A COMPUTER ANALYSIS OF SERIAL INTERACTIONS IN SPACED RESPONDING¹

BERNARD WEISS, VICTOR G. LATIES, LOUIS SIEGEL,
AND DAVID GOLDSTEIN

UNIVERSITY OF ROCHESTER SCHOOL OF MEDICINE AND DENTISTRY,
ROCHESTER, NEW YORK

Serial dependencies in interresponse times were studied by means of a digital computer. In monkeys exposed to a DRL 20-sec schedule of reinforcement, serial interactions appeared at all stages of training. Early in training the serial effects consisted of trains of relatively long interresponse times interspersed among trains of relatively short ones. Later on, the serial effects appeared to be characterized by a tendency to drift up and down in long wavelength periods around the minimum interval required for reinforcement. After training to a point at which most interresponse times produced reinforcement, serial effects of a still more subtle nature appeared. These effects were made apparent by autocorrelation and power spectrum methods and consisted of both long-term and extremely short-term fluctuations in interresponse times.

When exposed to a schedule of reinforcement that reinforces only those interresponse times (IRTs) exceeding a specified duration—the DRL schedule—organisms adjust to such requirements by lengthening the interval between responses.

Various aspects of how organisms space responses in time on the DRL schedule of reinforcement have been studied but little is known about the apparent serial interaction in IRT duration noted by a number of authors. One instance of this interaction is a tendency for reinforced IRTs to occur in clumps (e.g., Kelleher, Fry, and Cook, 1959; Farmer and Schoenfeld, 1964). Another instance is the tendency for closely-spaced responses to be emitted in bursts, especially after IRTs just short of the duration required for reinforcement (e.g., Sidman, 1956). Recently, Ferraro, Schoenfeld, and Snapper (1965) reported significant first-order dependencies in rat performance; that is, successive IRTs were

positively correlated. In order to be able to offer a more detailed analysis of these rather subtle effects than has been possible in the past, the present study employed digital computers to acquire and process data.

METHOD

The subjects were three female *M. speciosa* monkeys, weighing about 3 kg each, which were maintained in primate restraining chairs (Foringer) enclosed in large chambers. The operanda were microswitch levers (Lehigh Valley monkey model) mounted on the chairs. About 30 g of force was needed to operate a lever. The levers were connected to the external sense lines of a LINC computer (Clark and Molnar, 1965) through Schmitt triggers and flip-flops. The LINC programmed the schedule and recorded all IRTs with a resolution of 40 msec. Each depression and release of a lever was recorded as a response and produced a 20 msec, 2500 cps tone in the chamber from a Sonalert (Electropac, Inc.), immediately after release. The release contingency was introduced to avoid prolonged holding of the levers. A white noise generator helped mask extraneous sounds.

Reinforcement occurred each time a response terminated an interresponse interval of 20 sec or more (DRL 20). The reinforcing stimulus consisted of 0.3 ml of a fruit drink delivered via a solenoid valve. A short metal nipple extending from the outlet port of the valve enabled the monkey to ingest the fruit

¹From the Department of Radiation Biology and Biophysics, the Center for Brain Research, and the Department of Pharmacology. A portion of this work was conducted at the Johns Hopkins University, supported in part by grants MH-03229, MH-08353, and MH-07498 from the National Institute of Mental Health. A portion was conducted at the University of Rochester, supported in part by grant MH-11752 from the National Institute of Mental Health and in part by a contract with the U. S. Atomic Energy Commission at the University of Rochester Atomic Energy Project. Reprints may be obtained from Bernard Weiss, Dept. of Radiation Biology and Biophysics, The University of Rochester School of Medicine and Dentistry, Rochester, New York 14620.

drink when it was ejected. Since the liquid flowed from the nipple when the valve opened, the monkeys tended to keep their mouths close to the openings and began to lick whenever the solenoid was actuated.

Subjects were allowed free access to food except during the experimental session but deprived of water for approximately 20 hr before a session. Three sessions per week were conducted, each session lasting 2 hr.

