

Classical conditioning as a nonstationary, multivariate time series analysis: A spreadsheet model

C. R. GALLISTEL

University of California, Los Angeles, California

The implementation of the Gallistel (1990) model of classical conditioning on a spreadsheet with matrix operations is described. The model estimates the Poisson rate of unconditioned stimulus (US) occurrence in the presence of each conditioned stimulus (CS). The computations embody three implicit principles: *additivity* (of the rates predicted by each CS), *provisional stationarity* (the rate predicted by a given CS has been constant over all the intervals when that CS was present), and *predictor minimization* (when more than one solution is possible, the model minimizes the number of CSs with a nonzero effect on US rate). The Kolmogorov-Smirnov statistic is used to test for nonstationarity. There are no free parameters in the learning model itself and only two parameters in the formally specified decision process, which translates what has been learned into conditioned responding. The model predicts a wide range of conditioning phenomena, notably: blocking, overshadowing, overprediction, predictive sufficiency, inhibitory conditioning, latent inhibition, the invariance in the rate of conditioning under scalar transformation of CS-US and US-US intervals, and the effects of partial reinforcement on acquisition and extinction.

Gallistel (1990, in press) describes a model of the classical conditioning process in which it is assumed that what the animal learns is the rate of unconditioned stimulus (US) occurrence to be expected in the presence of a conditioned stimulus (CS). The model predicts many experimental results that have posed problems for past or present associative models of conditioning, including: (1) the effect of partial reinforcement on the rate of acquisition and the rate of extinction (Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980); (2) the effect of the duty cycle (or ITI/ISI ratio) on the rate of acquisition and its lack of effect on the rate of extinction (Gibbon, Baldock, Locurto, Gold, & Terrace, 1977); (3) blocking and overshadowing (Kamin, 1969); (4) the blocking effect of background conditioning (Rescorla, 1968); (5) the effects of having the "background" USs signaled by another CS (Robbins & Rescorla, 1989); (6) inhibitory conditioning when the CS and US are explicitly unpaired; (7) the predictive sufficiency results of Wagner, Logan, Haberlandt, and Price (1968), in which the CS that accounts for more of the variance in US occurrence is the CS that gets conditioned; (8) inhibitory conditioning in overprediction experiments (Kremer, 1978); and (9) the noninhibitory retarding effect of a "latent inhibition" training phase on the rate of subsequent conditioning (Reiss & Wagner, 1972).

The model assumes that USs (brief shocks or small food pellets) are point events. It treats these point events as if they were generated by processes obeying Poisson statistics. The following are fundamental assumptions of the model: (1) It is the rate of US occurrence in the presence of a CS that is relevant to conditioning, not the probability of US occurrence. (2) The effects of different CSs on the rate of US occurrence are assumed to be additive until the data show that the rate predicted by two CSs acting concurrently is not the sum of the rates predicted when each CS acts in isolation. (3) The influence of a CS on the rate of US occurrence is assumed to be stationary until the data show statistically significant deviations from stationarity. (Experimental extinction is an example of nonstationarity, because the rate of US occurrence predicted by the CS prior to extinction training is not the same as the rate predicted during the extinction phase of training. The latent inhibition paradigm, in which there is a "preconditioning" training phase during which CS occurrence does not predict US occurrence followed by a conditioning phase in which it does, is another instance of nonstationarity.) (4) When the data are ambiguous about which CS predicts what fraction of the observed rate of US occurrence, the model gives the solution that minimizes the number of predictors, the number of CSs that have a nonzero effect on the rate of US occurrence. (5) The strength of the conditioned response is a sigmoidal function of the likelihood that the rate of US occurrence so far observed in the presence of a given CS is greater than would have been predicted in its absence.

This paper describes a spreadsheet implementation of the model. The implementation permits those without training in algebra to derive from the model the results

I am grateful to Tom Wickens for discussions that led to the algorithm for testing stationarity. The costs of creating this spreadsheet and preparing the manuscript were partially covered by NSF Grant BNS89-96246. Correspondence should be addressed to the author at the Department of Psychology, 405 Hilgard Ave., University of California, Los Angeles, CA 90024-1563.

that it predicts for novel training protocols—the predicted results of experiments not yet performed. The implementation also permits detailed scrutiny of the intermediate stages of computation, facilitating an understanding of why the model predicts what it predicts. The implementation described here is done in Excel, but it should be possible to implement it in any spreadsheet program that supports matrix operations (the calculation of determinants, matrix inversion, and matrix multiplication).

STRUCTURE OF THE SPREADSHEET

The layout of the spreadsheet is shown in schematic form in Figure 1. It has five sections. Section 1 contains the columns into which the training protocol for an experiment is entered. Section 2 contains the accumulation columns. Here, the model computes the running totals for the amount of time that each CS and each pairwise combination of CSs has been present and the total number of US occurrences in the presence of each CS. The computation of the rates of US occurrence predicted by each CS depends solely on these temporal and numerical totals. Section 3 contains the matrix operations and the condition-checking operations for computing these rates. Section 4 computes the likelihood that a given CS has had a non-zero effect on the rate of US occurrence, and finally, the subjective confidence that the animal accords to the hypothesis that a given CS affects the rate of US occurrence. This confidence determines the strength of the conditioned response. Section 5 computes the expected and observed number of USs to be imputed to the influence of each CS at the time of each event, as well as the Kolmogorov-Smirnov statistic, which is the difference between the expected and observed numbers of occurrence normalized by (divided by) *n*, the total number of observations. It also computes a crude approximation to the likelihood that this normalized deviation is greater than is to be expected by chance—a crude index of the likelihood that the series has not been stationary.

