

Chapter 9

Time series analysis

9.1 Introduction

The correlation between ON and OFF time has been considered in the previous chapters in terms of four correlation variables (viz., $\text{Corr}(X, Y)$, $\text{Corr}(X, Y_{-1})$, $\text{Corr} D(X, Y)$, and $\text{Corr} D(X, Y_{-1})$). However, there are other approaches to determining the correlation between two variables. This chapter examines the correlation between ON and OFF time by the use of time series methods.

Three properties of shuttling behaviour suggest that the use of time series methods may provide information on the nature of the underlying neural processes (the ‘black box’ generating mechanism). Firstly, under CRF stimulation and at a reasonably high level of intensity, the behaviour of the self-stimulating rat may be considered to be under the control of the stimulation. Secondly, the behaviour of the rat in such a situation may be assumed to reflect the activity in the neural substrate that is being stimulated. Thirdly, the behaviour may be completely described across the trial period by an ordered series of observations that describes the behaviour as being either ON or OFF. If a time series can be regarded as the realisation of a physical process, then some of the properties of that series must reflect the underlying activity that generated that series (Gregson, 1983).

Another reason for an investigation of the inter-relations in the ON and OFF time data is that classical statistical analyses require independence, or, at least zero correlation, among the relevant observations. The existence of autocorrelation in data has been shown to lead to considerable bias in variance estimates and confidence intervals (Gottman, 1981; Gregson, 1983). The statistical analyses of the experiments reported here were generally based on group data. In turn, the data for each

animal in a group were based on mean values obtained over a reasonably long period of observation; and the observation periods themselves were reasonably spaced. This procedure and the fact that no sequence effects could be found in the data in Experiment Ia, suggests that data obtained in this way are essentially independent.

However, statistical calculations based on individual ON and OFF times from within a particular trial may not be expected to show independence. Regressions or F-values calculated on a rat's times may be useful descriptive statistics but the reference of these statistics to probability levels, which have been derived on the assumption of independence, may be hazardous if even moderate autocorrelation is present (Gottman, 1981; Gregson, 1983). Biological systems, and, in particular, neurophysiological systems, may be expected to contain significant memory components which would invalidate the assumption of independence between successive observations. The ANOVAs calculated for individual trials in Experiment III may be of doubtful use if it can be shown that significant autocorrelation exists between adjacent observations. Inspection of the raw data associated with the calculated F-values in those analyses, suggested that the ANOVAs were quite accurate in identifying consistent patterns of responding, but a more formal analysis is required.

The theoretical basis for time series analysis will not be considered in detail since extensive reviews of the time series literature exist elsewhere (e.g., Anderson, 1976, 1980; Bennett, 1979; Cryer, 1986; Gregson, 1983). The present discussion will outline some elementary features only. The standard, time-domain, approach to time series analysis (Box & Jenkins, 1976) is a statistical procedure that uses the information available from the autocorrelation and cross-correlation functions to infer the type, and order, of linear equation that fits the data best. Time series methods assume that the system responsible for the generation of that time series can be represented by a 'black box', although model fitting and systems analysis methods might eventually allow a 'peek' inside that system's 'black box' (Marmarelis & Marmarelis, 1978, p. 377).

9.2 ARMA models

The time-domain approach of Box and Jenkins (1976) traditionally requires the value of a variable to be observed at repeated intervals constantly spaced in time. However, a series of values may be analysed with time-domain methods as long as there exists a set of observations ordered along some dimension (Anderson, 1976). For example,

morphology changes along a line of latitude or river-height loss along a stream bed may be satisfactory data sets for time series analysis (Bennett, 1979). For the purposes of examining the relationship between ON and OFF time, the values of these variables will be considered to be an ordered series of observations occurring along some, event-related, unidimensional axis. The underlying process which generated the values is assumed to be continuous and to continue indefinitely.

Two uses of the Box-Jenkins methods will be considered here. These are a univariate analysis of some CROSS time data and a transfer function (or, input-output) analysis of the ON and OFF times that constituted those CROSS times.

9.2.1 Univariate output series.

The ON and OFF times within a particular trial may be combined to give a CROSS time (i.e., $CROSS_i = ON_i + OFF_i$, where the OFF_i is the OFF time immediately following ON_i), so that the resultant series is a univariate set of numbers which describes the behavioural output in terms of the time from one initiation to the next. The input to such a series is assumed to be unknown. Obviously, the 'distant' input is known in that the intensity, pulse width and interpulse interval of the electrical stimulus are all known precisely; however, the immediate input to the neural system that directly results in the shuttling behaviour may be a conversion of those stimulus parameters into some new set of neurophysiologically interpretable signals which are assumed to be unknown. The 'second stage set' of neurons (Gallistel et al., 1981) may be the recipients of an unknown code of inputs (such as action potentials or 'packets' of neurotransmitter). The nature of the relevant signals will be assumed to be unknown although considerable progress has been made in determining the speed of conduction and the type of transmitter. The VTM has been suggested as a possible location for the second stage neurons (Shizgal et al., 1980) and has been implicated as a central region in ICSS models (e.g., Routtenberg, 1968; Wise & Bozarth, 1984).

The study of a single output series comprises, essentially, the iterative fitting of members of the family of mixed autoregressive and moving-average models known as ARMA (p, q) models, to arrive, empirically, at the most parsimonious mathematical representation of that series (Bennett, 1979; Cryer, 1986; Gregson, 1983; Tong, 1980). The autoregressive AR(p) terms may be thought of as the reliance of the present system output on previous outputs (i.e., memory of past responses), where the number of significant terms or order of this reliance is denoted by p . The moving-average

MA(q) terms reflect the reliance of the present output on previous inputs; that is, the effect of previous (unknown) inputs which have not dissipated from the system and are thus still contributing to the present output.

The set of unknown input signals is assumed, quite arbitrarily, to be an independent, normally distributed random variable (or 'white noise' process) denoted as $\{a_i\}$ with zero mean and constant variance σ^2 . The form of the general ARMA (p, q) model may then be written as

$$Z_i = \phi_1 Z_{i-1} + \phi_2 Z_{i-2} + \dots + \phi_p Z_{i-p} + a_i + \theta_1 a_{i-1} + \dots + \theta_q a_{i-q} \quad (1)$$

The above model is sometimes referred to as a 'shocks' model (Bennett, 1979) in which the current value of the process Z_i is made up of the current shock a_i plus the effects of q previous shocks a_{i-1}, \dots, a_{i-q} , plus memory terms incorporating the last p values of the series Z_{i-1}, \dots, Z_{i-p} . The physical interpretation of the ARMA model is that the $\{Z_i\}$ output sequence is dependent upon an unknown set of variables, which can be reproduced by the moving average terms plus internal memory (i.e., autoregressive terms, — Bennett, 1979). The 'shocks' are stored in internal memory and have effects which persist after the moment of their occurrence. Thus, although the shocks are considered random, their effects are not.

The identification of memory terms in the process would indicate a dependence of the current CROSS time on previous CROSS times, whereas the identification of moving average terms would indicate the importance of previous 'barrages' of impulses or white noise shocks in the determination of the current CROSS time value.

By considering an ordered set of CROSS times, it may be possible to model or 'capture' the essential properties of the unknown stimulus input that 'drives' the behaviour observed in the shuttlebox. Obviously, an identification and estimation procedure that arrives at an adequate descriptive model does not imply that the input process is exactly as has been modelled. All that can be said is that the essential properties of the input process have been included in that, most parsimonious, model. In other words, the empirically, best-fitting model may not be directly interpretable in terms of neurophysiological components. However, from the process of ARMA modelling, hypotheses may be generated concerning the nature of the reliance of the CROSS time series on its previous values and on the unknown components represented by the a_i 's. These hypotheses may then be tested by further experimentation.

9.2.2 Input-output analysis.

The second approach within a time series framework to be considered here is to regard the OFF time series as the output from the ON time series. The relationship between the individual ON and OFF times may be examined in more detail and over a greater range of lags and leads than has been considered with the within-trial correlation variables (i.e., as used in Chapters 5–7). Each ON time might be more correlated with an OFF time further removed than the succeeding, or preceding, OFF time. The cross-correlation function (CCF) summarises the relationship between ON time and the OFF time series. Bennett (1979) suggests that ‘When a series of stochastic inputs is involved, the cross-correlation function provides an estimate of the impulse response averaged over the length of the series.’ (p. 135). Or, in another way, it is possible to determine the system response that develops from a typical input.

Method

For the purposes of this analysis, nine 10-minute trials were taken from the M-Series data in Experiment Ia for closer examination. These trials are indicated in Table A.1. Appendix A, and were taken from the second highest level of intensity (i.e., Level 6) used in that experiment.

The particular trials chosen for analysis represented periods during which a fairly high number of crosses occurred (approximately 100) and were also periods in which observation notes had indicated an intuitive set of “good data”. Crossing rate was high, crossing behaviour appeared smooth, consistent and unhindered, and the times recorded showed a consistent order of magnitude over the period of observation. Also, three data sets from each of the ascending, descending and random sequences were chosen. These nine trials will be considered ‘unstabilised’ trials since each trial was both preceded and followed by a trial in which a different intensity prevailed.

In addition, four trials from the R-Series of animals were used for analysis since these data represented long periods of observation (20 minutes) after the behaviour had ‘stabilised’. The R-Series animals had been through screening and training procedures (as described in Chapter 4), and through 15-20 5-minute shuttling trials in which various pulse width/interpulse interval combinations were compared for stability of crossing. Different PW/IPI combinations of 1/8, 3/4, 5/0 msec (i.e., still 100 Hz) were traded-off against intensity to produce a moderately high, to high, rate of crossing (i.e., similar to but somewhat less than the intensity level used with

the M-Series animals). (This experiment is not reported in further detail — no differences could be found for the different PW/IPI combinations. Later experiments i.e., Experiments II and III therefore adopted a central PW/IPI value of 3/4 msec — also see Section 4.3.2).

In summary, the 14 data sets represented, approximately, the same level of ICSS behaviour under 100 Hz stimulation and CRF, but in some cases, the electrical charge was delivered by different pulse amplitude/pulse width dimensions.

Statistical analysis

The analysis was conducted using the BMDP program BMDP2T. This program conducts a standard Box-Jenkins time-domain analysis. Autocorrelation functions (ACFs), partial autocorrelation functions (PACFs), model fitting, parameter estimation, filtering, cross-correlation functions (CCFs) and transfer function identification operations may be carried out with the aid of this program. The raw data may be found in Appendix F.

Results

Figures 44– 53 depict graphically the main results for the 14 data sets used in this analysis. The data for three of the M-Series animals and all four R-Series animals are shown in some detail. Only the cross-correlation functions are shown for the remaining six M-Series data sets. Strictly speaking, the time values should be drawn as discrete columns at each event, but, for ease on computer use, connecting lines are used. A feature of data obtained from shuttling behaviour is the tendency for linear trend, or wandering, to occur in the data over 10-minute or longer observation periods. Therefore, all data sets have been differenced once to remove this source of nonstationarity (Box & Jenkins, 1976; Experiment Ia).