No preliminary training, not even familiarization with the fruit drink dispenser, was given, and the DRL 20 reinforcement schedule was in effect from the very beginning. This procedure sought to prevent, as much as possible, any strengthening of tendencies to emit responses in bursts.

RESULTS

Serial interactions early in training

Despite the absence of conventional preliminary training, the monkeys began to respond almost immediately. Figure 1 is a display of the first 256 IRTs emitted by monkey M 11. Succession is from left to right, and IRT duration is represented by height. The upper dotted line represents an IRT duration of 20 sec. The lower one represents the median IRT and is equivalent to 4 sec. A substantial rate of responding developed rather quickly, as it did in the other two monkeys.

The apparent tendency of IRTs to occur in homogeneous groups, shown by Fig. 1, suggested that serial effects were present as early as the first session. The tendency toward

Fig. 1. Duration of the first 256 IRTs emitted by M 11 during the first session on DRL 20. Succession is from left to right and duration is represented by height. The upper dotted line represents an IRT of 20 sec, the lower one an IRT of 4 sec, the median value for this group of IRTs.

grouping may be displayed simply in the following way. Assign the value 0 to each IRT below the median value. Assign the value 1 to each IRT above the median. Then plot the frequency of occurrence of various sequences of 0's and 1's. Such a display is presented in Fig. 2. The top row represents sequences of four IRTs; reading from left to right, these are equivalent to the binary form of the decimal numbers 0 to 15. The height of the leftmost bar thus indicates the relative incidence of sequences equal to 0000 (four IRTs in a row of less than median length), the next bar the relative incidence of sequences equal to 0001, the next the relative incidence of sequences equal to 0010, and so on. The rightmost bar represents sequences equal to 1111. The middle row of bars represents sequences of three IRTs (the decimal numbers 0 to 7), and the bottom row sequences of two IRTs, or a 2 by 2 transition probability matrix, with the bars from left to right equivalent to 00, 01, 10, and 11. It is clear that the most frequent sequences were those containing IRTs in the same category since a random shuffle of sequence produced bars of equal height (*cf.*, Weiss and Laties, 1965). Thus, even in the first session, serial effects could be observed. The other monkeys displayed a similar tendency for IRTs to be grouped according to duration.

Serial interactions during intermediate stages of training

By the eighth session, approximately half of the IRTs produced by M 11 were long

Fig. 2. Sequence display showing the relative incidence of sequential combinations of IRTs above and below the median IRT. Top, sequences of four; middle, sequences of three; bottom, sequences of two. $N = 256$.

enough to be reinforced. Definite serial effects appeared during this session; these effects are visible in Fig. 3. This figure shows 128 successive IRTs, succession proceeding from left to right, with height representing duration.

The sequential relationships in the data are unmistakable. The IRTs drifted up and down in long wavelength periods, producing extended trains of reinforced and unreinforced responses. Similar effects were observed in data from the other monkeys.

Serial interactions late in training

With further exposure to the schedule contingencies, the incidence of unreinforced responses gradually fell. Figure 4 shows three IRT distributions, one for each monkey, during the later phases of the experiment. The variability from bin to bin is much greater than that normally observed on conventional IRT histograms, but the resolution in the present case is 160 msec, which provides much more detail in the IRT histogram than normally seen.

All three monkeys emitted a substantial proportion of reinforced IRTs, but M 11 was considerably more efficient and less variable than the other two. During Session 30, for example, M 11 was reinforced 233 times out of the first 256 responses; only 23 scattered responses were unreinforced, and these were still close to the minimum value of 20 sec. A portion of the cumulative record for this session appears in Fig. 5 and displays this efficiency quite clearly. If serial effects are present here, they should be rather subtle compared to those observed on earlier ses-

Fig. 3. Duration of 128 successive IRTs from Session 8 of monkey M 11. Succession is from left to right and duration is represented by height. The dotted line represents a duration of 20 sec.

