JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1968, 11, 251-258 NUMBER 3 (MAY)

# *THE STRUCTURE OF RESPONSE RATE*[1]

DAVID R. WILLIAMS

UNIVERSITY OF PENNSYLVANIA

Interresponse time distributions of the terminal rate under fixed-interval and fixed-ratio schedules were examined, using data from three rats in each case. By means of a sequential analysis, the overall interresponse time distributions were separated into orderly components. Consideration of the component distributions suggested that multiple determinants of rate act in succession, not simultaneously, and that probability of reinforcement has an important effect on the probability of occurrence of interresponse times.

---

Rate of responding is the basic datum of the free-operant situation. The selection of an appropriate quantitative description is important, complicated by the tendency of any measure to lose or distort some characteristics of the original performance. For example, a description of overall rate (responses per hour) conceals information concerning the pattern of responding between successive reinforcements. To the extent that such features of performance are irrelevant to a given analysis, this suppression of information may even be an advantage; until a situation is well understood, however, there is a troublesome possibility that an important aspect of performance will be hidden by a particular quantitative description.

Detailed investigations of response rate often use frequency distributions of the time between successive responses (interresponse time, or IRT, distributions) as a means of quantitatively describing rate in detail. The main results of these investigations, and the use of IRT distributions, have recently been discussed by Morse (1966). Although IRT distributions reveal characteristics of rate which are obscured in a simple statement of mean rate, they conceal sequential patterning of the IRTs from which they are compiled. The present paper reports an analysis of sequential patterns of responding found in the rates maintained by two conventional schedules of reinforcement. Its purpose was to consider whether the sequential organization of IRTs might contain useful information for the investigation of response rate.

The analysis was carried out on responding maintained by fixed-interval and fixed-ratio schedules. Although stable performance under both of these schedules is characterized by marked changes in rate during the interval between reinforcements, a steady rate (the "terminal rate") is typically found immediately before reinforcement (Ferster and Skinner, 1957). The present analysis was carried out only on responding at the terminal rate; responding earlier in the interreinforcement interval was not included.

Interval and ratio schedules were both used because they provide a contrast in the way reinforcements are allotted to IRTs of different durations. As Skinner (1938) pointed out, if all IRTs are not of precisely equal duration, an interval schedule will provide reinforcement with a higher probability after longer IRTs than after shorter ones. Revusky (1962) demonstrated this property of interval schedules mathematically, and Ferster and Skinner (1957) have shown experimentally that this property enters into the determination of response rate. Under ratio schedules, by contrast, the relative frequency of reinforcement of IRTs is proportional to their relative frequency of occurrence: the reinforcer does not occur differentially following IRTs of different durations. These characteristics of interval

[1]Preparation of this report was supported in part by a grant, NSF GB 5418, from the National Science Foundation. The experimental work was carried out at Yale University. Reprints may be obtained from the author, Dept. of Psychology, College Hall, University of Pennsylvania, Philadelphia, Pa. 19104.

and ratio schedules, and the related experimental work, are reviewed in detail by Morse (1966).

## METHOD

### *Subjects*

Six adult female albino rats were maintained at 80% of their free-feeding weight. They had no previous experimental history.

### *Apparatus*

The experimental chamber was enclosed in a sound-attenuating icebox in which a fan and a buzzing relay provided a constant masking noise. The experimental chamber measured 9 in. square and was 7 in. high. A standard Gerbrands lever projected 0.6 in. into the cage at a height 3.5 in. above the floor. The lever required 15 g force for operation. A tray into which single 45-mg Noyes pellets were delivered was placed on an adjacent wall, 5 in. from the center of the lever. The chamber was illuminated by a 7.5-w red bulb. Reinforcement was contingent on release of the lever and a press was defined as including both depression of the lever and its release. An IRT was defined as the elapsed time between successive releases of the lever. A relay made an audible click with each release of the lever. Water was continuously available from a bottle on the wall opposite the lever.

In the fixed-interval experiment, all programming and recording were done automatically. All IRTs were recorded successively in 0.5-sec class intervals on a 20-pen Esterline-Angus recorder, and later transferred to punched cards for analysis. In the fixed-ratio experiment, the final five IRTs of each ratio were registered on Standard Electric timers, and immediately written down to the nearest 0.1 sec, in order. These IRTs were subsequently transferred to punched cards for analysis. Throughout these experiments, performance was also monitored on a cumulative recorder.