Univariate analysis of CROSS time data

Examples of the original series, the differenced series (which is then the series under analysis), and the respective ACFs and PACFs, are shown in Figures 44– 47 for six animals (M01, M12, M16, R09, R24, R49).

The essential feature shown by these data series is the strong suggestion of an MA(1) process (Anderson, 1976; Bennett, 1979; Cryer, 1986). An MA(1) process

(Anderson, 1976; Box & Jenkins, 1976) is indicated by a strong negative autocorrelation at lag $k = 1$ only, combined with a (comparatively) decaying set of partial autocorrelations for the first three or four terms.

For both data sets, the MA(1) process often needed additional, higher-order terms, or occasionally AR(1) terms, to empirically reduce the residuals from model fitting to a 'white noise' process. There were no obvious differences in the form of the ACF and PACF between the 'stabilised' and 'non-stabilised' groups. The ACFs and PACFs for the ON time data alone closely resembled those for CROSS time.

If the first order moving average model were to be considered as representative of the differenced series then the original series becomes an integrated moving average model, or IMA(1,1). Cryer (1986) has described the IMA(1,1) model thus: "we may usefully think of $Z_{[t]}$ [i.e., CROSS times], as being the equally weighted accumulation of a large number of white noise values" (p. 90). This model describes a nonstationary, 'wandering' process (Anderson, 1976).

The ACFs and PACFs provided little evidence for significant memory terms in the differenced series. For first-order autoregressive models the ACF declines gradually while the PACF tends to truncate (i.e., the converse of the MA(1) model). The presence of an AR(1) term (in addition to the MA(1) term), is indicated in the ACFs by a positive autocorrelation at lag $k = 2$. For example, the CROSS time data for M01 (Figure 45) and for R09 (Figure 47). However, these were not significant.

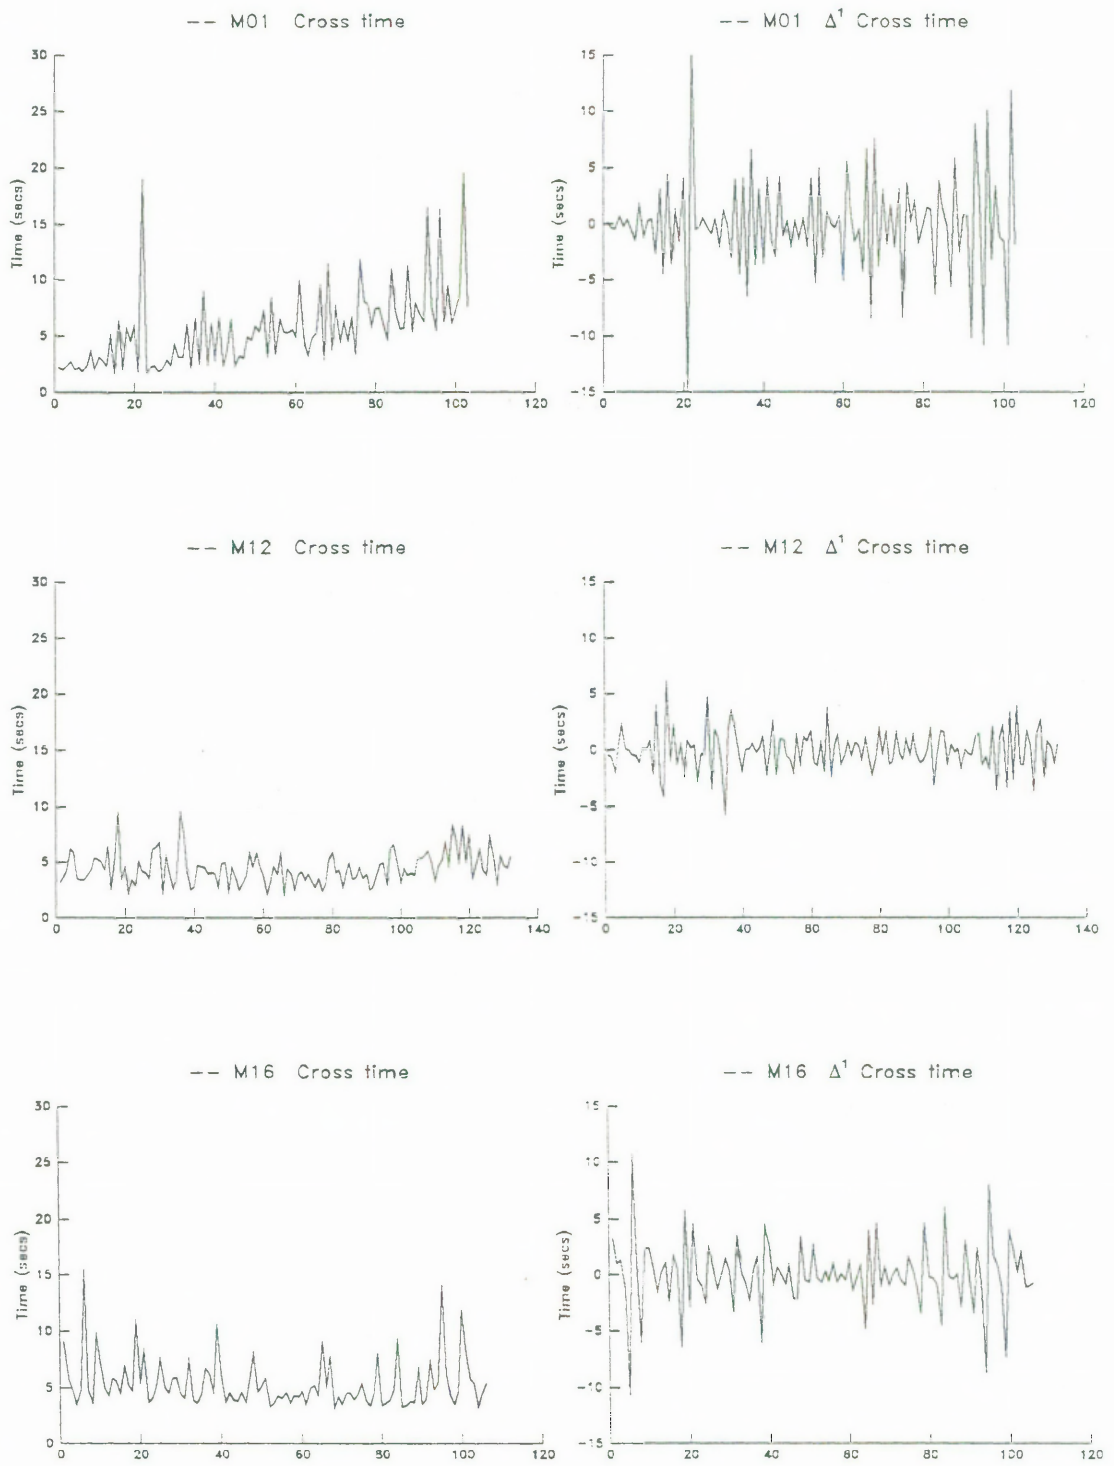


Figure 44: CROSS times and first differences (Δ^1) for M01, M12, and M16.




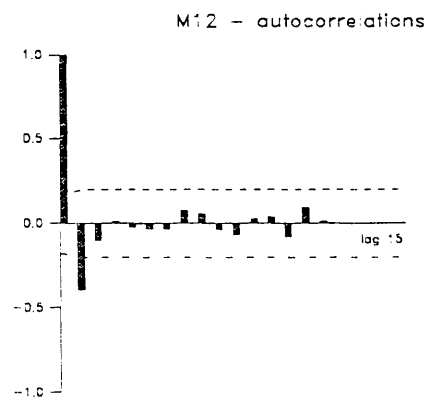
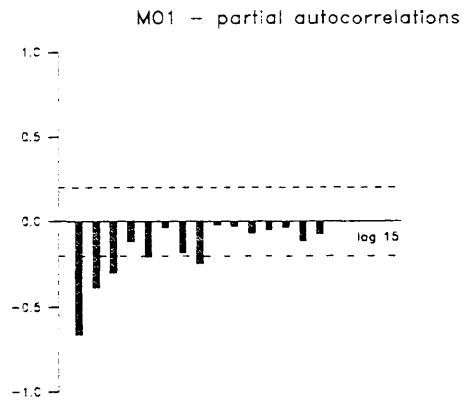
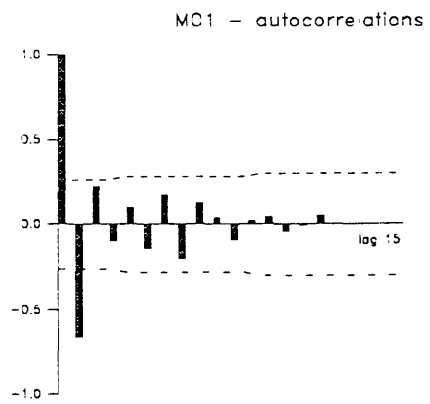


Figure 45: ACFs and PACFs for the CROSS time data in Figure 44. Dotted lines indicate confidence limits.

