The Stanford Institute for Mathematical Studies in the Social Sciences

APPLIED MATHEMATICS AND STATISTICS LABORATORIES
STANFORD UNIVERSITY

Reprint No. 33

Application of Stimulus Sampling Theory To Situations Involving Social Pressure

PATRICK SUPPES AND FRANKLIN KRASNE

Reprinted from
Psychological Review
Vol. 68, No. 1, 1961

Psychological Review
1961, Vol. 68, No. 1, 46–59

APPLICATION OF STIMULUS SAMPLING THEORY TO SITUATIONS INVOLVING SOCIAL PRESSURE[1]

PATRICK SUPPES AND FRANKLIN KRASNE

Stanford University

Much contemporary social psychological theory appears to be theory of group behavior qua group. While group oriented concepts may be transiently useful in classifying and systematizing the vast experimental literature of social psychology, we hold that ultimately group behavior can be explained entirely in terms of the behavior of the individuals who constitute the group. In other words we regard the theory of social behavior as a highly important special case of the general theory of individual behavior. Furthermore, to our minds recent quantitative formulations of stimulus-response-reinforcement theory provide excellent conceptual tools for effecting such a subsumption.

This is not to claim that we are prepared to give a detailed stimulus-response-reinforcement analysis of every experimental situation now considered important by social psychologists, but it is to claim that the general lines of such an analysis are clear and in fact may be given in detail, as we shall see for a representative example, for a rather large class of social interaction situations. We conceive of an interaction situation as one in which each member of a group (potentially) provides stimuli and reinforcements for every other member of the group with the behavior of each member entirely explicable on an individual basis given the sequence of stimuli and reinforcements impinging on him. Examples of small group experiments analyzed from this point of view are Hays and Bush (1954), Atkinson and Suppes (1958, 1959), Burke (1959), and Suppes and Atkinson (1960).

Here we wish to consider the theory of social comparison processes, an area of social psychology usually spoken of in terms of frames of reference, social pressure, group norms, and the like, whose experimental work is both especially provocative and particularly well suited to precise experimentation.

Using strictly the notions of stimulus, response, and reinforcement, it is natural to construe the kinds of social situations with which we are concerned as classical or "almost classical" discrimination experiments. The problem, of course, is to identify in each experiment just what is to be considered as stimulus, what as response, and especially difficult what as reinforcement. Once the identifications have been made it is easy enough to assume that a response reinforced in some stimulus situation will have an increased tendency to be emitted on future occasions of that situation.

In the social situation objective stimuli and the behavior of other members of a group combine to form the relevant stimulus situations for each subject. Response classes are in principle arbitrary, and if an experiment is well structured for the subject, the experimenter will be easily led to a "natural" classification of responses.

[1] This research was supported by the Rockefeller Foundation and the Group Psychology Branch of the Office of Naval Research. This paper will be included in the Reprint Series of the Stanford Institute for Mathematical Studies in the Social Sciences.

Considerable difficulty arises in making reinforcement identifications, because subjects bring with them to an experiment a large number of covert verbal responses having secondary reinforcing properties. Consequently, we are simply forced to limit the number of reinforcers which we wish to recognize as important and for analytical purposes to ignore the rest. We shall suppose first that social support per se is reinforcing; although admittedly there will frequently be factors working against it, we will consider them explicitly when they seem important. The proprioceptive stimuli produced by an overlearned response have consistently preceded reinforcement, and therefore the making of an overlearned response is self-reinforcing via the mechanism of secondary reinforcement; thus, we shall assume secondly that there is some automatic (secondary) reinforcement for making responses which have been well overlearned in everyday experience. Lastly, there may be experimenter-controlled rewards of a less ambiguous nature such as money payoffs, prizes, and primary drive reduction.

As a specific example of the type of analysis which we are suggesting we shall study in detail in the remainder of this paper an experimental situation which is similar in certain respects to the classic experiments of Sherif (1935) and in other respects to those of Asch (1956). The theoretical ideas which have been described will be embedded in a stochastic model of the general type described by Suppes and Atkinson (1960), which is a variant of the stimulus sampling theory of Estes and Burke.

In the experiment to be described subjects were required to make a choice on the basis of an objective but slightly ambiguous stimulus situation; in particular they were asked to indicate which of two lines they thought was longer. This choice was followed by an indication of what the subjects were instructed to believe was the correct answer.