### *Procedure*

Preliminary training was identical for all subjects. Two days of magazine training, during which one pellet per minute was delivered for 60 min, were administered with the lever removed. On the next two days, the lever was inserted, and 1 hr of training on continuous reinforcement (CRF) was provided. On the first day, all subjects responded within 30 min, and no shaping was required.

After this pretraining, three subjects were placed directly on FI 4, and no further changes in procedure were made for the remainder of the experiment. The other three subjects were rewarded on a small fixed ratio, the value of which was increased to FR 60 over the next seven days. Once FR 60 was reached, no further changes in procedure were made.

The data were drawn from Sessions 71 to 80 on FI 4, and from Sessions 61 to 65 on FR 60. On FI 4, data from the first five intervals of each day, and IRTs from intervals which included fewer than 15 presses, were excluded from analysis. Approximately 15% of all intervals were discarded for these reasons.

## RESULTS

### *Fixed Interval*

Throughout the sessions which were analyzed, a well-developed fixed-interval scallop was in evidence, and no systematic changes in performance were noted on inspection of cumulative records, or on inspection of IRT distributions from the first and second halves of the recording period. Before the main analysis, distributions of reinforced IRTs, of IRTs 1-back from reinforcement, 2-back, *etc.* were constructed for the final 12 IRTs in each interval. For R-13 and R-14, the IRT distributions for responses 4-back through 11-back from the reinforced response did not differ in any systematic fashion as determined by inspection, nor did a $\chi^2$ test indicate any reliable differences ($P > 0.5$ for R-13; $P > 0.1$ for R-14). Accordingly, IRTs 4- through 11-back from the reinforced response were taken as representative of the terminal rate, and analyzed further. In the case of R-15, the distributions 4- through 7-back appeared similar and did not differ reliably according to a $\chi^2$ test ($P > 0.2$), although they did differ from the distributions 8- through 11-back. Further analysis in this case was carried out only on IRTs 4- through 7-back from the reinforced response.[2]

---

[2]The exclusion of data, as described here, is a conservative procedure; it might reduce the reliability of the findings, or possibly obscure some relationship among IRTs. It will not, however, artificially introduce false relationships among IRTs.

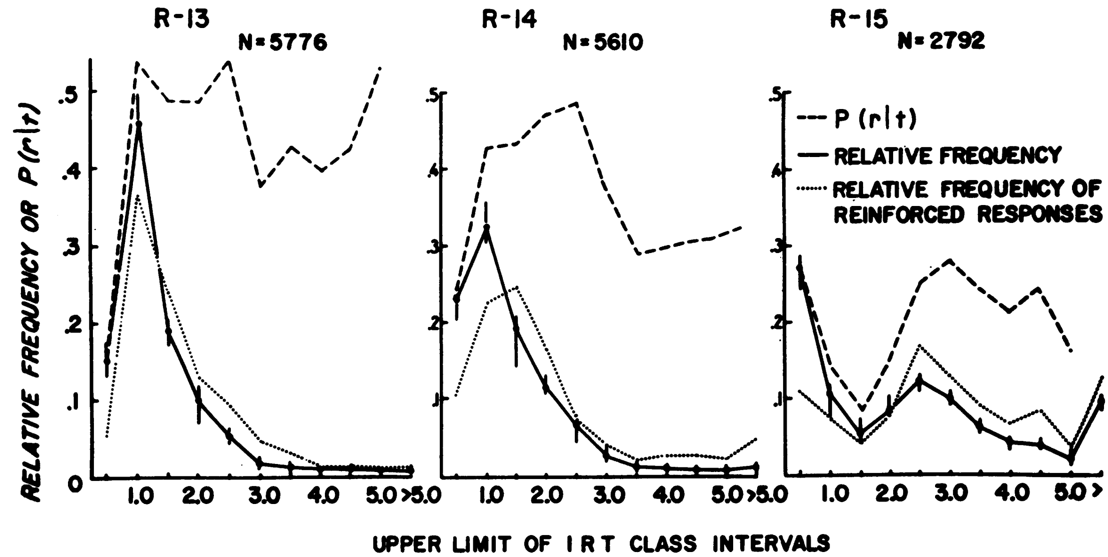


Fig. 1. Relative frequency and P(r|t) for each subject at the terminal rate on FI 4; the number of IRTs included is indicated in each panel. The relative frequency of reinforced IRTs is also shown.