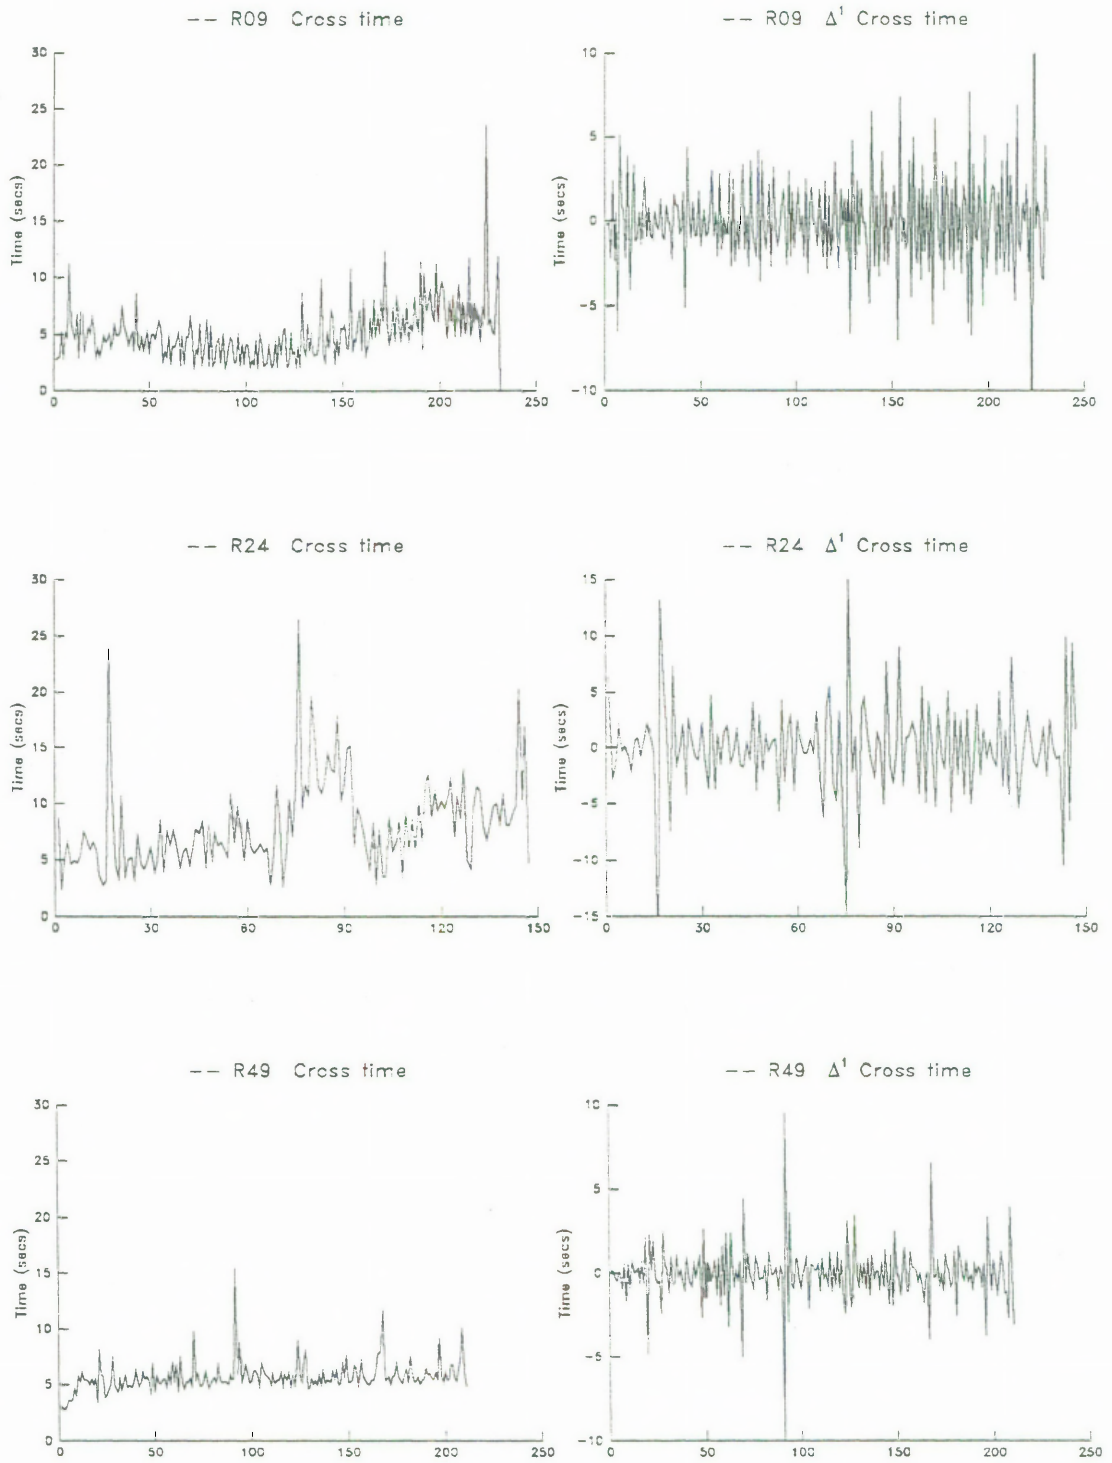


Figure 46: CROSS time and first differences (Δ^1) for R09, R24, and R49.

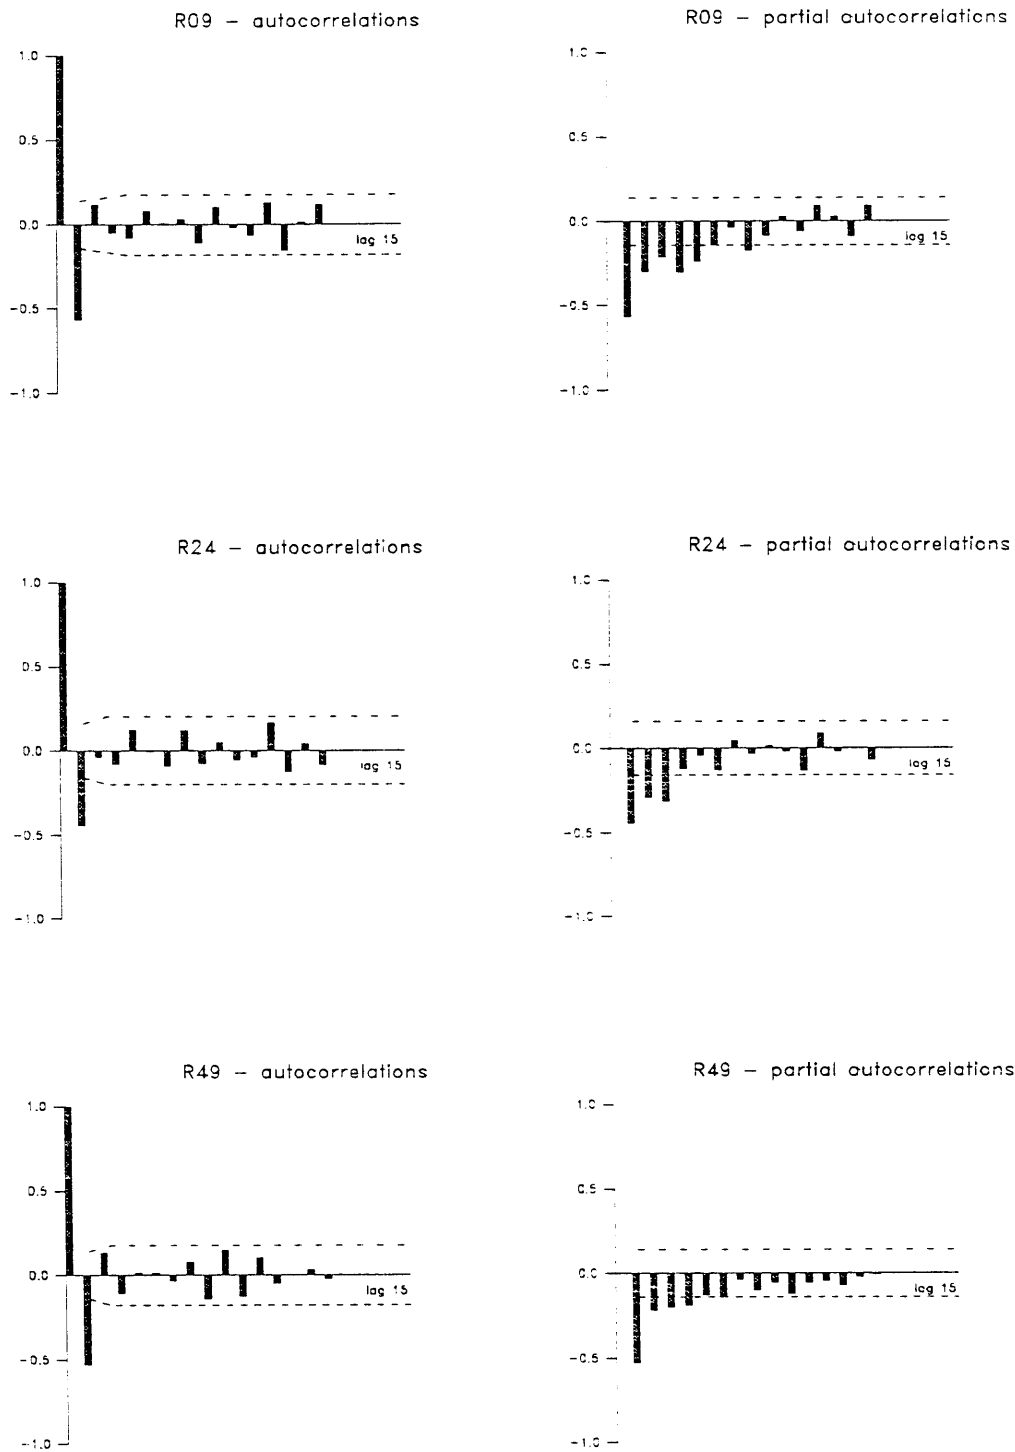


Figure 47: ACFs and PACFs for the CROSS time data in Figure 46. Dotted lines indicate confidence limits.

Input-output analysis

The CCFs for all data sets are shown in Figures 48– 52. The empirical, best-fitting model for ON time ('best-fitting' was defined as least residual mean square with no more than three terms of which one was MA(1)), was used to filter the differenced OFF time data. Residuals from the fitting of the model to the ON time data (called 'pre-whitening') and the residuals from the filtered OFF time data were then used as the basis for calculating the cross-correlation function (Anderson, 1976; Bennett, 1979; Box & Jenkins, 1976). The pre-whitening procedure is necessary to reduce the input series to a white noise process (Bennett, 1979; Gregson, 1983) and the filtering is to remove the correlation in the output series due to the correlation in the input series. The output series may then be considered as the result of a random input signal plus the effect of the neural process itself. Note that the ARMA model used for pre-whitening was not derived from the ACFs and PACFs provided in Figures 44– 47, as these were for the combined CROSS time data. The ON (and OFF) time data showed a very similar ACF and PACF structure to the combined CROSS time data (i.e., MA(1)) but there was often evidence of the effect of longer lags in the ON time series in particular (see Table 21 for filters used).

Figure 53 gives examples of the ON time series when that series is smoothed by the averaging of successive observations (i.e., high-frequency components are removed — Gottman, 1981). Differencing removes low-frequency components and amplifies high-frequency components, thus this figure shows more clearly the appearance of low frequency cycles (Gottman, 1981).

For the M-Series animals, there appeared to be some similarities with regard to strong cross-correlations at lag $k = -1$ (M11, M12, M18, M19, M26, perhaps M23). In these particular examples at least, ON time appears most associated with the preceding OFF time. Note, however, that for M11 the correlation was in the opposite direction to the others. Also to be noted is the lack of gradual decline in the correlations. Despite the correlation between ON time and the preceding OFF time being strong, there was no significant correlation with earlier OFF times (i.e., at lags $k = -2$ or less) or following OFF times (i.e., at lags $k = 0$ and greater). Also, the significant correlations tended to be near the 'typical' ON time and not randomly dispersed across the ± 10 range of the CCF.

Examination of the correlations calculated without pre-whitening and filtering (i.e., the variables $\text{Corr}(X, Y)$, $\text{Corr}(X, Y_{-1})$, $\text{Corr } D(X, Y)$, $\text{Corr } D(X, Y_{-1})$ —

Table 21: ARMA models

ARMA models used to ‘pre-whiten’ ON time and filter OFF time. The terms in the models were derived by an empirical, ‘best-fitting’ procedure and were restricted to no more than three terms of which one was MA(1). Higher order MA(q) terms and AR(p) terms used in the filters were not necessarily significant. The least residual mean square was used as the criterion for ‘best-fit’.