Fig. 4. IRT histograms for three monkeys on DRL 20. Resolution = 160 msec. Ordinate = proportion of IRTs in each bin; abscissa = IRT duration in seconds. The number of hours of practice preceding these data was 78 for M 8, 46 for M 10, and 54 for M 11. $N = 1024$ per histogram, and encompasses the next four sessions.

sions. And, in fact, when one inspects a display of successive IRTs, as shown in Fig. 6, no decisive tendencies toward serial interaction can be detected. Figure 6, which corresponds to the cumulative record in Fig. 5, is most notable for its relative lack of variability. The apparent freedom from complex patterns of drift, however, makes this kind of data suitable for more advanced methods of time series analysis than were possible with the data from earlier sessions. These are discussed in detail below.

Fig. 5. The cumulative record of 128 responses made during 58 min of Session 30 (M 11); cf., Fig. 6. Slash marks indicate reinforcement.

Autocorrelation analysis

The purpose of time series analysis is to minimize the contribution of random fluctuations to the serial dependencies, periodic or otherwise, that may be present in the series. One method is autocorrelation analysis (Abelson, 1953; Weiss, Coleman, and Green, 1955; Blackman and Tukey, 1958; Walter, 1963).

Assume a time series, made up of observations $X_1, X_2, X_3, \dots, X_{N-1}, X_N$, where the subscript denotes the order of occurrence of each observation. This series is only one sample derived from a population of potential samples of N observations each. No conclusions could be reached about the population from this single sample were it not for the assumption of *stationarity*. A time series that can be shifted along the time axis without changing its statistical properties, *i.e.*, its mean, variance, *etc.*, is said to be stationary, and permits the assumption that time has the effect only of ordering the observations. A time series with a trend is not stationary. Given a stationary time series, it can then be asked

whether $X_1, X_2, X_3, \dots, X_N$ are independent or whether a particular observation depends on those preceding it. Autocorrelation analysis summarizes the dependencies by computing the covariance of X_{i+L} and X_i , where L represents the shift, or *lag*, along the time axis. When $L = 0$, the correlation is perfect. When $L = 1$, the correlation between the series $X_1, X_2, \dots, X_{N-1}$ and the series $X_2, X_3, \dots, X_N$ is computed. Since correlation is of samples drawn from the time series which are shifted relative to one another by an amount (L) along the time axis, the amount of correlation is expected to be somewhat less.

The autocovariances, which some authors call the autocorrelogram, are first computed. The autocovariance for lag L is defined as

$$AC(L) = \frac{1}{N-L} \left[\frac{\sum_{i=1}^{N-L} X_i X_{i+L} - \frac{\sum_{i=1}^{N-L} X_i \sum_{i=1+L}^N X_i}{N-L}}{N-L} \right]$$

Since the autocovariance for $L = 0$ is the variance of the original series, autocorrelation coefficients, $A(L)$, may be computed similar to the usual correlation coefficient by dividing each autocovariance $AC(L)$ by $AC(0)$. Plotting this normalized autocovariance against lag gives the autocorrelogram.

Autocorrelograms based on 256 consecutive IRTs each from sessions 27, 28, 29, and 30 appear in Fig. 7. In order to make certain that these series were stationary, any linear trends present were removed. This transformation was accomplished by subtracting from each datum the corresponding ordinate value of the least squares line for the entire data set. The data were then readjusted so that the mean value was zero. The assumption of stationarity following these corrections was tested by taking a typical series (Session 30) that displayed a slight but significant linear trend and then computing four autocorrelo-

Fig. 6. 128 successive IRTs from Session 30 of M 11. Duration is represented as height, and succession is from left to right. The superimposed dotted line represents an IRT of 20 sec. The leftmost IRT is the one terminated by the first response shown on the cumulative record of Fig. 5.