The distributions representing the terminal rate are shown in Fig. 1 by the heavy lines, which indicate the relative frequency of IRTs in 0.5-sec class intervals. The vertical bars through each data point extend from the maximum to the minimum relative frequency of IRTs in each class interval of the sub-distributions (4-back, 5-back, *etc.*). The small range of variability among the sub-distributions indicates that there were at most negligible changes of rate during the period immediately before reinforcement.

The dotted lines of Fig. 1 represent the distribution of reinforced IRTs. Although there are distinct similarities between these distributions and the distributions characterizing the terminal rate, the reinforced distribution contains relatively fewer IRTs in the early class intervals, and relatively more in the later ones. A $\chi^2$ test conducted separately for each subject indicated that the two distributions differed reliably ($P < 0.001$ in every case). When the solid and dotted curves are compared, the tendency of interval schedules to reinforce longer IRTs selectively becomes evident, and the magnitude of the effect is readily seen.[3]

The dashed line of Fig. 1 is an alternative representation of the terminal rate distribution, and indicates the conditional probability of response in the various class intervals. This measure, introduced by Anger (1956) as "IRTs /Op", is denoted here as P(r|t), namely the probability of a response given that t sec have elapsed since the previous response. It is useful because the frequency of IRTs in a particular class interval depends not only on the probability of a response (*i.e.*, a press) in the class interval itself, but also on the number of times that the class interval was reached. For example, the number of 3-sec IRTs in a session depends not only on the tendency to respond after a pause of 3-sec duration, but also on the number of times a 3-sec pause took place. The conditional measure, P(r|t), indicates the tendency to respond in a given class interval, regardless of the number of occasions when a pause of sufficient duration provided an opportunity to respond there. Thus, P(r|t) can be large (near unity) even when a small number of IRTs were actually of duration t. P(r|t) complements the relative frequency measure because it separates the tendency to respond at t, considered alone, from the tendency to respond at t relative to all other times.

[3]The distributions of IRTs 1- through 3-back from reinforcement were excluded from the terminal rate distribution because of sequential dependencies in the data (see below). When successive IRTs are related, the over-representation of some IRT value in the reinforced distribution implies a related distortion in the immediately prior distributions. Again, the exclusion is conservative.

A somewhat different relationship among IRTs in the terminal rate samples is evident when P(r|t) is used. In the case of R-14, for example, P(r|t) rises throughout the first five class intervals, even though relative frequency falls after the second. This rising trend in P(r|t) embraces 93% of all IRTs included in the sample, and indicates that the tendency to respond, as successive class intervals were reached, increased throughout this period. The P(r|t) curve for R-13 does not show any falling trend of major proportions until after the fifth class interval, although the relative frequency curve falls off after the second class interval. In this case, the first five class intervals include 95% of all IRTs. The relative frequency and P(r|t) curves for R-15 are much more nearly similar in overall shape: both appear bimodal, peaking in the first class interval and again in a later one.

In order to analyze the terminal response rates for the presence of sequential patterns, a series of IRT distributions separated according to the duration of the preceding IRT was compiled for each rat. Thus, separate distributions were developed for IRTs following a 0.5-sec IRT, following a 1.0-sec IRT, *etc.* If there were no sequential dependencies (patterns) in the terminal rate, these new distributions would be the same except for sampling error: for example, the IRT distribution following a short IRT would resemble the IRT distribution following a long IRT. If relationships among successive IRTs were present, however, they would be reflected by consistent differences among the separated IRT distributions. Using a $\chi^2$ test, the hypothesis that the separated distributions were of the same form was rejected ($P < 0.001$ for each subject).

The separated components of the overall distribution are shown in Fig. 2, in terms of

