It simply comes down to this. If a woman with a higher degree is really career oriented and we think she will spend all or at least most of her time working, we'll hire her; if not we won't.

Another commented:

In general, length of service is less for women at the upper levels than for men but there are many exceptions. We hired some Ph.D. women 15 years ago who are still with us. We hired several real fireball Ph.D. men 2 years ago, and they have already left us.

One of the persistent needs reported by women chemists who have left the labor force for the "3 M's" is for part-time employment (*2*). Almost no part-time job opportunities exist in these 65 laboratories, nor are they likely to become available in the future. About 80 percent of the firms said they never hired women (or men) on a part-time basis. The three reasons they most commonly gave were that (i) space and equipment were limited; (ii) most jobs require continuity and cannot be broken into small parts; (iii) too many tasks require continuous supervision, consultation, or teamwork to make part-time work practicable. Yet there were a few notable exceptions. In one large, well-known firm employing about 50 women with B.S. degrees in chemistry, eight of the women were on part-time assignment, working either three 8-hour days a week or 4 hours a day, 5 days a week. This part-time program, resulting from a special need, had been in operation for 5 years with satisfactory results.

What is the employment outlook in research laboratories for women with training in chemistry? Over three-fourths of the firms said "very favorable." With the smaller firms excluded, the favorable rating was virtually 100 percent. It appears to me that this favorable outlook by laboratory directors rests on two assumptions: (i) the total number of research jobs will expand as rapidly in the next 5 years as in the past 5 years, during which laboratories employed an ever-increasing proportion of the nation's rapidly expanding professional, technical, and scientific force (*3*); (ii) beginning jobs will be increasingly available to women in the large-scale laboratories as men continue to upgrade their training and enter chemistry above the B.S. level (*4*).

JOHN B. PARRISH
College of Commerce and Business Administration, University of Illinois, Urbana

References and Notes

1. The selected laboratories are those that were largest in professional laboratory personnel among the 141 members of the Pharmaceutical Manufacturers' Association as of 9 April 1962. Personnel data from *Industrial Laboratories of the United States* (National Research Council, Washington, D.C., ed. 12, 1960).
2. J. B. Parrish, *J. Chem. Educ.* **41**, 506 (1964).
3. ———, *Res. Develop.* **15**, 32 (1964).
4. Three indicative studies may be cited. A report on 1163 women who received B.S. degrees in chemistry in 1934–39 revealed only 2 percent employed in industrial research [E. L. French, *J. Chem. Educ.* **16**, 576 (1939)]. My forthcoming study of women with B.S. degrees in chemistry from a number of large Midwestern universities in 1958–63 will report 28 percent employed in industrial research laboratories. The U.S. Department of Labor, in "Manpower Resources in Chemistry and Chemical Engineering," *U.S. Dept. Labor Bull. No. 1132* (1953), p. 17, reported that 70 percent of all employed women chemists 40 years of age or over had at least the M.S. degree, and that only 25 percent of those under 40 had at least the M.S. degree.

1 March 1965

Reinforcement Schedule Generated by an On-line Digital Computer

Abstract. *A LINC digital computer was used to generate an autoregressive schedule of reinforcement. On such a schedule the probability of reinforcement is a function of the similarity in duration of the intervals between successive responses. A detailed analysis of the data obtained from monkeys on this schedule demonstrated two distinct tendencies in their behavior: a tendency for periodic response and a tendency for serial dependence between successive interresponse times.*

We now know a great deal about how to shape and maintain operant behavior by the process of differential reinforcement. The most elegant and searching manifestations of this process have been provided by experiments on schedules of reinforcement (*1*). The reason is that reinforcement schedules permit one to make an extremely significant transformation; they make it possible to convert complex patterns of stimuli and responses into serial and temporal ones, much as the nervous system is thought to transform information into temporally coded patterns (*2*).

Digital computers add new dimensions to the study of reinforcement schedules. They enable us to fabricate schedules that are impossible to program by conventional methods and then facilitate detailed observations of the fine structure of behavior.