Section 1: Training Protocols

The model takes as input the experimental protocols, which specify the times at which the conditioning events occur (Figure 2, left three columns). The time at which

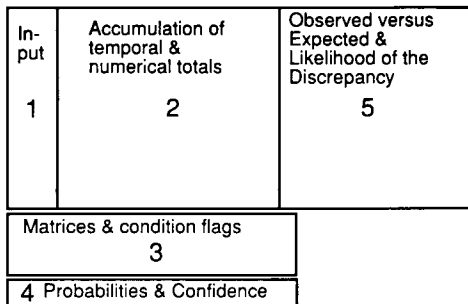


Figure 1. Schematic showing the layout of the subsections of the spreadsheet and their functions.

Event Time	Totals:	60	10	10	5	5		
	C1	C2	US	TC1	TC2	TL2	NC1	NC2
0	1	0		60	10	10	5	5
10	1	1		50	10	10	5	5
12	1	0	1	48	8	8	5	5
22	1	1		38	8	8	4	4
24	1	0	1	36	6	6	4	4
34	1	1		26	6	6	3	3
36	1	0	1	24	4	4	3	3
46	1	1		14	4	4	2	2
48	1	0	1	12	2	2	2	2
58	1	1		2	2	2	1	1
60	1	0	1	0	0	0	1	1

Figure 2. A sample of the input columns and the accumulation columns. The first five trials of a simple conditioning protocol with a 2-min tone (C2) and a 10-min intertrial interval (CS offset to next onset). Only the columns for the two CSs (the tone and the background = C1) are shown. Note that the total time that the two CSs have been on together (TL2) is the same as the total time that C2 has been on, because the background (C1) is present when the tone is on as well as when it is not.

an event occurs is entered in the first column (headed Event Times). An event is the onset of one or more CSs, the offset of one or more CSs, or the occurrence of a US. If events coincide in time—for example, if two CSs come on at the same time, or if a US occurs at CS offset—only one time is entered. The four protocol columns to the right of the event time columns are for up to four CSs (only two are shown in Figure 2). A 1 is entered in the column for a given CS if that CS was present at the time of the event. If the event was the onset of that CS, a 1 is entered. If the event was the offset of the CS, a 0 is entered, or, better, the cell is left blank. Excel treats blank cells as containing 0s, so it is not necessary to actually enter a 0.¹ The protocol is easier to read if one uses blanks in place of 0s. Because of these conventions regarding the recording of CS presence at the onset and offset transitions, the occurrence of a CS is often indicated by a single 1, marking the onset of the CS. This happens whenever the next event after the onset of the CS is its offset—that is, whenever no other events intervene between the onset and offset of the CS. In entering and reading the protocols, it is important to remember that in such cases, the CS was present from its onset time (given in the row marked by the 1) up until its offset time (given in the next row down).

The sixth protocol column is for the US (fourth column in Figure 2). A 1 in this column indicates the occurrence of the US at the indicated event time. Unlike CSs, USs do not last from the onset time to the offset time indicated in the next row down. USs are point events; their onsets and offsets coincide.

In the absence of an event time, cells in the Event Time column should contain the =N/A() formula.² When no number is entered in a cell containing this formula, the cell displays #N/A, which indicates to computational cells elsewhere on the sheet that the sequence of events does not extend down far enough to fill that cell. This prevents cells with computational formulas that read event time cells from reading a blank cell as an event time of 0.

Before entering a new protocol, the CS and US columns should be made blank, and the Event Time column should be filled with the =N(A) formula. This is done with an initialization macro, which is run before one enters a new experimental protocol.

Section 2: Accumulation Columns

These columns accumulate durations and numbers of occurrences retroactively from the most recent event; hence, the higher in the column one looks, the greater the accumulated total (Figure 2, five columns on the right). The entry at the top of the TC1 column gives the accumulated interval over which CS₁ has been present since the start of training (or, more precisely, since the time indicated in the first entry under Event Times); the entry at the top of the T1.2 gives the accumulated interval over which CS₁ and CS₂ have been simultaneously present; the entry at the top of the NC1 column gives the number of USs that have occurred in the presence of CS₁; and so on. As the protocol lengthens, the totals at the tops of these columns—and all the intervening subtotals from the lowest (most recent) entry in a column up to the top—increase. The formula entered in each cell of a single TC column is

$$=IF(AND(NOT(ISNA($A8)),B7=1),$A8-$A7+H8,H8).$$

(The \$ is the column- or row-freezing symbol in Excel. It forestalls a relational change in the column or row referred to, when the formula is copied into another column or row.)

The formula for the pairwise CS columns (e.g., the column that records the time that CS₁ and CS₂ have been on together) is

$$=IF(AND(NOT(ISNA($A8)),AND($B7=1,C7=1)),$A8-$A7+L8,L8).$$

The formula for the retroactive accumulation of the number of USs that have occurred in the presence of CS₁ is

$$=IF(AND(B6=1,$F7=1),1+R8,R8).$$

For these formulas to work, the bottommost cell in each accumulation column must contain a 0 rather than the accumulation formula contained in all the cells above it.

The formulas in row 5, the next row but one above the uppermost row of accumulation formulas, contains the current total time for each CS, each pairwise combination of CSs, and each N. This total is equal to the topmost accumulation in that column (the accumulation in row 7) plus whatever initial totals have been entered in row 4 of the column. The ability to insert initial totals for the accumulations enables one to bypass protocol composition altogether. One can often compute from the design of the experiment what these accumulations must be after a given number of trials. These accumulations determine the state of conditioning of the animal (its conditional rate estimates and the *p* values attached to them). If one can easily compute these accumulations from the design of the experiment and one does not want to follow the course of conditioning, there is no need to compose the protocol. One simply enters the accumulations in the shaded cells in row 4.