Subject	Model
M01	MA(1) – AR(1)
M08	MA(1,9)
M11	MA(1,4)
M12	MA(1,7)
M16	MA(1,2)
M18	MA(1,4)
M19	MA(1,10) – AR(1)
M23	MA(1,12) – AR(1)
M26	MA(1,9)
R09	MA(1,4) – AR(1)
R24	MA(1)
R49	MA(1) – AR(1)
R48	MA(1,2) – AR(7)
R48 High	MA(1,14) – AR(1)

— MA(1,4), for example, may also be written as MA(1) + MA(4).

Table 22: Relationship between CCF and correlation variables

Within-trial correlation variables as they were calculated in Experiment Ia (i.e., from Tables A.9–A.12, Appendix A). Data for R-Series animals is also provided.

Within-trial correlation variables				
Subject	Corr (X, Y)	Corr (X, Y_{-1})	Corr $D(X, Y)$	Corr $D(X, Y_{-1})$
M01	0.522	0.148	0.416	-0.323
M08	-0.052	0.262	-0.274	0.245
M11	0.059	-0.017	0.215	-0.012
M12	-0.085	-0.040	-0.305	-0.002
M16	-0.003	-0.053	0.083	-0.105
M18	0.038	0.055	-0.207	-0.034
M19	-0.132	-0.128	-0.296	-0.035
M23	-0.223	-0.024	-0.318	0.079
M26	0.036	0.033	-0.309	-0.074
R09	0.413	0.278	0.174	-0.174
R24	0.226	0.156	0.046	-0.071
R49	0.266	0.156	0.107	-0.100
R48	0.679	0.508	0.250	-0.135
R48 High	0.622	0.413	0.428	-0.367

Tables A.9–A.12, Appendix A), indicated that the best predictor for the strong correlation at lag at $k = -1$, was Corr $D(X, Y)$ (non-parametric correlation of -0.86 with $k = -1$ and 0.90 with $k = 0$). Pre-whitening and filtering methods can be laborious and consuming on computer processing time, so the important characteristics of the CCF may be estimated more quickly by using Corr $D(X, Y)$.

Rat M01 showed a strong cross-correlation at lag $k = 0$, that is, the immediately following OFF time. Still others showed no apparent relationship (M08, M16) although there was some suggestion of correlation at lag $k = +2$ in the M16 data and at both $k = -1$ and $k = +1$ for M08. The lack of a significant correlation with M16 might have been a reflection of the poor resolution in the OFF time data. The magnitude of the correlation coefficient is partly a function of the variability of the two measures (Roscoe, 1969).

The longer stabilised R-Series data showed little evidence for consistent CCF patterns. R09, R24, and R48 showed a number of significant correlations but these did not conform to any obvious pattern. These results might indicate that stabilisation

produces data in which the energy is spread over a greater portion of the series. However, the one example provided in Figure 51 for a rat at two intensity levels does not support this. The CCF at the 10% higher intensity did not provide any evidence for a more concentrated energy spectrum. The significant correlation at lag $k = 0$ remained, as did the significant correlation at $k = -4$, but other significant correlations were evident at the higher intensity (i.e., at lags $k = -7$ and leads $k = -1$ and $+10$).

The results indicate the importance of an IMA(1,1) model in ICSS behaviour, and, as a consequence, the lack of significant effects of recent CROSS times on the value of the current CROSS time. AR(p) terms were not strongly evident in the ACF and PACF structure. Because the same or very similar results were obtained when ON times was considered rather than CROSS times, the same conclusions may be made for ON time. That is, the reward system, does not show any strong capacity to remember its own past. Each ON time or CROSS time appears to be a function of the current stimulus value and is not modulated by the gradually dissipating effects of previous durations. This suggests that the use of individual ANOVAs, as in Experiment III (and based on 6-cross blocks), should produce reasonably valid probability estimates as autocorrelation does not extend beyond adjacent times. Apart from the negative autocorrelation at lag $k = 1$ in the ACFs for the differenced series (and which is largely induced by the differencing operation), in no case was there evidence for a significant correlation at other lags. The computer generation of a set of random numbers (from uniform or normal distributions) which are then differenced and analysed as above, produces ACFs and PACFs similar to those produced by shuttling behaviour.

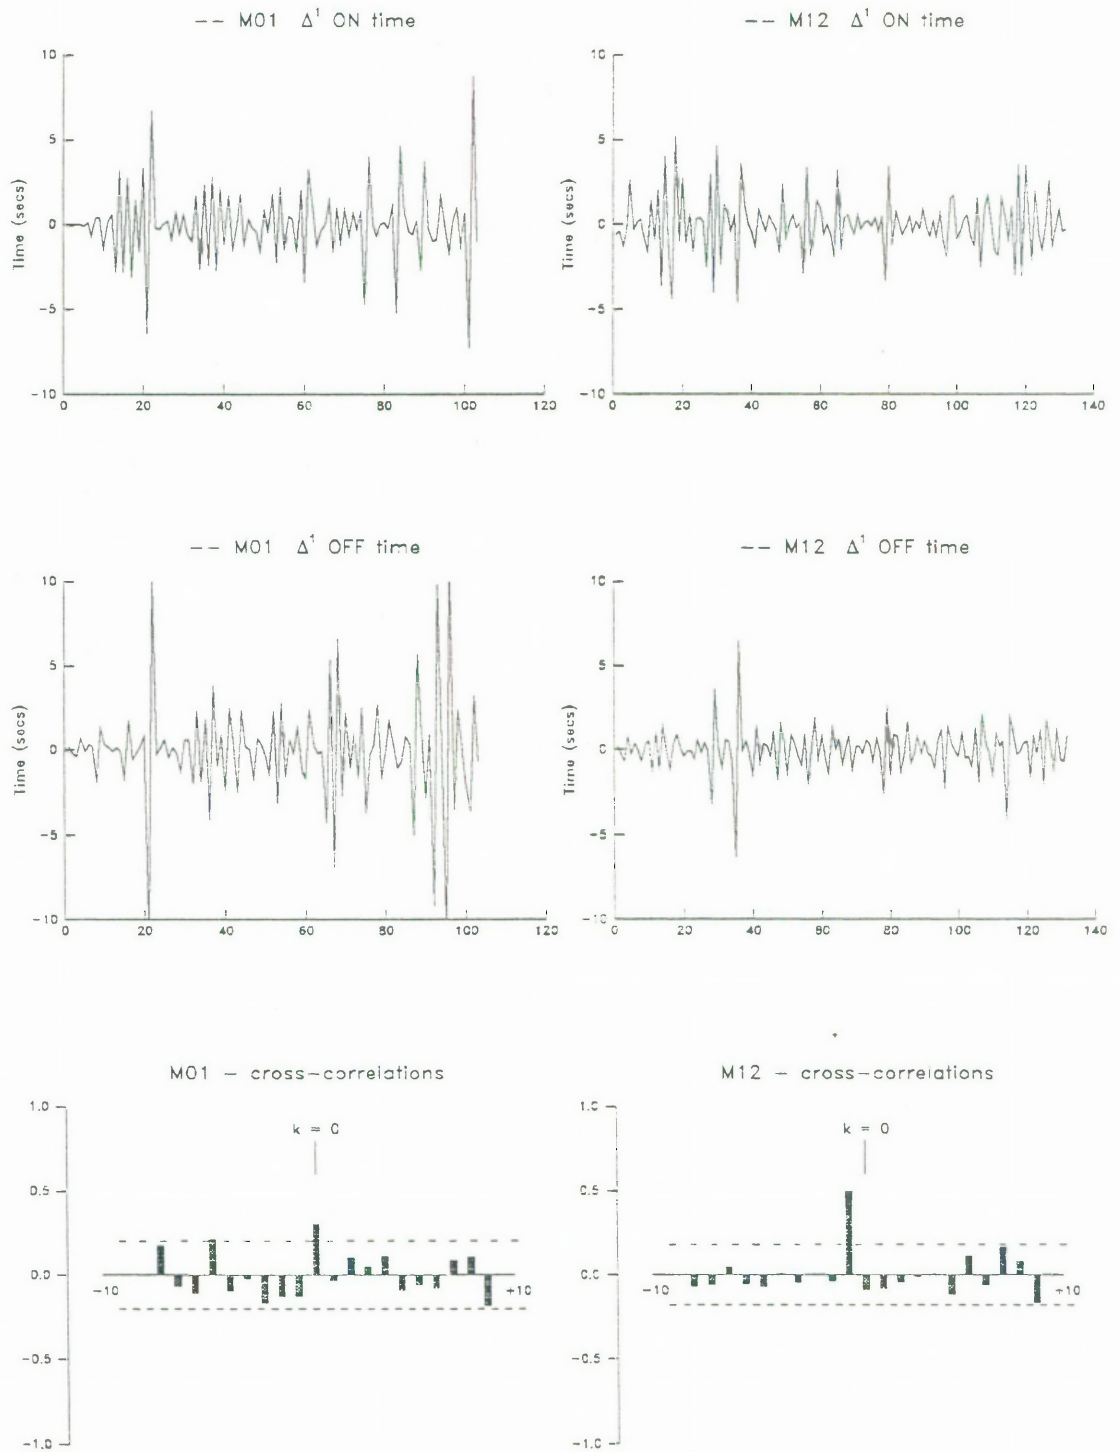


Figure 48: Input \Rightarrow output relationship for M01 and M12. Dotted lines indicate confidence limits i.e. twice the standard deviation.