An example that reveals the potential of computers is the problem of differentially reinforcing low variability in response rate. Although with conventional equipment we can devise schedules which produce low variability—for example, "Differential reinforcement of low rate with limited hold" (*1*)—the low variability is only a by-product of the contingencies. In this report we describe a computer-generated schedule which controls variability directly, and we demonstrate how a detailed analysis of the data reveals the multiple effects of the schedule. Both the programming and the analysis were performed by a small, high-speed LINC computer (*3*).

Our "autoregressive reinforcement schedule" is a stochastic schedule that takes its name from a category of time series in which successive observations depend on a function of the previous term or terms plus a random additive error (*4*). It differentially reinforces low variability in response rate by promoting consistency in the intervals between responses, or the interresponse times. It specifies that the probability of reinforcement of the response that terminates an interresponse time I_i, depends on the similarity of I_i and the previous interresponse time, I_{i-1}. The closer the similarity, the greater the probability of reinforcement. The organism's own behavior then becomes the basis of the response-reinforcement correlation.

To determine whether the conditions for reinforcement are met, the program first computes the quotient of the two successive interresponse times, always placing the larger value in the numerator. The quotient corresponds to a number, equivalent to a certain *p*-value, in a table stored in the computer memory. A random number is then generated and compared with the table entry. If the table entry is greater, reinforcement is programmed. The function relating probability of reinforcement to the quotient appears in Fig. 1. Probability ranges from 0 to 1.0 on the ordinate and the quotient ranges from 1.0 to 1.4 on the abscissa. If $I_i = I_{i-1}$, $p = 1.0$. If the quotient exceeds 1.4, $p = 0$. A quotient of approximately 1.05 is equivalent to $p = 0.5$.

Three female monkeys, one *Macaca speciosa* and two *M. nemestrina*, were placed in primate restraining chairs and subjected to the schedule. A response was defined as depression of a lever far enough to actuate a microswitch. A fruit drink (0.1 ml) was used as reinforcement. Each experimental session lasted 90 minutes. During preliminary training sessions, each press

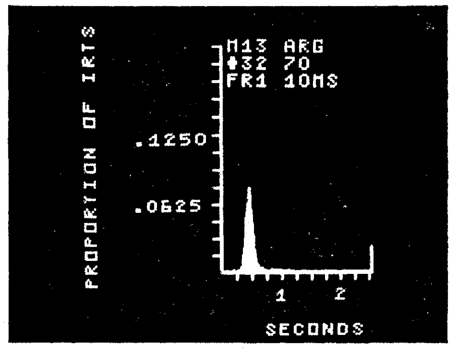

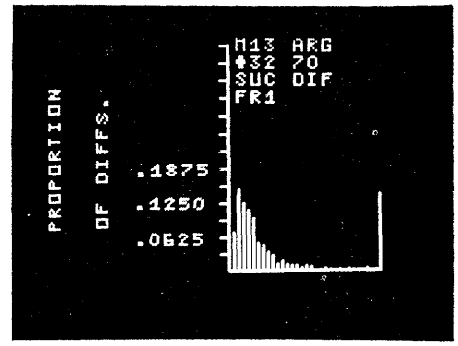

Fig. 2 (left). Distribution of interresponse times on the autoregressive reinforcement schedule. Abscissa markers are placed at 250-msec intervals. $N = 2048$. Fig. 3 (right). Distribution of differences between successive interresponse times. Successive bars represent differences of 0, 10, 20, and 30 msec, and so forth. The last category represents all differences greater than 320 msec. $N = 1032$.

of the lever was reinforced. The monkeys were then subjected to the autoregressive reinforcement schedule and their behavior rapidly came under its control. Figure 2 shows a histogram of 2048 successive interresponse times for one monkey. This histogram is a photograph from the LINC oscilloscope, as are the succeeding figures. The distribution appears Gaussian in shape and rather narrow, the modal interresponse time being about 450 msec. The data for this monkey were intermediate in variability, compared to those for the other two monkeys. Such a narrow distribution demonstrated clearly that the contingencies of the autoregressive reinforcement schedule had considerable effect on the consistency of the interresponse times. Figure 3 shows the distribution of the differences between successive interresponse times; starting at the left, each bar represents differences between successive times of 0, 10, 20, 30 msec, and so forth. The modal difference was 10 msec, or 1 clock pulse of resolution.

