R.C. Atkinson (Ed.), Studies in Mathematical Psychology.
Stanford: Stanford University Press, 1964, pp. 358-379.

The Stanford Institute for Mathematical Studies in the Social Sciences

VENTURA HALL
STANFORD UNIVERSITY

Reprint No. 76

Empirical Comparison of Models for a Continuum of Responses with Noncontingent Bimodal Reinforcement

PATRICK SUPPES AND HENRY ROUANET
MICHAEL LEVINE, AND RAYMOND W. FRANKMANN

Reprinted from
Studies in Mathematical Psychology
1964

12

Empirical comparison of models for a continuum of responses with noncontingent bimodal reinforcement

PATRICK SUPPES, HENRY ROUANET, MICHAEL LEVINE, and RAYMOND W. FRANKMANN, *Institute for Mathematical Studies in the Social Sciences, Stanford University*

This paper is a direct extension of the earlier work reported by Suppes and Frankmann (1961). A general introduction to the problems of stochastic learning theory applied to a continuum of responses is given in the preceding paper in this volume by Suppes and Rouanet. The present experiment used the same circular apparatus as did Suppes and Frankmann and Suppes and Rouanet. We abstract briefly from the fuller description given by Suppes and Rouanet. The response x and the reinforcement y vary continuously along the circumference of the circle from 0 to 2π. In the Suppes and Frankmann study and also in the present study, the reinforcement is noncontingent, i.e., the probability distribution of reinforcement is independent of the subjects' responses and the same distribution is used on every trial. The particular distribution used by Suppes and Frankmann was a triangular distribution on the interval 0 to 2π. The present study has used the bimodal distribution constructed from the two equal triangular distributions on the interval, 0 to π, and π to 2π (see Fig. 2).

The aims of the present experiment have been twofold. The first has been to investigate the extent to which the common predictions of the linear and stimulus-sampling models for a continuum of responses will hold when the reinforcement distribution is no longer unimodal as it was in the case of the

This research has been supported by the United States Air Force Office of Scientific Research Directorate of Information Services, under Contract AF 49(638)–1037. Part of the computing related to the analysis of data has been supported by grant USPHS M-6154 from the National Institute of Mental Health. We are indebted to Jean Donio and to Sylvia Garfinkel for assistance in the analysis of data, and to Eleanor Williamsen for assistance in running subjects.

(Henry Rouanet is now at the University of Paris, and R. W. Frankmann is at the University of Illinois.)

Suppes and Frankmann study. Secondly, a larger number of trials was run in this study than in the Suppes and Frankmann study in order to provide an adequate amount of data at asymptote to test the differential sequential predictions of the linear and stimulus-sampling models. In the Suppes and Frankmann and the Suppes and Rouanet studies, no predictions are reported that differentiate the two kinds of models.

1. Summary of theoretical results

Using the notation of Suppes and Rouanet in the preceding paper, we may briefly summarize the theoretical results utilized in the analysis of data. These naturally fall into three categories.

First, there is the asymptotic response distribution whose density $r(x)$ is given by the equation

$$r(x) = \int_0^{2\pi} k(x; y) f(y)\, dy, \tag{1}$$

where $k(x; y)$ is the smearing density and $f(y)$ the noncontingent reinforcement density.

Second, there are the reinforcement-dependent statistics $P(X_n \mid Y_{n-1})$ and $P(X_n \mid Y_{n-1}, Y_{n-2})$. These statistics are derived in the preceding article. Their equations at asymptote are as follows:

$$P(X_n \mid Y_{n-1}) = (1-\theta) R(X_n) + \theta \frac{H(X_n, Y_{n-1})}{F(Y_{n-1})}, \tag{2}$$

$$P(X_n \mid Y_{n-1}, Y_{n-2}) = (1-\theta)^2 R(X_n) + \theta \frac{H(X_n, Y_{n-1})}{F(Y_{n-1})} + (1-\theta)\theta \frac{H(X_n, Y_{n-2})}{F(Y_{n-2})}, \tag{3}$$

where

$$R(X) = \int_X r(x)\, dx,$$

$$F(Y) = \int_Y f(y)\, dy,$$

$$H(X, Y) = \int_X \int_Y k(x; y) f(y)\, dx dy.$$

The important thing to note about these reinforcement-dependent statistics is that like the asymptotic response distribution, they are the same in the linear and stimulus-sampling models.

For the present experiment we also include the related conditional density

$$j(x_n \mid Y_{n-1}) = \int_{Y_{n-1}} j(x_n, y_{n-1})\, dy_{n-1} / F(Y_{n-1}),$$

which is given at asymptote by the equation

$$j(x_n \mid Y_{n-1}) = (1-\theta) r(x_n) + \theta \frac{H(x_n, Y_{n-1})}{F(Y_{n-1})}, \tag{4}$$

where

$$H(x, Y) = \int_Y k(x; y) f(y)\, dy .$$

Third, there are the asymptotic sequential predictions $P(X_n \mid Y_{n-1}, X_{n-1})$. Predictions of these sequential probabilities differentiate the linear and stimulus-sampling models. In the case of the linear model the result is the following:

$$(5) \qquad P(X_n \mid Y_{n-1}, X_{n-1}) = (1 - \theta) \frac{W(X_n, X_{n-1})}{R(X_{n-1})} + \theta \frac{H(X_n, Y_{n-1})}{F(Y_{n-1})},$$

where

$$W(X, X') = \frac{1}{2 - \theta}[2(1 - \theta)R(X)R(X') + \theta R(X, X')]$$

and

$$R(X, X') = \int_X \int_{X'} \int_0^{2\pi} k(x; y) k(x'; y) f(y)\, dx dx' dy .$$

For the N-element stimulus-sampling model that assumes that exactly one stimulus is sampled on every trial (i.e., the Estes pattern model), the expression is the following:

$$(6) \qquad P(X_n \mid Y_{n-1}, X_{n-1}) = \left(\frac{1}{N} - \frac{c}{N}\right) \frac{R(X_n, X_{n-1})}{R(X_{n-1})} + \frac{c}{N} \frac{H(X_n, Y_{n-1})}{F(Y_{n-1})} + \left(1 - \frac{1}{N}\right) R(X_n),$$

where c is the conditioning parameter. [In the case of statistics (2), (3), and (4), in terms of the pattern model $c/N = \theta$, i.e., the same expression holds in the pattern model if θ is replaced by c/N.] An important observation that can be made about Eq. (6) for the N-element model is that it may be written in the following form:

(6′)

$$P(X_n \mid Y_{n-1}, X_{n-1}) = \frac{1}{N}[\text{Pred. of one-element model}] + \left(1 - \frac{1}{N}\right) R(X_n).$$

In other words, the sequential prediction at asymptote is a linear combination of the predictions of the one-element model and the asymptotic response distribution. The rationale of this result is obvious. When the same element is sampled on trials n and $n - 1$ then effectively the one-element model may be used to make predictions. The probability of such an event, that is, of sampling the same element on both trials, is $1/N$. On the other hand, if the element sampled on trial n is not the same as the element sampled on trial $n - 1$, then the reinforcement on trial $n - 1$ as well as the actual response made has no direct effect on the response distribution on trial n, for that will be determined by the element sampled on that trial. In order that this intuitive interpretation not be taken too literally, however, it should be pointed out that the precise result given here does not hold except at asymptote. The reasons for this may be seen by examining the detailed derivation given in the Appendix.