The ability to enter initial values for the accumulations also enables one to pick up, so to speak, in the middle of training. One can follow the course of conditioning after previous training (e.g., after a phase of latent inhibition training, during which a CS is presented repeatedly without any US). The learning-relevant results of the previous training are given by the initial values of the accumulations. The training regime during the period for which initial accumulations are entered must be stationary; that is, the pattern of cooccurrence of CSs and the US must be constant throughout the phase of the experiment from which these initial totals come. Otherwise, the ensuing calculations are not valid.

Subsequent formulas are rendered more comprehensible by naming the totals—TC1, TC2, etc., for the accumulated duration of CS₁, CS₂, etc., and T1.2, T1.3, etc., for the accumulated duration of the pairwise combinations (the intervals where CS₁ and CS₂, CS₁ and CS₃, etc., have been on simultaneously), and NC1, NC2, etc., for the accumulated numbers of USs that have occurred in the presence of CS₁, CS₂, etc.

Section 3: Matrices, Rate Vectors, and Condition Flags

Figure 3 shows Section 3 of the spreadsheet. The coefficients of the unknowns in the system of four simultaneous equations that must be solved to obtain the corrected rate estimates are arrayed in the form of a 4×4 square array (matrix) of cells at the upper left. Since the temporal

Matrix				Corrected Rate Vector	Uncorrected Rate Vector
1	0.167	0	0	RC1= 0	URC1 0.0833
1	1	0	0	RC2= 0.5	URC2 0.5
#NUM!	#NUM!	1	#NUM!	DET12 0.83	URC3 #NUM!
#NUM!	#NUM!	#NUM!	1		URC4 #NUM!

Figure 3. Components of Section 3. The 4x4 matrix at the upper left has 0 or #NUM! coefficients in rows 3 and 4 and columns 3 and 4, because in the simple conditioning protocol used in this example, there is no CS, or CS_i; hence, the values of the corresponding coefficients are either 0 or undefined (#NUM!). The 2x2 submatrix is the one whose inverse is multiplied times the uncorrected rate vector <.08,0.5> to yield the corrected rate vector <0, 0.5>. DET12 is the determinant of the 2x2 matrix.

accumulations from which the coefficients are computed have been named—TC1, TC2, T1.2, and so on—the formulas for the coefficients in this 4x4 array are as follows (on the spreadsheet, each temporal ratio in the matrix below is preceded by the obligatory “=”, which indicates a computational formula in Excel)

	T1.2	T1.3	T1.4
1	T1	T1	T1
T1.2		T2.3	T2.4
T2	1	T2	T2
T1.3	T2.3		T3.4
T3	T3	1	T3
T1.4	T2.4	T3.4	
T4	T4	T4	1

The uncorrected rates, which are the inhomogeneous terms in the system of simultaneous equations, are computed in a column of four cells at the upper right (URC1, URC2, etc., in Figure 3). Again, since the numerical and temporal accumulations required to compute these uncorrected rate estimates have been named, the formulas in these cells directly reflect the algebra. The formula in the first cell is =NC1/TC1, in the second, =NC2/TC2, and so on. This column of four cells is the uncorrected rate vector.

The system of simultaneous equations is solved by inverting the temporal coefficient matrix and multiplying the uncorrected rate vector by the inverted matrix to obtain the corrected rate vector. This is accomplished by the array formula

$$\{=IF(DET4x4 < 0.001, "indeter", MMULT(MINVERSE(H95:K98), S95:S98))\}$$

—which is entered into a column of four cells to the right of the 4x4 array of cells containing the coefficients. The first part of this formula checks that the determinant of the matrix (named DET4x4) is substantially greater than 0. Because of rounding errors, the determinant of a matrix that has no inverse is not necessarily computed as exactly 0, which is why this condition checks for >.001 rather than >0. If the determinant is essentially 0, the values of the corrected rates in these four cells are entered as “indeter.” Otherwise, the formula enters the results of multiplying (MMULT) the inverse of the coefficient matrix [MINVERSE(H95:K98), where H95:K98 designates the 4x4 array of cells containing the coefficients] times the uncorrected rate vector (in cells S95:S98). The formula is enclosed in braces, as is required for array formulas in Excel.

This is the basic computation underlying the estimation of the rates, but there are many conditions under which the full 4x4 matrix will not yield an answer. For example, whenever there are less than four CSs, one or more of the simple temporal totals (for example, TC3 and TC4) will be 0, and the coefficients with these totals in their denominators will be undefined (their cells will display #DIV/0!). Alternatively, the coefficients may all be defined, but the determinant of the 4x4 matrix may be 0, in which case the computation yields no answer. These conditions are detected by statements that raise condition flags and by IF statements embedded in the computational

formulas. Thus, there is a cell that contains the formula for the determinant of the 4×4 matrix: $\{=MDETER(H95:K98)\}$. This value is named DET4×4 for transparency of reference in other formulas that use it.

When the full 4×4 matrix is not the correct matrix to use, the system uses a reduced version of the matrix, one of the 3×3 or 2×2 matrices that may be formed from the 4×4 matrix by omitting one column and one row or two columns and two rows. For example, if there are only three CSs, the correct matrix to use may be the 3×3 matrix formed by omitting the fourth row and fourth column of the original 4×4 matrix. It is not necessary to create a new array of cells for this particular reduced matrix, because the requisite array constitutes a contiguous subsection of the 4×4 array. Thus, the corrected rate vector calculated from this particular reduced matrix is computed by the following formula entered into a column of three cells farther to the right:

$$\{=IF(DET123 < 0.001, "indeter", MMULT(MINVERSE(H95:J97), S95:S97))\}.$$

This formula multiplies the uncorrected rate vector for the first three CSs (in cells S95:S97) by the inverse of the matrix formed by first three rows and first three columns of the original 4×4 matrix. Note that this formula first checks that the determinant of the reduced matrix is effectively greater than 0. This requires that there be a cell with the formula for the determinant of the reduced matrix. The formula in this cell is

$$\{=MDETER(H95:K98)\},$$

and the cell is named DET123 for transparency of reference.