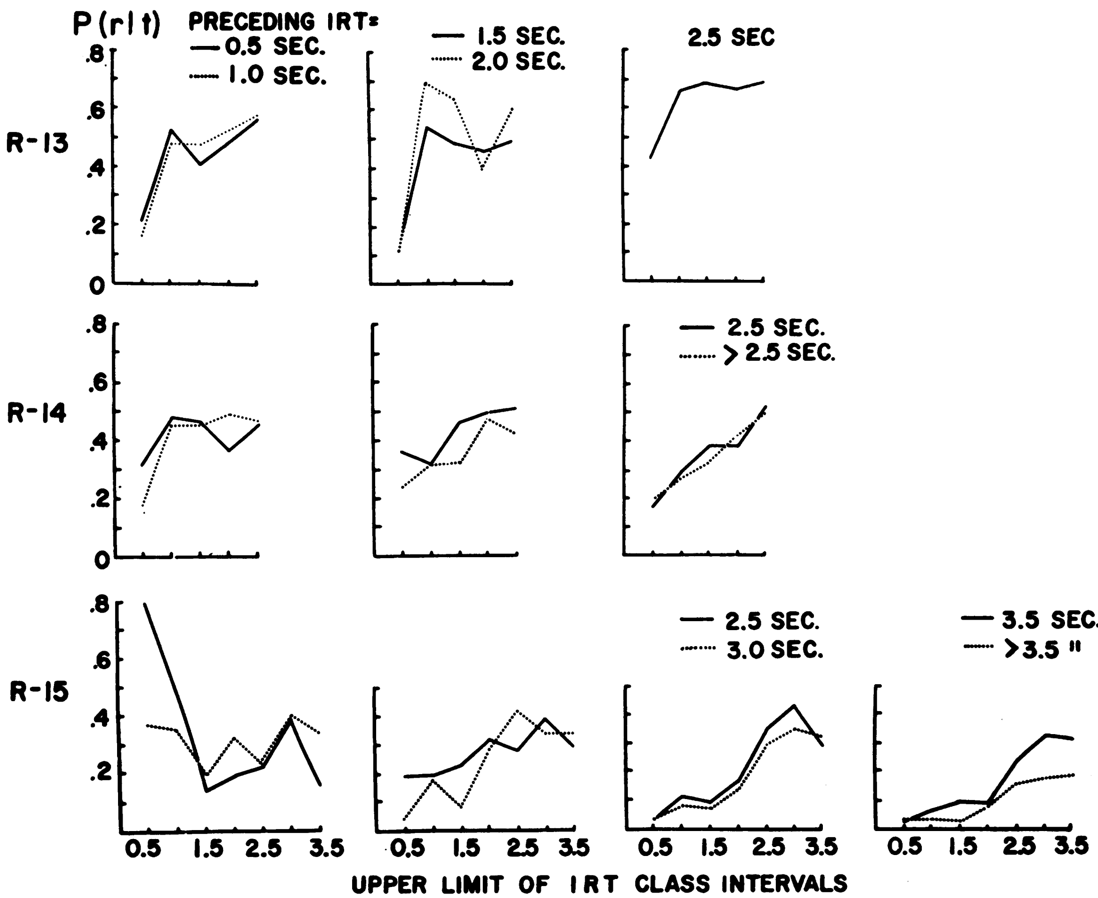


Fig. 2. Sequential analysis of terminal rate data presented in Fig. 1. The conditional probability of response is shown as a function of duration of the preceding IRT.

the conditional probability of response measure. The curves differ with respect to the length of the IRT which preceded them, and do not cover the very longest IRTs because there were too few of them to provide reliable estimates. Within each series of component distributions, an orderly progression of forms is evident as one compares curves following successively longer IRTs. In the case of R-15, for example, after an IRT of 0.5-sec duration, about 90% of the next IRTs are shorter than 1.0 sec (because these are conditional probability curves, the heights of these points do not indicate the number of IRTs included). As the duration of the preceding IRT increases, the tendency to respond with short IRTs diminishes until, following the longest IRTs, less than 10% of the IRTs are of short duration. Indeed, except following short IRTs, P(r|t) tends to increase steadily as a function of time since the last response. The curves for R-14 also show an orderly progression as the duration of the prior IRT increases, and also tend to become simple increasing functions following the longer IRTs. For this subject, P(r|t) is nearly constant after 0.5-sec IRTs, and develops a more and more positive slope following longer IRTs. A somewhat different pattern is found with R-13. An increasing function occurs only after IRTs of the shortest duration, and it appears that after a long IRT, R-13 emits a relatively short IRT, thus showing a pattern of "double presses" on these occasions.

Although an attempt was made to analyze patterns based on three rather than two IRTs, there were not sufficient data for this purpose. As pattern length increases, the number of observations needed to maintain a suitable number of cases increases exponentially. Even if only five class intervals are used, there are 25 separate distributions which need to be examined if two preceding responses are used rather than one.