2. Experimental method

Subjects. The subjects were 4 male and 26 female Stanford undergraduates. Each subject was paid $2.50 for the two-hour experimental session.

Apparatus. The general apparatus is the one described in Suppes and Frankmann (1961), but of the two circles, only the larger (5 feet in diameter) was used in the present study.

Procedure. The instructions to subjects, describing the experiment as a target prediction problem, were identical to those used in the Suppes and Frankmann study, with one exception. In the earlier study when the bar of red light was moved around the circumference of the circle at the beginning of the instructions, it was stopped at the top of the circle. In the present study it was stopped at the randomly selected physical position for the scale zero of a given subject.

As soon as questions had been answered by paraphrasing the instructions, 600 trials were run, with one interruption of about 3 minutes after the 300th trial. The average rate was 6 trials per minute.

Design. All subjects were run under the same experimental conditions. Thirty 600-trial reinforcement sequences were computed using the bimodal density shown in Fig. 2. By random choice, 30 equally spaced divisions of the circle, starting from an arbitrary physical zero, were assigned without repetition as scale zero points for the separate reinforcement sequences.

3. Results and discussion

The presentation of results has been organized into the following categories: conditional variance learning curve; estimation of parameter of smearing distribution; asymptotic response distribution; estimation of learning parameter; goodness of fit of reinforcement-dependent statistics; and model-differentiating sequential statistics.

Conditional variance learning curve. In previous studies we have found that the variance learning curve provides a good index of the rate of learning. In the case of a symmetric bimodal reinforcement distribution, such as was used in the present study, the variance is not a sensitive measure of learning. There is, however, a natural substitute, namely, the conditional variance, i.e., the variance of the conditional distribution restricted to half of the circle. Figure 1 shows the observed variance of the empirical conditional distribution on the interval 0 to π combined with that on π to 2π in blocks of fifty trials. It is evident that the variance is very close to an asymptote at the end of the first 100 trials and is essentially constant from trials 200 to 600. The asymptotic results that follow are based on the last 400 trials of the experiment.

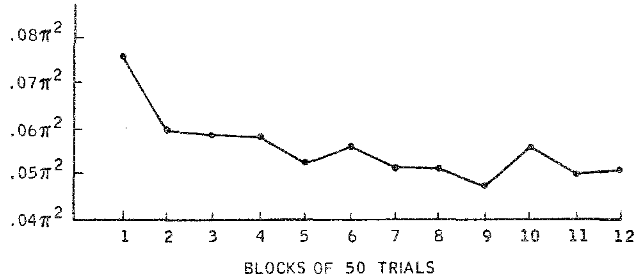

FIG. 1. Observed conditional variances of responses in blocks of 50 trials.

Estimation of smearing distribution. As in previous studies, we estimate the smearing distribution by using Eq. (1) for the asymptotic response distribution. However, the estimation is actually dependent upon using the method of moments, that is, upon estimating the single parameter of the smearing distribution $k(x;y)$ from the variance of the asymptotic response histogram. For the reasons just indicated, in the case of the symmetric bimodal distribution it is preferable to use the conditional variance. One result of the Suppes and Frankmann study was that the particular choice of the analytical form of the smearing distribution does not appreciably affect the shape of the asymptotic response distribution. The article by Suppes and Rouanet indicates that for a uniform reinforcement distribution restricted to a small arc of the circle, it is possible to discriminate between uniform and beta smearing distributions, but for the bimodal distribution of the present study it is undoubtedly the case that the results of the Suppes and Frankmann study are approximately true. Because of the considerably greater ease of making computations with a uniform smearing distribution than with a beta smearing distribution, the former has been chosen for use in the present study. With this choice of the smearing distribution the only parameter to be estimated is the half range a. To indicate the form of the equations for $r(x)$ as a function of a for the bimodal reinforcement and uniform smearing distributions, it is sufficient to give $r(x)$ for $0 \leq x \leq \pi/2$, because $r(x) + r(x + \pi/2)$ is a constant, or, intuitively speaking, because the other three parts of the curve have the same shape with a mirror reflection for two of them (see Fig. 2).

$$(7) \qquad r(x) = \begin{cases} \dfrac{1}{a\pi^2}(x^2 + a^2) & \text{for } 0 \leq x \leq a, \\ \dfrac{2x}{\pi^2} & \text{for } a \leq x \leq \dfrac{\pi}{2} - a, \\ \dfrac{1}{\pi} - \dfrac{1}{a\pi^2}\left[\left(\dfrac{\pi}{2} - x\right)^2 + a^2\right] & \text{for } \dfrac{\pi}{2} - a \leq x \leq \dfrac{\pi}{2}. \end{cases}$$

This expression and the others given below hold on the assumption that

$a < \pi/4$, which turns out not to be a real restriction, for the estimate of a falls definitely below this upper bound. It may be noted from Eqs. (7) that $r(x)$ coincides with the reinforcement density $f(y)$ for x between a and $\pi/2 - a$.

From (7) it is straightforward to derive that the asymptotic conditional variance (C.V.) on the interval 0 to π is given by the following expression as a function of a.

$$\text{C.V.} = \frac{\pi^2}{24} + \frac{a^2}{3} - \frac{a^3}{3\pi}. \tag{8}$$

From the natural symmetry of the circle, it is reasonable to combine the conditional variance computations for the two intervals $(0, \pi)$ and $(\pi, 2\pi)$ in order to use all the data. More precisely, for each response in the last 400 trials we computed its squared deviation about the mean of that one of the two intervals in which it fell, and then divided the sum of these deviations by 12,000, the number of responses, to obtain the empirical estimate of the conditional variance. As a partial behavioral check on this rather obvious assumption of symmetry, we tabulated, as shown in Table 1, the number of responses falling in each interval for each subject. The χ^2 on the null hypothesis that there are exactly 200 responses in each interval is also shown for each subject. The responses of subject 21 show a highly significant deviation from the null hypothesis of symmetry. When his $\chi^2 = 116.64$ is subtracted from the total, the resulting χ^2 is 37.44, which with 29 degrees of freedom is

TABLE 1

NUMBER OF RESPONSES IN THE LAST 400 TRIALS OF EACH SUBJECT IN THE INTERVALS $(0,\pi)$ AND $(\pi,2\pi)$

Subj.	$(0,\pi)$	$(\pi,2\pi)$	χ^2	Subj.	$(0,\pi)$	$(\pi,2\pi)$	χ^2
1	212	188	1.44	16	182	218	3.24
2	180	220	4.00	17	206	194	.36
3	206	194	.36	18	201	199	.01
4	195	205	.25	19	200	200	.00
5	194	206	.36	20	206	194	.36
6	212	188	1.44	21	92	308	116.64
7	220	180	4.00	22	187	213	1.69
8	198	202	.04	23	207	193	.49
9	189	211	1.21	24	223	177	5.29
10	208	192	.64	25	205	195	.25
11	189	211	1.21	26	192	208	.64
12	215	185	2.25	27	196	204	.16
13	200	200	.00	28	195	205	.25
14	186	214	1.96	29	219	181	3.61
15	193	207	.49	30	212	188	1.44

not significant at the .10 level. Data from subject 21 are retained in all the empirical computations given in the remainder of the paper, although some improvement of fit would have resulted from omitting his protocol.

