Can computers answer behavioral questions?*

BERNARD WEISS
University of Rochester, School of Medicine and Dentistry
Rochester, New York 14642

One of the most enduring, though time-consuming, pleasures of a scientific research career is showing visitors through the laboratory. Not only do visitors often help brighten those tedious, harrassing days we sometimes encounter, but by their grateful and puzzled expressions confirm the profundity of our work. I have found the presence of a computer in the laboratory, together with a short description of its operation vigorously embellished with technical jargon, a superb vehicle for eliciting such a response from visitors. I can think of only one other activity that is as reinforcing.

From time to time I have wondered, in the intervals between visitors (the IVI), whether computers in the behavior laboratory could serve some other function. During one such episode I conceived the project of a book comprised of chapters contributed by a variety of investigators who had survived the manic-depressive illness often associated with the installation of a computer. This project, I am glad to say, is finally coming to fruition. *Digital Computers in the Behavior Laboratory* (Weiss, in press) is in page proof, and I will be happy to take orders for it after this address. I also hope it forever eliminates papers which focus on pictures, say, of how the interface cabinet is painted to color-coordinate with the console. We have passed the stage of mass occupational therapy and should now proceed to use this superb technology to answer important questions about behavior. It has played such a role for me, I believe, as I will try to demonstrate.

TRANSITION STATES

When a stably performing organism is shifted from one reinforcement schedule to another, or from one schedule parameter to another, it takes some time for the performance to restablize according to some criterion such as rate. The intervening period, when performance is constantly changing, has been called a transition state. It formally is equivalent to the conventional learning paradigm except that the original and final baselines are more sharply defined and freer of extraneous variables.

It has always seemed to me more sensible to study this process in a schedule framework than in the conventional learning situation because of these factors, but it has not often been done. A major reason for the relative neglect of transition states is technological. I attribute a large part of it to the constraints of the earlier technology. If you plan to study well-modulated transition states, you need the ability to put the behavior under a microscope. Without computer technology, such a microanalysis is impossible. Let me illustrate by examining certain features of transitions from continuous reinforcement to spaced responding, to variable-interval performance, to fixed-ratio performance.

Spaced Responding

Many different kinds of organisms can be trained to separate successive responses (defined by some specified event such as a microswitch closure) by specified intervals of time (Weiss & Laties, 1967). Such spaced responding schedules, often called DRL for differential reinforcement of low rate (Ferster & Skinner, 1957), also have been and are still widely used for drug studies. Our interest in drug effects led us to examine more closely both the development and maintenance of spaced responding performance (Weiss, 1970a; Weiss, Laties, Siegel, & Goldstein, 1966).

The data I will show you came from three *M. nemestrina* female monkeys maintained in primate chairs and reinforced for successful responses by brief (60-msec) squirts of a fruit drink. The schedule required a minimum of 20 sec between successive responses to deliver reinforcement. No preliminary training was given with the lever. The spaced responding schedule was put into effect immediately. The experiment was controlled on-line by a LINC computer (Clark & Molnar, 1965). Three monkeys were studied at once.

Figure 1 is a three-dimensional plot which gives you a quick grasp of how Monkey M 11's performance developed over the first 30 sessions. These tracings were produced by a Calcomp digital plotter attached to the LINC. Each tracing represents a single experimental session. The x-axis represents successive sessions, the y-axis interresponse time (IRT) frequency, and the z-axis

*The preparation of this paper was supported, in part, by Grant MH-11752 from the National Institute of Mental Health, Grant NS-08048 from the National Institute of Neurological Diseases and Stroke, Grant GI-30097X from the RANN Program of the National Science Foundation, and in part by a contract with the U.S. Atomic Energy Commission at the University of Rochester Atomic Energy Project and has been assigned Report No. UR-3490-226.

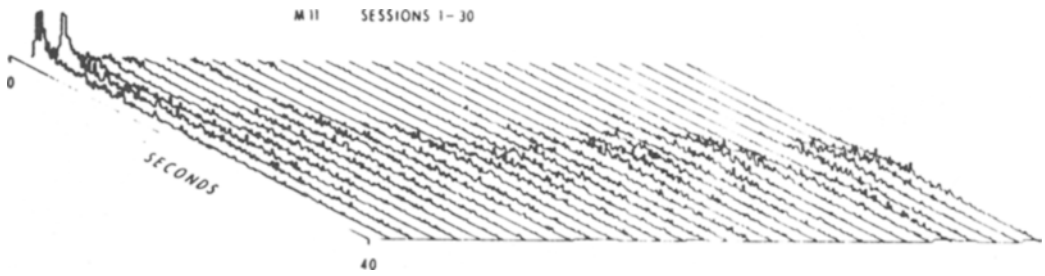

Fig. 1. Three-dimensional plot of the performance of Monkey M 11 over Sessions 1-30. Twenty seconds were required between successive responses for reinforcement. The x-axis represents successive sessions, the y-axis IRT frequency, and the z-axis the IRT in seconds (from 0 to 40). All longer IRTs fell in the last bin. Resolution is 40 msec.

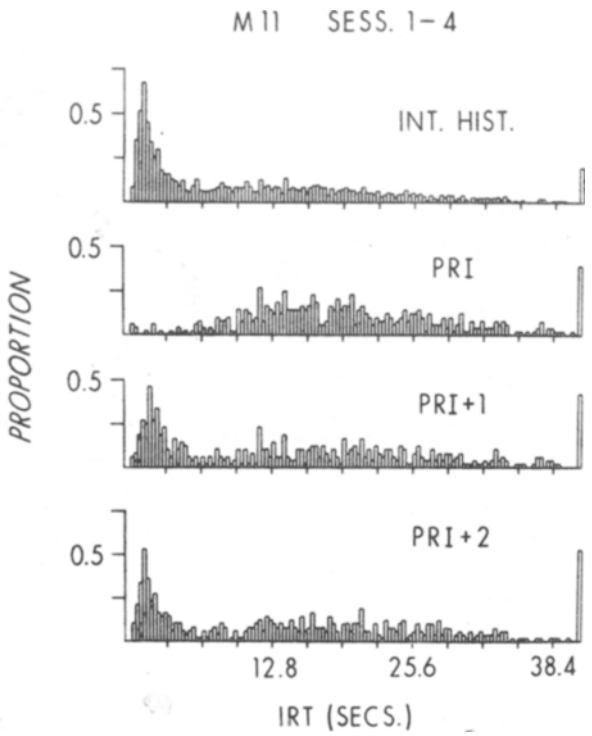

Fig. 2. Normalized histograms based on all IRTs (interval histograms), the interval directly between reinforcement and the first response afterward (PRI), the interval between the first postreinforcement response and the second (PRI + 1), and between the second postreinforcement response and the third (PRI + 2). These charts represent Sessions 1-4.

the IRT in seconds, with 40-msec resolution. Note the progression of adaptation to the schedule parameters. As training continued, the incidence of short IRTs fell and the incidence of IRTs clustered around 20 sec rose. Although this hardly constitutes a microanalysis, such a plot, by permitting an experimenter to view in one display a rather large section of an experimental history, can be a very useful guide in deciding upon further dimensions for study. It would be extremely tedious to produce, of course, without a computer-driven plotter.

Even though high rates persisted during the first few sessions, some IRTs exceeded the 20 sec required for reinforcement. It is illuminating to focus on what happened to the subsequent IRTs. Figure 2 is based on the first four sessions of the monkey whose history was drawn in Fig. 1. I have plotted here the conventional interval histogram, the postreinforcement interval (PRI) histogram, i.e., the interval between the occurrence of the response which produced the reinforcement and the next response, and histograms representing the interval between the first postreinforcement response and the second (labeled PRI + 1) and the interval between the second and third postreinforcement responses. Note that despite the short duration of the solenoid pulse which delivered the fluid (60 msec), the postreinforcement intervals tended to be longer than the interval histogram of all IRTs. Even PRI + 1 presents a somewhat different shape than does the interval histogram. Given, of course, that the PRI is longer than the bulk of other IRTs, it is

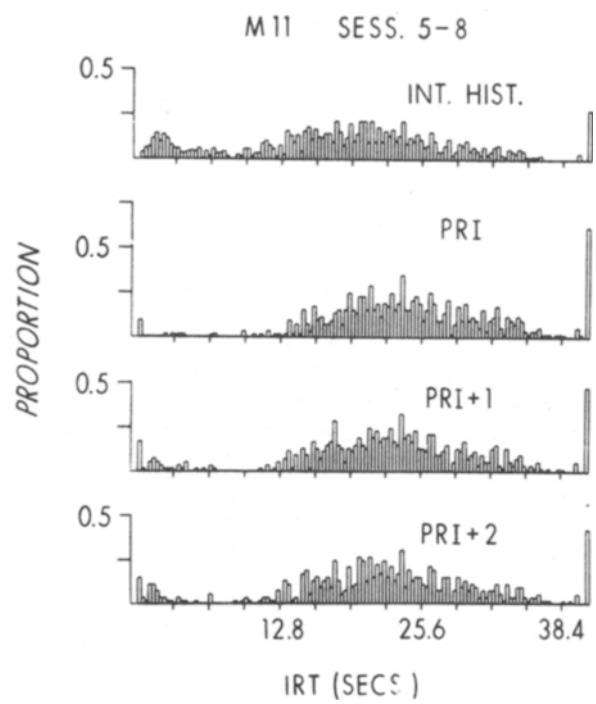

Fig. 3. Histograms depicting effects on performance for Sessions 5-8. See Fig. 2 for explanation.