The relatively narrow distribution of interresponse times in Fig. 2, and the distribution of successive differences in Fig. 3, suggest that the monkeys were responding almost periodically. The periodicity is demonstrated vividly by plotting a function called the Expectation Density, a term proposed by Huggins (*5*). This function gives the expected number of occurrences of an event within time interval t_i following a known occurrence at $t = 0$. It is computed by letting successive instances of the event—here a response—take on the value of $t = 0$. The distribution of later occurrences is given as an average over the ensemble of possible signals.

An Expectation Density plot is displayed in Fig. 4. The abscissa represents a total period of 2.5 seconds. The first peak lies at approximately 450 msec, which corresponds to the histogram for interresponse times in Fig. 2. The steep slopes of this peak indicate that, given an event at $t = 0$, it is most likely that the next event will occur about 450 msec later, within rather narrow limits. The second peak indicates that, given an event at $t = 0$, the second succeeding event is most likely to occur about 900 msec later. The fact that the second peak is flatter indicates that the distribution of times between every other response is more variable than the distribution of interresponse times, a conclusion difficult to reach merely by inspecting the appropriate scaled histograms. Since the peaking becomes less marked with distance from $t = 0$, and approaches a steady-state expectation of occurrence at the right-hand side of the plot, the process cannot be purely periodic. A pure periodic process would produce line spectra.

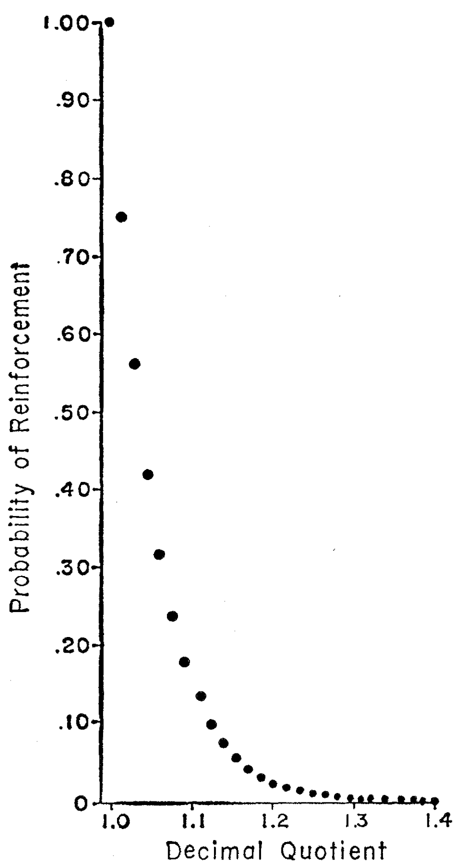

Fig. 1. Function relating probability of reinforcement to the quotient of two successive interresponse times.

Although the autoregressive reinforcement schedule generates nearly periodic responding, such a distribution in time might arise from a truly periodic process subjected to random perturbations. But, as McGill points out (*6*), such a process would be self-compensating; long intervals should tend to be followed by short ones and vice versa. The autoregressive reinforcement schedule, however, differentially reinforces just the reverse tendency. To determine the extent to which it maintains such behavior, we must examine the behavior as a time series —that is, we must look for serial dependencies.

One way to do this is to examine the data serially. In Fig. 5 we show a sequence of 128 interresponse times. Succession is from left to right and duration is represented as height. In order to smooth the data, we clipped the higher amplitudes and filtered the series once by a moving average (*7*) where $x_n = (x_{n-1} + 2x_n + x_{n+1})/4$. The smoothed function is not an unjust reflection of the raw data. It suggests that the interresponse times drift up and down about the mean and that the period of this drift is relatively long.

One other way to examine such dependencies is shown in Fig. 6, which shows the distribution of sequences of interresponse times above and below the median value. The plots are arranged as follows. The top set of traces represents sequences of four successive interresponse times; those times below the median are assigned the value 0. Those above are assigned the value 1. Then, reading from left to right, we show sequences equivalent to the binary representations of the decimal numbers 0 to 15, or 0000 to 1111. Thus, the trace at the extreme left represents the number of occurrences of the sequence 0000; the next, the number of occurrences of the sequence 0001; and the trace at the far right, the number of occurrences of the sequence 1111. The middle set of traces represents sequences of three, that is, binary representations of 000 to 111. The bottom set represents sequences of two interresponse times equivalent to 00, 01, 10, and 11 as one proceeds from left to right.

