


RESPONSE BURSTS ON DRL SCHEDULES

By



DENNIS WEBB

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Master of Arts

McMaster University

MASTER OF ARTS (1981)
(Psychology)

McMASTER UNIVERSITY
Hamilton, Ontario

TITLE: Response Bursts on DRL Schedules

AUTHOR: Dennis Wayne Webb, B.A. (University of Cincinnati)

SUPERVISOR: Dr. J. R. Platt

NUMBER OF PAGES: V, 40

ABSTRACT

The differential reinforcement of low rates schedule typically yields from the subject a distribution of interresponse times (IRTs) that is bimodal. One mode occurs at or near the reinforced value and a second one at IRT values less than 2.0 sec. The IRTs which comprise this latter mode have been termed "response bursts" (Sidman, 1956). Various explanations for the mode at short values have been offered: Blough (1963, 1966), arguing from pigeon data only, suggested that these short IRTs are due to variation in the topography of the pigeon's keypeck response, the double-peck being a variation of the peck. Kramer & Rilling (1970) noted that if the definition of what actually constitutes a response is not clear to the subject, then quick multiple responses might result. Finally, Sidman (1956) has viewed these sequences of rapidly-made responses as evidence of increased response strength at certain times.

A comparison of IRT distributions from mice, rats, and pigeons, however, shows such similarities in the patterning of bursts that Blough's suggestion of topographic variation in the pigeon's key peck cannot be a satisfactory explanation for the occurrence of these short IRTs. To test Kramer and Rilling's suggestion that lack of clear response definition may produce multiple responses, mice were trained on a DRL 18 sec schedule and given response feedback in the form of an auditory stimulus. The results of a series of such experiments indicated that while a relatively long-duration auditory stimulus (2.0 sec) would almost eliminate bursting, it did this by acting as an S^Δ

rather than as response feedback per se. The fact that a short-duration (0.2 sec) auditory stimulus did not suppress bursting showed that bursting was not the result of lack of response definition. Further analysis of the patterning of response bursts lent substance to Sidman's contention that response bursts are a reflection of momentary increase in response strength. Specifically, bursts were found to occur most often as time to reinforcement availability decreased; bursts never occurred immediately after responses which were reinforced (i.e., when time to reinforcement availability was greatest); and the length of a burst run (number of responses) was greater the closer it occurred to the time of reinforcement availability. It was concluded that Sidman's interpretation of the DRL schedule as a temporal discrimination task is consistent with the response burst data from the experiments described here.

ACKNOWLEDGEMENT

The author wishes to acknowledge the assistance of Dr. John R. Platt of McMaster University in supplying the computer software for running the experiments described in this paper. Dr. Platt also provided important guidance and timely encouragement to the author.

INTRODUCTION

The differential reinforcement of low rate (DRL) was conceived by Ferster and Skinner (1957, p.459) as a schedule that would insure that the response rate of the organism at the time of reinforcement was below a certain value. Their object was to confirm the effects on subsequent responding of the rate of response at the moment of reinforcement. "Practical considerations" however, led them to sample, as their measure of response rate, the individual interresponse time (IRT). While it is questionable whether, under this procedure, we can meaningfully speak of reinforcing a rate of responding in anything but the technical sense of the phrase, the procedure has survived as a reinforcement schedule in its own right.

The sampling of the single most recent IRT for the purpose of establishing "rate" on DRL has encouraged the analysis of DRL performance in terms of emitted IRTs. Attention has been focused on the distribution of IRTs on DRL (Kramer & Rilling, 1970), and the form of the characteristic distribution has provided ample material for discussions of underlying mechanisms.

The typical IRT distribution generated by subjects on DRL schedules is bimodal, with one mode centered approximately at the minimum reinforced (criterion) IRT value and a second mode at values less than 2.0 sec. Although the mode at the criterion IRT will follow this criterion if the latter is shifted (Conrad et al., 1958; Kelleher et al., 1959), the second mode

remains under 2.0 sec. It is not surprising that the criterion IRT mode on DRL adjusts to the value of the criterion IRT employed. The criterion IRT does, after all, define the temporal aspects of the schedule contingencies. However, the occurrence of a great number of IRTs less than 2.0 sec is somewhat surprising. The rapidly-occurring responses which result in these short IRTs have been called "response bursts" by Sidman (1956). He defined as a burst "any sequence of two or more responses in which no consecutive responses are separated by more than 2 sec."

Various hypotheses about the nature of these response bursts have been offered: Blough (1963; 1966) suggested that these short IRTs be considered topographical variation associated with the response. Kramer and Rilling (1970) put forward the idea that, due to insufficient response feedback, the definition of the required response may be somewhat ambiguous to the subject, a situation which might encourage multiple attempts. Finally, Sidman (1956) accepts response bursts on DRL as a reflection of response strength at any given moment.

This paper will discuss these three interpretations of response bursts on DRL. It will also describe experiments designed to closely examine the patterning of response bursts on DRL schedules and to test the validity of the explanations offered for their occurrence.

Bursts as Topography

In operant conditioning, the activity of the organism which is to constitute a "response" is defined by the experimenter and applied through the mechanical apparatus employed. This definition of the response necessarily includes the specification of certain threshold values of the topographical features of the response: the subject must be sufficiently close to the manipulandum, the excursion of the bar or key must exceed a certain value, the force applied cannot fall below a particular minimum, the duration of the response must be long enough to be sensed by the control equipment, and so on. The threshold values for these topographical features are typically such that they confront the subject with little difficulty and minimal physical effort. Consequently, experimenters have in general not concerned themselves with fluctuations in these features of responding, but have been satisfied to restrict their data to "criterion responses," whatever their topographic makeup. This neglect of topographic variation finds justification in the assumption that since only topography which results in a criterion response can be reinforced, only those values of topographical features which meet this criterion can be directly strengthened. As a result, rate of criterion response has endured as the predominant datum of interest.

However, some studies have examined response characteristics other than rate and have found that values of the response dimension of interest have systematically changed with the schedule

used. For example, Margulies (1961) found that the mean and variance of duration of rats' lever presses were much higher in operant level and extinction sessions than during continuous reinforcement. Eckerman and Lawson (1969) report greater variability in response locus in pigeons on various intermittent reinforcement schedules than during continuous reinforcement. Herrick (1965a, b) has shown that rats' mean lever displacement systematically decreases with ordinal number of press in a fixed ratio run, but is constant on variable ratio schedules. In the above studies response dimensions other than rate proved sensitive to changing contingencies which did not specify any change in criterion for those dimensions. Any response dimension which shows systematic variation under different conditions of reinforcement and motivation is potentially of interest in explaining behavior. On the other hand, dimensions which are insensitive to experimental manipulation are of little interest and are best ignored; or, if they affect rate of response (or some other datum of interest) simply by adding to its magnitude or variability, they may be filtered out by mechanical or statistical means.

The argument has been made by Blough (1963, 1966) that responses terminating short IRTs by pigeons on temporal schedules should be regarded as constituting a response dimension which is insensitive to schedule control, and is therefore best relegated to the status of topographical variation. Blough's analysis of these responses showed them to be variable and insensitive to schedule contingencies, and consequently of no interest in a

behavioral analysis. He sees them as essentially different in character from responses which terminate longer IRTs, since these latter mark off IRTs which are sensitive to schedule control.

In 1963 Blough reported data from pigeons working on VI, FR, DRL, and extinction schedules, all of which produced clusters of IRTs of about 0.4 sec. These short IRTs were also present during transitions between different schedules and on generalization tests. A similar finding was described by Blough in 1966, again with pigeons as subjects. In this case the schedule was one which reinforced least frequently emitted IRTs, relative to an idealized exponential distribution. Although changes in the rate parameter of the reference distribution were followed by corresponding changes in the pigeons' response rates, this relationship held only for IRTs greater than about 0.8 sec. The frequency of IRTs shorter than this value was unaffected by the reinforcement criterion. In addition, variability between subjects and session-to-session variability within subjects were much greater for these very short IRTs than for those greater than about 0.8 sec.

