platt, 1983). These are not really a new class of schedules, but rather the simultaneous application of two percentiles to shape responding in opposing directions on either side of an absolute target. One operates on responses above the target value and differentiates towards smaller values, whereas the other operates on responses below the target and differentiates towards higher values. These two percentiles drive responding tighter around the absolute value specified as target.

Thus, the percentile schedules provide a number of ways in which to differentiate responding while controlling variables not controllable under fixed-criterion procedures. Although admittedly more complex than fixed-criterion procedures, their complexity pays dividends in terms of experimental control. Under these procedures, the independent effects of defining as criterional different proportions of responses and the conditional probability or rate of reinforcement for criterional as well as noncriterional responses can each be examined, allowing quantitative models of the interactions to be posed and evaluated (Galbicka & Platt, in press; Scott & Platt, 1985).

RESPONSE DIFFERENTIATION: FUTURE DIRECTIONS

I have argued that the procedure Skinner (1938) used to demonstrate shaping insufficiently controlled variables important in quantifying differentiation. Percentile schedules provide the control needed to conduct such analysis, yet they remain an infrequently used analytic tool a decade since they were first described. This suggests that there remain questions and concerns about the operation of these procedures. Because continued reliance on the procedurally simpler fixed-criterion procedure will deter the analysis of response differentiation for the reasons cited above, it is important to consider possible sources of concern operating to impede adoption of percentile schedules by a wider range of behavior analysts.

The first obstacle to the adoption of percentile schedules is a methodological one; when percentile schedules were first described, the laboratory computers that make percentile schedules feasible were an anomaly. Fortunately, this is no longer the case. Behavior analysts have warmed to the laboratory computer, some admittedly slower than others. The computer industry has helped spur this change by providing more reliable, faster, cheaper, and friendlier hardware and software. As a result, even today's personal computers can handle the calculations necessary to program percentile schedules.

The second problem is a conceptual one. Many people express concern that the response criterion continually shifts under percentile schedules. This statement has the same effect as the statement "I tried reinforcement and it didn't work"; in both cases the problem has been misidentified. The former statement is as false a characterization of percentile schedules as the latter is of reinforcement. The criterion is consistent under percentile schedules; it always represents some proportion of the current distribution. What changes are the absolute values comprising the distribution (and hence the absolute value of the rank considered criterional), and it is to this that some people object. Yet even this does not seem adequately to represent the problem. Under alternative or interlocking schedules, for example, the criterion for reinforcement is not specifiable without reference to some or all of the behavior generated thereby, yet these procedures do not raise serious concerns.

The main conceptual distinction between percentile schedules and other reinforcement schedules, the fixed-criterion schedule in particular, is the referent used to determine whether a bit of behavior qualifies as criterional. With standard reinforcement schedules, the referent is absolute; all presses longer than *x* s, all latencies shorter than *y* s, all pecks on the ends of the response strip, and so on. With percentile schedules, the referent is a fixed proportion of a response distribution. The objection to this shift in perspective is either that such scaling involves the statistical control of response probability, anathema to the experimental analysis of behavior, or that such measurement is indirect and hence not characteristic of a natural science.

In responding to these concerns, it is essential that the consequences of these biases, originally responsible for their selection, clearly be understood. With respect to statistical analysis, Skinner notes:

> The statistical approach compensates for its lack of rigor at the stage of measurement by having recourse to statistical analysis, which the non-statistical approach in general avoids. ... The concepts established in the first case become a part of scientific knowledge only by virtue of statistical procedures, and their reference to the behavior of an individual is indirect. In the second case there is a simpler relation between a concept and its referent and a more immediate bearing upon the individual. It may be that the differences between the two approaches are transitory and that eventually a combination of the two will give us our best methods, but at the present time they are characterized by different and almost incompatible conceptions of a science of behavior. (1938, p. 443)

We have, in percentile schedules, an example of such a combination. Using a sample of the subject's own behavior to predict characteristics of future behavior is, after all, little more than the goal of experimental analysis. What is not being done is exactly what Skinner cautioned against, indirectly establishing concepts through reference to the behavior of a group. This is far more likely with fixed-criterion procedures, in which criteria are often established without recourse to an individual's behavior, and hence may depend on what works with most subjects, or what the next multiple of 10 is, or what is most easily programmable.