Intervals between responses that are close to one another tend to be similar in value; indeed, the most common sequences are those of interresponse times in the same category. This display also demonstrates that the interresponse times drift up and down around a central value, since the mean value remains relatively stable (Fig. 5). We can also determine departures from randomness in sequential observations by calculating the number of runs in a series. If a run is defined as a sequence of ones or zeros bounded by the opposite symbol or no symbol, the runs test (*8*) tells us whether, for a given number of observations, the number of runs obtained could have arisen by chance. In the present case, the number is lower than it would be if it arose by chance.

A way to examine serial relationships that makes greater use of the available resolution is to plot serial correlations, that is, the correlation of I_i and I_{i+L}, where L is the lag between successive interresponse times. A serial correlation plot for $L = 1$ is shown in Fig. 7. Despite the narrowness of the distribution, a definite positive relation is evident; the shorter interresponse times tend to be followed by short ones, and the longer interresponse times by long ones.

It is clear that serial dependencies do exist in our data. One way to assess their contribution to the temporal patterning of the behavior is to dismantle the original sequence of interresponse times and then reassemble it according to a random process (*2*). The results of such a random shuffle are given in Figs. 8 and 9.

Figure 8 shows the distribution of successive differences in the shuffled sequence. Compared to the original data (Fig. 3), the distribution is flatter and broader, indicating that the distribution in Fig. 3 does indeed arise from serial dependencies. Figure 9 shows an Expectation Density plot derived from the shuffled data. When this plot is compared with that of Fig. 4, it becomes apparent that the periodicity in the latter is due only to the overall distribution of interresponse times, and not to their serial relationship. In fact, eliminating the serial relationship makes the function slightly more periodic.

The behavior generated by an autoregressive reinforcement schedule appears, therefore, to be governed by two components. One is a tendency of the monkey to respond periodically. The second component, which may be viewed as superimposed upon the first, is a process more directly controlled by the schedule contingencies. It is reflected by the serial dependencies in the data. Both tendencies are responsible for the low variability observed.

We have displayed many features of the behavior generated by the autoregressive reinforcement schedule in order to demonstrate the ability of the digital computer to perform a detailed analysis of behavior. Although it is apparent that the various analyses and displays greatly overlap, each one views

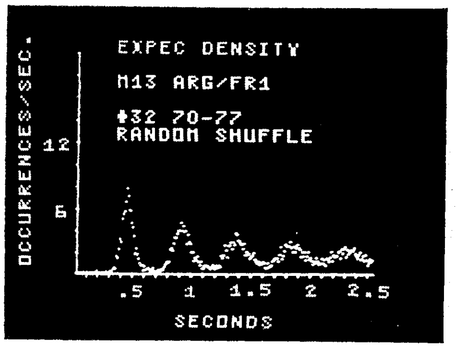

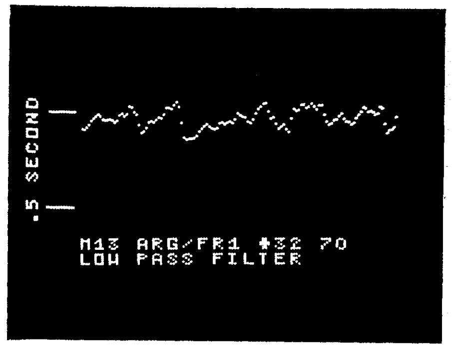

Fig. 4 (left). Expectation Density. Abscissa markers are placed at 100-msec intervals. Each ordinate marker represents a rate of two occurrences per second. $N = 2048$. Fig. 5 (right). Sequential distribution of 128 consecutive interresponse times. Time runs left to right, and duration is given by height. The bottom marker represents zero and the top marker represents 0.5 second.