*Fixed Ratio*

Owing to difficulties of recording IRTs at the high rates encountered under FR, it was possible to record only the last five IRTs in each fixed-ratio sequence of 60 responses. A more extended recording sequence, however desirable, is less important on fixed-ratio than it is on fixed-interval schedules: on fixed ratio, probability of reinforcement is independent of IRT duration, so there is no *a priori* reason to expect the distribution of reinforced IRTs to differ from the distribution characterizing the terminal rate. When the distribution of reinforced IRTs, the distributions 1-back, 2-back, *etc.* were compared, they did not differ reliably by $\chi^2$ ($P > 0.10$, $P > 0.25$, and $P > 0.50$, for R-16, R-17, and R-18, respectively). Therefore, these distributions were combined and used to represent the overall terminal rate.

The overall IRT distributions, and the related conditional probability curves are shown in Fig. 3, along with estimates of variability

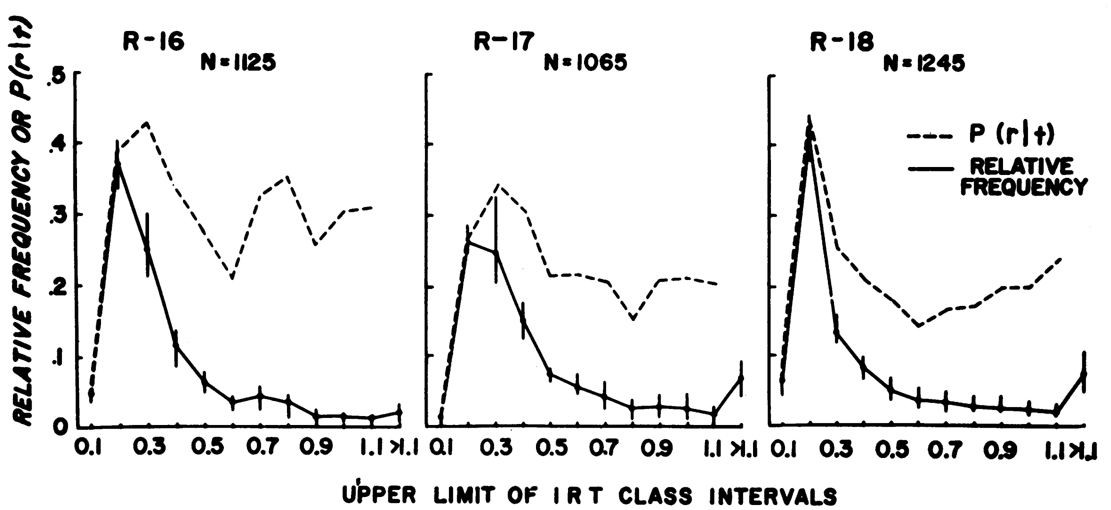


Fig. 3. Relative frequency and P(r|t) for each subject at the terminal rate on FR 60; the number of IRTs included for each rat are shown in the panels.

similar to those of Fig. 1. The greater variability reflects the fact that the sub-distributions are based on a smaller sample of performance. In comparing these curves with those displayed in Fig. 1, it should be noted that the abscissa here is in 0.1-sec units. In every case, conditional probability of response reaches a peak in the second or third class interval, and then declines.

The analysis of sequential patterns in these data was carried out in a manner similar to that described for the fixed-interval data. The heavy concentration of IRTs in the first few class intervals made it necessary to restrict to three the number of "preceding IRT" categories. As with the interval data, a marked heterogeneity was found among the resulting distributions. The hypothesis that they were of the same form was tested by $\chi^2$, and rejected ($P < 0.001$ for each subject).

The distributions that resulted are shown in Fig. 4. For R-18, $P(r|t)$ rises sharply to a peak at 0.2 sec following the shortest IRTs. Following the longest IRTs, $P(r|t)$ maintains a relatively constant value. The curve describing responding following IRTs of intermediate value (0.3 sec) lies between the other two functions. The analysis of responding by R-17 produced a similar pattern of results. In the case of R-16, the three distributions differed reliably from one another, but there was no obvious separation of the overall distribution into sharply peaked and relatively constant components.

## DISCUSSION

The terminal rates examined in the present study were "stable" in the ordinary sense of the term: neither the mean IRT (average rate) nor the overall distribution of IRTs changed during any part of the period examined. Because IRTs were sequentially patterned, however, the overall distribution did not reflect the momentary status of response tendencies. Thus, the probability of a short IRT might fluctuate from high following a short IRT to low following a long one, while the overall distribution would indicate only a moderate probability. By carrying out a sequential analysis, it was possible to separate the overall IRT distribution into component distributions which provided a more detailed description of momentary tendencies to respond.