Using the computation described, the empirical conditional variance is $.051684\pi^2$. On the basis of Eq. (8), this leads to the estimate $a^* = .1930\pi$, and this value is used in the remaining theoretical predictions reported.

Goodness of fit of asymptotic response distribution. Figure 2 presents the response histogram for the group of subjects in class intervals of $.05\pi$ for the last 400 trials. The predicted asymptotic response density $r(x)$, as well as the triangular reinforcement density $f(y)$, are also shown in the figure. It is evident from the most casual inspection that the response density $r(x)$ fits the empirical histogram of responses extremely well. In order to perform a statistical test of the apparent goodness of fit of the response density to the response histogram, it is necessary for the validity of the test to select observations that are statistically independent. It is well known that for stationary stochastic processes, observations sufficiently well spaced tend to become statistically independent. Because at asymptote the linear and stimulus-sampling models are stationary processes, it seems reasonable to select a figure for trial spacings and to test for independence. Following the procedure used in the Suppes and Frankmann study, we have selected every fifth observation. On a test of zero-order vs. first-order dependence on these observations with the interval 0 to 2π divided into the four subintervals of length $\pi/2$, the test of zero-order vs. first-order does not yield a significant χ^2 ($\chi^2 = 11.24$ with 9 df). Given this positive evidence of independence of every fifth response, we proceeded to use these 2400 responses to test goodness of fit of response density. As would be expected from Fig. 2, the χ^2 is not significant at the .05 level ($\chi^2 = 52.44$ with 38 df).

The results of this statistical test are similar to those obtained in the Suppes and Frankmann study for two groups that had a unimodal reinforcement distribution. That study and the present one, together with the discrimination study reported in the preceding paper, indicate that the linear

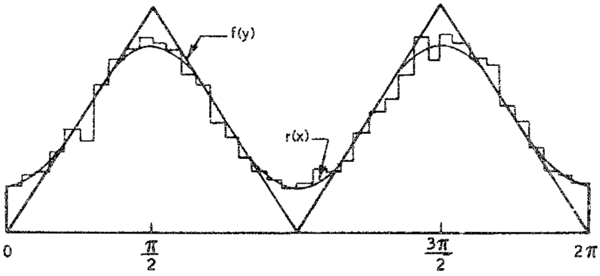

FIG. 2. Asymptotic response histogram with predicted asymptotic response and reinforcement densities.

and stimulus-sampling models for a continuum of responses are able to predict with good quantitative accuracy the asymptotic response distribution on the basis of estimating a single parameter, namely, the parameter of the smearing distribution.

Estimation of learning parameter. To analyze the goodness of fit of the reinforcement-dependent statistics, it is first necessary to estimate a second parameter. In the linear model this parameter is usually designated as the learning parameter θ. In the stimulus-sampling pattern models it is the learning parameter c/N, where c is the probability of conditioning the stimulus pattern sampled on each trial, and N is the number of stimulus patterns available for sampling. It should be noted that at the level of the reinforcement-dependent statistics it is not possible to make separate estimates of c and N but only to estimate the ratio c/N.

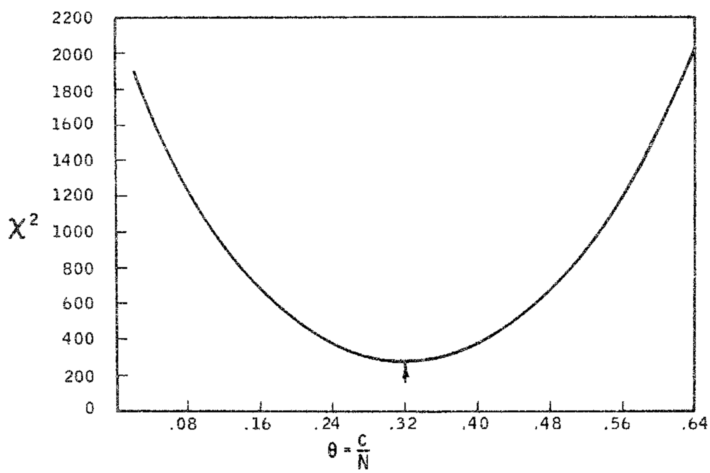

FIG. 3. Chi-square as a function of θ for the fit of $P(X_n \mid Y_{n-1})$.

For the bimodal reinforcement distribution the circle naturally divides into four symmetric arcs, each of length $\pi/2$. For these four divisions and the reinforcement statistic $P(X_n \mid Y_{n-1})$, we used a minimum χ^2 technique to estimate the learning parameter. The value obtained was $\theta^* = (c/N)^* = .32$. The entire χ^2 function for this statistic is shown in Fig. 3. As has proved to be the case in a number of learning experiments (see several examples in Suppes and Atkinson, 1960, as well as the corresponding figure in the preceding article), the exact value of θ (or c/N) is not too important for the goodness of fit. A value of $\theta^* \pm .02$ does not significantly affect the goodness of fit.

Goodness of fit of reinforcement-dependent statistics. The observed proportions and theoretical predictions for the statistic $P(X_n \mid Y_{n-1})$ are shown in Table 2. The qualitative agreement of the sixteen observed and predicted

TABLE 2

PREDICTION OF RESPONSE QUADRANT GIVEN QUADRANT OF THE PRECEDING REINFORCEMENT $[P(X_n \mid Y_{n-1})]$

(First number of each pair is observed proportion based on last 400 trials, and the second is the theoretical prediction)

Y_{n-1} \ X_n	1	2	3	4
1	.408	.268	.148	.176
	.428	.224	.170	.178
2	.293	.408	.149	.150
	.224	.428	.178	.170
3	.155	.151	.395	.299
	.170	.178	.428	.224
4	.147	.152	.278	.423
	.178	.170	.224	.428

probabilities is reasonably good. Yet, as is evident from Fig. 3, the χ^2 measure of the goodness of fit based on the 12,000 observations is in terms of conventional levels of significance not satisfactory ($\chi^2 = 285.14$, df $= 10$). As was remarked in the preceding paper, because of the large number of observations, the probability of errors of the second kind in these goodness of fit tests is essentially zero. Several different approaches may be used to provide additional measures of the goodness of fit. The most direct and natural statistical approach is to select a simple alternative hypothesis and equate errors of the first and second kind. However, rather elaborate calculations are required first to determine the noncentrality parameter and then to compute the integrals corresponding to the equating of the two kinds of errors. A simpler and intuitively appealing approach is to use the χ^2 computations to report the average percentage of error in the prediction. This computation is based on the following consideration.