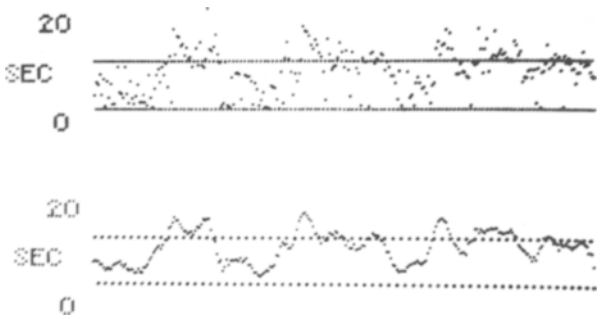

Fig. 4. A display showing 256 successive IRTs from Session 8, Monkey M 11. IRT duration is represented by height. Successive IRTs run from left to right. The bottom dotted line represents 0 sec. The upper one represents 20 sec, the minimum required for reinforcement. The upper section of the photograph shows the raw data. The lower portion shows the data after smoothing by moving average where $X_t = (X_{t-1} + 2X_t + X_{t+1})/4$.

more likely again to be followed by reinforcement. That is, reinforcements will tend to propagate, so to speak, which is one reason why a number of observers have seen trains of reinforced IRTs on spaced responding schedules.

By Sessions 5-8, as shown in Fig. 3, the conventional interval histogram shows control by the spaced responding requirement. Note again, however, that the early peak is missing in the PRI distribution and attenuated in the PRI + 1 distribution.

It seems to me that these findings give us a better understanding of the processes by which spaced responding is shaped. They suggest that the evolution of performance is not described properly simply by showing that longer intervals become more frequent. What seems to happen, as it were, is that one reinforcement leads to another. The entire character of the distribution changes.

Some additional features of spaced responding performance are most easily appreciated simply by displaying the raw data and subjecting them to various transformations. The upper part of Fig. 4 shows a train of 256 responses from Session 8 of Monkey M 11, photographed directly from the LINC oscilloscope. Note the fluctuations in IRT duration above and below the 20-sec minimum. A clearer view of this type of fluctuation is shown in the lower tracing, which was smoothed by a three-term moving average.

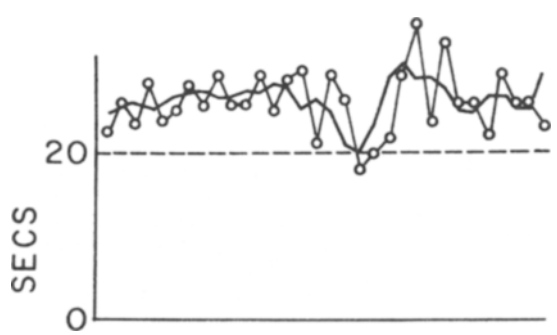

Fig. 5. IRTs (open circles) plotted against a three-term moving average. These are 32 successive IRTs from Session 30, Monkey M 11.

Smoothing is one component of a data manipulation program adapted by Louis Siegel of our laboratory from one devised earlier by A. J. Hance. It incorporates many such transformations and has proved extremely useful in what I call interactive data analysis. I use this term to describe the process of treating data from multiple perspectives, subjecting them to a variety of manipulations and burrowing, so to speak, within them. It represents a much more creative, improvisational, and demanding enterprise than does using the computer simply to print means and variances. I suppose I find it a more natural process because the experimental analysis of behavior emphasizes functional rather than statistical analysis.

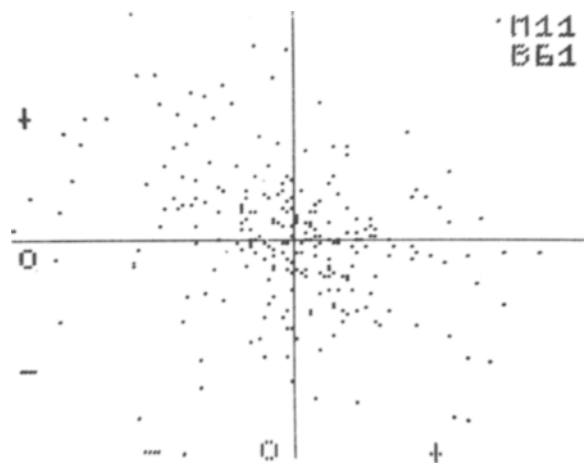

Fig. 6. Scatterplot showing successive IRT differences for Session 30, Monkey M 11 (see Fig. 5). The abscissa represents the difference between IRT_i and IRT_{i+1}, the ordinate being the difference between IRT_{i+1} and IRT_{i+2}. The scatterplot quadrants represent successions of −−, −+, +−, and ++ to show the relation between successive IRT differences. The plot comprises a range of ±10 sec.

Another example of what a microanalysis can reveal about spaced responding appears in Fig. 5. The solid, heavy line represents a three-term moving average of data from Session 30. The successive IRTs are given by the open circles. You can see here that imposed upon the overall trend is a fluctuation in IRT duration characterized mostly by alternation. The alternation can be revealed another way by plotting the values of the successive differences as in Fig. 6. The abscissa represents the difference between IRT_i and IRT_{i+1}, the ordinate being the difference between IRT_{i+1} and IRT_{i+2}. The scatterplot is divided into quadrants which represent successions of minus-minus, minus-plus, plus-minus, and plus-plus to indicate the relation between successive IRT differences. The scatterplot

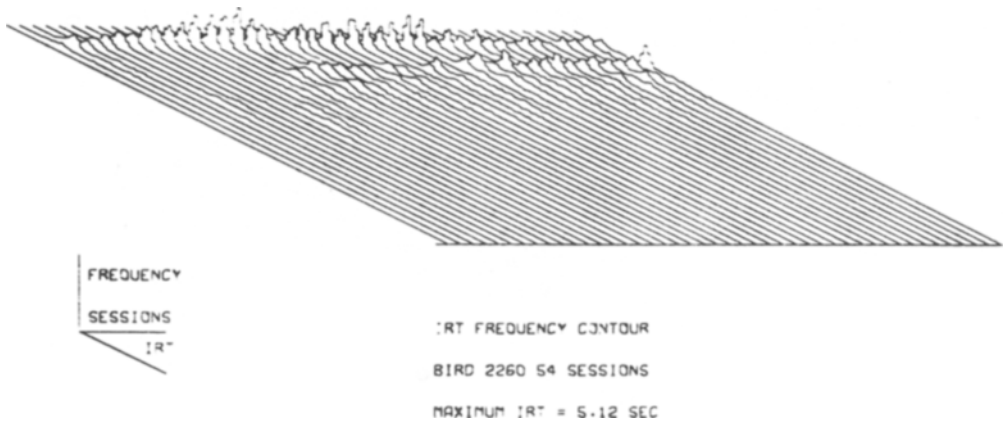

Fig. 7. Three-dimensional plot of the performance of Pigeon 2260 on VI 2.5 min over the first 54 sessions after the change from CRF.

indicates a negative correlation. That is, if Difference (i) is negative, Difference (i + 1) will tend to be positive. Again, this kind of analysis would not have been undertaken without the earlier interactive data analysis already described.