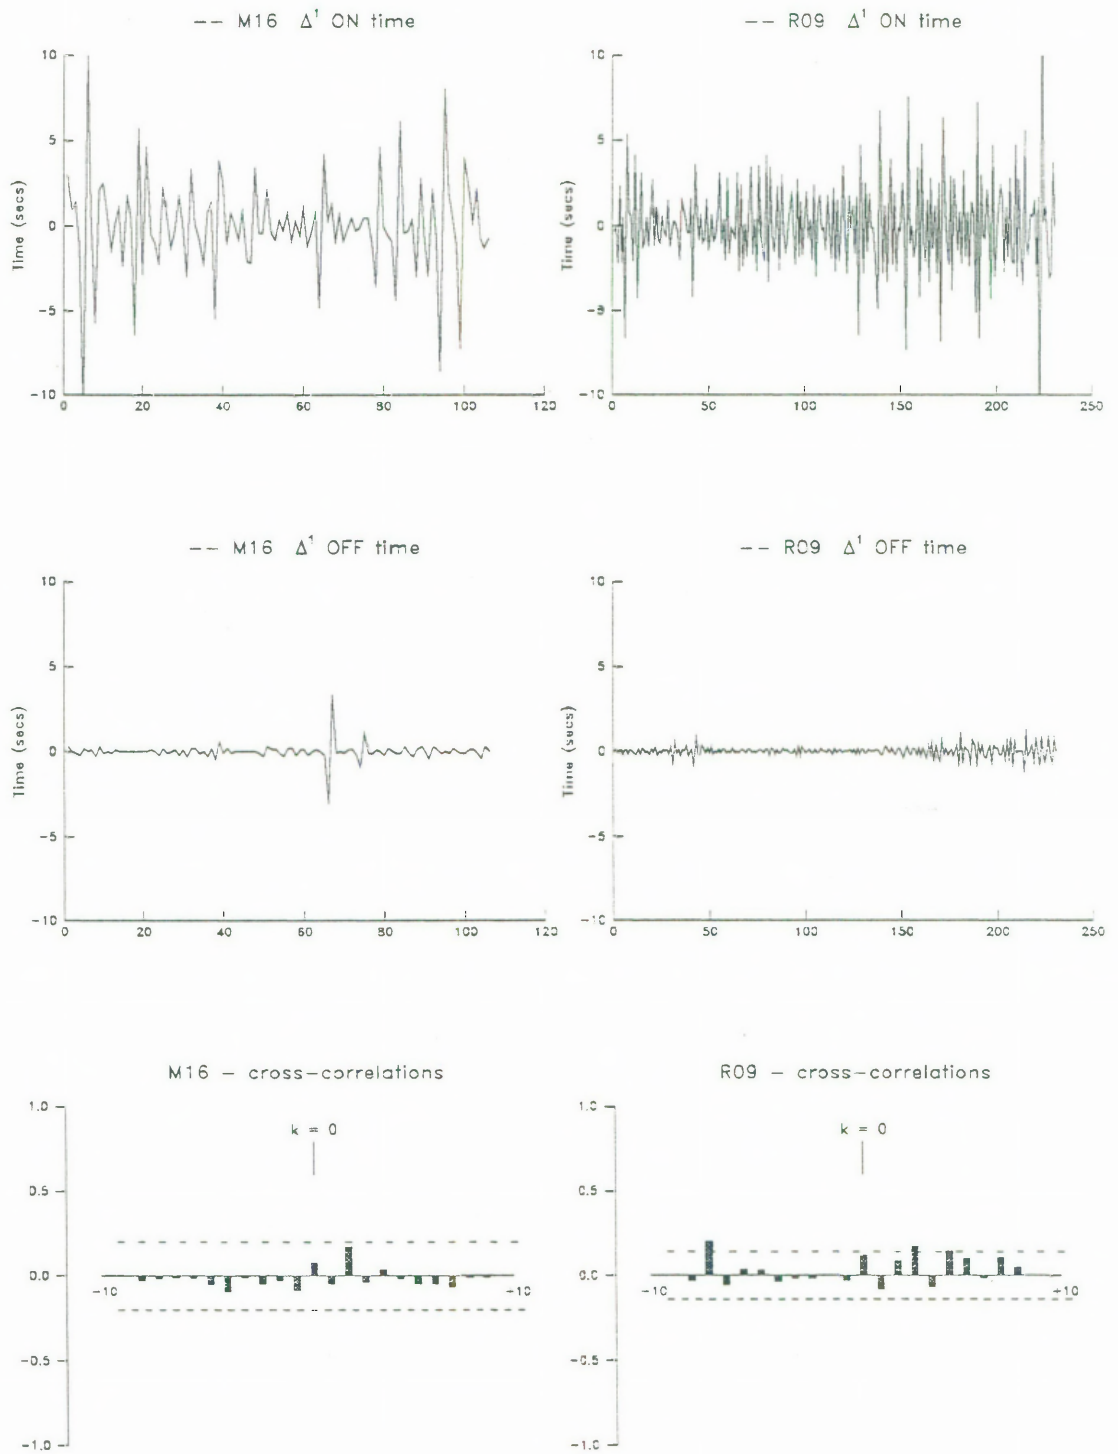


Figure 49: Input \Rightarrow output relationship for M16 and R09. Dotted lines indicate confidence limits i.e. twice the standard deviation.

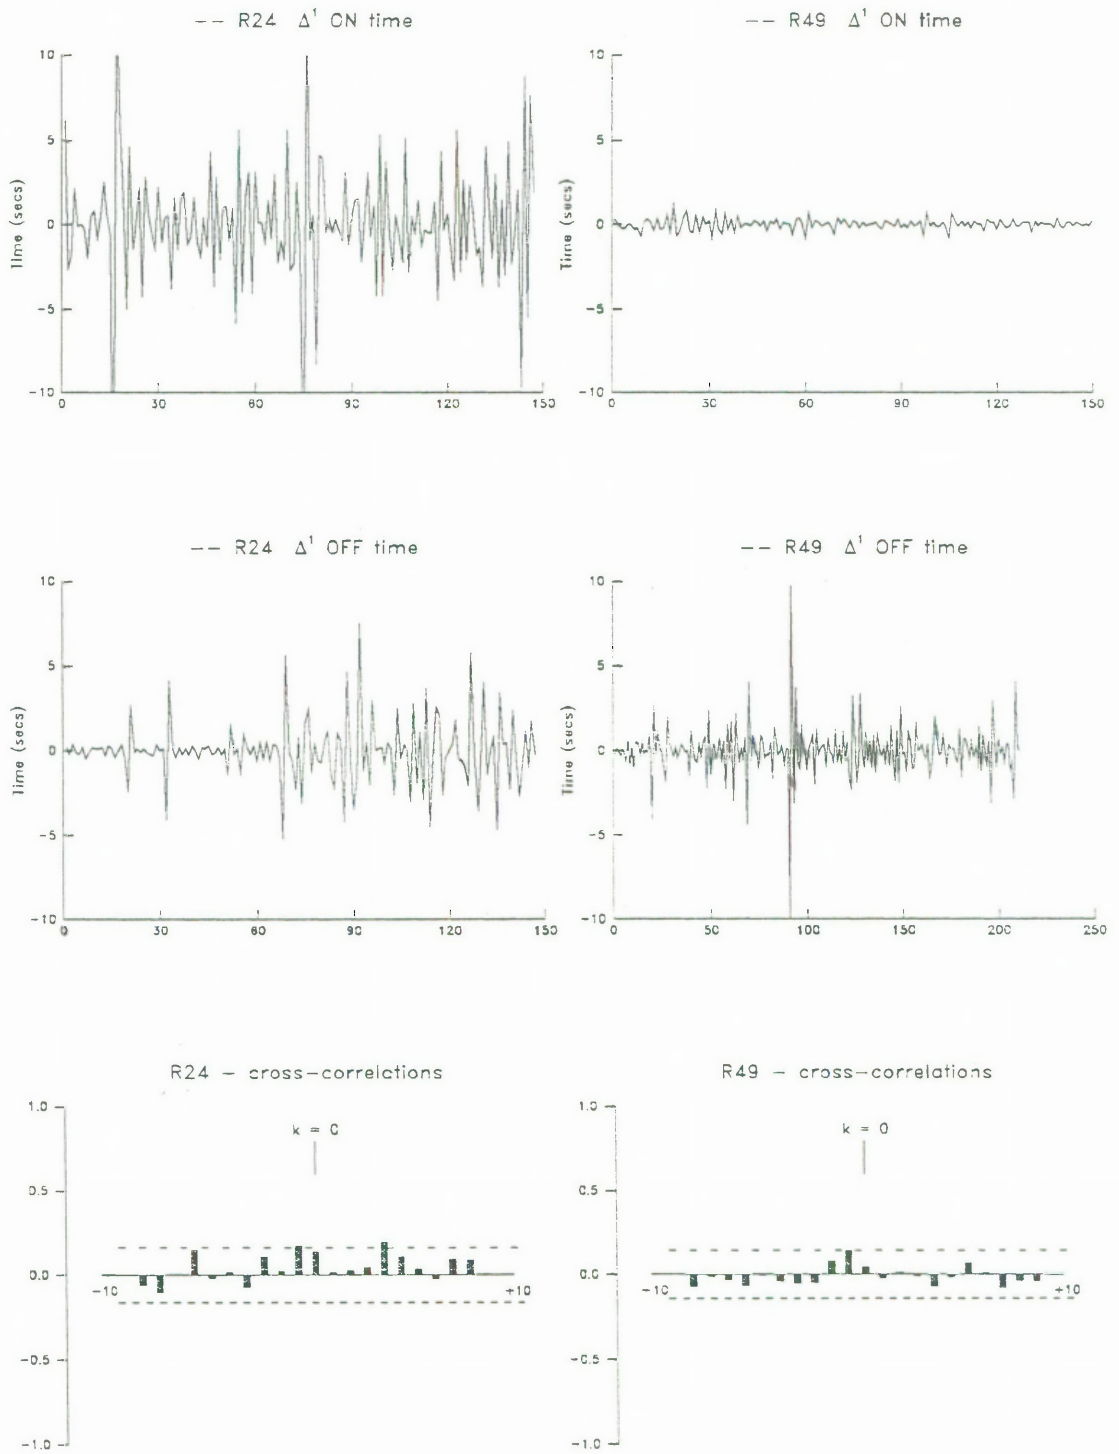


Figure 50: Input \Rightarrow output relationship for R24 and R49.