Similar formulas compute the corrected rate vector for the 2×2 matrix formed by omitting the last two rows and last two columns of the 4×4 matrix, and, of course, its determinant (DET12). These formulas are in the third of the three pairs of columns that appear to the right of the 4×4 matrix under the heading Corrected Rate Vectors.

Down six rows and back to the left edge of Section 3, there is a 3×3 matrix, whose columns are directly headed by TC1, TC3, and TC4. This is the reduced matrix obtained by omitting from the 4×4 matrix the second row and second column on the implicit assumption that CS₂ does not have predictive power. The rate predictions based on this assumption (the corrected rate vector) are given to the right of this matrix in rows labeled RC1=, RC3=, and RC4=. The formula for computing these predicted rates is *mutatis mutandis* the same as the formulas in the cases already discussed:

$$\{=IF(DET134 < 0.001, "indeter", MMULT(MINVERSE(H101:J103), S101:S103))\}.$$

In a complete model, there should be a third 3×3 matrix, the 1,2,4 matrix formed by dropping the third column and row on the assumption that CS₃ was irrelevant; but in practice no experiment has required this matrix, so it has not been put on the sheet. To the right of the 1,3,4 corrected rate vector is the computation when only CS₁ and CS₂ are presumed relevant. Finally, at the bottom of Section 3 is the matrix when only CS₁ and CS₄ are presumed to have predictive power, and to its right, the computation of the rates predicted under this assumption.

Flags. The flags are tests for various conditions, the outcome of which determines which matrix computation is used to predict the rates. FLAG4×4 is computed by

$$=IF(AND(NOT(ISERR(DET4×4)), DET4×4 > 0.001), 1, 0).$$

This flag "goes up" (has value = 1) whenever the value of DET4×4 is not an error value and the value of DET4×4 is greater than .001 (that is, not effectively 0). When FLAG4×4 is up, it indicates that the full 4×4 matrix is the correct matrix to use in computing the corrected rates of US occurrence. FLAG3×3 goes up whenever one of the two 3×3 matrices has a nonzero determinant. Its formula is

$$=IF(OR(DET123 > 0.001, DET134 > 0.001), 1, 0).$$

FLAG123 goes up if the solution yielded by the 1,2,3 matrix (the 3×3 matrix that drops the fourth CS) is as good as or better than (has as few nonzero predictors as or fewer nonzero predictors than) the solution yielded by the 1,3,4 matrix (the 3×3 matrix that drops the second CS). To decide whether one set of solutions is as good as or better than the other, it compares the sums of the absolute values of the estimated rates. In solutions that use more predictors (more CSs) than are necessary, the sum of the absolute values of the predicted rates is greater than it is in the solu-

tions that use the minimum number of predictors required to account for the data. The formula for FLAG123 is

$$=IF(AND(Flag4 \times 4 < 1, Flag3 \times 3 > 0, ABS123 \leq ABS134), 1, 0).$$

The formulas for ABS123 and ABS134 are as follows:

$$\{=IF(ISERR(SUM(ABS(O95:O97))), 1000, SUM(ABS(O95:O97)))\}$$

—and similarly for all the other absolute value variables. Note that the sum of the absolute values is set to a very high value (1,000, which is much higher than the rates will ever sum to) if the corrected rate vector whose values are being summed displays error values for its dimension values (which happens when the corresponding matrix has undefined coefficients or a zero determinant).

FLAG134 goes up if the 1,3,4 solution is as good as or better than the 1,2,3 solution. FLAG12 goes up if the 4 \times 4 and both 3 \times 3 matrices have 0 determinants and the 1,2 solution is as good as or better than the 1,3 and 1,4 solutions. Its formula is

$$=IF(AND(Flag4 \times 4 < 1, Flag3 \times 3 < 1, ABS12 \leq ABS13, ABS12 \leq ABS14, ISNUMBER(Q95)), 1, 0).$$

Cell Q95 has the value on dimension 1 of the 1,2 corrected rate vector. Under some now-forgotten condition, the check to verify that this value is a number proved necessary. FLAG13 behaves just as does FLAG12 (and has a similar formula), except that it is for the 1,3 solution, and similarly for FLAG14.

The final set of formulas in Section 3 uses the flags to pick out the best solution. The array formula

$$\{=IF(FLAG4 \times 4 > 0, M95:M98, '')\},$$

entered in a four-cell subcolumn (section of a column), puts the rates calculated by the 4 \times 4 matrix in that subcolumn if FLAG4 \times 4 is up. The formula

$$\{=IF(Flag123 > 0, O95:O97, '')\},$$

entered in the uppermost three cells in the adjacent subcolumn, fills in the rates calculated from the 1,2,3 matrix in those cells when those results are as good as or better than the 1,3,4 results, whereas in the lowermost cell of the subcolumn, the formula

$$=IF(Flag123 > 0, 0, '')$$

fills in a 0 for the effect of CS₄ (the missing CS in the 1,2,3 matrix). The formula

$$=IF(Flag134 > 0, O101, ''),$$

in the uppermost cell, and the array formula

$$\{=IF(Flag134 > 0, O102:O103, '')\},$$

in the lowermost two cells of the next subcolumn, fill in the rates from the 1,3,4 matrix when they are as good as or better than the 1,2,3 results, whereas the formula

$$=IF(Flag134 > 0, 0, '')$$

fills in a 0 for the effect of CS₂. Similar formulas fill in the next subcolumn when the 1,2 solution is as good as or better than the other two-CS solutions, the next subcolumn when the 1,3 solution is . . . , and the rightmost subcolumn when the 1,4 solution is . . .