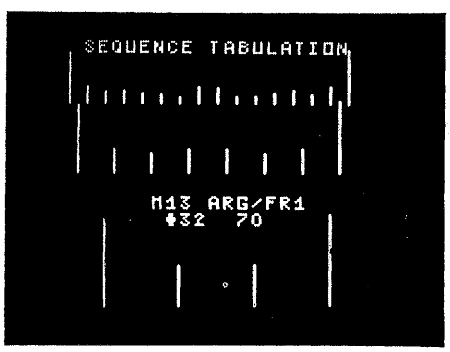

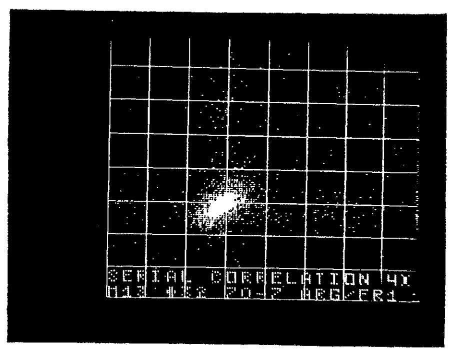

Fig. 6 (left). Tabulation of sequences of interresponse times above and below the median ($N = 512$). Top, sequences of 4; middle, sequences of 3; bottom, sequences of 2. Fig. 7 (right). Serial correlation plot (lag 1) based on 2048 consecutive interresponse times. Duration of I_i is given by the abscissa; duration of I_{i+1}, by the ordinate.

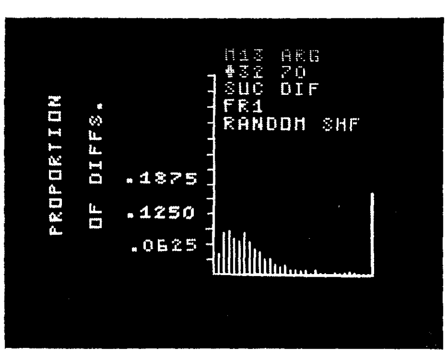

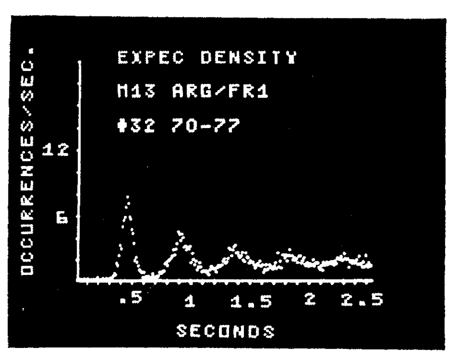

Fig. 8 (left). Distribution of successive differences after a random shuffle of order (compare Fig. 3). Fig. 9 (right). Expectation Density function after a random shuffle of order (compare Fig. 4).

the data from a slightly different point of vantage and contributes some unique information. It also is apparent that even a relationship between behavior and its consequences as conceptually simple as the autoregressive reinforcement schedule produces behavior whose finer and more critical details cannot be appreciated without the assistance of computer technology (*9*). Perhaps it is even more important to note that programming such a schedule would have been impossible with the facilities we now possess in most laboratories.

BERNARD WEISS
VICTOR G. LATIES

Department of Pharmacology and Experimental Therapeutics and *Department of Psychiatry, Johns Hopkins University School of Medicine, Baltimore, Maryland*

References and Notes

1. B. F. Skinner, *The Behavior of Organisms* (Appleton-Century-Crofts, New York, 1938); C. B. Ferster and B. F. Skinner, *Schedules of Reinforcement* (Appleton-Century-Crofts, New York, 1957).
2. G. F. Poggio and L. J. Viernstein, *J. Neurophysiol.* **27**, 517 (1964); R. W. Rodieck. N. Y. S. Kiang, G. L. Gerstein, *Biophys. J.* **2**, 351 (1962). We are indebted to Dr. Poggio for numerous enlightening comments.
3. W. A. Clark and C. E. Molnar, *Ann. N.Y. Acad. Sci.* **115**, art. 2, 653 (1964).
4. G. Tintner, *Econometrics* (Wiley, New York, 1952); M. G. Kendall, *The Advanced Theory of Statistics* (Griffin, London, 1946), vol. 2.
5. W. H. Huggins, *Proc. I.R.E. (Inst. Radio Engrs.)* **45**, 74 (1957).
6. W. J. McGill, *Psychometrika* **27**, 3 (1962).
7. R. W. Hamming, *Numerical Methods for Scientists and Engineers* (McGraw-Hill, New York, 1962).
8. P. G. Hoel, *Introduction to Mathematical Statistics* (Wiley, New York, ed. 2, 1954).
9. B. Weiss, *IRE (Inst. Radio Engrs.) Intern. Conv. Record* **9**, 103 (1962).
10. Supported in part by grants MH-03229, MH-08353, and MH-07498 from the National Institutes of Health. We thank Louis Siegel for technical assistance.