Variable-Interval Performance

The ability to store large amounts of data and to have them readily accessible has also enhanced our understanding of variable-interval performance. Variable-interval (VI) reinforcement schedules are often used as a baseline against which to measure the actions of various experimental treatments such as drugs or processes such as the conditioned emotional response. Ronald Wood, in our laboratory, has carefully examined the evolution of VI performance in pigeons using a schedule modeled after Anger's (1956), with a mean interval of 2.5 min. Figure 7 is a three-dimensional plot of the performance of one pigeon. I show it here to illustrate an important aspect of performance, namely, that when you measure IRTs grossly, you miss many of the important features of performance. The IRT distribution possesses a distinct multimodal character, and the relationships among the different modes change with duration of exposure to the schedule. An experimenter interested in how certain independent variables affect VI performance is compelled to examine these microfeatures of performance, because they may be under the control of different variables than previously assumed and may be differentially sensitive to different treatments.

Fixed-Ratio Acquisition

An even more compelling argument for the contribution of microanalysis to our understanding of schedule-controlled behavior comes from an analysis of fixed-ratio (FR) performance. As shown by Ferster and Skinner (1957) and others, a typical pattern during the transition from continuous (100%) reinforcement, or CRF, is a few ratios emitted at a high rate and without long IRTs, then a period during which there are many long IRTs, followed by a final period before stable FR performance develops during which the FR sequence shows some discernible patterning, as in Fig. 8.

The top tracing shows 10 ratios from one of nine pigeons in which we studied the transition from CRF to FR 30 (Gott & Weiss, 1972). The individual IRTs within a ratio are plotted from left to right for Ordinal

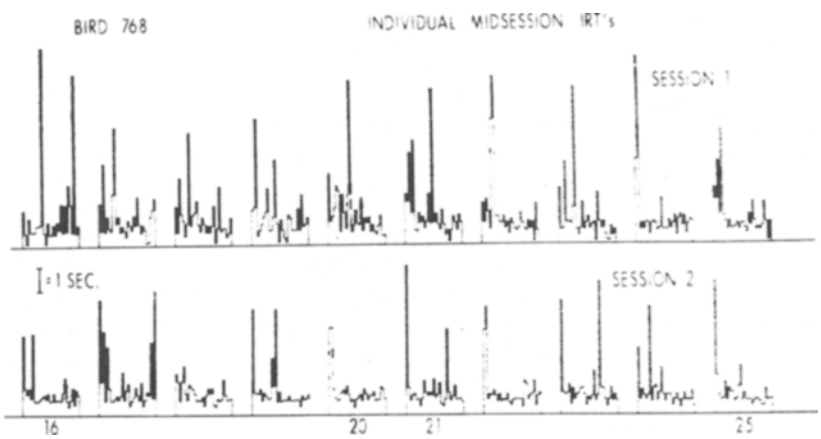

Fig. 8. Ten individual ratios from Session 1 and Session 2 after the transition from continuous reinforcement to FR 30, Pigeon 768. The numbers at the bottom show the ratio number within a session of 40 complete ratios. Each bar represents a single IRT within the ratio with successive ordinal IRTs running left to right.

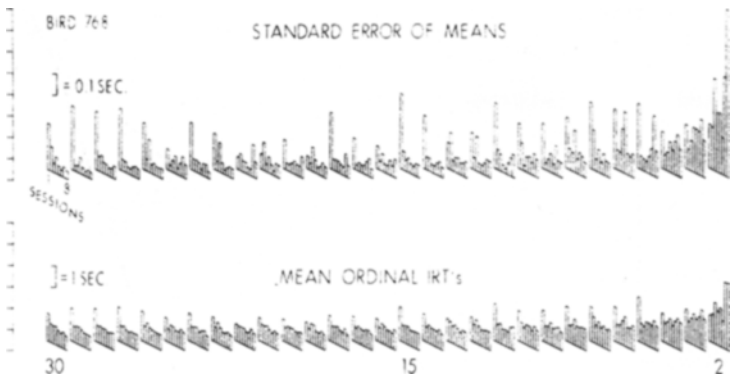

Fig. 9. Mean ordinal IRTs from IRT 2 to IRT 30 and associated standard errors of means for Bird 768, Sessions 1-8. Sessions run from 1 to 8 along the z-axis. IRTs run from 2 to 30 along the x-axis. IRT duration is represented by the y-axis.

Position 1 to Ordinal Position 30. Long IRTs appeared in all positions in Session 1 on FR 30. By Session 2 (40 ratios per session), a more distinct pattern had emerged. The first interval after reinforcement tended to be longer than the succeeding ones, while the rest of the IRTs tended to be relatively uniform.

A summary of this transition period is afforded by Fig. 9, which displays the means and standard errors of each of the IRTs from the 2nd to the 30th. In this pigeon, the IRTs early in the ratio increased while the later ones decreased in length from Session 1 to Session 8.

If one can simply examine the occurrence of IRTs greater than 1 sec long, which we have called *outliers*, a pattern emerges which accounts for a good deal of the stability of FR performance. Notice how the incidence of outliers on the first four sessions after the change from CRF to FR 30 appears greatest in the first few ordinal positions within the ratio (Fig. 10). It drops sharply after the postreinforcement interval, and from Ordinal Positions 6 to 30, especially during the later sessions, there is virtually no distinguishable difference in incidence.

Since the number of outliers decreased across the last four-fifths of the interval rather uniformly, we concluded that the terminal response (which directly preceded reinforcement) was not acting backward in time via a gradient, but was controlling some other aspect of the behavior.

That response topography is one such aspect is apparent if one observes birds responding during the transition period. Different varieties of responses are emitted. We have classified them as follows: a "nibble" is a response characterized by the bird standing close to the key with its beak opening and closing rapidly, especially near the edge of the opening through which the key is pecked. This topography typically produces trains of short IRTs shown in Fig. 11 at about 120 msec.

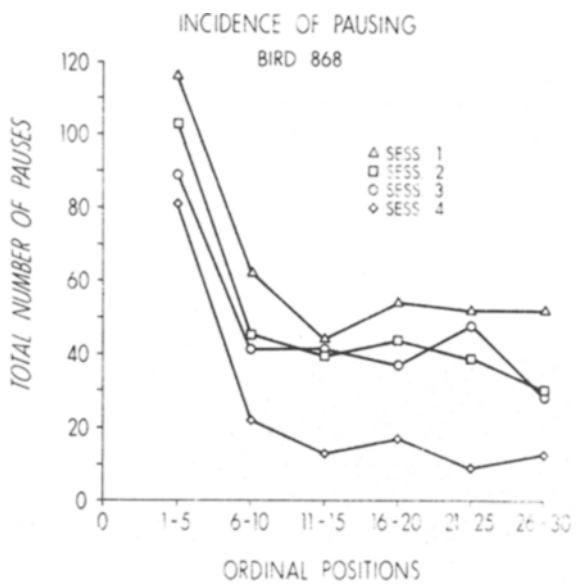

Fig. 10. Incidence of outliers (pauses > 1 sec) for Sessions 1-4, Pigeon 868. Ordinal position is represented along the abscissa as groups of five IRTs.

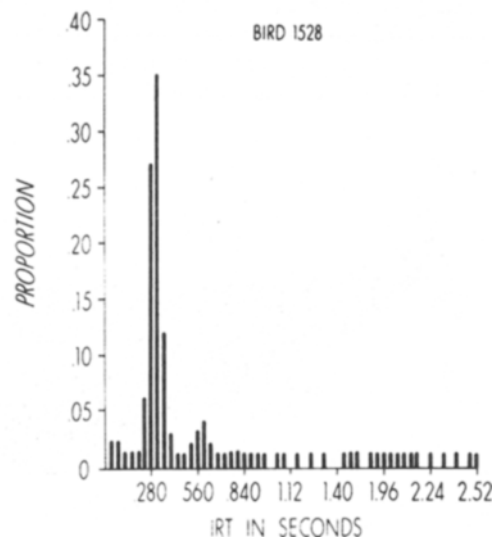

Fig. 11. IRT histogram showing the distribution of IRTs for Pigeon 1528. Session 8. IRT resolution was 40 msec.

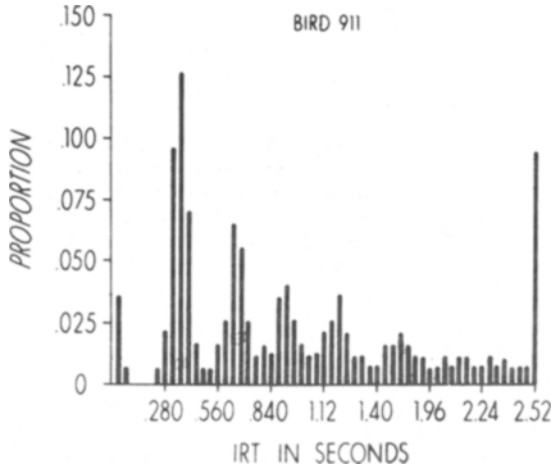

Fig. 12. IRT histogram for Bird 911 on Session 8 (see Fig. 11).