These conditionals establish the rates ascribed to each CS—except for the rate ascribed to the background (CS₁), which is adjusted by the following formula:

$$=IF(TC1=0, 0, IF(AND(TC2=0, TC3=0, TC4=0), NC1/TC1, IF(MAX(G112:L112) < 0.000001, 1/TC1, MAX(G112:L112))))).$$

The first conditional, IF (TC1=0), sets the background rate to 0 when the experiment has not yet begun—that is, when the animal has not yet been placed in the experimental environment for the first time. (This avoids unwanted error messages elsewhere.) The second conditional, IF (TC2=0 AND TC3=0 AND TC4=0), sets the background rate to the number of USs that have occurred, divided by the time during which the animal has been in the apparatus in cases in which there has been no CS other than the apparatus itself (pure background conditioning), because this condi-

tion yields no matrix with a nonzero determinant. The third conditional, IF MAX(rate for CS, computed by any matrix) IS EFFECTIVELY 0, assigns one occurrence of the US to the background in those cases in which the matrix computations impute no USs to the background. This is done to create an upper limit on the background rate, given the training so far. It is equivalent to assuming that a US imputable to the background will occur in the next instant. It is necessary to have an upper limit on the estimated background rate in order to test whether the rates observed in the presence of the other CSs are higher than would be expected, given the (unknown) background rate.

Finally, if all the preceding conditionals have failed, the rate assigned to the background is the rate computed by the matrix computations. This rate will always be the only nonzero entry in the designated subrow G112:L112, which is where the rate-selecting formulas specify the values of the rate imputed to the background. The same principle—that the rate estimate is the only nonzero estimate that the selection formulas will have placed in the subrow for a given CS—underlies the formulas that place whatever answer has emerged from the selection process in the “final answer” subcolumn. These formulas all have the following form:

$$=IF(MIN(G113:L113) < 0, MIN(G113:L113), MAX(G113:L113))$$

—which examines the subrow (G113:L113) for both values less than 0 and values greater than 0.

It is possible for a rate estimate with a negative sign to emerge from these calculations, and in fact, rate estimates with negative signs do emerge in the conditions that produce inhibitory conditioning. What a negative rate estimate means is that the CS in question is imputed with the power to reduce the rate of US occurrence. Obviously, when that rate would otherwise be zero, the effect of such a CS cannot be observed. It is perfectly possible to observe a negative (reducing) effect on a positive rate, but it is not possible to observe a negative rate.

Section 4: Probabilities and Confidence

The model assumes that in processing the training experience, the nervous system computes a value that plays the same role in the decision process leading to a conditioned response that the computed p value plays in the process of statistical hypothesis testing. From a statistical decision point of view, the problem confronting the animal is to decide whether the rate of US occurrence imputed to the joint effects of the background (CS_B) and another CS (CS_A) is improbably greater than or less than the rate imputed to the background alone. The probability, p , that the difference between two independent estimates of rate, λ_a and λ_b , based on observations of n_a and n_b occurrences during nonoverlapping intervals T_a and T_b , indicates a difference in the values of the underlying rate parameters is

$$p \left\{ F(2n_a, 2n_b) < \frac{n_a T_b}{n_b T_a} \right\}.$$

When applied to the present situation, this general formula becomes

$$p \left\{ F(2n_{CS}, 2n_B) < \frac{\lambda_{CS} + \lambda_B}{\lambda_B} \right\},$$

where n_{CS} is the number of USs that occurred while the CS_n and the background were both present, n_B is the number of USs that occurred when the background alone was present, λ_{CS} is the estimated rate of US occurrence due to the influence of CS_n, and λ_B is the estimated rate of US occurrence due to the background.

In implementing this calculation in the present context, one has to take into account that λ_{CS} may be negative and that it may be much greater than the effect of the background; hence, the formula above may yield an F that is negative or that is positive but significantly less than 1. Such eventualities cannot arise under more conventional circumstances where negative rates cannot exist. Thus, in applying this procedure one must use the absolute value of λ_{CS} . The test is whether the estimated rate ascribed to the background (which is always positive) plus the absolute value of the rate ascribed to the CS is greater than the background rate alone—given the number of independent observations on which the two rate estimates are assumed to be based. The computation of the F value in this test is implemented by

$$=IF(ABS(Rate2) < 0.001, 0.1, (ABS(Rate2) + Rate1) / Rate1),$$

and similarly for Rate3 and Rate4. Note that when the rate estimate is effectively 0 (absolute

value $< .001$), the F is arbitrarily given the value 0.1. This, together with the assignment of degrees of freedom, fixes the p values for the effects of CSs to which no effect on the rate of US occurrence is imputed.

The degrees of freedom for the denominator are calculated by

$$=IF(\text{Rate1}*(\text{TC1}-\text{TC2}) < 2, 2, 2*\text{Rate1}*(\text{TC1}-\text{TC2})),$$

and for the numerator by

$$=IF(\text{ABS}(\text{Rate2}) < 0.001, 2, 2*(\text{ABS}(\text{Rate2})+\text{Rate1})*\text{TC2}).$$

The formula above is for CS₂; the rate and time variables change appropriately in the CS₃ and CS₄ formulas. Notice that the degrees of freedom associated with each estimate are computed by multiplying the rate estimate times the interval over which that estimate is assumed to have been made. The actual number of USs observed in the assumed interval will often be different, because other CSs will also have been present. The system must estimate how many of the USs that occurred were due to the (CS_n+CS_B) combination over the interval when it was present (TC_n)—during some of which other CSs may also have acted—and it must estimate how many were due to the background alone over the interval when CS_n was not present (TC1-TC_n)—during some of which, again, other CSs may have acted.