Blough (1963) suggested that these response "bursts" are generally insensitive to schedule control and have a special character of their own:

"...certain of the responses that contribute to the "rate" on variable interval and DRL schedules are almost entirely controlled by prior responses; their probability does not vary with extinction, stimulus change, etc., except indirectly through changes in the probability of the responses on which they are dependent...It is possible that

the insensitive component in behavior on these schedules is benign, and applies only a multiplicative factor to the output of "sensitive responses." (Blough, 1963, p.245)

Blough (1966) concludes that short IRT responses are "perhaps best considered as part of the pigeon's response topography, the double peck being simply a variant of the peck." Changes in the frequency of these short IRTs (double pecks), then, become topographical drift.

The contention that pigeons' response bursts are not a sensitive feature of the behavior generated by temporal schedules raises questions about these same short IRTs produced by other species. Response bursts are a fairly reliable feature of responding on DRL schedules with a variety of species and responses (Sidman, 1956; Ferraro, Schonfeld, & Snapper, 1965; Staddon, 1965; Kramer & Rilling, 1969; Blough, 1966). If we are to ignore response bursts in pigeons on DRL, what about other species on this schedule?

Although explanations of bursting in pigeons may apply only to the pigeon as a subject, the general resemblance of pigeons' IRT distributions on DRL schedules to those of other species invites comparison of process. The following data and discussion will compare features of bursting on DRL schedules with pigeons, mice, and rats as subjects.

Figure 1 offers for comparison some representative data from pigeons, mice, and a rat. The pigeon and mouse data were obtained by the author; the rat data are reproduced from Sidman (1956):

The individual figures for the mice and pigeons resemble, in their general characteristics, those obtained from additional subjects of the same species run under the same conditions. Sidman (1956) assures that the figure for rat SD-20 is typical of those generated by his other rat subjects. (Here and elsewhere in this paper only the principal features of experiments performed by the author will be described. A detailed methodology will be found in the Appendix.) For each mouse, 10 consecutive days' data are combined (days 51-60, each session 45 min) to obtain a sufficiently large sample of bursts. The mice were trained to make a nose-poke response on a DRL 18 sec schedule for a milk reinforcer. They received no programmed feedback for their responses other than the delivery of the reinforcer when appropriate. Their IRT distributions were stable by day 50. (These mice were part of a larger experiment described below as Experiment I.) The pigeons were given 45 sessions of 26 min each on a DRL 8 sec schedule for a mixed grain reinforcer. Five consecutive days' data (days 41-45) are combined in Figure 1 for the pigeons. Sidman's rats bar-pressed on a DRL schedule which reinforced "responses spaced at least 20-21 sec apart." (This presumably means that IRTs of 20-21 sec were only sometimes reinforced; ones longer than 21 sec were always reinforced.) The reinforcer was water, and the figure combines data from five consecutive days.

For each subject in Figure 1 is plotted the probability that an emitted IRT will be followed by a burst. That is to say, each bar gives the probability that a response terminating an

IRT of a particular length will be followed by another response before 2.0 sec has elapsed (in the case of the rat and mice) or before 1.0 sec has elapsed (in the case of the pigeons). This measure is conveniently termed "conditional burst probability" because it conditionalizes the emission of a burst on the occurrence of various prior IRT lengths.

The first characteristic to be remarked upon, and one shared by all three species represented, is the contrast between burst probability following reinforced responses, and that following nonreinforced responses. In all cases, it is improbable that a response which is reinforced (i.e., one which terminates an IRT greater than 21 sec for the rat, 18 sec for the mice, and 8 sec for the pigeons) will be followed by a burst. Consequently, it can be seen that nearly all bursts follow nonreinforced responses. (Some bursts follow other bursts; this aspect of the data has been omitted here and will be treated shortly.) This similarity across species in burst probability after reinforced and nonreinforced responses is one bit of evidence against the notion that pigeons' short IRTs on DRL are of an essentially different character than those emitted by other species. It is the case that on DRL schedules pigeons, mice, and rats consistently emit only a single response if that response is reinforced and often emit multiple responses if reinforcement is not forthcoming for the first response.

Further examination of the mouse and rat data in Figure 1 reveals a consistent pattern of burst probability following nonreinforced responses: In general, the longer the previous

nonreinforced IRT, the greater is the probability of a burst occurring after it. This increase in burst probability with increasing length of previous IRT is fairly consistent in the rat and mouse data; it does not appear in the pigeon data in Figure 1, however. None of the pigeon subjects in this experiment showed any definite increase in burst probability with increasing non-reinforced IRT length.

From the data in Figure 1 it appears that the tendency to emit a burst on DRL is under the temporal control of the previous IRT length (at least in the case of rats and mice), and not merely under control of the previous response, as Blough suggests. That pigeons do not show the same tendency on DRL may indicate something essentially different about bursting in pigeons. On the other hand, it may merely reflect the pigeon's lack of sensitivity to the temporal aspects of the schedule. It is well documented that pigeons are unable to effectively adjust their IRT emission to DRL criteria which offer no difficulty to other species (Holz & Azrin, 1963; Holz et al., 1963; Reynolds, 1964a, b; Kramer & Rilling, 1969; Staddon, 1965). Since (a) efficient emission of reinforced IRTs, and (b) a positive correlation between nonreinforced IRT duration and burst probability, both require a sensitivity to the temporal demands of the schedule, the pigeon may very well lack this sensitivity.

One additional comparative aspect of bursting on DRL schedules is of interest here--the length of a run of bursts. So far, response bursts have been discussed as single events which sometimes follow other (longer) IRTs. But, in fact, when bursts

occur, they often occur in multiples. When a longer IRT is not reinforced, as Figure 1 shows, there is some probability that a burst will follow, but that burst is often followed by others, constituting a "burst run." Figure 2 shows, for the same mice and pigeons as in Figure 1 (and using the same data sample), the relative frequency of burst runs of various lengths. As can be seen, there are many instances in which only a single burst occurred (a burst run of length "one"). Longer runs, however, were not uncommon. In general, pigeons had fewer long burst runs than did mice. But it is the case that pigeons did emit many burst runs of length greater than one. These runs of multiple bursts and the shape of the run distribution they constitute provide another similarity between DRL bursts by mice and those by pigeons, eroding further the notion that short IRTs made by pigeons on DRL are essentially different than those emitted by other species on the same schedule. These runs also present a difficulty to the interpretation of the burst as a "double-peck" topography. If, as Blough (1963) suggests, response bursts in pigeons occur because of a double-peck variant of the single-peck response which is peculiar to the pigeon, then there should be no burst runs of length greater than one. Of course, strings of intermixed single and double pecks could form burst runs of any length. But, as Blough notes, these short IRTs by pigeons are of a fairly uniform duration, about 0.4 sec. To explain all these IRTs at about 0.4 sec, we would have to postulate that the IRT between two successive response units (whether the units themselves were single or double pecks)

be consistently about 0.4 sec, coincidentally the same length as the IRT within the double-peck unit. Thus, it appears more reasonable to regard these strings of responses as response runs in their own right and not as the result of special topographical considerations. Add to this the general similarity of burst runs in mice and pigeons, and it becomes plausible that a common process is involved.

In sum, a comparison of various features of response bursts on DRL schedules in mice, pigeons, and rats has shown general similarities. The only notable differences are: 1) the difference in the steepness of the mouse and pigeon burst run distributions, and 2) the lack of correlation between burst probability and length of previous nonreinforced IRT in the pigeon, a correlation which is present in both the mouse and rat data. But as mentioned, this lack of correlation may be a symptom of a general lack of temporal control by this schedule in the case of the pigeon. The general similarities among species here in features of bursting on DRL make doubtful an interpretation of response bursts in pigeons as different from those in other species. Further, the dependence of bursts upon nonreinforcement and length of previous IRT gives the response burst the status of a sensitive dependent variable.