The χ^2 statistic is given by the following equation:

$$\chi^2 = \sum_i n_i \left[\sum_j \frac{[\hat{p}_{ij} - p_{ij}(\theta)]^2}{p_{ij}(\theta)} \right], \tag{9}$$

where $\hat{p}_{ij}$ is the observed and $p_{ij}(\theta)$ the predicted transition probability, and n_i is the number of observations in the ith row. In computing χ^2 on a high-speed digital computer, it is convenient to print out for each row the quantity

$$\sum_j \frac{[\hat{p}_{ij} - p_{ij}(\theta)]^2}{p_{ij}(\theta)}. \tag{10}$$

Let us suppose now that the percentage of error of each prediction in the ith row is the same, i.e., we assume that approximately for each j,

$$e_i = \frac{|\hat{p}_{ij} - p_{ij}(0)|}{p_{ij}(0)}. \tag{11}$$

Then, combining (10) and (11), we have

$$\sum_j \frac{[\hat{p}_{ij} - p_{ij}(0)]^2}{p_{ij}(0)} = \sum_j \frac{e_i^2 p_{ij}(0)^2}{p_{ij}(0)} = e_i^2 \tag{12}$$

because $\sum_j p_{ij}(0) = 1$. The over-all average error we then define as

$$e = \frac{1}{m} \sum_{i=1}^{m} e_i. \tag{13}$$

We do not propose to give an exact statistical interpretation of e here, and we emphasize the approximate character of the computation, particularly as reflected in (11).[1] Nevertheless, we feel that reporting the value of e is useful, particularly to bring into perspective reported χ^2's for experiments with a very large number of observations. In the case of Table 2, $e = .15$. At this stage of investigation we are reasonably well satisfied with an "average" prediction error of 15 per cent.

Because of the symmetries in the reinforcement distribution, the theoretical predictions shown in Table 2 actually reduce to four different numbers. For example, for *all* of the models considered, when the reinforcement occurs in a particular quadrant and the prediction is the probability of a response in that same quadrant, the prediction is the same regardless of which of the four quadrants of the circle the reinforcement occurs in. Inspection of the observed proportions shown in Table 2 indicates that these symmetries are also reflected very well in the empirical data. The four numbers resulting from the fourfold division of the circle have a simple interpretation. First, there is the prediction that the response will be in the *same* quadrant as the reinforcement. Second, there is what we term the *line-adjacent* prediction that the response will be in the quadrant adjacent to the reinforcement quadrant in the direction of the point of maximum frequency of the reinforcement distribution. Third, there is what we term the *point-adjacent* prediction. In this case the prediction is that the response will be in the quadrant adjacent to the reinforcement quadrant but in the other main half of the reinforcement distribution. Finally, there is the case in which the prediction is for the quadrant that is separated from the reinforcement quadrant by one other quadrant. This we call the *alternate* case. This notation is made clear in Fig. 4. The letters S, L, A, and P (for *s*ame, *l*ine adjacent, *a*lternate and *p*oint adjacent) are placed in the quadrants of the circle on the hypothesis that

[1] A detailed investigation of various measures of error is now under way by the first author and Dr. Helen Chmura Kraemer. The use of e in the present paper is just intended to give a rough indication of fit. A first and immediate suggestion for change in the definition of e is to replace (13) by the sum weighted according to the row frequencies n_i.

the reinforcement event Y_{n-1} is the first quadrant on the left in the top figure, and that Y_{n-1} is the second quadrant in the bottom figure. The remoteness of the point-adjacent quadrant P in the top figure is deceptive because the figure does not show the periodicity of the circle. We call this the SLAP reduction of the 4×4 table.

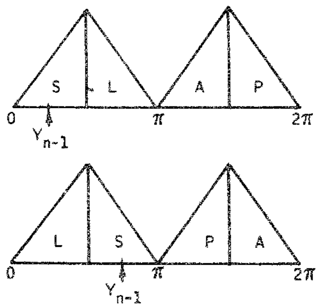

FIG. 4. Diagram illustrating the SLAP relations.

In order to check that this reduction of the data has been approximately guaranteed by the experimental design, a χ^2 test of goodness of fit was run with the four *empirical* SLAP numbers as the parameters. It is to be emphasized that this test in no way depends on the parameters of the models. We simply considered the 4×4 matrix of transition probabilities $P(X_n \mid Y_{n-1})$, and used as "theoretical" probabilities the four estimated SLAP numbers With 8 net degrees of freedom, the resulting $\chi^2 = 20.47$ is just significant at the conventional .01 level, but for the large number of observations the fit is excellent. In terms of the concept of "average" error introduced above, e is less than 4 per cent.

Applying the SLAP reduction to the observed data, we compare in Table 3 the observed and predicted probabilities for the four SLAP quadrants. The χ^2 measuring the goodness of fit is comparable ($\chi^2 = 261.02$) to that obtained for the unreduced data of Table 2.

In certain respects the kind of presentation provided by Table 2 and Table 3 belies the continuous character of the reinforcement and response distributions. A more detailed and instructive picture of the behavior of subjects with respect to the reinforcement-dependent statistics is given by the conditional densities corresponding to the discrete densities shown in Tables 2 and 3. We have used the symmetries implied by the SLAP test to construct a single conditional density $j(x_n \mid Y_{n-1})$. The observed histogram and predicted density are shown in Fig. 5. The figure is drawn for the SLAP data with Y_{n-1} always placed as the interval $(\pi, \frac{3}{2}\pi)$. The qualitative agreement between the observed histogram and the predicted density is quite good. The

TABLE 3

OBSERVED PROPORTIONS AND PREDICTIONS FOR SLAP REDUCTION OF $P(X_n \mid Y_{n-1})$

(Observed proportions based on last 400 trials)

	S	L	A	P
Obs.	.409	.284	.151	.156
Pred.	.428	.224	.170	.178

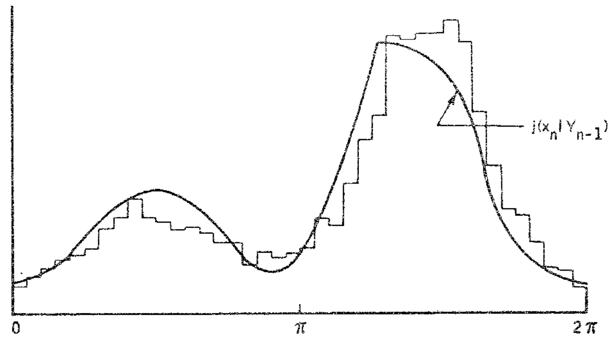

FIG. 5. Observed response histogram conditioned upon preceding reinforcement with corresponding predicted density $j(x_n \mid Y_{n-1})$.

two most important characteristics of the predicted response density that are well confirmed in the data should be noted. One is the asymmetry of the predicted distribution. The other, which is more subtle and much less likely to be predicted without a definite quantitative theory, is the secondary peak or wave in the density for the half of the circle in which a reinforcement did not occur. On non-theoretical intuitive grounds one natural prediction would be that the response density monotonically decreases away from the quadrant in which a reinforcement occurred. The prediction of the wave as opposed to a monotonically decreasing function is well supported by the empirical data.

We turn now to the second-order reinforcement statistics $P(X_n \mid Y_{n-1}, Y_{n-2})$. The qualitative agreement between the 64 observed and predicted probabilities for the four quadrants is reasonably good, but definite quantitative discrepancies do occur, as is reflected in the χ^2 that has a value of 1343.8 with 46 df. To examine these discrepancies more closely and yet reduce considerably the number of quantities to be considered, we have been able to apply the SLAP method of reduction to reduce the 16×4 table of second-order statistics to a 4×4 table. As in the case of Table 3, we first performed a χ^2 test of homogeneity to justify the SLAP reduction of the data. The results were not significant ($\chi^2 = 35.80$, df $= 32$, and $P > .30$).