7 January 1965

Observable Changes of Hypotheses under Positive Reinforcement

Abstract. *In mathematical models of concept learning it has consistently been assumed that positive reinforcement cannot lead to a change of the hypothesis determining the overt response. When hypotheses are experimentally identified and recorded along with positive and negative reinforcements of stimulus-response pairs, it can be shown that hypotheses may change after a positive reinforcement. Positive reinforcement has an information content for subjects that has not yet been adequately recognized in concept formation studies.*

Various stochastic theories of learning account for the classification or conceptualization of behavior in terms of hypothesis testing. The learning process is described as follows: the subject randomly samples from a population of hypotheses, and when a stimulus is presented he makes the response determined by the hypothesis sampled. In all hypothesis models proposed so far, it has been assumed that the same hypothesis is kept until information is received which infirms it (*1*). As a consequence, the probability is zero that a subject will shift from an incorrect hypothesis to the correct hypothesis on correct trials. The critical assumption of "no-change-if-no-error" has not been checked directly. Indeed, in most of the experiments designed to test these models of learning, it has not been possible to identify the hypotheses used on successive trials. Levine (*2*) attempted such identification by inferring hypotheses from sequences of responses given on blank trials (that is, without outcomes). However, since his procedure did not rule out the possibility that the hypotheses could change in the course of the test responses themselves, identification of the hypotheses could be impaired. This objection may be important, as suggested by the following results.

Three experiments were conducted with three different groups of college students. In experiment I, each subject was given a set of eight white cards; on each card was printed a string of three letters, each letter being either D or K, with all possible combinations represented once in the group of cards. The subject was requested to classify the cards correctly by placing them in two columns, headed by a pink and a blue label, respectively. After the subject had classified the cards once without any clue concerning the "correct" criterion, the experimenter replaced one of the white cards by a colored card, either pink or blue, having exactly the same pattern of letters as that on the white card removed. The subject was instructed to place this card in the group to which it belonged (identified by the color) and, furthermore, to make as many changes in the placement of the other cards, or none, as appeared necessary in order to achieve the correct classification of all cards. When the subject indicated that he was satisfied with his new classification, which might be the same as the previous one, he was given a second colored card. The presentation of each additional colored card started a new trial. The procedure was repeated until all white cards were replaced. As the subject kept all cards before him, the "correct" classification, defining the conceptual problem to be solved, was completely shown after eight trials. In this experiment the problem was the same for all subjects. It is described by the information cards which were given in this fixed order: DKD(blue), KKK(pink), DDD(pink), KDK(blue), DKK(pink), KDD(pink), DDK(blue), KKD(blue).

In experiment II, exactly the same cards and essentially the same procedure were used as in experiment I. But the subjects were given successively six problems to solve, with the same cards, ranging from the simplest (one-dimensional) problem to complex ones (disjunctive three-dimensional). The order of the problems and the sequence of information were constant.

In experiment III, the patterns to be classified included strings of 1, 2, 3, 4, and 5 letters, all 62 possible combinations of D's and K's being represented once. The strings were typed together on single sheets of paper, in fixed order for the individual, and in randomized order for the group of subjects, one sheet being used per trial. The patterns were to be classified by either circling or crossing the strings of letters. After each classification the subject selected one string of letters for which the experimenter indicated the "correct" response (circling or crossing). Each new sheet contained the past as well as the new information given by the experimenter. The single problem, given to all subjects, resembled that of experiment I: circling was the "correct" response requested for all strings of letters ending in DD or KK. Except for the differences in the experimental procedure, the instructions were similar to those given in experiments I and II.

Thus, the three experiments had the following common features: (i) on each trial the subject made a binary classification of the entire set of stimulus patterns (strings of letters); (ii) on each trial the classification of a single stimulus card was reinforced; (iii) the subject had before him a complete record of past information and was prevented from making any classification response inconsistent with this information.