The component distributions were not only different from the overall distribution, but they were related to each other in an orderly fashion. As the duration of the preceding IRT increased, the component distributions appeared to undergo a steady transformation of shape, from one form following the shortest IRT to another following the longest. This progression of forms is most clearly shown in the data from R-14 and R-15, where a relatively large number of preceding IRTs permits the sequence of progressively rising distributions to be seen in detail. Although based on a more limited number of component distributions, the data from R-17 and R-18 also

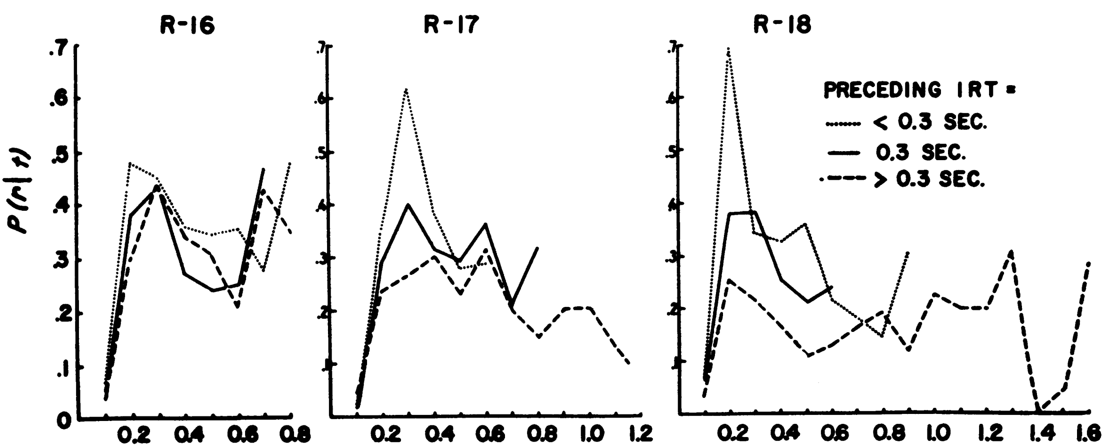


Fig. 4. Sequential analysis of terminal rate data presented in Fig. 3. The conditional probability of response is shown as a function of duration of preceding IRT.

demonstrate an orderly arrangement. In addition to the internal ordering, there was a relationship between the component distributions following longer IRTs and the schedule contingencies in force. Thus, following longer IRTs, the curves for R-14 and R-15 were increasing: a form of this sort is consistent with the tendency of interval schedules to reinforce longer IRTs with a higher probability than shorter IRTs. Under the ratio schedule, where reinforcement probability is independent of IRT duration, there were no generally increasing component curves. In this case, the curves were either sharply peaked (following the shortest IRTs), or relatively flat (following the longest IRTs). These observations suggest that the component distributions describe orderly aspects of response rate which are not apparent from inspection of the overall distribution.

The success of the sequential analysis in producing a series of orderly and distinctive component distributions suggests that the overall rate was a complex structure produced by two or more variables which operate in succession. As an illustration of how a sequential analysis would isolate the independent operation of successively controlling variables, suppose that an overall rate is controlled by two variables, one of which generates a preponderance of long IRTs, and the other, a preponderance of short IRTs. Under these circumstances, the occurrence of a long IRT would indicate temporary control by the first variable. If the first variable continued in effect for more than one IRT, the distribution of IRTs following a long IRT would largely reflect the isolated operation of this variable. Similarly, the occurrence of a short IRT would reflect temporary control by the second variable, and the subsequent IRT distribution would primarily reflect its isolated operation. In general, a sequential analysis would separate the effects of each variable, depending (a) on the extent to which some IRTs are predominantly related to each variable and thus "tag" its operation; and (b) on how long each variable remains in force without interruption. If IRTs of some value were produced equally by both variables, the distribution following an IRT of that length would be of a form intermediate between the two basic distributions. Applied to the present data, the sequential analysis generally seemed to produce two distinct distributions, and a series of intermediate forms. It seems reasonable to suppose, then, that the overall rate was an aggregated structure produced by several variables operating in succession.