The observed and predicted probabilities are shown in Table 4. A word of explanation is perhaps required about the use of the SLAP coordinates in the case of the second-order statistics. The row designations refer to the SLAP relations between the two reinforcements, and the column designations for the response X_n refer to the SLAP relation between the response quadrant X_n and the reinforcement Y_{n-1}.

Inspection of row P of Table 4 indicates what is perhaps the most serious discrepancy in the prediction of the second-order reinforcement statistics. The observed proportions for column L fit almost exactly the predicted probability for column P, and the observed proportions for column P fit almost exactly the predicted probability for column L. Analysis of Eq. (3)

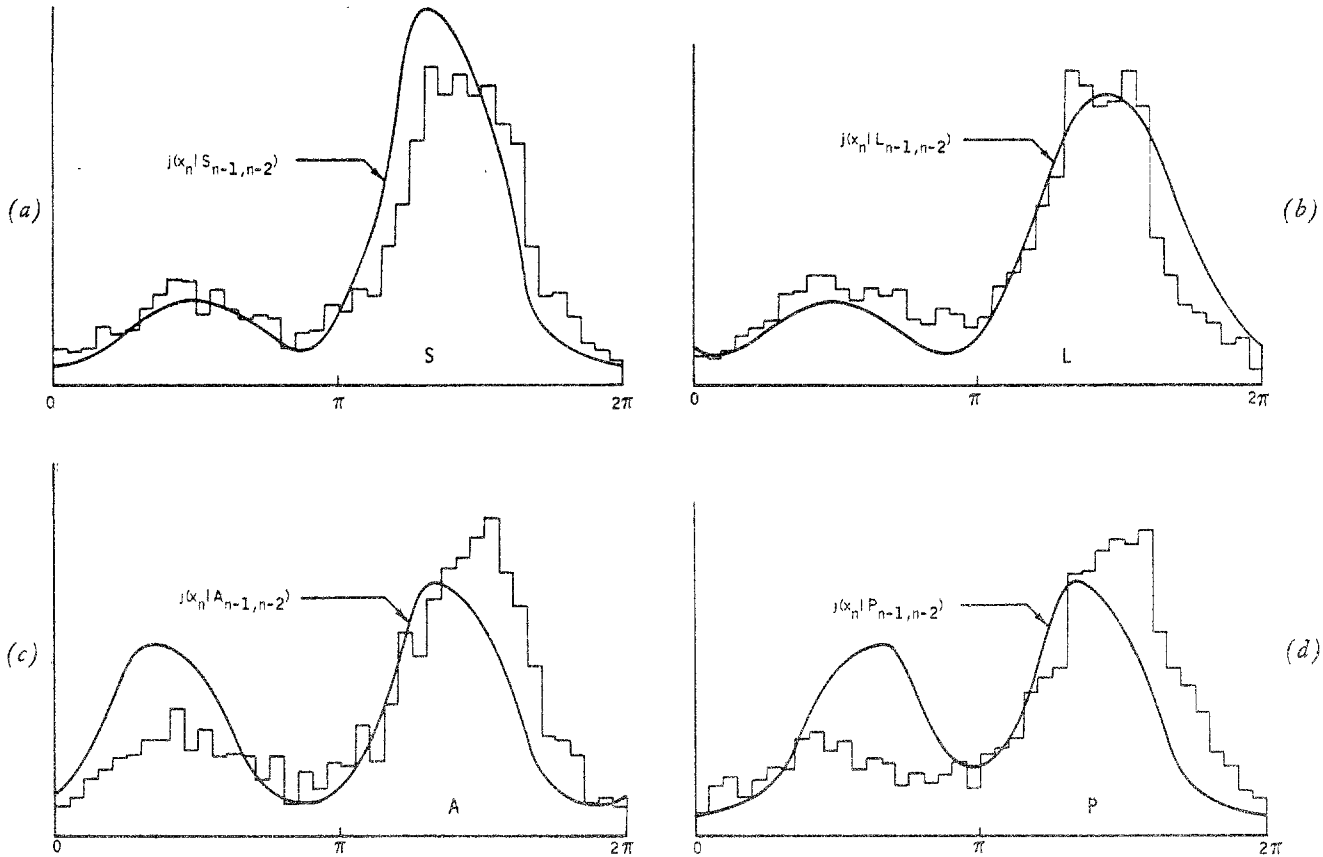

FIG. 6. Observed response histograms conditioned upon two preceding reinforcements with corresponding predicted densities.

TABLE 4

OBSERVED PROPORTIONS AND PREDICTIONS FOR SLAP REDUCTION OF SECOND-ORDER REINFORCEMENT STATISTICS $P(X_n \mid Y_{n-1}, Y_{n-2})$

(Observed proportions based on last 400 trials. First number of each pair is observed proportion)

Y_{n-2}, Y_{n-1} \ X_n	S	L	A	P
S	.433	.277	.146	.144
	.550	.206	.115	.129
L	.419	.258	.146	.177
	.410	.345	.121	.124
A	.394	.287	.168	.151
	.374	.175	.291	.160
P	.388	.314	.148	.150
	.380	.169	.152	.299

indicates that no simple change in the learning parameter could seriously improve these predictions. The difficulty is that the theory predicts much stronger effects for the second preceding reinforcement Y_{n-2} than in fact it seems to be having. We have the somewhat unexpected circumstance that if we compute the fit of the theoretical $P(X_n \mid Y_{n-1})$ to the second-order statistics given in Table 4, the fit is better than that of the predictions shown in the table. Adjustment of the parameter of the smearing distribution would improve the fit of the theoretical second-order reinforcement statistic to the observed data. It is evident from Eq. (3) that the value of the function $H(X_n, Y_{n-1})$ must increase considerably to account for the discrepancies in row P, but to increase this function in the case of column L a very substantial increase in the range of the smearing distribution is required. Such an adjustment would seriously disturb the excellent fit of the asymptotic response distribution, as shown in Fig. 2.

As would be expected from Table 4, when we turn to the conditional densities $j(x_n \mid Y_{n-1}, Y_{n-2})$, the worst fits are for the A and P cases of the SLAP reduction. It is in these two cases that Y_{n-1} and Y_{n-2} are most different, and from an empirical standpoint the theory errs in attaching too much importance to the reinforcing effects of Y_{n-2}, as is apparent from Fig. 6. The figures are drawn for the SLAP data with Y_{n-1} always placed as the interval $(\pi, \frac{3}{2}\pi)$. The notation $j(x_n \mid S_{n-1,\, n-2})$ in Fig. 6*a*, for instance, means that this is the conditional density for Y_{n-1} and Y_{n-2} being the same interval; similar SLAP definitions apply for $j(x_n \mid L_{n-1,\, n-2})$, $j(x_n \mid A_{n-1,\, n-2})$, and $j(x_n \mid P_{n-1,\, n-2})$. The large secondary wave in the predicted curves for the A and P cases (Figs. 6*c* and 6*d*) result from the weight given Y_{n-2} and are simply not substantiated by the data, although there is a small secondary wave in the empirical histogram of both figures. The relatively better fit of the S and L curves, where Y_{n-1} and Y_{n-2} are more nearly acting together, may be observed in Figs. 6*a* and 6*b*.