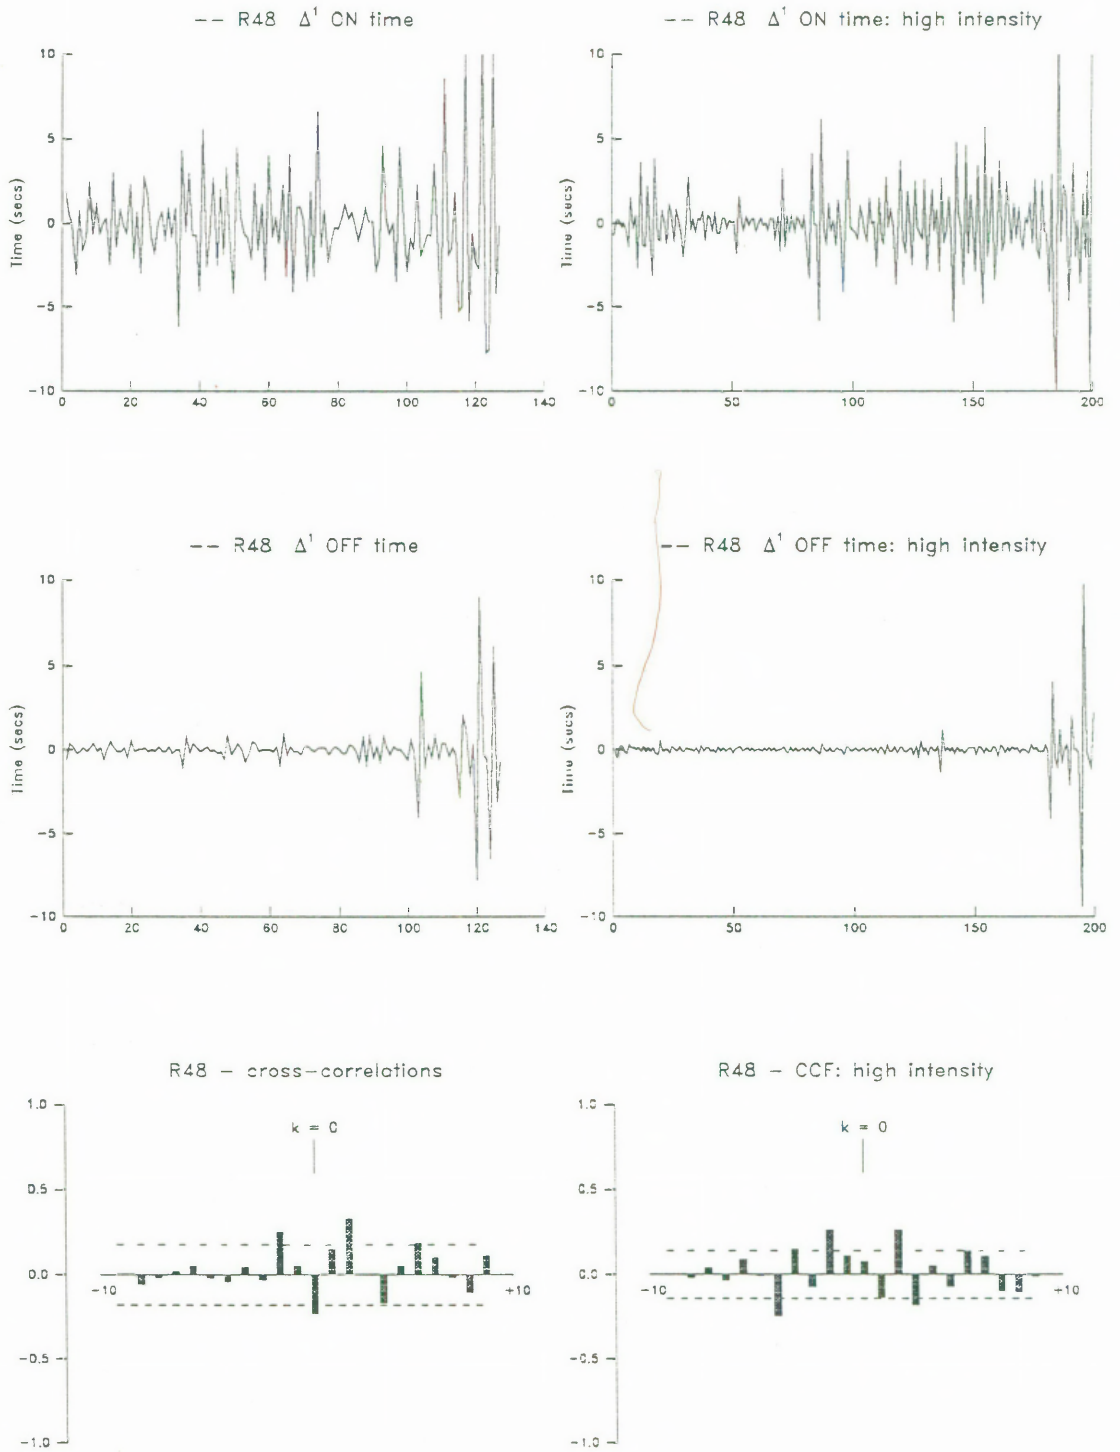


Figure 51: Input \Rightarrow output relationship for R48 at two intensities.

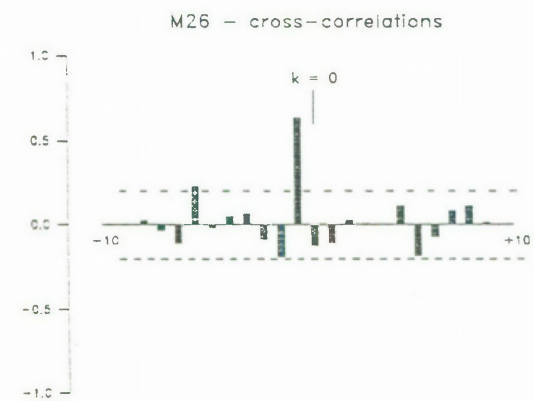
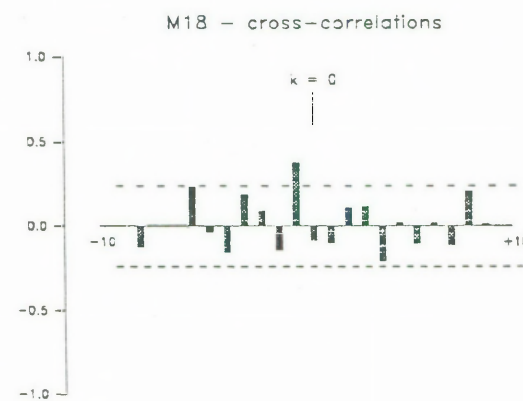

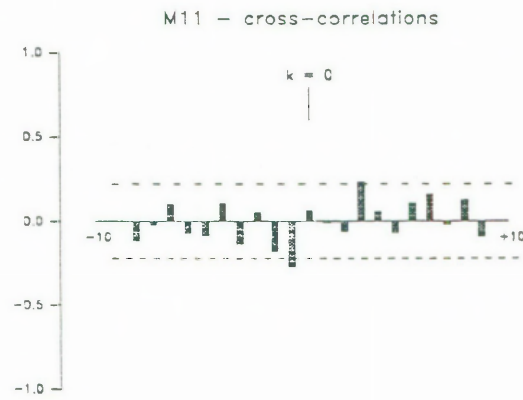




Figure 52: Remaining cross-correlation functions for M-Series data.

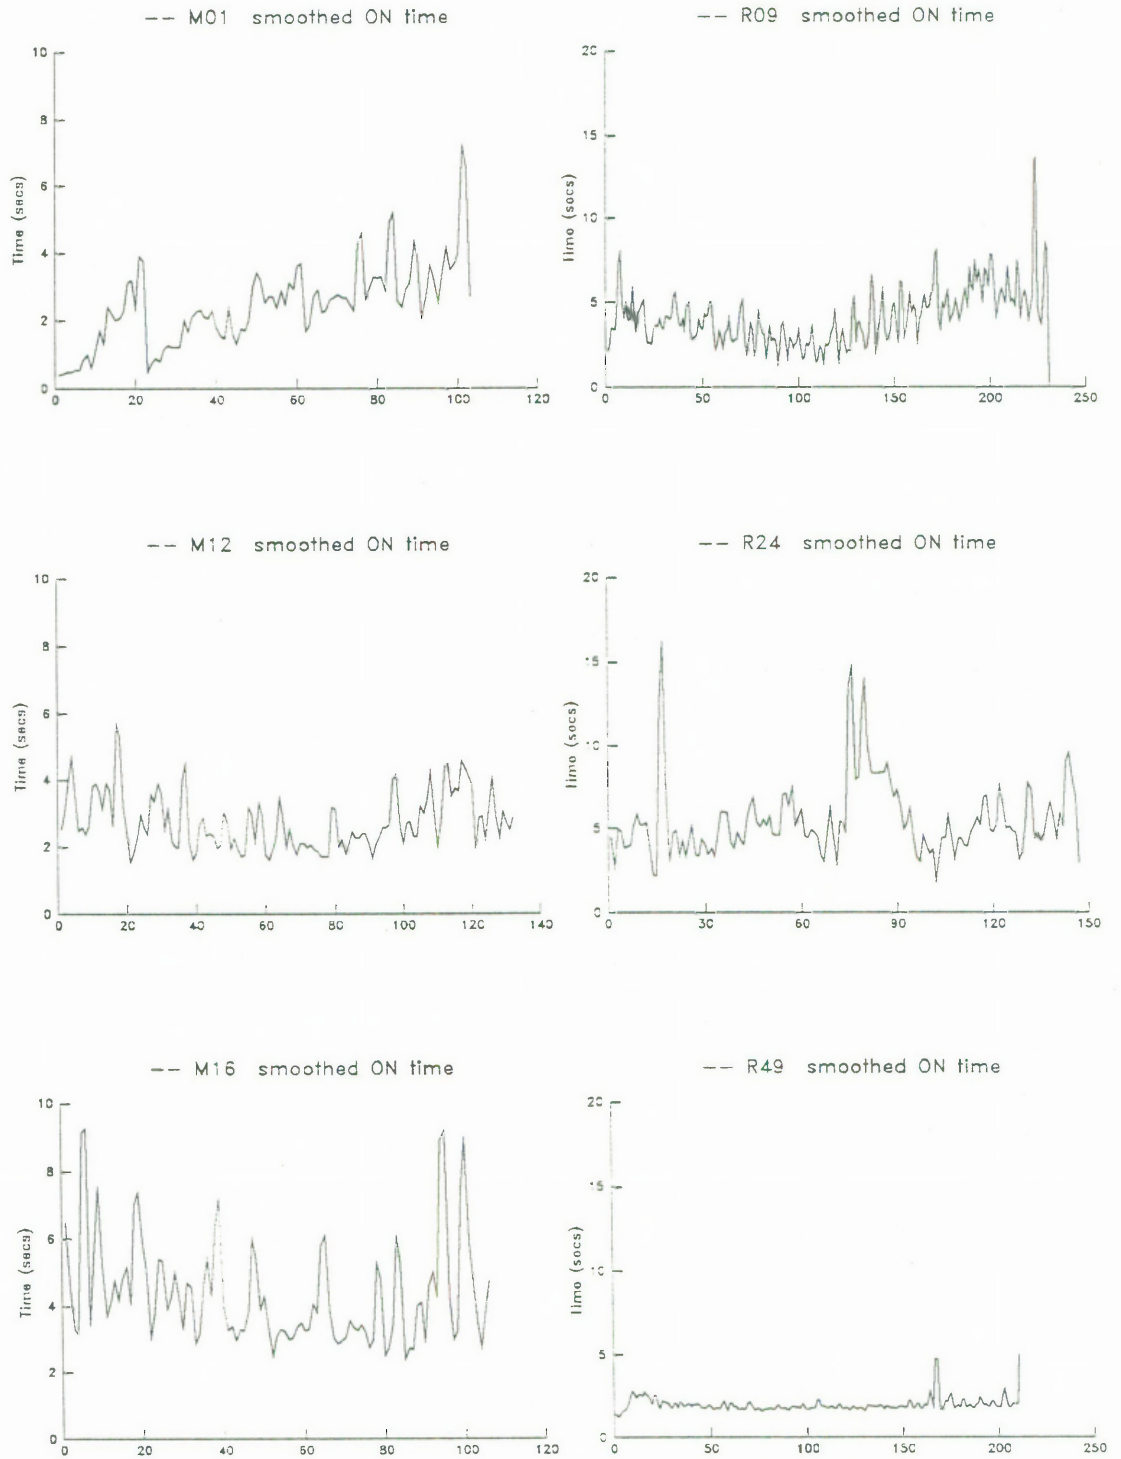


Figure 53: Smoothed ON time emphasising low frequency components.

9.3 Conclusions

The above analysis and discussion examined the relationship between OFF time and the duration of stimulation from a time series perspective. The main result was that at a constant intensity and frequency of stimulation, the sequence of ON times, OFF times, and CROSS times show little evidence for the effects of previous times on the current time. Each time appears to be generated largely on the basis of the current stimulus value only. This result demonstrates the stochastic nature of the generating mechanism and the dependence of the behaviour on the immediate 'shock' of the electrical stimulus. Differenced ON and OFF times may be considered randomly distributed, but the form of this distribution (e.g., normal, exponential) has yet to be determined. The indication that some rats do show highly specific correlations in some trials, and that these correlations tended to be associated with the 'typical' ON time, suggests that significant relationships can be present, but these might best be described as fragile. Further research is required into the factors that might determine when a rat's behaviour will show a strong correlation between ON and OFF time. The present analysis is also consistent with the phenomenon of rapid extinction often observed with ICSS behaviour on continuous reinforcement schedules. Because the behaviour is predominantly dependent on the current stimulation only, the behaviour should cease almost immediately once that stimulation is withdrawn.