Feedback For Responses

Kramer and Rilling (1970) note that while many DRL studies show a large number of response bursts (Sidman, 1956; Kramer, 1968; Kramer & Rilling, 1969), others show relatively few (Reynolds, 1964a, b; Carter & Bruno, 1968a; Kelleher et al., 1959). One aspect of the

study by Kelleher et al., however, may provide an explanation for this inconsistency. Those workers found that if their rats' lever presses were accompanied by a click from a mechanical relay, somewhat fewer bursts occurred. Rats trained with this auditory feedback for responses emitted fewer bursts than are often found on DRL schedules with rats as subjects, and the removal of this stimulus resulted in a small but consistent increase in the number of bursts emitted thereafter. Kramer and Rilling (1970) suggest that the removal of this feedback stimulus made criterion and noncriterion responses (those which did and did not operate the manipulandum, respectively) less unambiguous to the subject. Lack of clear feedback for a response may fail to inform the subject that a criterion response has actually been made, and so invite further attempts. Certainly, delivery of the reinforcer is a salient source of feedback for a response, and as shown in the previous section, bursts do not often occur after a reinforced response. (Of course, consumption of the reinforcer after reinforced responses will compete with any tendency toward further immediate responding.) Further support for the idea that feedback for responses reduces or eliminates bursting is the general finding in the DRL literature that little bursting occurs whenever some salient stimulus event accompanies a response. For example, Catania (1970), reinforcing latencies (ranging from 0.6 to 48.0 sec) in pigeons, obtained very few short latencies, presumably because each response initiated an intertrial interval. (When the length of this intertrial interval was shortened from 20 sec to 2.0 sec or less, some increase

in short latencies was seen, but Catania attributes this to the appreciable increase in reinforcement density.) Kramer and Rilling (1969) instituted a timeout following noncriterion IRTs on DRL and obtained very few short IRTs from their pigeons. They attributed the suppression of short IRTs to the punishing effect of the timeout and compared the effect to that obtained by using shock on the same schedule (Holz, Azrin, & Ulrich, 1963). But whereas Holz et al. found that the degree of suppression of short IRTs was in direct proportion to the intensity of the shock used, Kramer and Rilling found suppression independent of the duration of the timeout. This latter result makes it questionable whether the timeout acted as a punisher but leaves open the possibility that it controlled the emission of short IRTs by serving as response feedback.

In a situation where a criterion response is accompanied by no salient stimulus change (e.g., nonreinforced responses on DRL), the requirement for the response is not unambiguously defined for the subject. A sequence "noncriterion response--criterion response" may result in delivery of the reinforcer for the second response attempt; and if the subject does not discriminate that the first response attempt (i.e., the noncriterion response) was no response at all, then from the subject's perspective a double-response sequence (a short IRT) was reinforced. It is possible, then, that response bursts result from the inadvertent reinforcement of double responses, the first of which is not recorded by the control equipment, but is to the subject a genuine response.

To test this idea it was necessary to maximize and minimize

response feedback for different groups of subjects, with the expectation that the more salient the feedback, the less bursting would occur. To do this, a response was required which has as little inherent feedback as possible. The author had previously employed mice on various operant schedules using as the response a "nose-poke." The nose-poke is easy for the subject to make, is conveniently recorded as the interruption of a light beam falling on a photoresistor, and has very little natural feedback associated with it since the subject does not make physical contact with a manipulandum. Thus, the nose-poke satisfied the need for a response with minimum inherent feedback. To supply salient feedback to those subjects for whom feedback was to be maximized a Sonalert was used to produce a clearly audible tone. So the extremes of no feedback and very salient feedback were available to examine the role of feedback in response definition and to discover whether response bursts on DRL schedules are the result of the lack of such definition.

EXPERIMENT I

Sixteen naïve mice were liquid-deprived and learned to make the nose-poke response on a continuous reinforcement schedule, as described in the Appendix. The 16 mice were randomly divided into four groups of four mice each. A two-by-two factorial design was employed, with the two factors being (1) presence/absence of feedback, for (2) reinforced/nonreinforced responses. Specifically, the four group conditions (and their designations) were: feedback

for all responses (FEED-ALL), no feedback at all (NO-FEED), feedback for reinforced responses only (FEED-RF), and feedback for non-reinforced responses only (FEED-NONRF). The feedback for responses was a 2.0-sec tone from the Sonalert. After the final session of continuous reinforcement, the feedback conditions outlined above were instituted as the mice were placed on the DRL 18 sec schedule. These conditions remained in effect for 60 days. All the subjects had stabilized by about Day 50, and the subsequent 10 days' data were combined for each subject to obtain a large and reliable sample of performance.

The upper two graphs in Figure 3 show, for each subject in each group, the conditional probability of a burst occurring after a nonreinforced response (topmost graph) and after a reinforced response (second from top). The group designations are printed below each group's data, along with the experimental condition for the class of responses in question, and the subject numbers (to allow comparisons with other figures). Group means appear above each group's data.

The topmost graph shows the conditional probability of a burst following a nonreinforced response for all four groups of mice. The two groups that received no auditory feedback for nonreinforced responses (NO-FEED and FEED-RF) had a relatively high burst probability after nonreinforced responses. In contrast, the two groups that did receive auditory feedback for nonreinforced responses (FEED-ALL and FEED-NONRF) had relatively low burst probabilities after nonreinforced responses. The only exception

to this dichotomy is subject L1, whose burst probability resembles that of the subjects in the NO-FEED and FEED-RF groups. If the burst probabilities for FEED-ALL and FEED-NONRF are compared with those for NO-FEED and FEED-RF, the only overlap in the values occurs in the case of L1. From the topmost graph in Figure 3 it is clear that the presence of the auditory stimulus accompanying nonreinforced responses resulted in lower levels of bursting after these responses.

The next graph in Figure 3 (second from top) shows the probability of a burst following a reinforced response. Here all four groups had very low burst probabilities. It seems that the presence or absence of auditory feedback made little difference if the response was reinforced. Presumably, the delivery of the reinforcer itself was a salient enough event to render the auditory feedback redundant for those groups that received it. Groups NO-FEED and FEED-RF, however did exhibit a little more bursting than the other two groups. Since NO-FEED and FEED-RF did so much bursting after nonreinforced responses, there may have been some carryover of a tendency to emit series of responses even though the first response in the series activated the reinforcer mechanism. Or, lack of feedback for the great number of nonreinforced responses emitted by these two groups may have resulted in a more variable topography which occasionally was sensed by the photoresistor circuit as two responses.

The general conclusion to be drawn from the two upper graphs in Figure 3 is that feedback for responses on DRL will

greatly reduce the bursting that occurs after nonreinforced responses, but has no effect if it accompanies reinforced responses, probably because such feedback is redundant.

In light of these results, a further question may be asked: Was the effect of the auditory stimulus used in this experiment brought about by its role simply as feedback for responses, or did it serve as a discriminative stimulus for not responding (an S^{Δ})? An examination of the distribution of IRT lengths for IRTs of 0-2.0 sec (bursts) showed that most of these bursts are IRTs of about 0.5-1.0 sec. Thus, the auditory feedback used in this experiment, being 2.0 sec in duration, was longer than the burst IRTs. This fact would allow the auditory stimulus--as an S^{Δ} , since responses emitted in its presence were never reinforced--to effectively "cut off" an incipient series of responses (burst run) once the first response in the run activated the stimulus.

EXPERIMENT II

To determine whether the auditory stimulus reduced bursting by acting as an S^{Δ} , or through its role as feedback for responses, an experiment was performed, identical to the one just described except for two changes: (1) the auditory stimulus used was only 0.2 sec in duration, less than half the value of the shortest IRT bursts, and (2) the group of subjects receiving no feedback at all (NO-FEED) was omitted from this experiment. This latter change was made because groups NO-FEED and FEED-RF were redundant in light of the preceding experiment--they both would have received (1) no auditory stimulus for nonreinforced

responses, and (2) what in the above experiment proved to be an effective stimulus for reinforced responses, the reinforcer itself. All other conditions of training in this experiment were the same as in the previous one. Naive subjects were used (S1-S4 and S9-S16).

The two bottom graphs in Figure 3 give the results of this second experiment. Again the data are combined from Days 51-60. Examining first the burst probability after nonreinforced responses, it can be seen that there was some variability within groups, but little difference between groups. The three groups had comparable burst probabilities in spite of the differences in feedback conditions they were exposed to. From the bottom graph in Figure 3 it is again evident that presence or absence of feedback for reinforced responses made no difference; the delivery of the reinforcer was sufficient to prevent further responding.