Quite apart from the consideration of any models, it is clear from the near constancy of the empirical probabilities in the individual columns of Table 4 that the reinforcement Y_{n-2} two trials back had only a very slight effect on subjects' responses. This is particularly surprising in view of the relative success of Model II in the preceding paper, for Model II of that paper attaches equal weight to Y_{n-1} and Y_{n-2}. We have no real explanation of this discrepancy, but it is worth noting that in the experiment reported in the preceding article, the maximum difference between two reinforcements for a given discriminative stimulus is $\frac{2}{3}\pi$. In the present experiment, for cases A and P, Y_{n-1} and Y_{n-2} represent reinforcing events occuring in widely separated parts of the circle. Faced with this conflicting evidence of target location, subjects evidently rely on the last reinforcement almost entirely. In stimulus-response terms this suggests that the conditioning parameter c should be nearly 1. When we consider additional sequential statistics that permit a separate estimate of c and N this qualitative inference about c is well-confirmed, as we shall now show.

Model-differentiating statistics. It was remarked earlier that the asymptotic sequential predictions $P(X_n \mid Y_{n-1}, X_{n-1})$ may be used to differentiate the linear and stimulus-sampling models. As in the case of the reinforcement-dependent statistics, we performed the SLAP reduction defined the same as for $P(X_n \mid Y_{n-1}, Y_{n-2})$, with the response interval X_{n-1} replacing Y_{n-2}. When the SLAP reduction is tested against the 16×4 matrix of observed proportions, $\chi^2 = 56.40$, which is significant at the .01 level; e, the average error defined earlier, is about 6 per cent. This fit seemed good enough to justify using the reduction.

The observed and predicted probabilities for the linear and N-element stimulus-sampling models are shown in Table 5 in terms of SLAP coordinates. There are two striking things about this table. In the first place, as in the case of Table 4, the empirical data in a given column are nearly constant, indicating that subjects are only slightly influenced by their own immediately preceding responses in comparison with the preceding reinforcement. Secondly, the N-element model is able to reflect this empirical fact in a very direct and simple fashion. When $c = 1$, in the N-element model,

$$P(X_n \mid Y_{n-1}, X_{n-1}) = P(X_n \mid Y_{n-1}),$$

as may be easily seen from Eqs. (2) and (6) when it is noted once again that $\theta = c/N$ in Eq. (2). By then taking $1/N = c/N = .32$, as previously estimated from $P(X_n \mid Y_{n-1})$, the predictions of the N-element model shown in Table 5 are row invariant and simply those made earlier for $P(X_n \mid Y_{n-1})$. To check that these estimates were indeed the best possible for the N-element model, we computed a joint minimum χ^2 estimate of c and N. The

TABLE 5

OBSERVED PROPORTIONS AND PREDICTIONS FOR SLAP REDUCTION OF MODEL-DIFFERENTIATING STATISTICS $P(X_n \mid Y_{n-1}, X_{n-1})$

(Observed proportions based on last 400 trials. First number of each triple is observed proportion, the second is the prediction of the N-element pattern model, and the third is prediction of the linear model)

X_{n-1}, Y_{n-1} \ X_n	S	L	A	P
	.419	.278	.156	.147
S	.428	.224	.170	.178
	.492	.220	.138	.150
	.429	.264	.134	.173
L	.428	.224	.170	.178
	.424	.288	.142	.146
	.388	.311	.140	.161
A	.428	.224	.170	.178
	.396	.196	.234	.174
	.399	.284	.175	.142
P	.428	.224	.170	.178
	.401	.191	.166	.242

results were $c^* = 1$ and $1/N^* = .32$, showing that the fit of the N-element model could not be improved by a deviation from these values. The contour lines of the iso-χ^2 values for various values of the two parameters are shown in Fig. 7. The gradient near the minimum χ^2 is not very steep along either axis, and thus shows, as did Fig. 3, that the goodness of fit is not sensitive to the exact values of the parameters.

The value of the minimum χ^2 for the N-element model is 313.17, which is of course highly significant in terms of conventional levels of significance. The average error e is slightly less than 16 per cent.

The predictions of the linear model are shown as the third figure in each of the entries in Table 5. The χ^2 measuring the over-all fit of these predictions

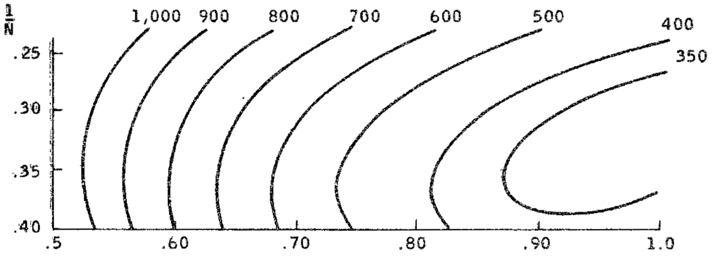

FIG. 7. Iso-χ^2 lines for χ^2 as a function of c and $1/N$.