Chapter 10

General discussion and conclusions

10.1 Main results

10.1.1 ON and OFF times

Self-stimulating rats respond to increased intensities of stimulation by selecting decreased durations of stimulation and decreased durations of no stimulation. As intensity is increased, OFF time tends to decrease at a faster rate than ON time, so that a greater proportion of the available time is spent with the stimulation ON. The charge accepted per initiation also decreases at higher intensities.

As a consequence of these concurrent changes the decrease in OFF time at higher intensities might result from either the increase in intensity or from the reduced duration of stimulation. Since the finding of decreased ON and OFF times represents one of the three characteristic features of shuttling behaviour that any model is required to explain, one aim of the present series of experiments was to differentiate the effects of intensity and duration on OFF time. Intensity might directly and independently decrease both ON and OFF time, or intensity might directly decrease ON time only, with the decrease in OFF time occurring only indirectly.

The results have shown that intensity is the principal determinant of OFF time. An increase in the intensity of stimulation produces significant decreases in OFF time that are largely independent of the decrease also seen with ON time. However, the results also show that at a relatively high intensity, an increase in duration of stimulation produces a small, but significant, increase in OFF time. In other words, there was also some support for the possibility that intensity has an indirect effect on OFF time. As intensity is increased over a central range, part of the decrease in OFF

time may be attributed to the concurrent decrease in the duration of stimulation. When intensity was fixed at a relatively high level, a fourteen-fold increase in the time over which that intensity prevailed resulted in an increase in OFF time from 1.67 secs to 2.71 secs. By contrast, a mean increase in intensity of 36% in Experiment Ia, resulted in a decrease in OFF time from 30.02 secs to 2.03 secs. These results indicate that although the duration of stimulation can have significant effects on OFF time, these effects are minor in comparison to the effects of intensity.

Experiment IIa was conducted at the one relatively high intensity level. The possibility has not been ruled out that duration of stimulation has a greater or even different effect at low to moderate levels of intensity. However, because the proportion of time ON was significantly greater at higher intensities, if duration was to be a major determinant of OFF time its influence was expected to have been greater at higher intensities.

Because an increase in intensity results in decreased OFF time, while an increase in the length of time over which the stimulation applies results in a slight increase in OFF time, these two aspects of electrical stimulation must reflect different neurophysiological properties. In particular, since OFF time, a time when no stimulation is received, is most significantly affected by the intensity of stimulation, a two-system basis for the ICSS neural substrate is supported. The changes seen with intensity would suggest that two neural systems were being activated during stimulation, both of these systems being predominantly dependent on the amplitude of the stimulating wave form.

The increase in mean OFF time that occurred when the duration of stimulation was increased might indicate that the neural reward system takes somewhat longer to recover after longer durations. According to the interpretation of the reward/aversion model followed here, the range of durations tested should not have produced significant aversion. Also, if OFF time inversely reflects the reward value provided by the intensity of stimulation, at a constant intensity OFF time should not change. The evidence is therefore more supportive of the proposal that some adaptation occurs in the neural substrate after prolonged stimulation. Perhaps the neurons that mediate the initial reward effect are still adapting at longer durations and hence take longer to recover. Alternatively, the increased OFF time might represent the time required for some recovery in the various autonomic effects of ICSS stimulation.

The question of which OFF time in a long series of ON/OFF times is most related to a particular set of stimulus parameters was not clearly answered. There was some

indication that an intensity increase on each third initiation resulted in a tendency for the OFF time immediately preceding the high intensity to decrease and the following OFF time to increase. This would indicate that OFF time is not simply elicited by the stimulation but results, at least to some extent, from qualitative comparisons being made between the different reward values provided by different intensities. The decreased OFF time before the high intensity might indicate an anticipatory effect and the increased post-high intensity OFF time a post-reinforcement pause effect. Since these effects are consistent with effects from natural rewards (e.g., Ferster & Skinner, 1957), the assumption that OFF time is inversely related to the intensity of stimulation may be validated. Pulse width had a similar but weaker effect to intensity.

When duration was increased on each third initiation there was also a tendency for the following OFF time to increase but indications of a decrease in the preceding OFF time were not so apparent. Only for those one or two cases in which the preceding OFF time did decrease might the longer duration be considered as providing an increase in reward value. For the remaining cases the longer OFF time might be more indicative of adaptation occurring in the neural substrate or recovery from autonomic activation.

Because of inconsistency in the manner in which animals adapted their behaviour to the stimulus contingencies the data can only be taken as suggestive. The inconsistency and general lack of effect of the stimulus changes might indicate an overriding influence of arousal or formication processes under continuous reinforcement that tends to blur the effect of temporary changes in stimulus conditions (Atrens, 1984). Perhaps the two low intensity conditions should have been even lower or an inter-block time-out introduced in order to reduce such effects. Liebman (1983) reports that post-reinforcement pauses in ICSS behaviour may only be apparent at moderate intensities, at higher intensities they tend to disappear. Also, other regions of the brain such as the lateral hypothalamus might be less influenced by arousal or formication than the ventral tegmental area.

10.1.2 Correlation data

The correlation between each ON time and the preceding, or succeeding, OFF time was examined throughout all experiments reported. There was no evidence to suggest

a consistent correlational structure (across animals or trials) in any of these experiments. Under some conditions, however, statistically significant effects did occur (although actual r s were low) and these will be listed and discussed briefly.

Firstly, Experiment I showed that the correlation between each ON time and the succeeding OFF time was significantly decreased by intensity (from 0.203 to 0.081). Experiment IIIb found a significant difference in the correlation between each ON time and the succeeding OFF time when intensity, rather than pulse width, was varied in a 2:1 format (-0.033 and 0.196 respectively). When the correlation between each ON time and the preceding, and succeeding, OFF time was averaged across all animals and conditions, significant positive correlations were found (Experiment Ia: 0.064 , 0.103 , respectively; Experiment Ib: 0.105 , 0.146 , respectively). However, in all these cases, if ON and OFF times were differenced once in order to remove linear trend from the data, the significant findings were eliminated. It has been argued here that the slight positive correlations are due to the tendency for ON and OFF times to increase over test periods of the order of 10 minutes.

The lack of significant change in the correlation between each ON time and the adjacent OFF times as a function of intensity (when the linear trend was accounted for) was in marked contrast to the changes seen with other behavioural measures. Mean ON time, mean OFF time, proportion of time ON, total time ON, mean charge, and total charge were all significantly altered as intensity was increased. The lack of effect of intensity on correlation indicates a high degree of independence between the (within-trial) ON and OFF times. The finding that the correlation between mean ON and OFF times was significantly positive at the four lowest levels tested suggests that relatively high intensities should be used in experiments that require independence between ON and OFF time when these are calculated from mean values.

Other significant results found with the correlation variables included a significant increase in the correlation between each ON time and the preceding OFF time (from -0.200 to 0.096) when OFF time was increased from 0.5 to 1.5 secs (Experiment IIb). Also, in the same experiment, a significant decrease occurred in the correlation between each ON time and the succeeding OFF time from the first to the third 30-cross trial (from 0.110 to -0.103). A 7% random variation in the timing of the onset of stimulation appeared sufficient to result in significant changes in the correlation between ON and OFF time: suggesting that the timing of the initiation response may be important in the determination of stimulation and no stimulation periods. The timing of the onset of stimulation might therefore be worthy of further investigation.

In Experiment IIIa a significant interaction was found between the constant ON and variable ON subgroups and the trials over which the experiment was conducted. The interpretation of this finding, however, is not readily apparent.

In general, the lack of consistent changes in the within-trial correlation variables described above suggests some caution in their interpretation until further experiments have been conducted. The overall evidence demonstrates that no strong or consistent correlation exists between each ON time and the adjacent OFF times, however, the presence of some significant findings suggests a weak relationship between the measures may still be possible. For a group of well-trained SSs (say 6–10) with electrodes in the VTM, tested for 3–4, 30-cross trials in a particular condition and at a moderate to high intensity, ON and OFF time are effectively independent whether calculated from mean values or from within-trial values. Atrens et al. (1983) found that the correlation between ON and OFF time declines rapidly over the first three crosses, from 0.354 to 0.099. The last correlation is consistent with the low positive values found throughout the present experiments (before differencing), suggesting that the decline does not continue past about the third cross.

In all three experiments, the correlation with the succeeding OFF time (0.103, 0.146, and 0.155) was slightly greater than the correlation with the preceding OFF (0.064, 0.105, and 0.114 for Experiments Ia, Ib, and IIIb, respectively) but these differences were not significant.

The relationship between ON and OFF time was also investigated from a time series perspective. The results again indicated that a consistent relationship between ON and OFF time (across animals or trials) does not occur. However, some animals in some trials did show highly significant correlations and these may have been more than just chance correlations. This is because they tended to occur near the 'typical' ON time and were not randomly dispersed across the ± 10 region of the cross-correlation function. Rather than a weak relationship existing between ON and OFF times, a fragile relationship might still occur but which is easily disrupted. The factors which determine if and when a significant correlation can develop between ON and OFF time need to be examined still further. Because prewhitening, filtering and model fitting methods can be very consuming on computer time it is suggested that the best indicator of a significant relationship within a particular trial, as estimated by the time series methods, is the correlation between each ON time and the succeeding OFF time for the differenced series.

Analysis of the autocorrelational structure of the series of ON times (and OFF

times) revealed that except for a tendency for linear trend (also described as ‘wandering’, or ‘drift’), there was no consistent autocorrelation. The neural system responsible for the generation of a series of ON times does not appear to be strongly influenced by previously generated ON times. The neural systems responsible for ON and OFF time appear to generate behavioural output in a stochastic manner independent of the stochastic output from the other system and independent of previous output.

In mathematical terms, a random variable observed over time is completely determined by its mean, variability, and autocorrelation function (e.g., Kedem, 1980). Because shuttling behaviour can be described in the same way, the essential information contained in the behaviour (or, in any particular response measure) may be obtained by measuring the relevant mean, standard deviation and autocorrelation function. These three measures contain independent information which may then be given interpretative meaning. It is suggested that more emphasis be given in further research to an assessment of variability and autocorrelation in shuttling behaviour.

The tendency for ON and OFF times to increase over a 10 minute period, or for baseline thresholds to shift over the period of an experimental regime, is not a desirable feature if the relationship between the times is important. Differencing is a statistical, post hoc method for dealing with measurement instability but implantation techniques that reduce this tendency such as the use of platinum electrodes rather than stainless steel (Bollinger & Gerrall, 1971) would constitute methodological improvements. Also, random variation in the intensity (and/or frequency) might lead to a truer response over a trial that is less confounded by the ability of the animal to predict the intensity of stimulation.