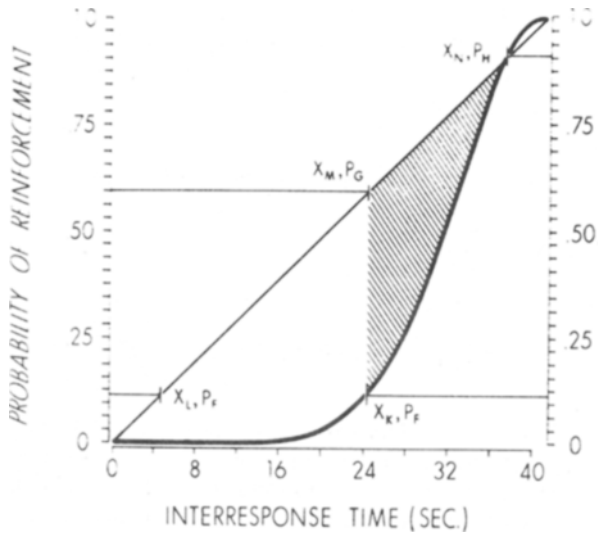

Fig. 13. Two functions representing stocastic reinforcement of waiting (SRW) schedules. The linear function makes available reinforcement once every 41 sec, on the average, provided that the mean response rate exceeds that value. The Gaussian function more closely simulates conventional DRL schedules. The shaded area emphasizes the increasing divergence of reinforcement probability as IRT is shortened.

The well-defined peak at about 300 msec seems to correspond to the conventional clean-cut keypeck. Notice, however, that at about 600 msec, there is a shorter peak which is more clearly seen in Fig. 12 from Bird 911. Here we see a succession of peaks. We call these "harmonics," and visual observations indicate that these arise from incomplete responses, i.e., feints toward the key or pecks which strike not the key but the wall nearby. The incidence of these harmonics falls with practice.

A microanalysis of the acquisition of FR performance, then, leads to the following conclusions. Fixed-ratio responding evolves via two processes. First, the frequency of very long IRTs, i.e., long breaks in responding, decreases sharply. Second, the topography is sharpened and made more precise. The proportion of harmonics goes down and the proportion of responses in the category correlated with a clean-cut keypeck goes up. Fixed-ratio performance develops as a shaping process in that efficient responding is differentially selected. Such a conclusion would not be possible if the data had not been recorded with a fair degree of temporal precision, if serial order had not been preserved, and if rapid manipulation and processing of large quantities of data had not been possible.

NEW RELATIONSHIPS

Morse (1966) pointed out that "simple" reinforcement schedules are schedules that are easy to instrument, not schedules that are simple to interpret. Almost any microanalysis documents the validity of that statement. Sometimes a complex contingency is more revealing of certain processes. For example, on a VI schedule, the probability of reinforcement rises as a function of the duration of the interval between

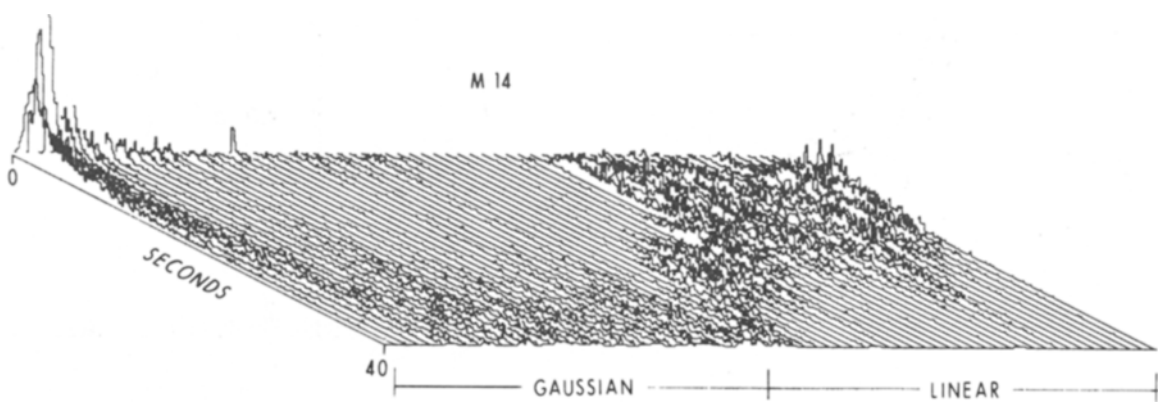

Fig. 14. Three-dimensional IRT histogram plot showing the changes in IRT distribution during the period of exposure to the Gaussian function and the subsequent reversal to the linear function previously employed. Sessions run along the abscissa. IRT frequency is represented by height, and the z-axis represents IRT duration from 0 to 40 sec.

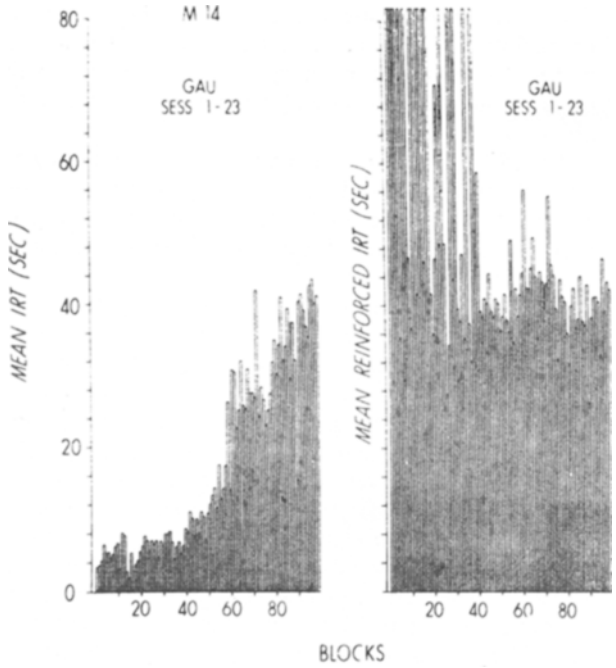

Fig. 15. Mean IRT and mean reinforced IRT in blocks of 100 IRTs for Sessions 1-23 on the Gaussian function for Monkey M 14.

responses. The longer an organism waits, the higher this probability. A more direct statement of this function can be provided by the computer, as shown on Fig. 13. We call this an SRW schedule for stochastic reinforcement of waiting (Weiss, 1970a). In this chart, you see both a linear and a Gaussian function. The linear function also corresponds to a variable ratio (VR). That is, if the average waiting time is 20 sec, the organism will be reinforced on a VR 2 schedule. Using a quasirandom number generator, such a schedule is very easy to program.

A more complex function, any function in fact, can be simulated, in essence, by plotting it in memory. The Gaussian function in Fig. 13 is close to a stochastic analog of the conventional spaced responding schedule. I have studied the transitions between these two functions, linear to Gaussian and back to the linear again, to determine whether or not exposure to the Gaussian function would produce any permanent change, and I employed the Gaussian rather than the simpler DRL schedule in order to stretch out the transition and to make its various stages more amenable to analysis. One example of the kind of change seen when moving from the Gaussian to the linear function is shown in Fig. 14, for Monkey M 14.

The first few sessions on the Gaussian function showed a peak at the short IRT values, corresponding to earlier performance on the linear function. As with the spaced responding monkeys, the distribution gradually drifted to the longer IRTs. When the linear function was reimposed, we saw a gradual drift to the shorter IRTs, but without fully returning to the pre-Gaussian features of the interval distribution. The early peak is noticeably absent. Even after approximately 90 sessions (145 h) on the linear function, there was no reappearance of the very short IRTs, suggesting a permanent change.

A detailed analysis shows another aspect of how the contingencies attained control. Figure 15 plots the mean IRT and mean reinforced IRT for each 100 IRTs. For the first several hundred IRTs, the mean reinforced IRTs were quite long, suggesting that they occurred as part of an extinction phenomenon, interrupting sequences of relatively short IRTs. Later, as mean IRT rose, mean reinforced IRT fell, a correlation that suggests the assumption of IRT control by schedule specifications.

The ability to perform arithmetic operations on-line confers still another dimension of control. Suppose, for example, that an experimenter is interested, as we have been, in controlling the variability of response rate or IRTs. Both spaced responding and pacing schedules (Ferster & Skinner, 1957) do so indirectly by prescribing IRT minima, maxima, or both. The question that we wanted to propose (Weiss & Laties, 1965), both for its intrinsic behavioral interest and because of its importance in behavioral pharmacology, was whether or not variability could be controlled by appropriate contingencies and how responsive it would be to drug effects.