The conclusion to be drawn from this second experiment is that the effect of the auditory stimulus in the first experiment was not due to the feedback it provided for nonreinforced responses. Rather, it apparently functioned as an S^{Δ} to suppress responding on those occasions when short IRTs would have occurred. Thus, it is possible to dismiss the notion that bursts on DRL schedules result from a lack of feedback qua feedback for nonreinforced responses.

The 2.0-sec tone which accompanied nonreinforced responses in Experiment I appears to be an S^{Δ} because: (a) responses made

in its presence were never reinforced, and (b) after some exposure to it the subjects did little or no responding in its presence. However, the tone did not occur randomly throughout the experimental session; rather, its onset always occurred as a response was made. Since the making of a response on DRL always initiates a new interval to be timed, the beginning of this new interval always coincided with the tone presentation and may itself have become a "temporal S Δ ." If, after repeated pairings of the tone with the beginning of the interval, the interval beginning takes on S Δ properties of its own, then the effectiveness of the tone as an S Δ may diminish or be entirely lacking. Although the presence of the 2.0-sec tone used in Experiment I was sufficient to suppress bursting, whether it was necessary, once training on the schedule had progressed for some time, for the continuance of this suppression is not assured. To determine whether the 2.0-sec tone remained the effective S Δ , or whether the beginning of the new interval eventually acquired this property, provided the rationale for an additional experiment.

EXPERIMENT III

The strategy here was to train mice on a DRL schedule with a 2.0-sec tone for all responses; after stabilization, the tone would be omitted for nonreinforced responses. If the tone was still an effective S Δ at the time of its removal, then some increase in bursting would result from its omission.

For this experiment four naive mice (T1-T4) were trained

on a DRL schedule, the criterion IRT of which was progressively lengthened over sessions. Days 1-7 were spent on DRL 4, 8-14 on DRL 8, and 15-30 on DRL 16. The purpose of this training regimen was to reduce the length of training required for stabilization at the highest value. Each increase in criterion IRT was made when no further increase in efficiency (defined as reinforcements per response) was seen at the current value. The apparatus, deprivation schedule, session length, reinforcer, and response acquisition were the same as in the previously-described experiments. A 2.0-sec tone accompanied every response during training, but beginning on Day 28 the tone was omitted for nonreinforced responses.

Figure 4 shows the probability of a burst following a nonreinforced response for each of the four mice at DRL 16 sec. Data from Days 25-27 are shown, since these were the final three days on which the tone was supplied for nonreinforced responses. Data for Days 28-30 show the burst probability for the period when the tone was omitted for nonreinforced responses. A comparison of the data obtained with and without the tone reveals that this stimulus was still required if bursting after a nonreinforced response was to be suppressed. Apparently, after 27 days of training with the tone accompanying each response, the tone was still functioning as an S Δ . However, if the range of burst probabilities after tone removal in Figure 4 (.11 to .34) is compared to the range for groups NO-FEED (.46 - .60) and FEED-RF (.41 - .68) at the top of Figure 3, it will be seen that the latter two ranges are considerably higher. (Like Days 28-30 for mice T1-T4, groups NO-FEED and FEED-RF

had no tone for nonreinforced responses.) These lower burst probabilities for T1-T4 in the absence of the tone may be the result of a division of S^{Δ} suppression between the tone itself and the IRT initiation which it always accompanied. Thus, elimination of the tone would only remove part of the effective S^{Δ} . On this account, subjects in groups NO-FEED and FEED-RF, never having had an external (auditory) S^{Δ} paired with IRT initiation (for nonreinforced responses), would not come under good control of the beginning of the IRT as an S^{Δ} . At any rate, a salient stimulus of some duration accompanying responses on DRL continues to function as an active S^{Δ} throughout training, and may impart some of its control to its temporal coincident.

RESPONSE BURSTS AND TEMPORAL CONTROL

The experiments previously described, in which a short (0.2-sec) auditory stimulus as feedback for responses failed to reduce response bursts to significantly lower levels than occurred when feedback was absent, demonstrate two things: First, the suggestion by Kramer and Rilling (1970) that bursts on DRL schedules may result from uncertainty by the subject as to what constitutes a response is untenable. Secondly, from an IRT-reinforcement point of view, response bursts cannot result from the inadvertent reinforcement of double-response sequences, the first response of which was not forceful enough to be sensed by the control equipment. For the use of feedback for responses allowed no discrepancy between the subject's response attempts and the apparatus' reading of those

attempts. Evidently, bursts are not to be explained by appealing to aspects of the response itself.

An alternative to the now untenable suggestions examined above is to regard response bursts on DRL schedules as a source of information about the processes at work on the schedule, rather than as an incidental phenomenon. The orderliness of response bursts on DRL does suggest that they are an integral part of these processes and a direct result of the temporal control the schedule exercises over responding. Some data comparing several features of responding on DRL will be presented to show the relationship between bursting and the temporal aspects of the schedule.

Figure 5 contains data from four mice (L5-L8) from Experiment I described earlier. These four subjects comprised the group that was given no explicit feedback for responses (NO-FEED). The schedule, it will be recalled, was DRL 18 sec. The data combine Days 51-60 for each subject. Since this group of subjects received no auditory stimulus for responses, no suppression of response bursts occurred and the emission of these short IRTs was free to take its course.

Three graphs are shown for each subject in Figure 5. The top row shows the conditional IRT probability (the hazard function for the IRT distribution) for the subjects. This measure, described by Anger (1956), gives the probability that a response will occur (i.e., that an IRT will terminate) t seconds after the previous response, given that it has not occurred before t seconds. Thus, the probability of any particular IRT is conditionalized

upon the number of opportunities for that IRT to occur. This measure corrects for the fact that short IRTs of necessity have more opportunities to occur than longer ones. The second row of figures displays, for the same subjects, the conditional probability that a burst will occur, given that a particular IRT has just been emitted. Finally, the bottom row shows the mean length of the burst runs following particular IRTs. A burst run is here defined as a sequence of responses, the interposed IRTs of which are 2.0 sec or shorter. A burst run, then, is simply the responses comprising a run of bursts and includes the response that terminated the previous IRT. This measure gives the mean number of responses, occurring in rapid succession, which followed "waiting times" longer than 2:0 sec.

The graphs of conditional IRT probability in Figure 5 show good temporal control by the DRL 18 sec schedule. Although frequent short IRTs in the range 0-2.0 or 0-4.0 sec are emitted by all subjects, the higher IRT probabilities are at longer values approaching the 18-sec criterion. The lowest response probability occurs from about 2.0-14.0 sec after the previous responses (provided, of course, no response occurs before 2.0 sec has elapsed). That few intermediate length IRTs (2.0-14.0 sec) occur shows that on DRL subjects emit long and short IRTs in nonrandom sequences. Some of these sequences are described by the other graphs in Figure 5.

As an example, the second row of graphs shows the probability of a burst following various other IRTs. As described in the first section of this paper, and depicted here again, bursts occur almost

exclusively after nonreinforced IRTs, rarely after reinforced IRTs. Another consistent feature of the conditional burst probability in Figure 5 is that the burst probability steadily increases from about 6.0-18.0 sec. This rise in burst probability with increasing length of previous IRT parallels the rise in IRT probability with time since the last response (top row of Figure 5). Thus, the longer the subject has waited after his last response, the more probable a response becomes, and also the more probable a burst will follow that response (provided the reinforcer is not presented). It appears, then, that response bursts are controlled by the same variable that controls the responses terminating the longer IRTs on DRL schedules--namely, the passage of time.

As described earlier in this paper, bursts often occur in multiple, or runs. The bottom row in Figure 5 describes the number of successive bursts occurring after various periods of no responding. As explained above, this includes the response which terminates the waiting period. These graphs resemble, in their general characteristics, those for conditional burst probability immediately above them. It can be seen that the mean number of responses emitted in quick succession depend upon the length of the preceding period of no responding. The longer the subject has waited before responding, the more responding he does. Reinforced responses, however, were only rarely followed by any additional responses.