10.1.3 Models of ICSS shuttling behaviour

Shuttling behaviour may be described in terms of three characteristic features. These are: that rats terminate apparently rewarding brain stimulation; that once the stimulation has been terminated, the stimulation is also readily re-initiated; and that the self-selected duration of stimulation and the self-selected duration of no stimulation both decrease as the intensity of stimulation is increased. These features were considered to be the minimum conditions that any model of the behaviour must explain. The present study examined two models on this basis. Both the reward/aversion model and the reciprocal inhibition model require modification to explain the results

obtained.

Reward/aversion model

The hypothesis that a self-stimulating rat learns to terminate the stimulation because of impending aversiveness may explain the remaining two features if it is also assumed that the rat terminates whenever a fixed level of aversiveness is reached and that OFF time inversely reflects the reward value obtained on initiation.

Because an experienced self-stimulator normally terminates before aversion has accumulated to punishing levels, the use of test durations less than or equal to the preferred duration should not produce punishing levels of aversion. Also, if intensity alone determines the reward value obtained on initiation, OFF time should not change if intensity remains constant. Atrens et al. (1983) have suggested that even though very long durations do become aversive, the initial reward effect adapts and it is this adaptation that determines termination under continuous reinforcement. The suggestion from Experiments IIa and IIIa that some adaptation does occur in the neural substrate would support the proposal that the initial reward effect adapts. Whether this adaptation alone is sufficient to lead to termination, however, is not likely in view of the feature that ON time decreases as intensity increases. This is because at higher intensities the reward effect should take longer to adapt, adapt to a higher level, or remain unchanged if independent. This reasoning rests on the assumption that self-stimulating rats normally terminate well before significant aversion has accumulated, and the evidence for this has only been indirectly derived (Atrens et al., 1983). A more direct comparison between preferred duration under CRF and the duration at which the stimulation becomes aversive under partial schedules needs to be determined.

The present evidence suggests that the best interpretation of the reward/aversion model would include the property that the initial reward effect begins to adapt from the moment of initiation while aversion starts to accumulate (Atrens et al., 1983). Rate of adaptation must be considerably less than rate of accumulation of aversion. When the difference between the two quantities becomes less than some criterion, a signal is produced which motivates the animal to terminate the stimulation. In order to account for the independence of ON and OFF times the model might incorporate stochastic functions for the rate of accumulation of aversion and the reward value obtained on initiation. These functions might represent the normal biological variation

inherent in the neural substrate.

Reciprocal inhibition model

The reciprocal inhibition model describes a neural switching mechanism dependent on cross-inhibitory coupling between two neurons or between two systems of neurons. The periods or cycles of dominance of one system over the other have been shown to be dependent on the amplitude or frequency of the incoming pulses. As the intensity or frequency of the incoming pulses is increased, the cycles of dominance decrease (Ludlow, 1976, 1980; Reiss, 1962). The RI model is therefore capable of explaining the basic features of shuttling behaviour as described previously.

The results of Experiment IIa and IIb in which either ON or OFF time was controlled do not support the RI model. If ON and OFF time reflected the periods of dominance of one neural system over the other, and both were linked in a reciprocally inhibitory manner, the same changes should have been observed whether ON or OFF time was controlled. When ON time was increased OFF time was significantly increased; however, a similar increase in OFF time had no effect on ON time. To the extent that these results indicate adaptation in the neural substrate during ON time, the RI model may be applicable. However, the RI model needs to be modified to account for the lack of a similar effect when OFF time was controlled. One possibility might be to introduce a third neuron capable of gating, or modifying in some way, the inhibition in one direction.

Also, an RI model predicts that no changes should occur in ON or OFF time as a function of the degree of experience with the stimulus parameters. The interaction of the two neural systems and the cycling periods of dominance are presumed to be determined by the parameters of stimulation only, and should not be altered by increased experience. The evidence from the within-trial studies (Experiment III) suggesting that OFF time can change with experience indicates an element of operant control over OFF time.

Although OFF time might have changed with greater experience, ON time, in general, remained unchanged. This was also true for the case in which longer ON times occurred on high intensity stimulations. This would indicate that ON time is relatively invariant and more strongly dependent on the parameters of stimulation than is OFF time. Therefore, in well-trained self-stimulators the termination response may be approximated by a process of elicitation. This interpretation would

therefore support Fibiger's (1978) claim that the termination response was a respondent response (elicited by the preceding stimulation) while the initiation response was an operant response (reinforced by the consequences of the response).

It is unlikely that the processes that determine OFF time can be represented in exactly the same way as those that determine ON time. OFF time is more likely to reflect the sum activity from several feedback and/or reverberatory circuits that function in ways different to the circuits activated during direct stimulation.

The relation between the proposed models and the known neural substrate also needs to be considered. The three features may be determined by the second stage neural network (Gallistel et al., 1981) only after the direct stimulation has been transformed into a more neurophysiologically interpretable set of signals, or the directly stimulated substrate may be involved (Skelton & Shizgal, 1980). For an electrical stimulus of a given strength, the number of impulses arriving at the second stage network might be sufficient to activate a mechanism that produces the three features and the directly stimulated substrate would play no role. However, because adaptation probably does occur in the directly stimulated substrate, because the reward effect appears to be almost immediate, and because of a major element of elicitation, the involvement of the directly stimulated substrate in determining ON time is implicated. Because evidence implicating involvement of second stage networks has mostly come from lever pressing or alley running measures (which are more correlated with OFF time), second stage networks might be more applicable to the determination of OFF time.

Finally, the two models need not be completely exclusive because the point of switch-over in an inhibitory mechanism might be accompanied by aversion at some, if not all, ICSS sites.

A speculative explanation of shuttling behaviour in the VTM region might include the following (a modification of Wise & Bozarth, 1984). Initiation of stimulation rapidly excites the directly stimulated substrate which, in turn, stimulates dopaminergic and/or opioid neurons in a post-synaptic network (exactly how this might happen is not clear). The release of dopamine is responsible for the reward effect, but the accumulation of dopamine in the synaptic cleft also self-inhibits and/or collaterally inhibits (German et al., 1980) those same dopaminergic neurons. Combined with the gradually adapting directly stimulated substrate (and therefore less production of dopamine), inhibition of the dopamine neurons leads to a rapid decline in the reward effect. An increase in intensity would accelerate this process (in

a way similar to the RI model). Higher intensities lead to a more rapid initial release of dopamine and are therefore more rewarding. After the stimulation has been terminated, the directly stimulated substrate recovers its responsiveness, the excess dopamine is taken up by the terminals and the process is ready to start again. The gradual increase in ON and OFF times over a period of ICSS interaction might be indicative of a gradual depletion of transmitter reserves because of a less than complete reuptake processes.

Intensity must also activate a second system during stimulation which is a major contributor to the time an SS leaves the stimulation OFF. The output from this second system is uncorrelated with the output from the first. Perhaps opioid neurons performing different functions, or different types of dopamine neurons (Creese et al., 1983; Joyce, 1983), or dopamine neurons with more remote connections (e.g., nucleus accumbens) are stimulated. The suggestion from Skelton and Shizgal (1980) that two different populations of directly stimulated neurons are responsible for ON and OFF reponding may be the most appropriate explanation for the present findings. Both populations are directly activated during stimulation but one population responds more slowly. If the different populations have different spatial distributions but similar excitability characteristics, the results of moveable electrode studies which tend to show gradual changes in responsivity as the electrode is lowered (e.g., Corbett & Wise, 1980), may be more easily accommodated. Memory for the initial rewarding effect must also return during OFF time.

Obviously, the above is very speculative, particularly with the regard to the putative role of dopamine. Also, whether the above speculation actually accounts for the basic features of shuttling behaviour is difficult to determine without a simulation approach. Formalisation of the proposed models into mathematical relationships and subsequent computer simulations would help assess their suitability. Also, the effect of different assumptions (e.g., different spatial densities of the two systems) could be assessed by their effect on compatibility with anatomical and experimental information.

10.2 Conclusions and further research

The experiments reported here examined the relationship between ON and OFF time and the intensity of stimulation. The results indicated that no consistent correlation existed between the within-trial times at any level of intensity. The results also

established that the most significant determinant of OFF time at a relatively high intensity is the intensity of stimulation itself rather than the duration over which the stimulation is applied. An increase in the duration of stimulation produces a slight but significant increase in OFF time whereas an increase in intensity produces large, and highly significant, decreases in OFF time. These results support the conclusion that two neural systems are activated during stimulation; both of these systems are most dependent on the amplitude of the stimulating pulses and both systems generate stochastic output independent of the stochastic output produced by the other and independent of previous output.

Because OFF time is largely a function of intensity and not the duration, and because it was suggested (Experiment I) that crossing rate and charge interacted over some number of crosses to maintain a constant total charge, further research might investigate the effect of random variation in the intensity of stimulation rather than the constant intensity usually employed. ON and OFF time might then be correlated directly with intensity and might more clearly reveal the relationships among these variables. (Variation in frequency would not alter the number of neurons involved and might therefore be a further improvement.) If variation in the intensity of stimulation over a trial followed predetermined probabilistic statements (e.g., as a time series), analysis of the resulting behaviour might be in terms of how the behaviour follows these probabilities (Weiss, 1962; Weiss, Laties, Siegel, & Goldstein, 1966).

The use of shuttling behaviour under continuous reinforcement as a measure of an experimental manipulation (especially drug studies), would be improved by the establishment of parametric standards at which the behaviour has definable properties. The use of 100 Hz frequency already seems to be an accepted standard, and monopolar, cathodal pulses of 0.2 msec duration may be a better standard than used here. However, the present research suggests that a reasonable standard for intensity would be an intensity 5–10 % above the point on the rate-intensity function for each rat that best estimates the greatest rate of change in crossing rate. For grouped data, initiation rates increase linearly over mid-range intensities, but because correlations between mean values may not be independent at lower levels, the intensities used in studies that assume independence between mean times should be selected from the higher end of this range.

Alternatively, a standard mean and variability might be specified for the ON and OFF times for each rat. Schmitt et al. (1981) have shown that ON and OFF times are separately manipulable by varying pulse width /interpulse combinations so that

a unique frequency may be found for each rat that approximately equates baseline times.

In terms of measuring the behaviour, the present research suggests that the cross-correlation function provides the most comprehensive way of assessing the relationship between within-trial ON and OFF times. However, further research might examine the use of binary time series methods of measurement (Kedem, 1980; Keenan, 1982a,b). Comparability of baseline measures across laboratories might be improved by quoting proportion of time ON and mean charge accepted in addition to mean ON and OFF times.

Investigation of the behaviour that results from electrical stimulation of the brain's reward system may eventually allow some insight into how and why an animal responds to the various rewards and motivating circumstances in its environment. Although ICSS shuttling behaviour under continuous reinforcement may have interpretative difficulties (Atrens, 1984), and may not even represent a unitary phenomenon (Phillips, 1984; Redgrave, Dean, & Andrews, 1981), it remains a powerful and intriguing behaviour worthy of continued investigation.