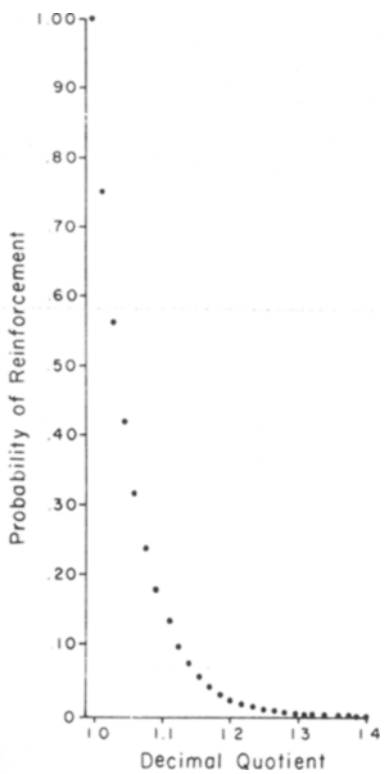

Fig. 16. The autoregressive schedule function. Two successive IRTs are compared by computing the quotient of the larger over the smaller. The closer the quotient lies to 1.0, the higher the probability of reinforcement. This function is stored as a table in memory.

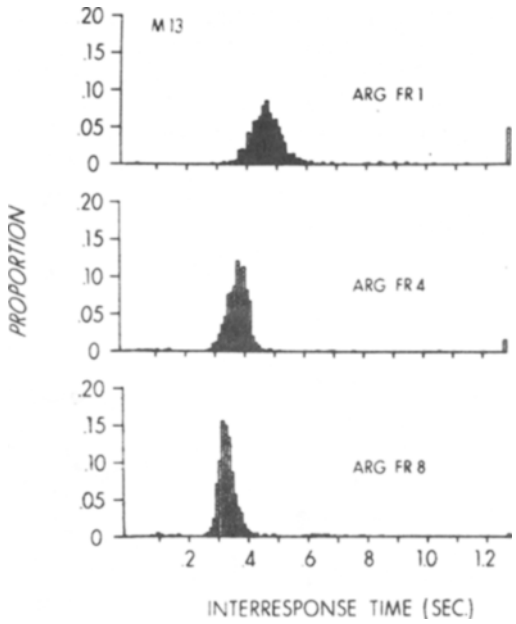

Fig. 17. IRT distributions generated by Monkey M 13 under three variations of the autoregressive reinforcement schedule. FR 4 and FR 8 refer to conditions in which every fourth and eighth response eligible for reinforcement was followed by the delivery of juice.

The contingency we employed is based on serial variability. We called it the autoregressive reinforcement schedule after the time series model of serial correlation. It is diagrammed in Fig. 16. Each time a response is made and terminates an IRT, that IRT is compared to the previous one by computing the quotient of the two. The closer the quotient lies to 1.0, the higher the probability of reinforcement. The function diagrammed in Fig. 16 resides in memory and corresponds to the range of numbers generated by a pseudorandom number generator.

The IRT distributions generated by this reinforcement schedule are shown in Fig. 17. These interval histograms are based on 10-msec resolution. The topmost histogram shows a condition we call FR 1, in which every response eligible for reinforcement was accompanied by the delivery of juice to a monkey (*M. speciosa*). With FR 4 and FR 8, where every fourth and eighth response eligible for reinforcement was followed by juice, the distribution narrowed, variability becoming reduced.

This was one of the methods that the monkeys used to solve the problem posed by the contingencies, namely, behaving very much like an oscillator. In addition, however, they also introduced a considerable measure of serial correlation. As shown on Fig. 18, they tended to drift up and down around a stable mean value in a manner reminiscent of spaced responding performance. Figure 18 also illustrates some aspects of our GENDIS program which make such an analysis easy to perform.

Finally, I want to discuss the application of computers to the control of adjusting or titration schedules in which some parameter of the situation is changed in accordance with the organism's behavior. In previous studies with drugs (Weiss & Laties, 1964), we found response duration to be very sensitive, in dogs, to combinations of amphetamine and barbiturates, a phenomenon that we wished to explore at greater length with the computer. For this purpose, I devised a

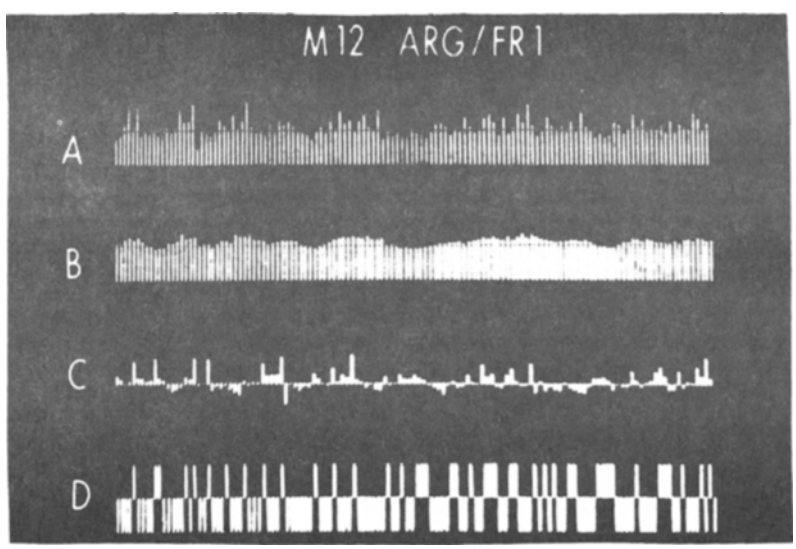

Fig. 18. Photographs from the oscilloscope display of the LINC computer showing features of performance on ARG FR 1 from Monkey M 12. A represents successive IRTs. B shows the same sequence after smoothing by a three-term moving average. C shows successive differences between IRTs. D shows the direction of deviations from the mean IRT. These displays indicate that serial correlation was a significant feature of the performance.

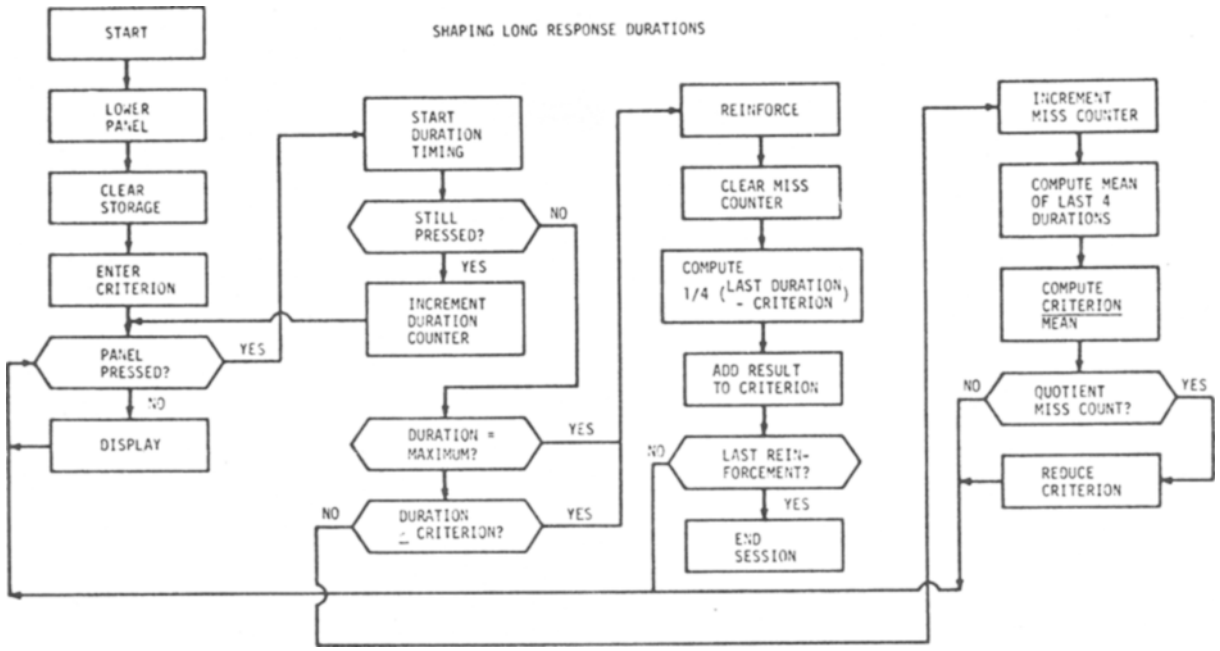

Fig. 19. Flow chart of a program (SHAPE) that shapes long response durations. Dogs were trained to press a panel with the snout. If the response succeeded, a response criterion was raised by one-fourth of the difference between the old criterion and the response duration which exceeded it. Under certain conditions, the criterion could be lowered, but lowering took place more quickly the closer the mean of the last four response durations was to the criterion.