Fig. 7. Normalized autocorrelograms, reading left to right, for Sessions 27, 28, 29, and 30, respectively, of monkey M11. Each correlogram is based on the first 256 IRTs of the session. The solid lines represent the empirical values. The dashed lines represent correlograms based on the assumption that the serial dependencies reflect a Markov process (see text).

grams based on 64 observations each; that is, the original series of 256 observations was sliced into four sections. The resulting values confirmed the assumption of stationarity because the four autocorrelograms were quite similar and the variances differed only slightly. A failure to check for stationarity may produce spuriously high autocorrelation values (*c.f.*, Ferraro *et al.*, 1965).

Figure 7 (solid line) shows the values for $L = 0$ to $L = 10$. For $L = 0$, $A(L) = 1.0$, since the series is perfectly correlated with itself. For successive observations, $L = 1$. According to the formulas given in Tintner (1952, p. 243) the 5% level of significance would be close to .099 and the 1% level close to .124. For all four series, $A(1)$ is significant at the 1% level.

If the serial interactions arose only from a correlation between adjacent observations, the relationship most frequently examined, then we would be dealing with a Markov process (*c.f.*, Abelson, 1953; Weiss *et al.*, 1955); that is, a process whose change depends only on the probabilities associated with its present state, being completely independent of its history preceding that state. For a Markov process, the autocorrelation for any given L is related to the autocorrelation for $L = 1$ by: $A(L) = [A(1)]^L$. Using this relationship, the first 10 autocorrelations for a Markov process were calculated and are compared in Fig. 7 (dashed line) with the corresponding empirical autocorrelations. The close correspondence of

the empirical and derived correlograms, taking into account the inevitable presence of statistical variability, suggested that we could be dealing with a Markov process and that the first-order relationships were the only significant ones. Partly because of the deviations from the model, however, and partly because of the striking periodicities in Fig. 3 and similar displays, we decided to pursue the analysis of serial dependencies with another method, spectral analysis.

Spectral analysis

The autocorrelogram is perhaps most useful when a signal is dominated by a single periodic fluctuation obscured by a random process. Where the underlying process is more complex, the sampling constraints of the autocorrelation function may make interpretation difficult. These constraints arise from the lack of independence in the autocorrelation coefficients. For example, if $A(1)$ is high, then $A(2)$ will also tend to be high for most functions. Sampling peculiarities in the first few lags, moreover, intrude artificially into the later lags.

One way to bypass these constraints is to compute the cosine transform of the autocorrelation function by a technique known as spectral analysis (Abelson, 1953; Weiss *et al.*, 1955; Blackman and Tukey, 1958; Walter, 1963). The basis of this technique is that any stationary time series can be described as a sum of an infinite number of cosine-wave

components of particular frequencies and amplitudes. In more common statistical terminology, the contribution to the total variance of the time series of each component frequency (the *spectral density*) can be determined and then a *power* (or *variance*) spectrum showing the distribution of power over frequency can be plotted. The utility of plotting the power spectrum can be appreciated by noting that the autocorrelation function of a sine wave is itself a sine wave, but that the power spectrum shows one dominant frequency; i.e., the power spectrum is plotted against frequency, not time.

Following Abelson's (1953) excellent exposition, these statements can be developed in the following way. Assume a time series of observations $\{X_i\}$. Spectral analysis considers that the fluctuations in the series are accounted for by a sum of cosinusoidal oscillations of varying frequencies, all frequencies being admissible. Then

$$\{X_i\} = \sum_{j=1}^{\infty} \alpha_j \cos(\omega_j t_i + \phi_j)$$

where,

$\{X_i\}$ is the population of X 's at any single value of i ,

α_j is the amplitude of the j^{th} component wave,

ω_j is the frequency of the j^{th} component wave,

ϕ_j is the phase (varying randomly) of the j^{th} component wave.

The total variance of the time series population is given by

$$\sigma_x^2 = \sum_{j=1}^{\infty} \frac{1}{2} \alpha_j^2$$

Thus, each component wave of frequency ω_j contributes to the total variance as the square of its amplitude.