The picture that emerges from the three measures displayed in Figure 5 is one of increasing response strength with time since

previous response. And on DRL, time since previous response is synonymous with the time into the interval to be timed. The longer the subject has waited since his last response: (a) the more likely responding becomes, and (b) the greater the quantity of responding that occurs. The dichotomy between long and short IRTs seen in the top row of Figure 5 is the result of this pattern of no responding (long IRTs) and subsequent emission of multiple responses (short IRTs).

Although bursts are generally defined as IRTs of 0-2.0 sec in duration, following Sidman (1956), this distinction is somewhat arbitrary. The least probable IRT length varies slightly from subject to subject, and all those IRTs shorter than this length may best be regarded as constituting bursts. It will be noticed in Figure 5 that for some subjects bursts and burst runs are more frequent after 2.0-6.0 sec IRTs than after 6.0-10.0 sec IRTs. In these cases, the 2.0-6.0 sec IRTs behave as bursts--they occur intermixed with the 0-2.0 sec IRT sequences. All IRTs shorter than the value of the lowest IRT probability in the distribution appear to be functionally homogeneous.

DISCUSSION

The previous sections have made clear that response bursts on DRL schedules are not the result of uncertainty by the subject as to the definition of the response. The provision of a brief but clear feedback stimulus for responses did not produce a reduction in bursting below the levels exhibited by subjects not given such feedback. It was seen, however, that by lengthening the stimulus

to 2.0 sec a significant reduction in bursting could be obtained, and that the effect of this longer stimulus was that of an S^{Δ} . The removal of the stimulus after extensive training caused an immediate increase in response bursts in all subjects--evidence that it was still actively functioning as an S^{Δ} . It has also been shown that short IRTs do not occur randomly on DRL, but are sensitive to environmental events as well as to the temporal dimension. These results bear upon the interpretation of performance on DRL in several ways.

Implicit in the classification of the auditory stimulus in Experiment I as an S^{Δ} , and that in Experiment II as merely feedback is a functional distinction between the two. The procedural difference was the duration of the two stimuli, but the fact that the two durations had very different effects (or were effective or not effective) on responding prompts the making of a functional distinction. The distinction suggested here is that since both the long and short duration stimuli were certainly of a character to be heard by the subject, differing only in duration, then both can properly be called feedback. But since only the longer duration stimulus came to suppress responding which occurred (or on subsequent occasions would have occurred) in its presence, it can also properly be seen as an S^{Δ} . This distinction raises questions about why the duration of the stimulus determined its effectiveness. It was mentioned earlier that the short IRTs on DRL emitted by the mice were of values of about 0.5-1.0 sec. The effect of a stimulus in suppressing these short IRTs may require that the stimulus be present

when the second response of the pair forming the IRT is emitted. Put more generally, if a series of responses in rapid succession occurs on certain occasions, a stimulus which accompanies a response must be of a duration to overlap the succeeding response if it is to be capable of becoming a discriminative stimulus. This distinction between feedback and discriminative stimulus is not incompatible with the functional definition of a discriminative stimulus, which specifies that the stimulus comes to control responses occurring in its presence.

This distinction ignores the question of why the short (0.2 sec) stimulus used in these experiments might not also serve as an effective S^{Δ} by virtue of a stimulus trace persisting after the presentation of the stimulus itself. It may be that any S^{Δ} effects of such a trace were small and were overshadowed by the excitation associated with the temporal stimulus at the time of responding.

As mentioned in the Introduction, it is difficult to conceive of a response rate being reinforced on DRL in any more than the strictly technical sense of the phrase. The DRL schedule insures that no responding has occurred prior to the reinforced response for a period of time equal to the criterion IRT. Assuming that the effect of reinforcement on behavior extends back over a definite period of time preceding the reinforcer delivery, this time period is filled on DRL with other behavior. The effects of the reinforcement must extend to the single response which produces it and possibly to other behavior which is not being measured (Laties et al., 1965, 1969). It is not easy to believe,

however, that the initiation of an IRT of, say, 18 sec or more is directly reinforced by an event which occurs more than 18 sec later. We must distinguish here between the procedure of reinforcing an IRT (which certainly has been shown to have the effect of making the reinforced IRT values more likely to occur, Kuch & Platt, 1976) and the mechanism by which those reinforced IRTs become more frequent. The argument just made addresses the mechanism of IRT reinforcement and suggests that presentation of the reinforcer in the presence of the temporal stimulus directly strengthens (or makes more probable) the key-peck or nose-poke which terminates the IRT, but not the actual interresponse time per se. The production of particular IRT values, then, is a matter of stimulus control by the temporal dimension. The data describing response bursts presented in this paper demonstrate how the DRL schedule produces long IRTs by selectively making reinforcement available for responses occurring at various times. DRL essentially confronts the subject with a temporal discrimination task.

Sidman (1956) described the DRL schedule as a temporal discrimination task on which response strength fluctuated from moment to moment. Procedurally, this description is consonant with the schedule contingencies: With response occurrence and reinforcer presentation as reference points, the temporal stimulus changes value with the passage of time. Responses occurring in the presence of S^{Δ} (the first t seconds after the most recent response or reinforcer delivery-- t being equal to the IRT requirement) are never reinforced; responses occurring in the presence of S^D

(any time after t seconds has elapsed) are always reinforced. Unlike a typical discrimination experiment, however, DRL does not allow the presentation of particular values of the temporal stimulus at particular times. The subject, by responding on DRL, determines when the temporal stimulus will again begin at its lowest value.

The conceptualization of the DRL schedule as a temporal discrimination task is most compatible with the data presented here. It was seen that as the time since the last response or reinforcement progresses (i.e., the temporal stimulus changes) not only does the probability of a response increase, but the number of responses emitted by the subject in quick succession before he begins "waiting" again also increases. That is, the response rate associated with the different values of the temporal stimulus increases as those values approach the value in the presence of which responses will be reinforced. This only refers to nonreinforced "trials," however, since we have seen that reinforced trials almost always end with a single response (the response which produces the reinforcer). Response strength is presumably just as high at time $t + 1.0$ sec as it is at $t - 1.0$ sec, but the presentation of the reinforcer for a response at $t + 1.0$ sec radically changes the prevailing stimulus situation. In short, we have two different kinds of trials, and on those on which the reinforcer is not presented we have the opportunity to see response strength reflected in response rate. If, for example, on selected trials the reinforcer were omitted for a response made at time $t + 1.0$ sec, undoubtedly multiple

rapid responses would occur, as they do at time $t - 1.0$ sec. Such responses do not occur on DRL schedules because delivery of the reinforcer effectively cuts them off. The preceding analysis of the DRL schedule as a procedure that provides the subject with a temporal discrimination task elevates the response burst to the status of primary datum of interest, since it indicates the response strength at any given time. This is in marked contrast to its customary role as an incidental schedule phenomenon which requires explaining apart from the processes at work on the schedule.

The DRL schedule is peculiar among temporal schedules in that it provides very predictable periods of reinforcement availability and nonavailability. This makes it possible for the subject to become sensitive to the reference points for the temporal aspects of the schedule and hence to the temporal aspects themselves. Where reinforcement is less predictable on the basis of the passage of time (e.g., VI schedules) temporal patterning of responses as on DRL will not occur; instead, more molar processes such as overall rate of reinforcement may play a large role in determining response rate. But where time provides the subject with information about reinforcement availability, he will almost certainly make use of that information. Bursting on DRL shows what the subject knows about time.

APPENDIX

Subjects

In the experiments designated I, II, and III naive male white laboratory mice (CD-1, Charles Rivers) were used. Sixteen subjects served in Experiment I, 12 in Experiment II, and 4 in Experiment III. Each subject was approximately 60 days old at the beginning of the experiment in which it served. An additional experiment employed four adult, experimentally-naive White Carneaux pigeons obtained from Palmetto Pigeon Plant, maintained throughout on mixed grain at 80% of their free-feeding weights.