To fix a p value for the effects of CSs to which no effects have been imputed, the degrees of freedom for these CSs are set to 2 by the conditionals in the formulas above. This reflects the implicit assumption that although no US has so far been observed that could be imputed to that CS, one such will be observed in the next instant. Without such an assumption, the rate estimate for such a CS is undefined; hence, one cannot give a likelihood for the hypothesis that the CS affects the rate of US occurrence. It is essential to assign some likelihood, because the model assumes that the animal's response or lack of response to the CS is based on this likelihood. As already noted, this assumption is also essential in the computation of a likelihood for any CS when no effect can be imputed to the background. In such cases, the background rate, which establishes what is to be expected on the basis of the null hypothesis, is undefined. The assignment of one occurrence to the background fixes the rate for the background at the highest value consistent with the observations made thus far (namely, that an interval of length TC1 has elapsed without the occurrence of a US imputable to CS₁).

The p values are computed through reference to an external, commercially available (Heizer Software) worksheet called *Statfuns*, containing a function called *Fapprox*, which returns the area under the F distribution for a given F value and degrees of freedom. The formula for the p value for the effect of CS₂ (pval2) is

$$=\text{Statfuns!Fapprox}(\text{Fval2}, \text{DeFr2}, \text{DeFr1}).$$

The Fval and the DeFr variables change in the formulas for the CS₃ and CS₄. The complement of a p value (pcompl) is $1-p$. The likelihood ratio for the hypothesis that a CS_n has an effect is $p/(1-p)$ or pval/pcompl.

Confidence. The decision process that translates the likelihood ratio into a conditioned response to the CS (the confidence function) is assumed to have the characteristics of the conventional decision process in hypothesis testing statistics. It is assumed to be a sigmoidal function of the likelihood ratio, whose value usually does not rise appreciably above zero until this ratio is fairly high (e.g., 10:1 to 100:1). Just how high depends on the perceived payoff matrix—that is, on the perceived costs and benefits of true positives (responding when the CS really does have an effect) versus false positives (responding when it really does not), true negatives (not responding when it really does not), and false negatives (not responding when it really does). The confidence function used in the model is a two-parameter Weibull function, implemented by

$$=1-2^{-(\text{alpha}*pval2/\text{pcompl2})^s}.$$

The α parameter plays the role of the critical value in a statistical hypothesis testing procedure. The smaller the value of α , the more conservative the animal is—that is, the higher the likelihood ratio must be before the animal will show any appreciable tendency to respond to the CS. The s parameter determines what might be called the decisiveness parameter. The higher the value of s , the steeper the transition from no responding to responding as the p value approaches the value of α .

Event	Totals:	CS2		D/n	conf. λ
		obs	exp		
		3.84	4		
Tim	C1 C2 US				changed
	1	3.84	4	0.04	0.00
10	1 1	3.84	4	0.04	0.00
12	1 1	2.88	4	0.16	0.00
22	1 1	2.88	4	0.16	0.00
24	1 1	1.92	3	0.27	0.00
34	1 1	1.92	3	0.27	0.00
36	1 1	0.96	3	0.39	0.01
46	1 1	0.96	3	0.39	0.01
48	1 1	0.00	2	0.50	0.13
58	1 1	0.00	2	0.50	0.13
60	1 1	0.00	2	0.38	0.01
70	1 1	0.00	2	0.38	0.01
72	1 1	0.00	1	0.25	0.00
82	1 1	0.00	1	0.25	0.00
84	1 1	0.00	1	0.13	0.00
94	1 1	0.00	1	0.13	0.00
96	1 1				

Figure 4. Section 5 of the spreadsheet: the expected and observed cumulative occurrences ascribed to each CS, the normalized discrepancy, and the confidence that the normalized discrepancy indicates nonstationarity. In this example, four conditioning trials have been followed by four extinction trials.

Section 5: The Test for Stationarity

Figure 4 shows the fifth section, which tests whether the effects of the CSs on the rate of US occurrence have been stationary—that is, whether they have been constant throughout the training. To test for stationarity, the model computes the retrospectively accumulated expected number of US occurrences ascribable to a given CS and compares this to the retrospectively accumulated numbers of USs whose occurrence has in fact been ascribed to the CS. By retrospectively accumulated number or intervals, I mean the number or temporal interval accumulated starting from the latest event and looking backward in time over each earlier event, so that the accumulation grows as one looks farther back in time.

The retrospectively accumulated number of expected US occurrences for a given CS is the retrospectively accumulated time during which the CS has been present, multiplied by the rate ascribed to the CS. The retrospectively accumulated time that CS₁ has been present is in column G. Therefore, in the top row (row 7), where the greatest accumulations are, the retrospectively accumulated number of USs expected from the influence of CS₁ is given by

$$=Rate1*G7.$$

The formula in each row below this has the row number incremented by 1. The formula for the other CSs is the same, except for the appropriate change in the rate variable (Rate2, etc.).

The computation of the retrospectively accumulated total for the number of USs whose occurrence was ascribable to a given CS (the “observed” number of USs for that CS) is more complex. The principle is that when a US occurs, credit for its occurrence must be prorated among the active, positive CSs in proportion to the rate estimates associated with them. An *active* CS is a CS that is present when the US occurs. A *positive* CS is one that has a positive effect on the rate of US occurrence. Similarly, when a US fails to occur because of the influence of one or more active, negative CSs, the failures to occur (the negative occurrences) must be prorated among the active, negative CSs. A *negative* CS is one that has a negative effect on the rate of US occurrence (reduces the rate).