program whose flow chart is shown in Fig. 19. The dogs were trained to press a panel with their snouts. This behavior was maintained by the delivery of dry dog food. After further training and parameter exploration, they were subjected to the requirements shown in this diagram. An experimental session began with a predetermined criterion of .25 sec. Any response duration above the criterion produced delivery of food and elevated the criterion by 25% of the difference between the old criterion and the new response duration which had exceeded it. This proportion was chosen empirically. If the criterion were raised too quickly, the behavior would be lost, as it is in any kind of shaping procedure where the change in criterion makes no contact with the behavior. At the same time, it was necessary, if behavior began to deteriorate even under these conditions, to recapture it. The following procedure was employed for this purpose. A running mean was kept of the last four durations. The decision about whether or not to reduce the criterion was made on the basis of how closely the mean approached the criterion. The closer it was to the criterion, the fewer the number of unreinforced responses required to reduce the criteria. This was done in order to prevent the animals from emitting long chains of short duration responses and bringing the criterion down very rapidly.

The kind of behavior produced is shown in Fig. 20. By looking at the bars, you can see a gradual increase in the duration of successive responses, a discrimination that the dogs eventually began to make with rather surprising precision. The criterion, which is represented by the line woven through the bars, tends to be exceeded by only a relatively small amount. When this behavior stabilized, it became possible for us to study the drug interactions we had seen earlier with much greater precision.

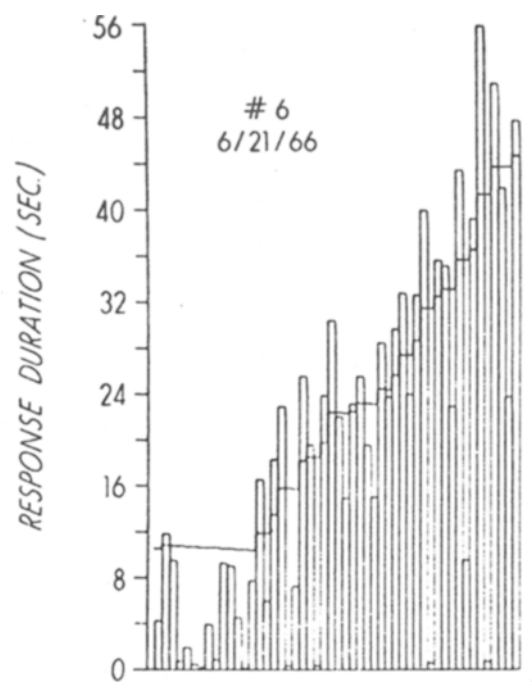

Fig. 20. Performance by Dog 6 on the SHAPE program (see Fig. 19). Individual bars represent successive response durations. The line woven through the bars represents the criterion. Note that after the first few durations, the criterion dropped slowly until long duration responding resumed.

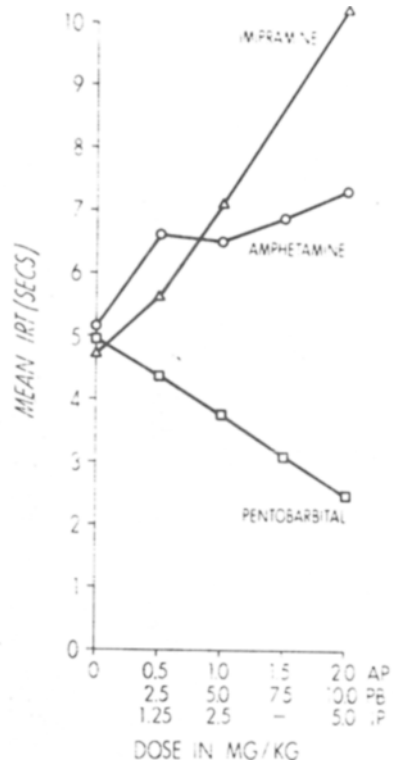

Fig. 21. Effects of the three drugs shown on the duration of the postreinforcement interval on FR 30. The points are means based on nine pigeons.

CHEMICAL INFLUENCES ON BEHAVIOR

The kind of microanalysis and microcontrol that I have been discussing so far has found a useful role in studies of how chemicals affect behavior. Let me turn my attention first to studies of fixed-ratio performance. One of the reasons for undertaking an evaluation of FR behavior is that it often has seemed resistant to the effects of drugs (e.g., Bignami & Gatti, 1969). It exerts such powerful control that even animals grossly ataxic from high doses of barbiturates may continue to respond at relatively high rates. For instance, on multiple FI FR schedules, drugs such as pentobarbital can virtually eliminate the FI component without grossly altering the FR rate (Herrnstein & Morse, 1956). In recent years, moreover, many investigators have examined drug effects in the context of the rate dependency hypothesis (Kelleher & Morse, 1969), which states that the effects of a drug on behavior will depend upon the baseline rate of responding independent of reinforcement schedule. By subjecting FR performance to a fine-grained analysis, other aspects of drug response become apparent (Weiss & Gott, 1972). Figure 21 shows the effects of three drugs on the interval following reinforcement, before the first response of the ratio. Notice that the administration of amphetamine to pigeons tends to lengthen, while the administration of pentobarbital tends to shorten this period, even at relatively high doses.

Some sample plots of patterning through the ratio are seen in Figs. 22 and 23. In Fig. 22, you can see that a dose of 0.5 mg/kg of d,l-amphetamine sulfate to this bird tended to lengthen not only the interval directly after reinforcement, but the next few IRTs as well. Sodium pentobarbital reduced these first few IRTs and exerted a similar effect on the rest of the ratio.

Other aspects of the data are even more interesting. Figure 24 compares two interval histograms, one for a control injection and one for 1.0 mg/kg of amphetamine. A number of changes took place after the drug. The main peak shifted to the right, the number of very short IRTs (the nibbles) declined, and the harmonic peaks increased in size, indicating more aborted or misdirected pecks. In contrast, the histogram for

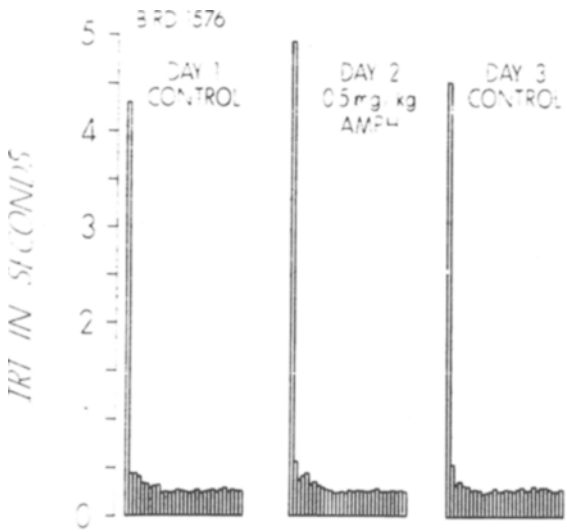

Fig. 22. IRT patterning on FR 30 in response to a dose of 0.5 mg/kg d,l-amphetamine sulfate.

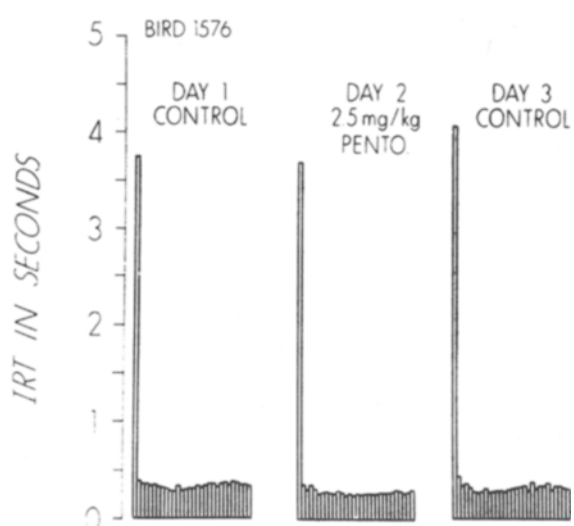

Fig. 23. IRT patterning on FR 30 in response to a dose of 0.5 mg/kg sodium pentobarbital.