Apparatus

Four identical operant chambers for the mice were constructed entirely of Plexiglas with length 20 cm, width 10 cm, height 20 cm. All sides were transparent, except the floor and front panel, which were painted flat black. The front panel contained a circular aperture 2.2 cm in diameter. A response consisted of a nose-poke into the aperture, which interrupted a light beam falling on a photoresistor. Condensed milk was used as the reinforcer and was delivered by a liquid dipper through a hole in the floor 1.4 cm in diameter centered 4 cm from the front panel and midway between the sides of the chamber. The dipper remained in the down position when not delivering the reinforcer. Reinforcement consisted of raising the dipper for 2.0 sec, although the mice consumed the 0.01 cc milk delivery in considerably less time. A houselight was mounted on the front

panel 8 cm from the ceiling and remained on during the experimental session. For some subjects a Sonalert was sounded at 70 db SPL (C scale): it was mounted on the chamber ceiling and was powered by a 28 V dc power supply through a 27,000 ohm resistor. The resulting auditory stimulus was loud enough to cause some mice to "freeze" when it was first presented, but this response was not seen after the first few presentations. The chamber was contained within a sound-attenuating housing upon which was mounted an exhaust fan for ventilation. Chamber events and data collection were effected by a PDP8/e computer located in an adjoining room.

For the pigeons four identical, three-key Lehigh Valley pigeon chambers were employed. Only the center key was used. This key was mounted 19 cm above the floor and was transilluminated by a No. 1820 incandescent bulb operated at 28 V dc. Three sec access to mixed grain was the reinforcer. A PDP8/e computer controlled the contingencies and recorded data.

Procedure

Upon receipt from the breeding farm the mice were given one week free access to food and water. They were then gradually put under water deprivation until, at the end of 30 days, one half hour access to water maintained stable weights. Food was always available in the home cages. At this point, weights were at about 80-85% of ad lib values. All mice were then given a one-hour experimental session with every nose-poke reinforced (CRF). For reinforcement, the liquid dipper was raised for 5 sec. No response

feedback (via the Sonalert) was given at this time. Two subjects in Experiment I (L3 and L15) failed to acquire the response during this initial session, as evidenced by very few nose-pokes. An additional one-hour session (identical to the first) was provided these two subjects on the following day. Both acquired the response during this session. On the following day all subjects were exposed to the same conditions (one-hour session, every response reinforced, no auditory feedback), but with the liquid dipper access reduced to 2.0 sec for each reinforcement. Beginning the next day, the mice were exposed to a DRL schedule. For Experiments I and II the DRL criterion was 18 sec. For Experiment III the DRL value was 4 sec for Days 1-7, 8 sec for Days 8-14, and 16 sec for Days 15-30. From the first day of DRL training subjects received auditory feedback for their responses from the Sonalert. The details concerning this feedback are described in the appropriate places in the text. The experimental sessions were 45 min long and were run six days a week. After each session the mice were given 30 min access to water in their home cages.

The pigeons were autoshaped (Brown & Jenkins, 1968) and then placed on a DRL 8 sec schedule for 45 days. Sessions were 26 min long and were run six days a week. The birds were fed a small amount of mixed grain shortly after the session to maintain their weights at about 80% of free-feeding.