With respect to direct versus indirect measurement, the important consequence of this bias is the replicability of experiments. Without scales that are absolute, values vary with each measurement operation, and hence neither procedures nor results are repeatable. But how should experiments with respect to behavior be replicated, when the baselines established in different subjects and/or at different times vary before the experiment is even begun? It is not particularly strange in such situations to normalize each subject's behavior before proceeding. This, for example, is done in behavioral pharmacology, where drug effects are presented relative to the baseline rate of responding. It is no less questionable to normalize the workings of a procedure relative to behavior. To allow historical accident to determine the paradigms, and thus define the science, seems particularly odd for a science that has always judged itself in terms of the prediction and control its methodology afforded. To the extent data generated under percentile schedules are as orderly as those obtained under fixed-criterion procedures, the question will provide its own answer.

This leads to the final problem, an empirical one. The law of effect as demonstrated in percentile schedules is often at odds with the law of effect characterized by current quantitative formulations of behavior. The latter almost all specify relations between molar response and reinforcement rates. Percentile schedules demonstrate the effectiveness of reinforcement in differentiating responding when gross reinforcement rate changes are absent, and even when overall reinforcement rate changes in a manner counter to molar models (e.g., shaping longer IRTs with a constant probability of reinforcement implies a *decrease* in overall reinforcement rate). None of the molar quantitative formulations have attempted to incorporate these data; most have simply ignored the results.

It is perhaps unfortunate that a procedure has been correlated with a theoretical stance emphasizing local changes in reinforcement frequency at a time when the prediction and control of response rate have been dominated by quantitative models emphasizing molar relations. It is unfortunate because the procedure has no inherent bias. It can provide information on local and molar relations independently (Galbicka & Platt, 1986), something other scheduling arrangements cannot so easily do. It is not surprising that the outcome of such independent manipulation is that both classes of variables affect response rate. What is surprising is that, once isolated, the independent effects of molar variables can be greater than effects observed under the procedures typically used to study them. That is, changing reinforcement rate has an increasingly greater effect on response rate as local reinforcement contingencies for IRT duration are eliminated. In that variable- or random-interval schedules both are characterized by relatively strong local reinforcement contingencies, they tend to underestimate the effect reinforcement rate has on response rate. Yet, precisely these arrangements have provided the basis for the molar models of responding. It seems that percentile schedules, because they allow local and molar effects of reinforcement to be isolated, are ideally suited for the analysis of both sets of variables. So pervasive a discrepancy as undermatching may even be experimentally manipulable by examining interactions of local and molar variables in concurrent situations (Galbicka & Platt, 1986).

The development of a science mirrors the operation of percentile schedules. What defines an appropriate paradigm at any point in the evolution of a scientific discipline is not a fixed set of procedures, but rather how effectively extraneous variables known to affect the subject matter are controlled without constraining the range of independent variable values possible. The number of extraneous variables that must be controlled increases as more is learned about the subject matter, but the criterion, judged relative to the distribution of possible confounds, is always to control them all. As such, methodologies must be refined continuously to increase our precision of control as more of the variables known to affect the subject matter are discovered. We know that response differentiation involves reinforcement and extinction of criterional and noncriterional responses. Percentile schedules provide a means of independently assessing the effects of each of these, across response classes, dimensions, and subjects that may or may not differ under baseline. The analysis Skinner outlined 50 years ago can finally be accomplished. Now if only the increased prediction and control this new methodology affords will differentiate our responding as scientists to include percentile and related paradigms.

REFERENCES

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

Anger, D. (1956). The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, **52**, 145–161.

Antonitus, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, **42**, 273–281.

Blough, D. S. (1966). The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, **9**, 581–591.

Catania, A. C. (1970). Reinforcement schedules and psychophysical judgments: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1–42). New York: Appleton-Century-Crofts.

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

Eckerman, D. A., Hienz, R. D., Stern, S., & Kowlowitz, V. (1980). Shaping the location of a pigeon's peck: Effect of rate and size of shaping steps. *Journal of the Experimental Analysis of Behavior,* **33,** 299–310.

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

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

Galbicka, G., & Branch, M. N. (1981). Selective punishment of interresponse times. *Journal of the Experimental Analysis of Behavior,* **35,** 311–322.