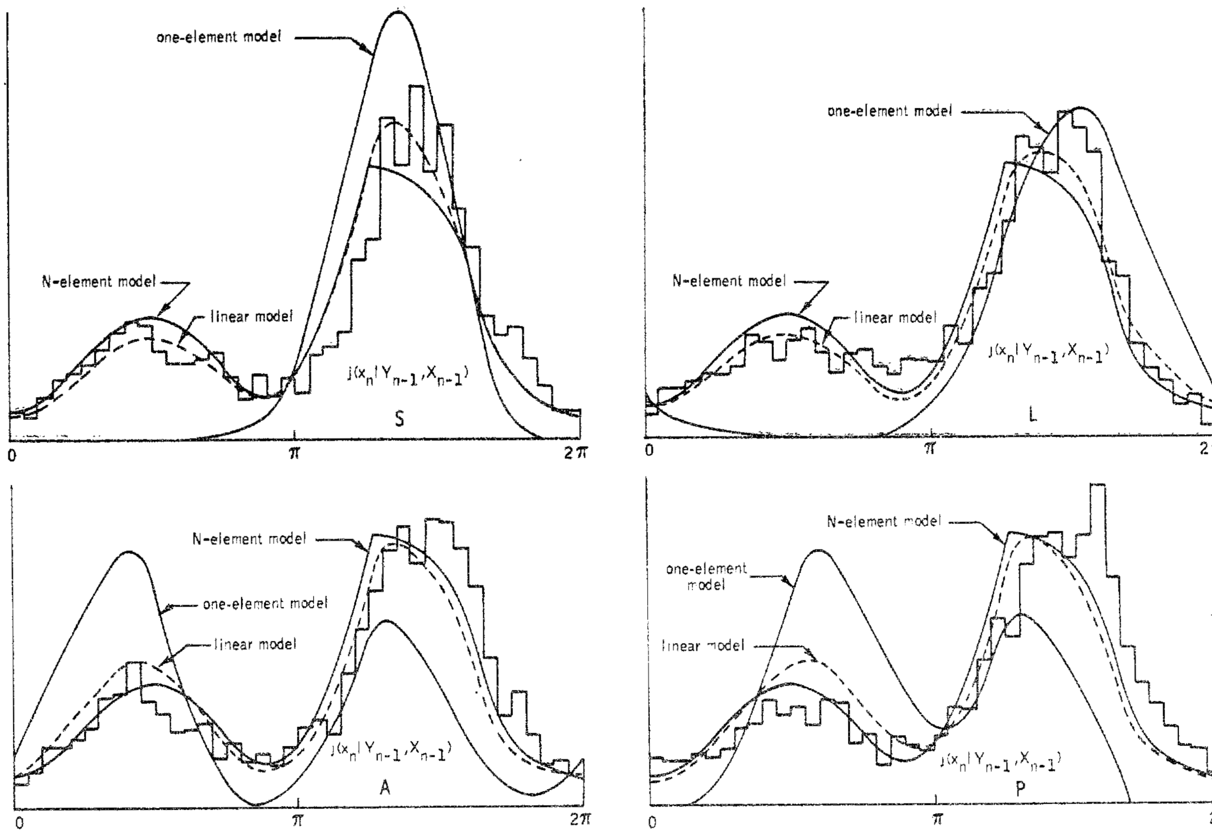

FIG. 8. Observed response histogram conditioned upon preceding reinforcement and response with corresponding densities predicted by various models.

is 626.80, which is twice as large as that obtained for the N-element model. As would be anticipated from remarks already made and as is also evident from inspection of Table 5, the linear-model fit is good for rows S and L, where the reinforcement and previous response are close together, but is much worse for rows A and P, where Y_{n-1} and X_{n-1} are separated. In terms of the obtained χ^2's the evidence definitely favors the N-element pattern model over the linear model.[2]

On the other hand, the one-element stimulus-sampling model fits the data of Table 5 very badly ($\chi^2 = 10124.1$). The difficulties here are those that beset the one-element model in the analysis of sequential statistics in discrete choice data: repetition of a reinforced response is predicted with too high a probability (Suppes and Atkinson, 1960). Thus in the discrete choice experiment, $P(X_{i,n} \mid Y_{i,n-1}, X_{i,n-1}) = 1$, and for the present experiment in SLAP coordinates, $P(S_n \mid S_{n-1,n-1}) = .764$, but the observed proportion is .419.

The observed histograms as well as the conditional densities $j(x_n \mid Y_{n-1}, X_{n-1})$ of the linear, N-element, and one-element models are shown in Fig. 8. SLAP coordinates are used to reduce the number of densities to four. Thus, when Y_{n-1} and X_{n-1} are the same quadrant, we represent the conditional density as $j(x_n \mid S_{n-1,\, n-1})$, and similarly for the other three cases L, A, and P. We selected the one-element model (with $c^* = .32$) as well as the better fitting N-element model (with $c^* = 1.00$ and $1/N^* = .32$) to exhibit in these figures for two reasons. First, it is instructive to see just where the large discrepancies occur in its fit to the data, and, second, the conditional densities of the N-element model have in fact already been shown in Fig. 5, for as already remarked, with $c = 1$, $P(X_n \mid Y_{n-1}) = P(X_n \mid Y_{n-1}, X_{n-1})$, and thus there is just the one density of Fig. 5 to represent all four SLAP conditions.

The four figures bear out the earlier remarks about the linear model, that is, that the fit for conditions S and L is definitely better than for A and P.

It is interesting to note that in two of the four conditions, namely, S and L, the one-element model does not predict the secondary peak in the observed histogram. The prediction of the outstanding qualitative features of various continuous conditional response densities is probably the most important task at the present stage of development for the kind of models considered in this paper. We are ourselves encouraged by the relatively good job done by the N-element model. It is also apparent that the consideration of such densities provides a relatively sharp tool for discriminating between models. This is not surprising, for as the sequence of conditional events is increased in length—for example, as in going from $j(x_n \mid Y_{n-1})$ to $j(x_n \mid Y_{n-1}, X_{n-1})$—a closer approximation to a sufficient statistic for all the models is reached.

---

[2] Results of the same order of magnitude hold when the two models are tested against the unreduced 16×4 matrix: $\chi^2 = 369.96$ for the N-element model; $\chi^2 = 657.58$ for the linear model.

4. Summary

Statistical learning theory has been generalized to produce models for the study of learning phenomena in which both responses and reinforcements are measured on a continuous scale. This paper reports another empirical test of these models.

Subjects were instructed to predict, in a sequence of 600 trials, the position of a point target on the circumference of a circle. In fact, the correct positions of the target—the reinforcements—were independent samples from a fixed distribution having a symmetric, bimodal density.

All of the models considered predict the same asymptotic unconditional response distribution and conditional response distribution whenever the conditioning event specifies only reinforcements. The correspondence between the observed and predicted asymptotic response distribution is strikingly close. Predictions of the conditional densities range from fair to excellent. One unanticipated finding (namely, a secondary peak in the conditional response densities) is well described by the models.

The models differ in their predictions of conditional response distributions if the conditioning event specifies a preceding response. On the basis of these statistics, one easily rejects the one element model. Both the N-element and linear models predict the outstanding qualitative characteristics of these conditional distributions, but the N-element model, being better able to emphasize the effects of the most recent past, more precisely describes the data.

Appendix

In this Appendix we derive the asymptotic expressions for $P(X_n \mid Y_{n-1}, X_{n-1})$ in the N-element stimulus-sampling model. The axioms as well as the notations used are those of Suppes (1960).

To obtain the desired result, we first need to derive the joint distribution at asymptote of the smearing parameters z and z' of two distinct stimulus elements s and s'. We begin by deriving a recursion for $g_n(z, z')$. We first note that if neither s nor s' is sampled on trial $n - 1$, then $g_n = g_{n-1}$. The probability of this event is $(1 - 2/N)$. Second, if either element is sampled but conditioning is not effective, then also $g_n = g_{n-1}$. The probability of this event is $2/N(1 - c)$. The two remaining cases to be dealt with arise when either s or s' is sampled and conditioning is effective. For definiteness suppose s is sampled. We then have with probability c/N:

$$\text{(A-1)} \qquad g_n(z, z') = \iiint \delta((z, z') - (y_{n-1}, z'_{n-1}))f(y_{n-1}) \, g_{n-1}(z_{n-1}, z'_{n-1}) \, dy_{n-1}dz_{n-1}dz'_{n-1}.$$

Integrating out z_{n-1} and applying the usual technique to the Dirac delta function δ, we obtain

$$g_n(z, z') = f(z)g_{n-1}(z').$$

Combining these various results, we obtain then the following recursion:

$$g_n(z, z') = \left(1 - \frac{2}{N}\right)g_{n-1}(z, z') + \frac{2}{N}(1 - c)g_{n-1}(z, z') + \frac{c}{N}f(z)g_{n-1}(z') + \frac{c}{N}g_{n-1}(z)f(z').$$

As $n \to \infty$, writing g for both the joint and marginal distributions we have, after a little simplification,

$$2g(z, z') = f(z)g(z') + g(z)f(z'). \tag{A-2}$$

Integrating both sides of this expression with respect to z', we obtain that

$$2g(z) = f(z) + g(z)$$

and thus

$$g(z) = f(z),$$

and by a similar argument

$$g(z') = f(z'),$$

whence from (A-2) we conclude that

$$g(z, z') = f(z)f(z'), \tag{A-3}$$

the result we need.