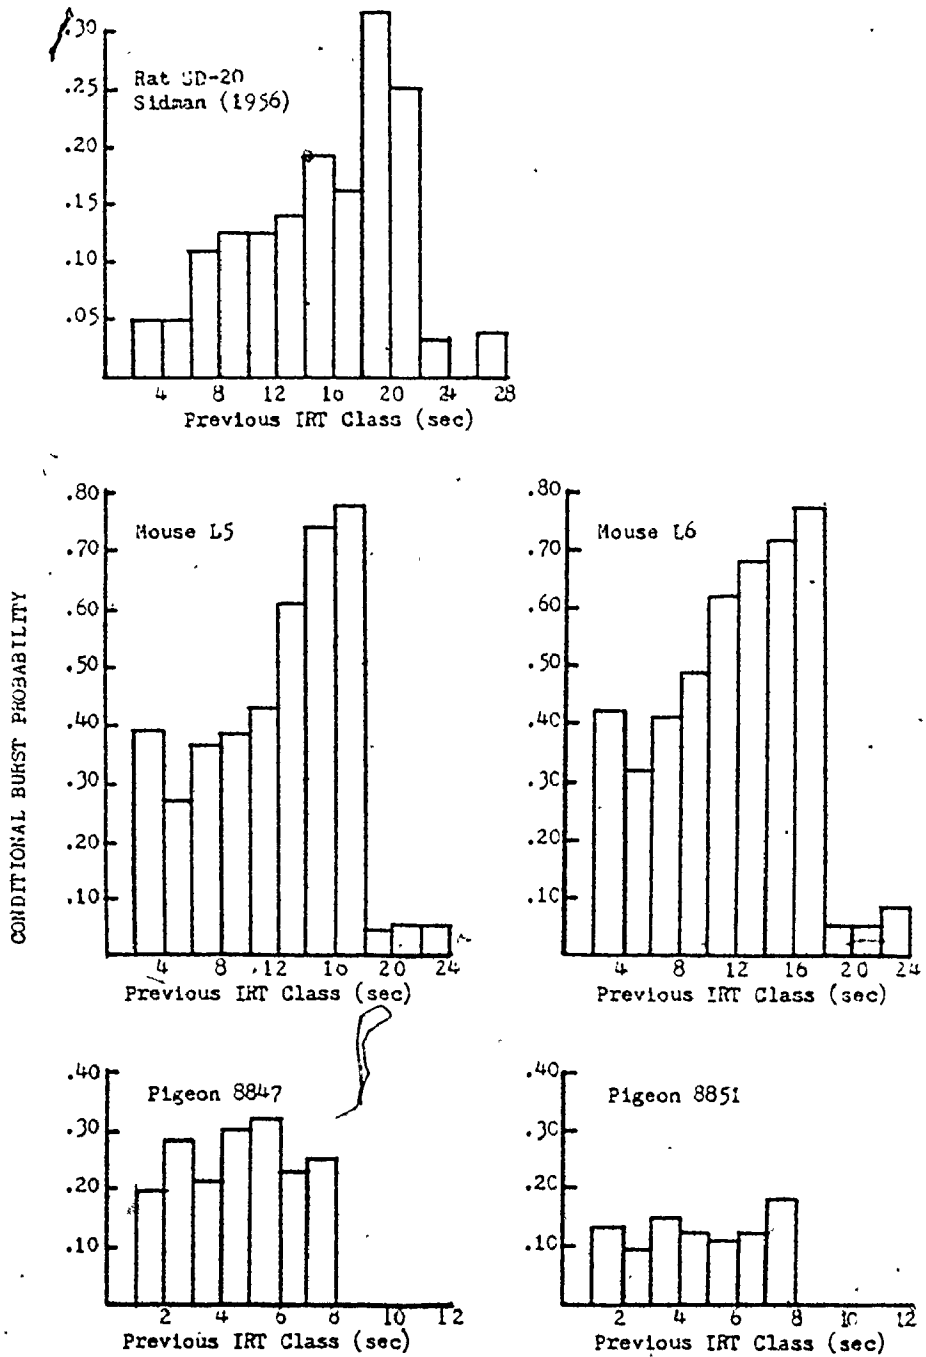


Figure 1

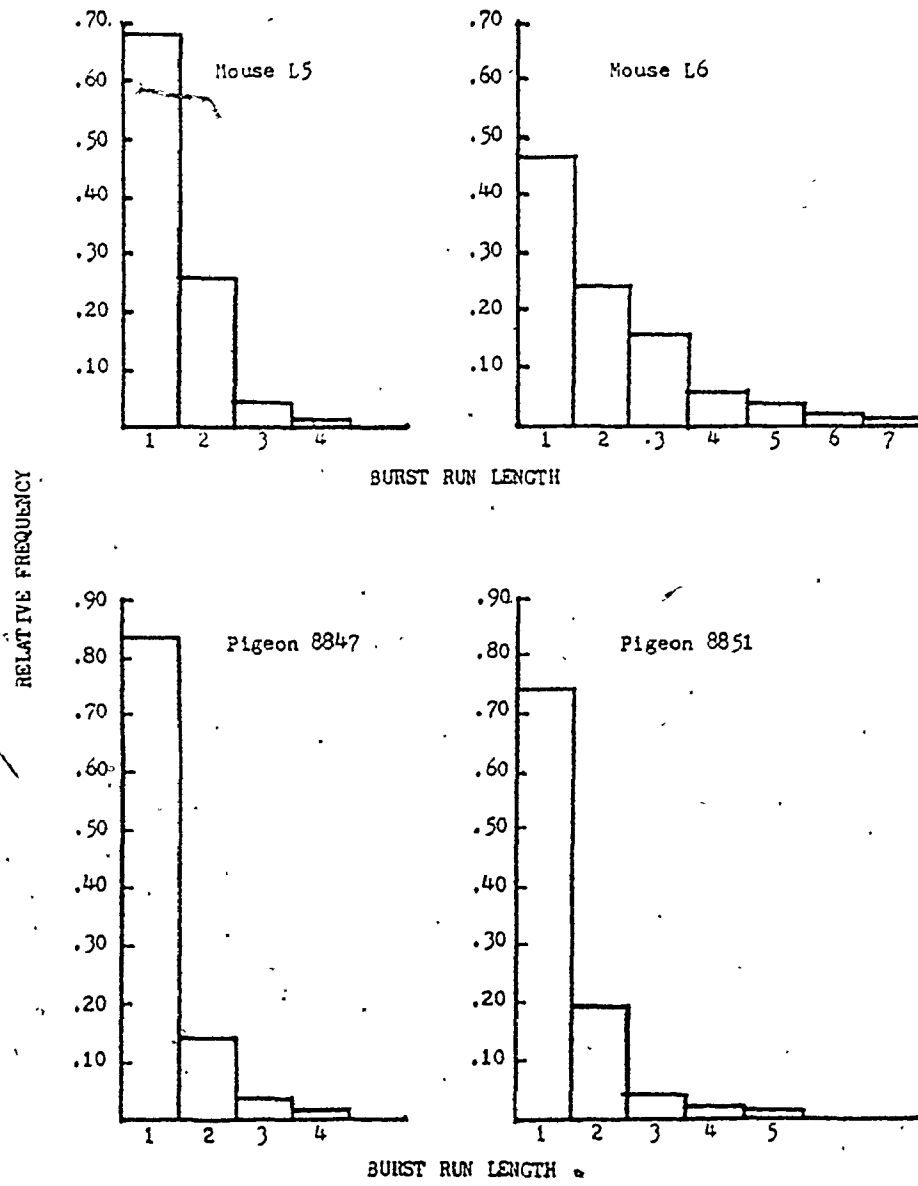


Figure 2

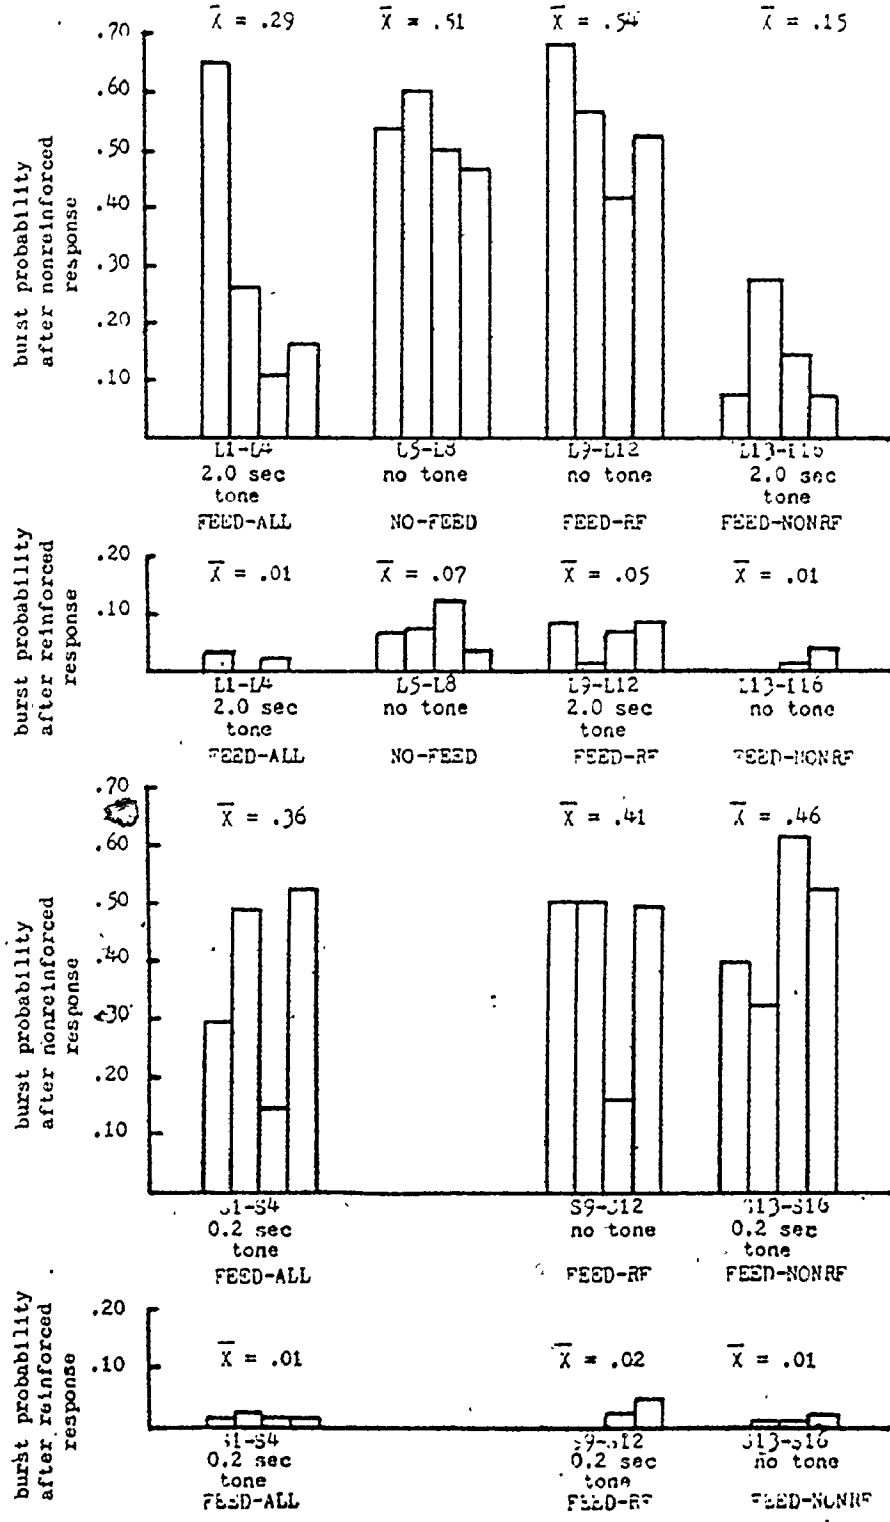


Figure 3


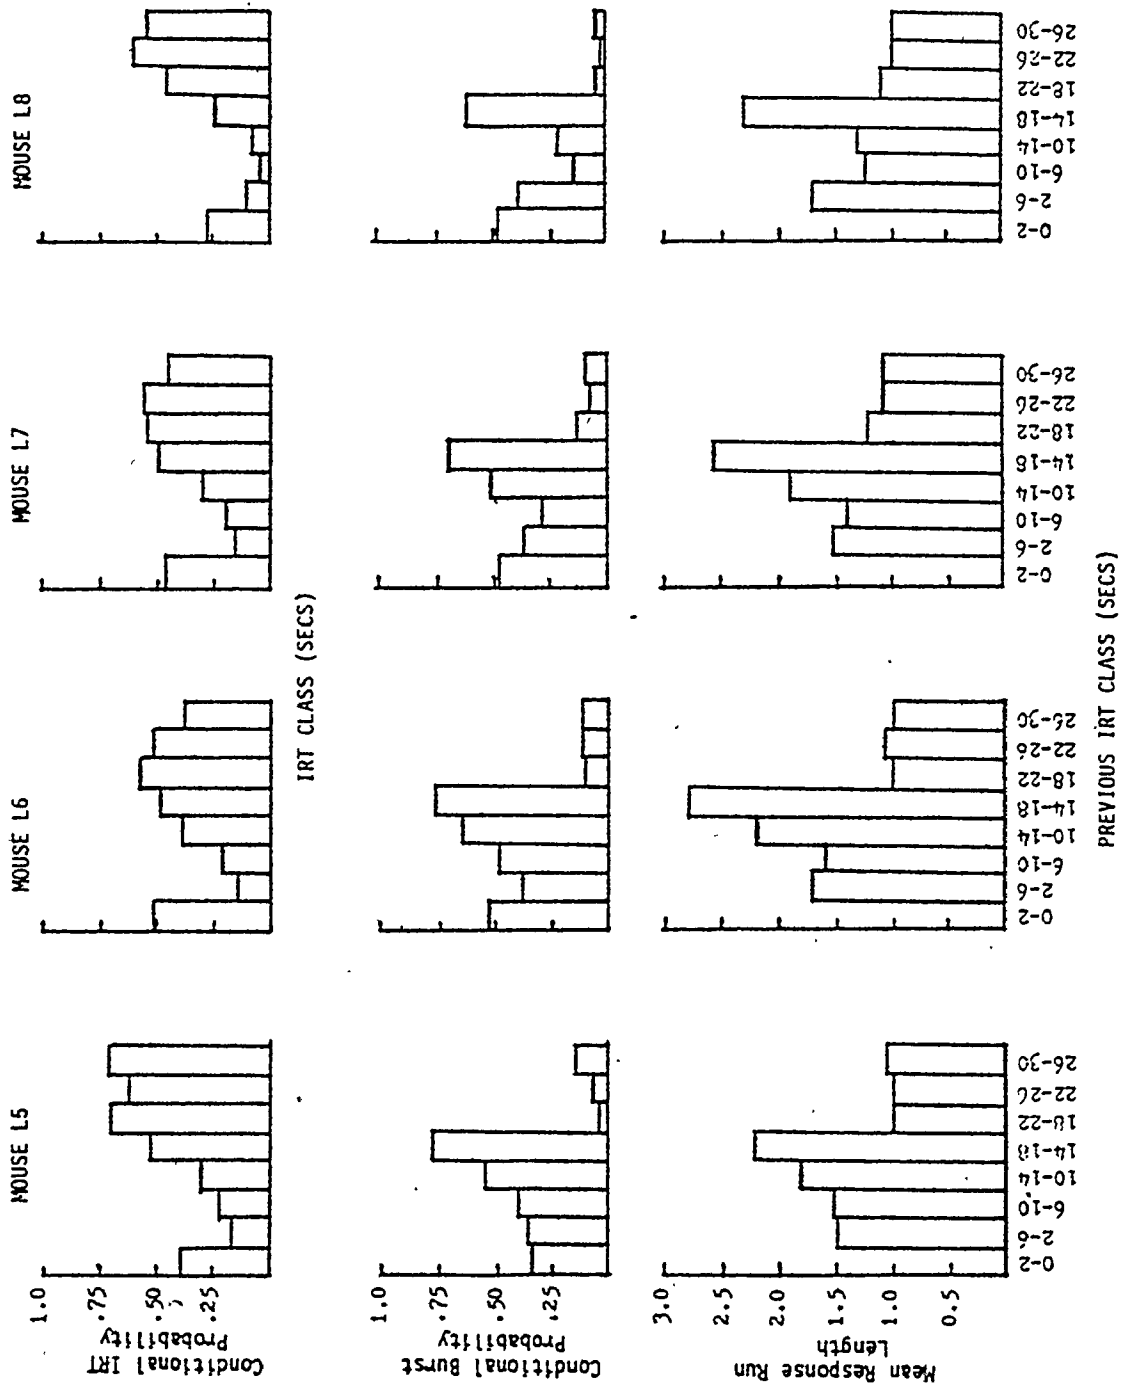


Figure 4

Figure 5



REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. J. exp. Psychol., 1956, 52, 145-161.
- Blough, D.S. Interresponse time as a function of continuous variables: A new method and some data. J. exp. Anal. Behav., 6, 1963, 237-246.
- Blough, D.S. The reinforcement of least frequent interresponse time. J. exp. Anal. Behav., 1966, 9, 581-592.
- Brown, P. and Jenkins, H. Autoshaping of the pigeon's key peck. J. exp. Anal. Behav., 1968, 11, 1-8.
- Carter, D.E. and Bruno, J.J. Extinction and reconditioning of behavior generated by a DRL contingency of reinforcement. Psychonomic Science, 1968, 11, 19-20 (a).
- Catania, A.C. Reinforcement schedules and psychophysical judgments: a study of some temporal properties of behavior. In W.N. Schonfeld (Ed.), The Theory of Reinforcement Schedules. New York: Appleton-Century-Crofts, 1970, Pp. 1-42.
- Conrad, D.G., Sidman, M., & Herrnstein, R.J. The effects of deprivation upon temporally-spaced responding. J. exp. Anal. Behav., 1958, 1, 59-65.
- Eckerman, D.A. and Lanson, R.N. Variability of response locus for pigeon's responding under continuous reinforcement, intermittent reinforcement, and extinction. J. exp. Anal. Behav., 1969, 12, 73-80.
- Ferraro, D.P., Schonfeld, W.N., & Snapper, A.G. Sequential response effects in the white rat during conditioning and extinction on a DRL schedule. J. exp. Anal. Behav., 1965, 8, 255-260.
- Ferster, C.B. and Skinner, B.F. Schedules of Reinforcement. New York: Appleton-Century-Crofts, 1957, p. 459.
- Herrick, R.M. Lever displacement under fixed-ratio schedules and subsequent extinction. J. Comp. Physiol. Psych., 1965, Vol. 59, No. 2, 263-270.
- Herrick, R.M. and Bromberger, R.A. Lever displacement under variable ratio schedules and subsequent extinction. J. Comp. Physiol. Psych., 1965, Vol. 59, No. 3, 392-398.