Power spectra, or autospectrograms as they are sometimes called, were computed according to the formula

$$S(I) = 4 \left[\frac{1}{2} \{AC(0) + AC(M) \cos(I\pi)\} + \sum_{L=1}^{M-1} AC(L) \cos\left(\frac{IL\pi}{M}\right) \right]$$

where,

$S(I)$ is the spectral density at frequency I ,
 M is the number of lags used in the computations,

L is the lag,

AC is the autocovariance.

In the present instance, 25 lags were computed since the highest lag is recommended to be about 10% of the total series (Blackman and Tukey, 1958).

The "raw" power spectrum was then smoothed by a procedure called hamming, which widens and lowers the peaks. A second hamming was applied to the present data simply for ease of plotting.

The hamming procedure is carried out as follows:

$$= (I)H \begin{cases} 0.54 S(I) + 0.23 S(I-1) + 0.23 S(I+1) & I = 1, 2, \dots, M-1 \\ 0.54 S(0) + 0.46 S(1) & I = 0 \\ 0.54 S(M) + 0.46 S(M-1) & I = M \end{cases}$$

where,

$H(I)$ is the spectral density at I after hamming,

$S(I)$ is the spectral density at I before hamming,

M is the number of lags.

Spectral density plots corresponding to the sessions whose autocorrelograms appear in Fig. 7 are shown in Fig. 8. The abscissa is designated as frequency and reflects the periodicity of the fluctuation in the time series. A peak at a frequency of .25 means a fluctuation with a period or cycle length of four observations. The highest possible frequency is 0.5, or a fluctuation with a period of two observations, which is the practical limit.

Despite the differences in the total variance of each time series, which corresponds to the integral of the power spectrum, the patterns associated with each session are quite similar. The dominant peak lies at a long cycle length, because the dominant relationship in the time series, as shown in Fig. 7, is a first-order one. In addition to this early peak, however, all four time series contain other, less prominent, peaks, at shorter cycle lengths. One lies in an abscissa position that corresponds to a period of about 2.3 observations, or a frequency of about 0.5-0.3. Another less prominent peak appears at a point which corresponds to a

Fig. 8. Power spectra, reading left to right, for Sessions 27, 28, 29, and 30, respectively, of monkey M 11. Each spectrum is based on the first 256 IRTs of the session. The abscissa, denoted frequency, is the reciprocal of cycle length, so that the ordinate corresponding to 0.2, for example, represents the power (variance) attributable to a fluctuation with a period of five observations.

period of about 4-5 cycles, a frequency of about 0.25-0.2. These secondary peaks indicate that other dependencies besides the first-order one do exist, but in diminished amplitude, which partly accounts for their absence from the autocorrelograms.

Two checks on the reliability of the spectra in Fig. 8 were performed. In one, the data from each session were subjected to a random shuffle of serial order and the spectra recomputed. These showed markedly disparate patterns giving no hint at all of any systematic variation. In the second, power spectra were computed from the four truncated series of 64 observations each, described earlier. The consistency of the spectra again confirmed the assumption of stationarity and also supported the conclusion that their most prominent features were reliable ones.

The power spectra, therefore, indicate that although these time series are dominated by first-order dependencies which lead to fluctuations of long period, they also contain components of much shorter period. The latter may be interpreted as reflections of sharp fluctuations in IRT duration. A visual indication of the long-term drift can be observed by sighting along the cumulative record in Fig. 5, holding the figure parallel to the plane of vision. A fluctuation with a period of slightly more than one cycle can be detected in this manner.

DISCUSSION

A variety of evidence demonstrates serial dependencies in interresponse times at all stages of training on the DRL schedule of reinforcement. Such dependencies do not arise from specified contingencies as they do on other schedules studied by means of a computer (Weiss and Laties, 1965). They arise from contingencies over which the experimenter has exerted no deliberate control. This also appears to be true of Sidman avoidance (Wertheim, 1965).