From a social psychological point of view the subject plus the experimenter form a dyad which is the simplest case of a small group. According to the social orientation the experimenter exerts (social) "pressure" on the subject to modify his choices in the direction which the experimenter calls correct. In this sense there are certain similarities to Asch's (1956) study; however, since the subject does not find out what the experimenter considers as correct until after his own response has been made, one might suppose, following Sherif (1935), that the effect of the social pressure imposed by the experimenter is to modify the subject's frame of reference on future stimulus presentations.

Our approach to the problem will be to treat the experimental situation as a stimulus discrimination experiment. The situation is complicated slightly in that subjects have strongly overlearned a relevant visual discrimination prior to this experiment; we acknowledge the strong effect of this past learning by positing a secondary reinforcer which may operate to accentuate or attenuate the effect of the experimenter-controlled reinforcer of social support.

It is the conflict between past learning and the experimenter-controlled reinforcer which generates social pressure that does not exist in the usual discrimination experiment. The subject holds Opinion A (here interpreted as a response probability) and another person (in this case the experimenter), who has for one reason or another the power to reinforce the subject, holds B, a potentially different opinion. That the person giving the second response is the experimenter is irrelevant. It could equally well have been another subject, and by the same token we could easily pass from the dyad to a larger group. This would involve a more laborious but strictly analogous treatment.

THE EXPERIMENT AND ITS THEORY

On each of a sequence of trials, a pair of lines, one slightly longer than the other, was projected for a few seconds on a screen. The subjects were asked to record on answer sheets which of the two lines (labeled Line 1 and Line 2) they thought was longer. They were then told (what they had been instructed to believe would be) the correct answer; however, it was in reality correct only on a randomly chosen subset of the trials.

In order to describe the situation more precisely, let us introduce some notation:

S_1 = Event of projecting a line pair of which Line 1 is longer
S_2 = Event of projecting a line pair of which Line 2 is longer
A_1 = Response of subject on answer sheet indicating that he thinks Line 1 is longer
A_2 = Response of subject on answer sheet indicating that he thinks Line 2 is longer
E_1 = Reinforcing event of experimenter saying, "Line 1 is longer."
E_2 = Reinforcing event of experimenter saying, "Line 2 is longer."
Ω_1 = Secondary reinforcement of A_1 response
Ω_2 = Secondary reinforcement of A_2 response
$\pi_1 = P(E_1|S_1)$
$\pi_2 = P(E_2|S_2)$
$\gamma = P(S_1)$
δ = Probability that a subject makes a correct discrimination in a control experiment where no information is being given him regarding the correctness of his responses

We will be primarily concerned with the behavior of various conditional probabilities of response as a function of π_1, π_2, and δ.

On any trial first a stimulus event (S_1 or S_2), then a response (A_1 or A_2), and then a reinforcing event (E_1 or E_2) occurs. In addition we introduce automatic secondary reinforcing events, Ω_1 and Ω_2, which reinforce the "perceptually correct" response. Thus, on any trial of an experiment the sequence of events may be indicated by the string of symbols,

$$C \rightarrow S \rightarrow A \rightarrow \Omega \rightarrow E \rightarrow C'$$

where C and C' represent conditioning before and after the trial and where we say that a response is *conditioned* to a stimulus event if the response is elicited as a result of that stimulus event. At all times exactly one response is conditioned to a particular stimulus event.

Independently of what response was actually made on a trial we will say that both E_k and Ω_k *reinforce* A_k. If a reinforcement is *effective*, then the reinforced response becomes conditioned to the stimulus event occurring on the trial. If no reinforcement is effective, then conditioning remains unchanged. It will be assumed that exactly one of the following occurs on every trial: social support is effective, secondary self-reinforcement is effective, neither is effective. Since these events are mutually exclusive and exhaustive, we define

$$\theta_1 = P\,(A_k \text{ is effectively reinforced by } \Omega_k | \Omega_k, E_j)$$

$$\theta_2 = P\,(A_j \text{ is effectively reinforced by } E_j | \Omega_k, E_j)$$

$$1 - \theta_1 - \theta_2 = P\,(\text{No reinforcement is effective } | \Omega_k, E_j)$$

Consider the subsequence of trials on which S_1 and S_2, respectively, occur. Since the conditioning of a response to a stimulus event can be affected only if that stimulus event occurs, we can treat these two subsequences as separate and independent simple learning situations (simple learning as opposed to discrimination learning). In particular it is easily seen that each learning process is characterized under our assumptions as a Markov chain with A_1 and A_2 as its states (that is to say on any trial either A_1 or A_2 is conditioned to S_i).