We now begin with the joint density $j(x_n, y_{n-1}, x'_{n-1})$ and make the usual expansion in order to apply the axioms. The integrals are taken over the entire interval of responses, and the summations are taken over the entire relevant set of events:

$$j(x_n, y_{n-1}, x'_{n-1}) = \sum_s \int_z \sum_i \sum_{s'} \int_{z'} \cdot j(x_n, s_n, z_n, e_{i,n-1}, y_{n-1}, x'_{n-1}, s'_{n-1}, z'_{n-1})\, dz_n dz'_{n-1}, \tag{A-4}$$

where s_n is the event of sampling stimulus s on trial n, z_n is the N-dimensional vector of parameters of the smearing distributions of the N stimuli on trial n, $e_{i,n}$ is the event of conditioning being effective or not on trial $n - 1$ ($i = 1$ if effective; $i = 0$ if not), s'_{n-1} is the event of sampling stimulus s' on trial $n - 1$, and z'_{n-1} is the vector of parameters on trial $n - 1$. Subsequently we use the notation $z_{s,n}$ for the parameter of the smearing distribution of stimulus s on trial n, i.e., $z_{s,n}$ is the component for stimulus s of the vector z_n. The primes have been introduced in order that we may drop the trial subscripts to simplify the rather lengthy expressions arising in the derivation. We now conditionalize the joint density on the right-hand side of (A-4) and use the relations

$$k_s(x; z_s) = j(x \mid s, z, e_i, y, x', s', z'), \qquad \frac{1}{N} = j(s \mid z, \cdots, z'),$$

$$j(e_i) = j(e_i \mid y, x', s', z'), \qquad f(y) = j(y \mid x', s', z'),$$

$$k_{s'}(x'; z'_{s'}) = j(x' \mid s', z'), \qquad \frac{1}{N} = j(s' \mid z'),$$

to obtain

$$\text{(A-5)} \qquad j(x, y, x') = \sum_{s} \int_{z} \sum \sum_{s'} \int_{z'} k_s(x; z_s) \frac{1}{N} \cdot j(z \mid e_i, y, x', s', z') j(f_i) f(y) \cdot k_{s'}(x'; z'_{s'}) \frac{1}{N} j(z') \, dz dz'.$$

We now consider two cases, according to whether $s = s'$ or $s \neq s'$, and we also may integrate out all of the vector z except the component z_s and all of the vector z' except $z'_{s'}$. If $s = s'$, then

$$\delta(z_{s'} - y) = j(z_{s'} \mid e_i, y, x', s', z') \qquad \text{if } i = 1,$$
$$\delta(z_{s'} - z'_{s'}) = j(z_{s'} \mid e_i, y, x', s', z') \qquad \text{if } i = 0,$$
$$c = j(e_i),$$
$$1 - c = j(e_0),$$

and if $s \neq s'$,

$$g_n(z'_s \mid z'_{s'}) = j(z_s \mid e_i, y, x', s', z').$$

If $s = s'$ or $s \neq s'$, then

$$g_n(z'_{s'}) = j(z'_{s'}).$$

Applying these results to (A-5), we then have

$$\text{(A-6)} \quad j(x, y, x') = \frac{1}{N^2} \Big\{ \sum_{s'} \int_{z_{s'}} \int_{z'_{s'}} k_{s'}(x; z_{s'}) [c\delta(z_{s'} - y) + (1 - c)\delta(z_{s'} - z'_{s'})] \cdot f(y) k_{s'}(x'; z'_{s'}) g(z'_{s'}) \, dz_{s'} dz'_s$$
$$+ \sum_{s \neq s'} \int_{z_s} \sum_{s'} \int_{z'_{s'}} k_s(x; z_s) g(z_s \mid z'_{s'}) f(y) k_{s'}(x'; z'_{s'}) g(z'_{s'}) \, dz_s dz'_{s'} \Big\}.$$

Equation (A-6) holds for every n. We now apply the asymptotic result expressed by (A-3) to derive an expression holding *only* at asymptote. To apply (A-3) we observe that $g(z_s, z'_{s'}) = g(z_s \mid z'_{s'}) g(z'_{s'})$. We may also note that on the assumption that the forms of the smearing distributions are the same for all stimuli, we may perform the indicated summations in (A-6) and drop the subscripts. At the same time we perform the integration involving the Dirac delta functions and obtain as a summary result

$$\text{(A-7)} \qquad j(x, y, x') = \frac{1}{N^2} [Nc \textstyle\int k(x; y) f(y) k(x'; z') f(z') \, dz'$$
$$+ N(1 - c) \textstyle\int k(x; z') f(y) k(x'; z') f(z') \, dz'$$
$$+ (N^2 - N) \textstyle\iint k(x; z) f(z) f(y) k(x'; z') f(z') \, dz dz'].$$

Eq. (A-7) brings us close to the desired result. We note that in terms of the definitions given in connection with Eqs. (4) and (5),

$$h(x, y) = k(x; y) f(y)$$
$$r(x') = \textstyle\int k(x'; z') f(z') \, dz'$$
$$r(x, x') = \textstyle\int k(x; z') k(x'; z') f(z') \, dz',$$

whence by substitution in (A-7) we get

$$j(x, y, x') = \frac{c}{N}h(x, y)r(x') + \frac{(1-c)}{N}r(x, x')f(y) + \left(1 - \frac{1}{N}\right)r(x)r(x')f(y). \tag{A-8}$$

Equation (6) follows immediately from (A-8) by integration over X_n, Y_{n-1}, and X_{n-1} and appropriate conditionalization. Thus

$$P(X_n, Y_{n-1}, X_{n-1}) = \frac{c}{N}H(X_n, Y_{n-1})R(X_{n-1}) + \frac{1-c}{N}R(X_n, X_{n-1})F(Y_{n-1}) + \left(1 - \frac{1}{N}\right)R(X_n)R(X_{n-1})F(Y_{n-1}),$$

and

$$P(X_n \mid Y_{n-1}, X_{n-1}) = \frac{c}{N}\frac{H(X_n, Y_{n-1})}{F(Y_{n-1})} + \frac{(1-c)}{N}\frac{R(X_n, X_{n-1})}{R(X_{n-1})} + \left(1 - \frac{1}{N}\right)R(X_n),$$

which is precisely (6).

The derivation of Eq. (5) for the linear model is a much simpler matter and may be obtained quite directly by extension of results in Suppes (1959) and in the Appendix of Suppes (1960).

REFERENCES

Suppes, P. A linear model for a continuum of responses. In R. R. Bush and W. K. Estes (Eds.), *Studies in mathematical learning theory*. Stanford, Calif.: Stanford Univer. Press, 1959. Pp. 400–414.

Suppes, P. Stimulus-sampling theory for a continuum of responses. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical methods in the social sciences*, 1959. Stanford, Calif.: Stanford Univer. Press, 1960. Pp. 348–365.

Suppes, P., and Atkinson, R. C. *Markov learning models for multiperson interactions*. Stanford, Calif.: Stanford Univer. Press, 1960.

Suppes, P., and Frankmann, R. W. Test of stimulus sampling theory for a continuum of responses with unimodal noncontingent determinate reinforcement. *J. exp. Psychol.*, 1961, 61, 122–132.