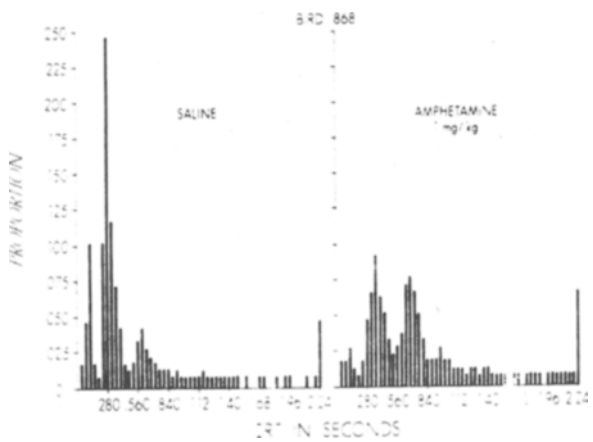

Fig. 24. IRT histogram on FR 30 after 1 mg/kg of amphetamine.

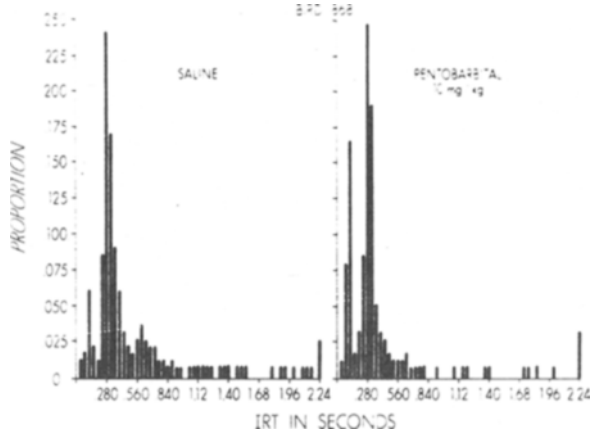

Fig. 25. IRT histogram on FR 30 after 10 mg/kg of pentobarbital.

10 mg/kg of sodium pentobarbital, as shown on Fig. 25, reveals a shift of the main peak to the left, an increase in nibbles, and a decrease in the harmonic components, indicating more efficient responding. Although much of the data can be adequately described by response rate changes, a closer examination of what underlies these changes leads one to consider what processes these drugs might be affecting. For example, they may be influencing stimulus control in the sense that they modify the selectivity of response topography, which would be an important dimension of drug action and one, given current research in neuropharmacology, that might more easily be related to basic neurochemical processes in the brain (Weiss & Laties, 1969).

The microanalysis of FR performance has also helped us to study the toxicology of methylmercury. Methylmercury is a potent CNS poison and, in the last 2 years, has been the object of much concern in the United States because of its presence in fish. In order to determine some adequate dose-response relationships in the context of the Rochester mercury research project, we began treating trained pigeons with various dose levels of methylmercury chloride. Methylmercury has an extremely steep dose-response curve, and one may see a sudden change taken place in an animal from one day to the next in overt neurological symptoms. We are interested in more sensitive indices of impairment so that we can more precisely determine the brain concentrations and distributions at which behavioral changes begin to appear. Figure 26 describes the changes in FR performance that appeared before this pigeon showed overt toxic symptoms (Evans & Kostyniak, 1972). The main change was a lengthening of the postreinforcement time and a gradual, though slight, extension from session to session in IRT length.

A precise temporal analysis of VI performance can also help illuminate some aspects of how chemicals interact with behavior. Even simple IRT distributions can be helpful. Figure 27 shows the changes induced in one pigeon's behavior by 10 mg/kg of imipramine

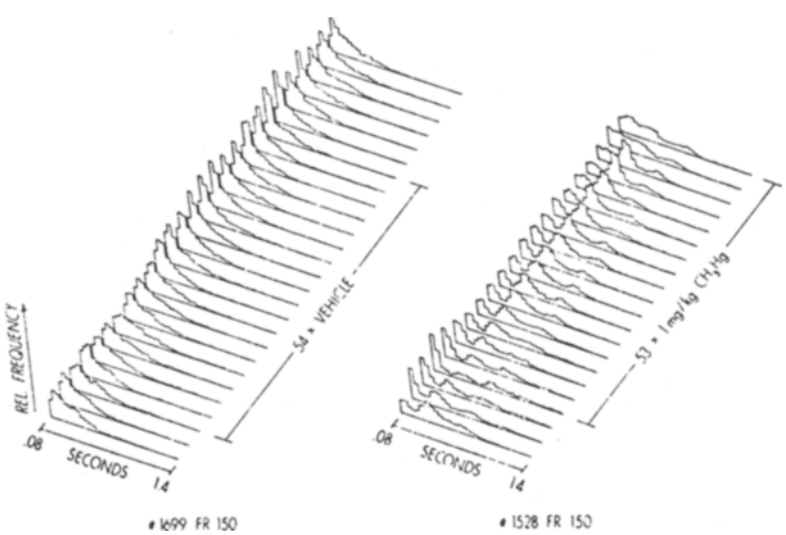

Fig. 26. Effects of methylmercury chloride on FR 150 performance. Pigeon 1699 was given the sodium carbonate vehicle. Pigeon 1528 was given 53 administrations of 1 mg/kg methylmercury five times per week. Selective IRT histograms from this period show a diminution of the early peak of the distribution and, toward the end, a flattening of the distribution.

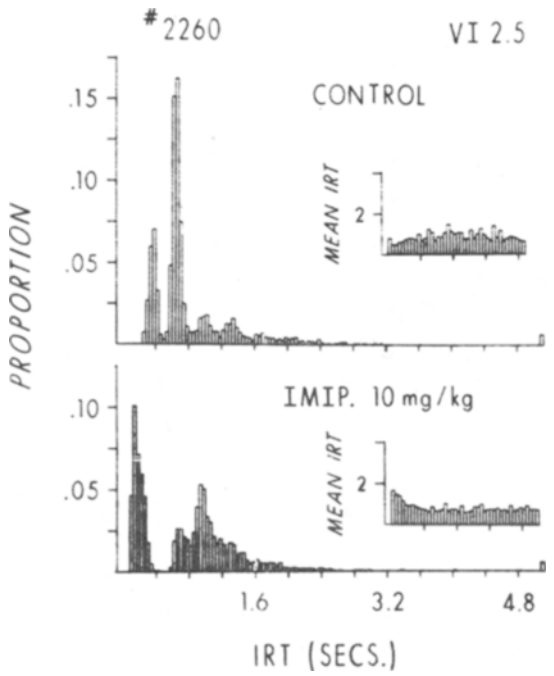

Fig. 27. Interval histogram under VI 2.5 during control conditions and after 10 mg/kg of imipramine hydrochloride. The insets show mean IRT per 100 IRTs.

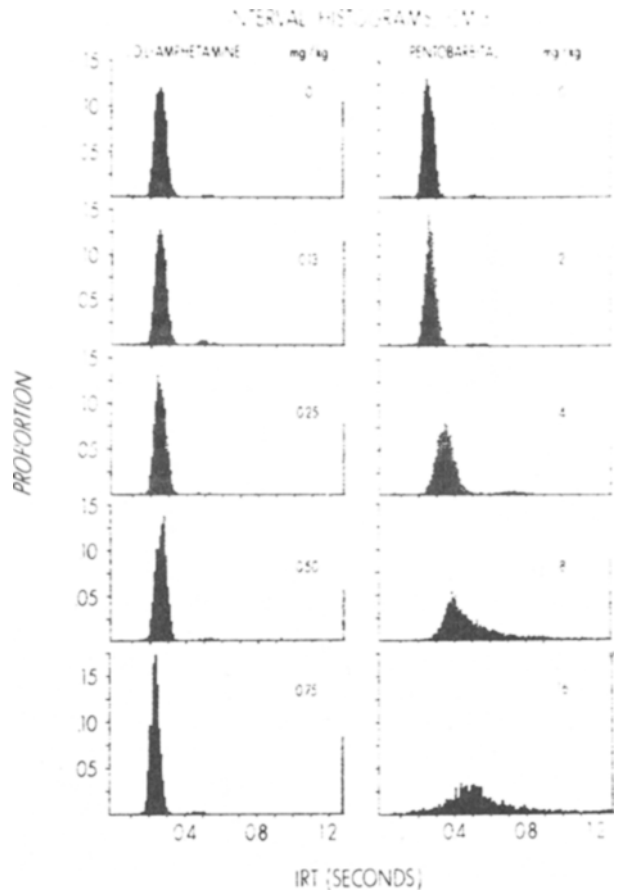

Fig. 28. Interval histograms on the ARG schedule, Monkey CM-3, after varying doses of amphetamine and after pentobarbital. IRTs were recorded with a resolution of 10 msec.