- Holz, W.C., Azrin, N.H., & Ulrich, R.E. Punishment of temporally-spaced responding. J. exp. Anal. Behav., 1963, 6, 281-290.
- Holz, W.C. and Azrin, N.H. A comparison of several procedures for eliminating behavior. J. exp. Anal. Behav., 1963, 6, 399-406.
- Kelleher, D.T., Fry, W. & Cook, L. Inter-response time distribution as a function of differential reinforcement of temporally-spaced responses. J. exp. Anal. Behav., 1959, 2, 91-106.
- Kramer, T.J. Effects of timeout on spaced responding in pigeons. Unpublished Master's thesis, Michigan State University, 1968.
- Kramer, T.J. and Rilling, M. Effects of timeout on spaced responding in pigeons. J. exp. Anal. Behav., 1969, 12, 283-288.
- Kramer, T.J. and Rilling, M. Differential reinforcement of low rates: A selective critique. Psychological Bulletin, 1970, 74, 225-252.
- Kuch, D.O. and Platt, J.R. Reinforcement rate and interresponse time differentiation. J. exp. Anal. Behav., 1976, 26, 471-486.
- Laties, V.G., Weiss, B., Clark, R.L., Reynolds, M.D. Overt "mediating" behavior during temporally spaced responding. J. exp. Anal. Behav., 1965, 8, 107-116.
- Laties, V.G., Weiss, B., & Weiss, A.B. Further observations on overt "mediating" behavior and the discrimination of time. J. exp. Anal. Behav., 1969, 12, 43-57.
- Margulies, S. Response duration in operant level, regular reinforcement, and extinction. J. exp. Anal. Behav., 1961, 4, 317-321.
- Mintz, D.E. Force of response during ratio performance. Science, 1962, 138, 516-517.
- Reynolds, G.S. Accurate and rapid reconditioning of spaced responding. J. exp. Anal. Behav., 1964, 7, 273-276.
- Reynolds, G.S. Temporally spaced responding in pigeons: Development and effects of deprivation and extinction. J. exp. Anal. Behav., 1964, 7, 415-421.
- Sidman, M. Time discrimination and behavior interaction in a free operant situation. J. Comp. Physiol. Psychol., 1956, 49, 469-473.
- Staddon, J.E.R. Some properties of spaced responding in pigeons. J. exp. Anal. Behav., 1965, 8, 19-28.