Galbicka, G., & Platt, J. R. (1984). Interresponse-time punishment: A basis for shock-maintained behavior. *Journal of the Experimental Analysis of Behavior,* **41,** 291–308.

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

Galbicka, G., & Platt, J. R. (in press). Response–reinforcer contingency and spatially defined operants: Testing an invariance property of phi. *Journal of the Experimental Analysis of Behavior.*

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. (1963). Lever displacement during continuous reinforcement and during a discrimination. *Journal of Comparative and Physiological Psychology,* **56,** 700–707.

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

Herrick, R. M. (1965). Lever displacement under a fixed-ratio schedule and subsequent extinction. *Journal of Comparative and Physiological Psychology,* **59,** 263–270.

Herrick, R. M., & Karnow, P. (1962). A displacement-sensing, constant-torque response lever. *Journal of the Experimental Analysis of Behavior,* **5,** 461–462.

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

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

Kuch, 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.

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

Lane, H. (1964). Differential reinforcement of vocal duration. *Journal of the Experimental Analysis of Behavior,* **7,** 107–115.

Lane, H., Kopp, J., Sheppard, W., Anderson, T., & Carlson, D. (1967). Acquisition, maintenance, and retention in the differential reinforcement of vocal duration. *Journal of Experimental Psychology,* **74** (2, Whole No. 635).

McDowell, J. J, & Wixted, J. T. (1986). Variable-ratio schedules as variable-interval schedules with linear feedback loops. *Journal of the Experimental Analysis of Behavior,* **46,** 315–329.

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

Millenson, J. R., & Hurwitz, H. M. B. (1961). Some temporal and sequential properties of behavior during conditioning and extinction. *Journal of the Experimental Analysis of Behavior,* **4,** 97–106.

Morris, C. J. (1987). The operant conditioning of response variability: Free-operant versus discrete-response procedures. *Journal of the Experimental Analysis of Behavior,* **47,** 273–277.

Notterman, J. M., & Mintz, D. E. (1965). *Dynamics of response.* New York: Wiley.

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

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.

Platt, J. R. (1984). Motivational and response factors in temporal differentiation. In J. Gibbon & L. Allan (Eds.), *Timing and time perception* (pp. 200–210). New York: New York Academy of Sciences.

Platt, J. R., Kuch, D. O., & Bitgood, S. C. (1973). Rats' lever-press durations as psychophysical judgements of time. *Journal of the Experimental Analysis of Behavior,* **19,** 239–250.

Richardson, W. K. (1973). A test of the effectiveness of the differential-reinforcement-of-low-rate schedule. *Journal of the Experimental Analysis of Behavior,* **20,** 385–391.

Richardson, W. K., & Loughead, T. E. (1974). Behavior under large values of the differential-reinforcement-of-low-rate schedule. *Journal of the Experimental Analysis of Behavior,* **22,** 121–129.

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

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

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

Silberberg, A., & Ziriax, J. M. (1985). Molecular maximizing characterizes choice on Vaughan's (1981) procedure. *Journal of the Experimental Analysis of Behavior,* **43,** 83–96.

Sizemore, O. J., & Maxwell, F. R. (1983). Differential punishment of other behavior: Effects of pause-dependent shock on variable-interval performance. *Animal Learning & Behavior,* **11,** 266–274.

Sizemore, O. J., & Maxwell, F. R. (1985). Selective punishment of interresponse times: The roles of shock intensity and scheduling. *Journal of the Experimental Analysis of Behavior,* **44,** 355–366.

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

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

Stebbins, W. C., & Lanson, R. N. (1962). Response latency as a function of reinforcement schedule. *Journal of the Experimental Analysis of Behavior,* **5,** 299–304.

Vaughan, W., Jr. (1981). Melioration, matching, and maximizing. *Journal of the Experimental Analysis of Behavior,* **36,** 141–149.

Vaughan, W., Jr. (1987). Reply to Silberberg and Ziriax. *Journal of the Experimental Analysis of Behavior,* **48,** 333–340.

Webster, J. R. (1976). *An experimental analysis of response differentiation.* Unpublished doctoral dissertation, McMaster University, Hamilton, Ontario.

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