Three intermediate quantities are used in this bookkeeping exercise. The first consists of the *accountably missing* USs for a given interevent interval. The accountably missing USs are those whose absence can be explained by the presence of one or more CSs to which a negative effect on rate has been imputed. The top-row (row 7) formula for computing the missing USs is

$$=(Q8-Q9-V7+V8-(X7-X8 > 0)*(X7-X8)-(Z7-Z8 > 0)* (Z7-Z8)-(AB7-AB8 > 0)*(AB7-AB8)).$$

The quantity $Q8 - Q9 - V7 + V8$ is the increment in the number of USs ($Q8 - Q9$) minus the increment in the number expected because of CS_1 ($V7 - V8$). From this, one must further subtract the increment in the number expected because of CS_2 , if CS_2 was present over the interevent interval—that is, if $(X7 - X8 > 0)$. (When a conditional is true in Excel, it has numerical value 1; when false, it has numerical value 0.) Similarly, one must subtract the increments expected from CS_3 and CS_4 if they were present. The negative (or zero) quantity thus arrived at represents the number of expected USs that are missing from the interevent interval. This quantity (In Cell AL7) then figures in the conditional formula

$$=IF(AND(AL7 < -0.1, AN7 < -0.1), AL7, 0),$$

which determines whether the missing USs are accountably missing—that is, missing because of the influence of a negative CS. If there really is no missing US (in which case, $AL7 > -0.1$), or if there were no negative CSs present in that interval (in which case $AN7 > -0.1$), then this formula has the value 0.

The other two intermediate quantities are the denominators for prorating the positive and negative US occurrences among the active positive and active negative CSs, respectively. For the positive CSs, the denominator formula is

$$=Rate1 * \$B7 + (Rate2 > 0) * Rate2 * \$C7 + (Rate3 > 0) * Rate3 * \$D7 + (Rate4 > 0) * Rate4 * \$E7.$$

Rate 1 is always positive (because the model assumes one US occurrence ascribable to it, when none have in fact been observed); hence, it requires no test for positivity. Each of the other rates must be tested for positivity (by the expressions in parentheses). A rate contributes to the denominator only if the corresponding CS is present, because $\$B7$, $\$C7$, $\$D7$, and $\$E7$ are positive only if CS_1 , CS_2 , CS_3 , and CS_4 are present, respectively. Since Rate1 is always at least slightly positive, and since CS_1 is always present, this denominator is never 0. The formula for this denominator is in the AM column.

The denominator for prorating among the negative CSs is computed by

$$=(Rate2 < 0) * Rate2 * \$C7 + (Rate3 < 0) * Rate3 * \$D7 + (Rate4 < 0) * Rate4 * \$E7.$$

This will be zero when no negative CSs are present, which causes the values of prorating formulas to be undefined, so this denominator, which is in the AO column, is examined by the following conditional

$$=IF(AO7 < 0, AO7, 0.01).$$

This IF statement leaves the denominator as is when it is less than 0, but sets its value to 0.01 when it is 0, to avoid division by 0 in the prorating formulas. This conditional is in the adjacent AI column, and the value in this column is actually used in the prorating formulas.

The prorating formula for the USs to be credited to CS_1 is

$$=IF(\$B7=0, 0, Rate1 * (\$F8 - AK7) / AM7 + U8).$$

The cell $\$F8$ contains the increment in the number of USs over the interevent interval. The column AK7 contains the (negatively signed) number of accountably missing USs; hence $(\$F8 - AK7)$ is the number of US occurrences to be prorated. The proportion of this number that is credited to CS_1 is determined by the ratio of Rate1 to the denominator (in column AM7), which is the sum of the rates of the active, positive CSs. The prorating formula for CS_1 is the simplest because CS_1 always has a positive effect on rate. The prorating formula for CS_2 , which like all the higher CSs may have either a positive or a negative effect, must provide for both possibilities. It is

$$=IF(\$B7=0, 0, (Rate2 > 0) * Rate2 * \$C7 * (\$F8 - AK7) / AM7 \\ + (Rate2 < 0) * Rate2 * \$C7 * AK7 / AN7 + W8).$$

The first half is just as for CS_1 , except for $\$C7$, which has the value 1 if CS_2 was present during the interevent interval and 0 otherwise. (Thus, CS_2 gets a share only if it has been present.) The second half of the formula, beginning with the test $(Rate2 < 0)$, prorates the accountably missing USs among the active negative CSs.

The statistic that tests for stationarity is the Kolmogorov-Smirnov statistic, which is the difference between the cumulative observed and cumulative expected number of occurrences, divided by the total number of expected occurrences. The statistic is computed by

$$=IF(ABS(V\$5) > 0.01, ABS(U7-V7)/ABS(V\$5), 0).$$

The quantity in V\$5 is the number of USs expected from the influence of that CS over the entire training period (including the initial or unprotocolled phase). The quantity (U7-V7) is the difference between the cumulative observed and cumulative expected USs. The IF conditional ensures that the Kolmogorov-Smirnov statistic is 0 when there is less than 1 US expected (from that CS) over the entire training period. The critical value of the Kolmogorov-Smirnov statistic for the $p < .02$ level is given by the formula

$$=IF(V5 < 2, 0.99, IF(V5 < 3, 0.9, IF(V5 < 4, 0.78, 1.35 * V5 ^ -0.474398))).$$

This is a purely empirical formula that yields critical values approximating those given in the standard tables of critical values for the Kolmogorov-Smirnov statistic. By comparing the normalized discrepancy (D/n) with this value, one can judge whether the discrepancy approaches statistical significance. This comparison is done crudely by the formula in the adjacent column headed conf lambda changed:

$$=1-50^-((AC7/\$AD\$3)^10).$$

When the value of this expression is very close to or greater than 1, it indicates that the observed discrepancy is highly unlikely; that is, one can conclude for nonstationarity with some confidence. The event time at which this value peaks indicates the time at which the rate of US occurrence for that CS changed. A formula giving the p value as a function of the normalized discrepancy and the n would be highly desirable in place of this formula. I would appreciate hearing from anyone who knows a nonrecursive formula for calculating these p values.