The complex structure of rate suggested by this interpretation is in line with recent analyses by Morse (1966) and Norman (1966) regarding the composition of response rate. Morse stated as his central thesis that "most schedule-controlled responding results from the joint operation of the differential reinforcement of IRTs and the generalized effects of reinforcement to strengthen responding". The component distributions produced by the sequential analysis were in line with this point of view. Components following longer IRTs seemed to reflect the way interval and ratio schedules reinforce IRTs of different lengths. Other, more sharply peaked, component distributions were localized around particular IRT values, and may reflect a tendency to press at a rate determined directly by the reinforcement of pressing itself, independent of any selective effect on IRTs. In this regard, it is interesting to note that Blough (1966) found short IRTs to be relatively insensitive to differential reinforcement procedures, whereas long IRTs were readily influenced.

The identification of the specific variables which determine the composition of response rate requires a great deal of further experimental analysis. If such variables operate in succession, as the present results seem to indicate, statistical "compromise" rules may prove unnecessary, and a detailed analysis of response rate would be considerably simplified. If validated in further work, the concept of successive control could provide a useful method for isolating and identifying the variables which underlie response rate.

In the fixed-interval portion of this study, $P(r|t)$ curves over the first five or six class intervals often tended to increase. The overall curves of R-13 and R-14, and the curves following longer IRTs for R-14 and R-15 illustrate this trend. The increase is consistent with Skinner's early notion that the emission of IRTs on interval schedules is related to their probability of reinforcement, which is higher for longer IRTs. Anger (1956) examined $P(r|t)$ curves of rats pressing on a variable-interval schedule, and found decreasing functions. As a result, he rejected the hypothesis that probability of reinforcement controls the emission of IRTs of different durations. On the basis of a similarity of shape between hourly rate of reinforcement and the P(r|t) curves, Anger explored the hourly rate variable, and found it to be effective. Although hourly rate of reinforcement was confounded with probability of reinforcement in these later studies, the probability variable had already been dismissed, and Anger concluded that the effectiveness of hourly rate had been demonstrated.

In comparing Anger's results with those of the present study, many procedural differences are apparent, *e.g.*, the use of variable- rather than fixed-interval reinforcement. There are also major differences, however, with regard to data analysis. In the Anger study, IRTs were recorded in 4-sec class intervals, with the result that 50 to 80% of all IRTs fell into the first class interval. Because these IRTs could not be differentiated, trends within this major part of Anger's data could not be ascertained. If the overall IRT distribution reflects the operation of more than one variable (however they combine), it is possible that some parts of an overall distribution are more sensitive to the effects of a particular variable than are other parts. It is possible, then, that increasing P(r|t) functions, consistent with the probability-of-reinforcement formulation, were also present in Anger's data, but were obscured by the relatively large class intervals. A second difference in data treatment lies in the sequential analysis, which was not carried out with Anger's data. It is impossible to ascertain whether component distributions consistent with the probability-of-reinforcement hypothesis were contained in Anger's overall results. While these differences in data analysis preclude a close comparison of Anger's and the present results, they may account for Anger's failure to find evidence that probability of reinforcement influences the IRT distribution. Although a final answer must rest with further experimentation and analysis, the present results strongly suggest that the issue be considered further.

A report by Kintsch (1965) describes statistically reliable sequential dependencies in VI and VR responding, but apparently no orderly component patterns were observed. Kintsch noted that the training his rats received was "limited". Perhaps at the early stage of training Kintsch observed, separate sequential control had not been well established. Had training been continued, more distinctive patterns corresponding to those observed in this study might have emerged.

Further development of the implications of this study demands a broader empirical base. It is important, for example, to separate the influence of a particular apparatus or manipulandum from the more general effects of reinforcement schedules. Should further work continue to reveal a sequentially patterned rate structure of the kind encountered here, the concept of successive control, and the methods it suggests, may prove valuable in analyzing the several factors that control response rate.

## REFERENCES

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

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

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

Kintsch, W. Frequency distribution of interresponse times during VI and VR reinforcement. *Journal of the Experimental Analysis of Behavior*, 1965, **8**, 347-352.

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

Norman, M. Frank. An approach to free-responding on schedules that prescribe reinforcement probability as a function of interresponse time. *Journal of Mathematical Psychology*, 1966, **3** (2), 235-268.

Revusky, S. Mathematical analysis of the durations of reinforced inter-response times during variable interval reinforcement. *Psychometrika*, 1962, **27**, 307-314.

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

*Received 2 May 1966.*