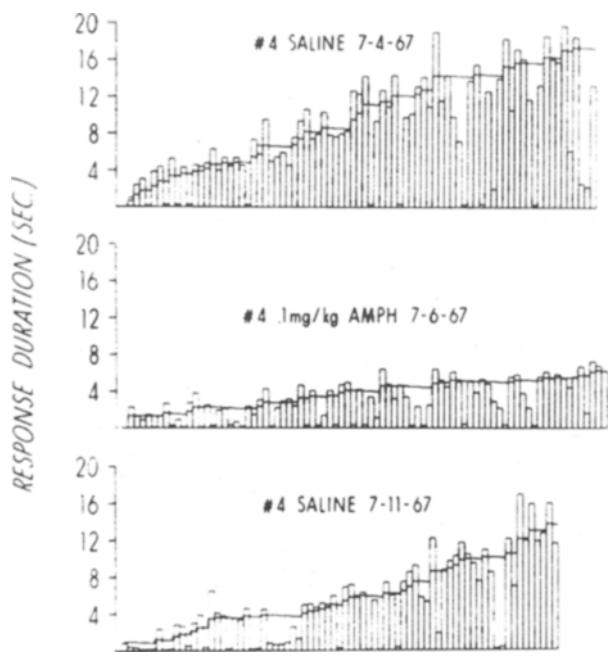

Fig. 29. Performance of Dog 4 on the shape schedule after saline and after 0.1 mg/kg of amphetamine. Individual bars represent response duration. The continuous line represents the criterion.

hydrochloride. The insets show mean IRT through the session in blocks of 100 IRTs.

The mean rate is somewhat lower under control conditions than after imipramine, but the whole character of the IRT distribution is different after the drug, which I consider a more important finding. Note that the first peak is shifted left, which probably accounts for the higher rate, but also that this first peak is now dominant and that there appears to be a distinct separation between the relatively short IRTs and the others. This suggests to me that this pigeon, under drug, emitted rapid trains of pecks separated by relatively long IRTs. Other features of the two distributions, such as the location and relative amplitude of the subsequent peaks, also indicate a significant change in temporal patterning. I maintain that these kinds of changes reveal more about basic drug-behavior mechanisms that do simple rate changes; they represent one reason I see the computer as an important tool in behavioral pharmacology and toxicology.

An equally cogent argument not only for microanalysis but for microcontrol comes from drug studies with the ARG schedule. Gage (1970) administered a range of doses of d-amphetamine and

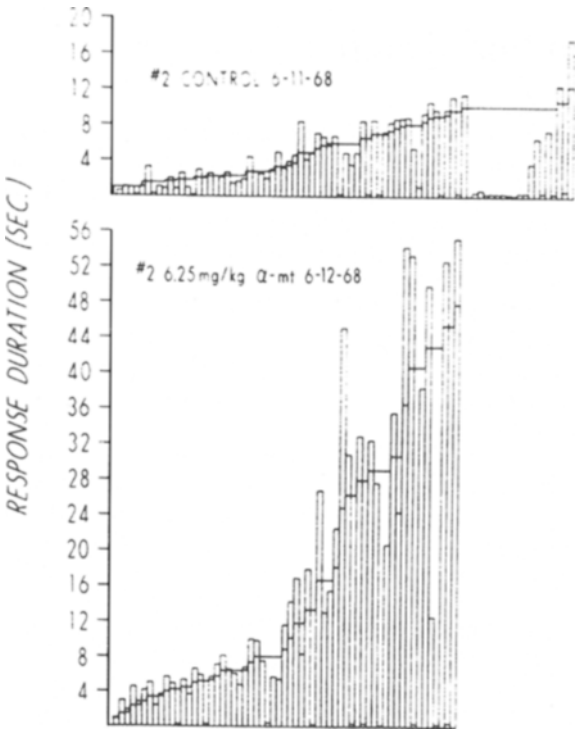

Fig. 30. Effects on response duration of 6.25 mg/kg of alpha-methyltyrosine in Dog 2.

pentobarbital to monkeys maintained by this schedule and obtained the results shown in Fig. 28. Although this obviously is a high-rate performance, in contrast to FR, on which amphetamine tends to lower and pentobarbital tends to raise rates, responding after amphetamine was more uniform and was emitted at a higher rate than after pentobarbital.

The shaping schedule with the dogs has also proven useful from a pharmacological point of view. We found the schedule extremely sensitive to amphetamine (Weiss, 1970b), with clear-cut effects appearing at 0.1 mg/kg (see Fig. 29), a dose within the range used, say, for the treatment of obesity. Similar effects occurred after barbiturates and with the amphetamine-barbiturate combination. In order to determine whether some common neurochemical processes underlay these effects of the two different drugs, we administered the compound alpha-methyltyrosine, which inhibits the rate-limiting step in the synthesis of the catecholamines. We began with the relatively high doses used by most investigators and found effects which caused us to continue to lower dose level until we reached 3.1 mg/kg, which still produced a lengthening of response duration, an effect we had not seen previously with any other drug. This effect caused us to induce Dr. Alfred Heller, from the University of Chicago, to examine the neurochemical consequences of these low doses; he discovered that even at dose levels as low as 6.25 mg/kg, significant changes appear to take place in rat brain levels of catecholamines. Figure 30 shows the effect of this dose level on dog performance.

CONCLUSIONS

I have tried to document my assertion that computer technology has made a contribution to our understanding of behavior and that it can be used both to ask and answer important questions about behavior and about behavioral pharmacology. Obviously, the technology is going to be developed further, be made available more cheaply, and made easier for us to use. I am happy to see such developments, but we surely do not have to await the perfect system in order to do significant research on behavior.

REFERENCES

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

Bignami, G., & Gatti, G. L. Analysis of drug effects on multiple fixed-ratio 33 fixed-interval 5 min in pigeons. Psychopharmacologia, 1969, 15, 310-332.

Clark, W. A., & Molnar, C. E. A description of the LINC. In R. W. Stacy and B. D. Waxman (Eds.), *Computers in biomedical research*. Vol. 2. New York: Academic Press, 1965. Chap. 2.

Evans, H. L., & Kostyniak, P. J. Effects of chronic methylmercury on behavior and tissue mercury levels in the pigeon. Federation Proceedings, 1972, 31, 561. (Abstract)

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

Gage, M. I. Amphetamine and pentobarbital effects on interresponse time distributions of monkeys reinforced for low sequential variability of interresponse times. Unpublished doctoral dissertation, University of Rochester, 1970.

Gott, C. T., & Weiss, B. The development of fixed-ratio performance under the influence of ribonucleic acid. Journal of the Experimental Analysis of Behavior, 1972, 18, 481-497.

Herrnstein, R. J., & Morse, W. H. Selective action of pentobarbital on component behaviors of a reinforcement schedule. Science, 1956, 124, 367-368.

Kelleher, R. T., & Morse, W. H. Determinants of the behavioral effects of drugs. In D. H. Tedeschi and R. E. Tedeschi (Eds.), *Importance of fundamental principles in drug evaluation*. New York: Raven Press, 1969. Pp. 383-405.

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.

Weiss, B. The fine structure of operant behavior during transition states. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970a. Pp. 277-311.

Weiss, B. Amphetamine and the temporal structure of behavior. In E. Costa and S. Garattini (Eds.), *Amphetamine and related compounds*. New York: Raven Press, 1970b. Pp. 797-812.

Weiss, B. (Ed.) *Digital computers in the behavior laboratory*. New York: Appleton-Century-Crofts, in press.

Weiss, B., & Gott, C. T. A microanalysis of drug effects on fixed-ratio performance in pigeons. Journal of Pharmacology & Experimental Therapeutics, 1972, 180, 189-202.

Weiss, B., & Laties, V. G. Effects of amphetamine, chlorpromazine, pentobarbital, and ethanol on operant response duration. Journal of Pharmacology & Experimental Therapeutics, 1964, 144, 17-23.

Weiss, B., & Laties, V. G. Reinforcement schedule generated by an on-line digital computer. Science, 1965, 148, 658-661.

Weiss, B., & Laties, V. G. Comparative pharmacology of drugs affecting behavior. Federation Proceedings, 1967, 26, 1146-1156.

Weiss, B., & Laties, V. G. Behavioral pharmacology and toxicology. Annual Review of Pharmacology, 1969, 9, 297-326.

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