DISCUSSION

A striking feature of this model is the absence of parameters, free or otherwise, in the portions of the model that compute what the animal learns. There are no parameters, because the model assumes that the animal learns the objective statistical facts of the training situations to which it is exposed. The only parameters in the model occur in the confidence function, which translates what has been learned (the rate of US occurrence predicted by a CS and the likelihood that the CS really has an effect on US rate) into an observable conditioned response. The explicit specification of the performance function, which translates what has been learned into a conditioned response, is another noteworthy feature of the model.

Among the noteworthy predictions of the model is that, like the model of Gibbon and Balsam (1981), it predicts the profound effect of the duty cycle on the rate of conditioning. More particularly, it predicts that the rate of conditioning—that is, the number of reinforced trials to reach a conditioning criterion—will be inversely proportionate to the duty cycle, which it is (Gallistel, in press; Gibbon et al., 1977). It also predicts that the rate of acquisition under partial reinforcement will be constant if the normalized strength of responding is plotted as a function of the number of rewards received (rather than as a function of the number of trials) and that the rate of extinction will be constant when the normalized strength of responding is plotted as a function of the number of expected rewards that have been omitted rather than as a function of the number of extinction trials (Gibbon et al., 1980). The model predicts quantitatively the number of trials to extinction after both continuous and partial reinforcement (Gallistel, in press).

The model predicts “latent inhibition”; that is, it predicts that repeated occurrences of the CS alone prior to the conditioning phase in which the US is introduced will retard conditioning. In the model, this retarding effect is not “inhibitory” in nature; it does not involve the animal’s learning that the CS predicts a reduction in the rate of US occurrence. Thus, the model predicts that the retardation of conditioning by prior CS exposure will be observed whether the subsequent conditioning phase involves excitatory or inhibitory conditioning, which is what has been found (Reiss & Wagner, 1972).

The model also predicts blocking and overshadowing (Kamin, 1967, 1969). It predicts that raising the rate at which USs occur when only the background is present to match the rate of occurrence when an intermittent CS (e.g., a tone = CS₂) is present will block conditioning to the intermittent CS (Rescorla, 1968). It also predicts that signaling the “background” USs with another CS will eliminate the blocking effect of the “background” USs (Goddard & Jenkins, 1987; Rescorla, 1984).

None of the predictions above depends on parametric assumptions. Thus, the model accounts for more experimental results than do other formalized models of the classical conditioning process, while making many fewer ad hoc parametric assumptions.

REFERENCES

- GALLISTEL, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.
- GALLISTEL, C. R. (in press). Classical conditioning as an adaptive specialization: A computational model. In D. L. Medin (Ed.), *The psychology of learning and motivation*. New York: Academic Press.
- GIBBON, J., BALDOCK, M. D., LOCURTO, C. M., GOLD, L., & TERRACE, H. S. (1977). Trial and intertrial durations in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*, 264-284.
- GIBBON, J., & BALSAM, P. (1981). Spreading associations in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Ed.), *Autoshaping and conditioning theory* (pp. 219-253). New York: Academic Press.
- GIBBON, J., FARRELL, L., LOCURTO, C. M., DUNCAN, H. J., & TERRACE, H. S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning & Behavior*, *8*, 45-59.
- GODDARD, M. J., & JENKINS, H. M. (1987). Effect of signaling extra unconditioned stimuli on autoshaping. *Animal Learning & Behavior*, *15*, 40-46.
- KAMIN, L. J. (1967). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the Prediction of Behavior: Aversive stimulation* (pp. 9-33). Miami, FL: University of Miami Press.
- KAMIN, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Ed.), *Punishment and aversive behavior* (pp. 276-296). New York: Appleton-Century-Crofts.
- KREMER, E. F. (1978). The Rescorla-Wagner model: Losses in associative strength in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 22-36.
- REISS, S., & WAGNER, A. R. (1972). CS habituation produces a "latent inhibition effect" but no active "conditioned inhibition." *Learning & Motivation*, *3*, 237-245.
- RESCORLA, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative & Physiological Psychology*, *66*, 1-5.
- RESCORLA, R. A. (1984). Signaling intertrial shocks attenuates their negative effect on conditioned suppression. *Bulletin of the Psychonomic Society*, *22*, 225-228.
- ROBBINS, S. J., & RESCORLA, R. A. (1989). The role of signaled periods of nonreinforcement in responding on a random schedule in autoshaping. *Animal Learning & Behavior*, *17*, 304-310.
- WAGNER, A. R., LOGAN, F. A., HABERLANDT, K., & PRICE, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, *76*, 171-180.

NOTES

1. Caution is urged with respect to implementing the model in other spreadsheets. Some distinguish between empty cells and cells containing 0s.—*Editor*
2. Some spreadsheets will interpret expressions such as =N/A() and =A1+A2 as text rather than formulas. Usually this problem may be overcome with expressions of the form +N/A() and +A1+A2 (which will work also in Excel).—*Editor*