Early in training, the serial interactions consisted of trains of longer IRTs alternating with trains of relatively short IRTs. With further training, the tendency toward alternation changed into a long wavelength drift, characterized by sequences of reinforced IRTs followed by sequences of unreinforced—but not particularly short—IRTs. This kind of drift can produce the grouping of reinforcements observed by Kelleher *et al.* (1959) and Farmer and Schoenfeld (1964). After training to a point at which only a small proportion of IRTs were too short for reinforcement, serial effects of a more subtle nature appeared.

A tendency for reinforced IRTs to occur in groups may not be difficult to explain. For example, one could posit that the after-effects of ingestion provide discriminative

stimuli which help make a sufficiently long pause more probable (*cf.*, Farmer and Schoenfeld, 1964). But the subtlety of the effects shown in the later sessions for M 11, in which nearly every IRT was reinforced, demands a more complex explanation. We would like to suggest that one possibility is provided by the chains of overt behavior first reported by Wilson and Keller (1953) that can be shown to control the pauses between responses (*e.g.*, Laties, Weiss, Clark, and Reynolds, 1965). Since such mediating chains are maintained in strength only by adventitious relations to reinforcement, their topography is subject to drift; the duration of a particular chain is certainly variable. A drifting mediational chain could produce periodic fluctuations in IRT duration even in a series whose mean is above the minimum interval required for reinforcement.

Other mechanisms could be advanced, of course. Perhaps this is the point: a detailed analysis of even a well-studied schedule may engender new questions about how the behavior is controlled and maintained.

REFERENCES

- Abelson, R. P. Spectral analysis and the study of individual differences in the performance of routine, repetitive tasks. Princeton, N. J.: Educ. Test. Serv., 1953.
- Blackman, R. B. and Tukey, J. W. *The measurement of power spectra*. New York: Dover, 1958.
- Clark, W. A. and Molnar, C. E. A description of the LINC. In Stacy, R. W. and Waxman, B. D. (Eds.) *Computers in biomedical research*, Vol. 2. New York: Academic Press, 1965, Ch. 2.
- Farmer, J. and Schoenfeld, W. N. Effects of a DRL contingency added to a fixed-interval schedule. *J. exp. Anal. Behav.*, 1964, 7, 391-399.
- Ferraro, D. P., Schoenfeld, W. N., and Snapper, A. G. Sequential response effects in the white rat during conditioning and extinction on a DRL schedule. *J. exp. Anal. Behav.*, 1965, 8, 255-260.
- Hamming, R. W. *Numerical methods for scientists and engineers*. New York: McGraw-Hill, 1962.
- Kelleher, R. T., Fry, W., and Cook, L. Interresponse time distribution as a function of differential reinforcement of temporally spaced responses. *J. exp. Anal. Behav.*, 1959, 2, 91-106.
- Laties, V. G., Weiss, B., Clark, R. L., and Reynolds, M. D. Overt "mediating" behavior during temporally spaced responding. *J. exp. Anal. Behav.*, 1965, 8, 107-116.
- Sidman, M. Time discrimination and behavioral interaction in a free operant situation. *J. comp. physiol. Psychol.*, 1956, 49, 469-473.
- Tintner, G. *Econometrics*. New York: Wiley, 1952.
- Walter, D. O. Spectral analysis for electroencephalograms: mathematical determination of neurophysiological relationships from records of limited duration. *Exp. Neurol.*, 1963, 8, 155-181.
- Weiss, B. and Laties, V. G. Reinforcement schedule generated by an on-line digital computer. *Science*, 1965, 148, 658-661.
- Weiss, B., Coleman, P. D., and Green, R. F. A stochastic model for time-ordered dependencies in continuous scale repetitive judgments. *J. exp. Psychol.*, 1955, 50, 237-244.
- Wertheim, G. A. Some sequential aspects of IRTs emitted during Sidman-avoidance behavior in the white rat. *J. exp. Anal. Behav.*, 1965, 8, 9-15.

Received March 8, 1965