As an example of the method used in deriving the transition matrices for these chains, we shall analyze in terms of a tree of logical possibilities a trial on which S_1 occurred. If the subject ended the last S_1 trial conditioned to A_1, then the relevant tree is

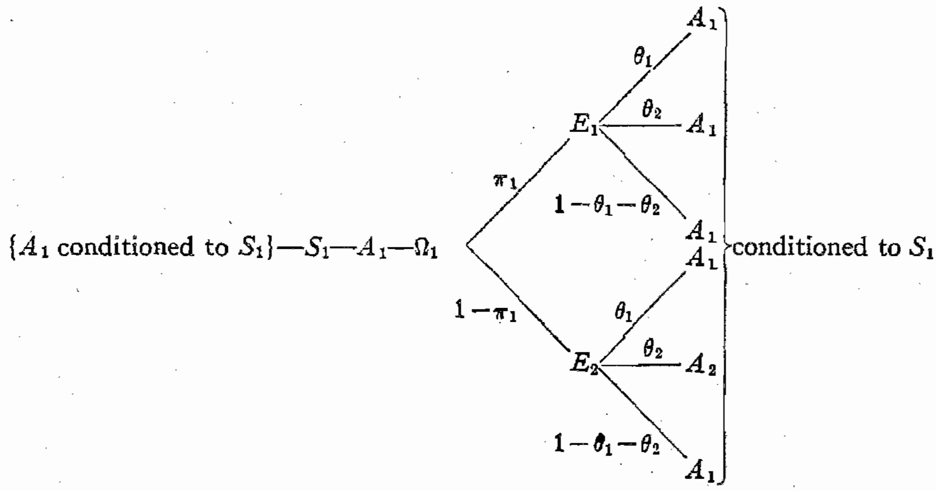

and if the subject ended the last S_1 trial with A_2 conditioned to S_1 we have

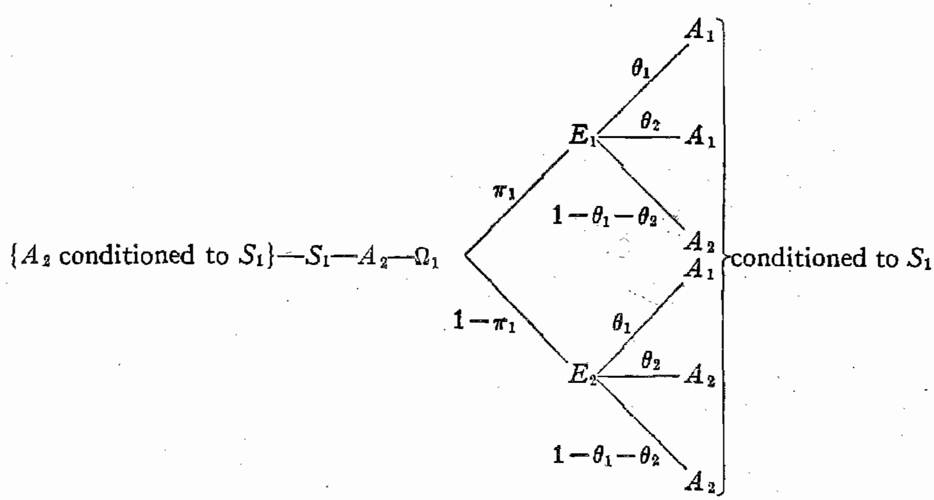

Using similar methods we arrive at the following Markov chains for the S_1 and S_2 learning processes, respectively. On S_1 type trials we identify two states, 1 and 2, referring to A_1 and A_2, respectively, as being conditioned to S_1; this gives the transition matrix

$$\begin{bmatrix} 1-(1-\pi_1)\theta_2 & (1-\pi_1)\theta_2 \\ \theta_1+\pi_1\theta_2 & 1-\theta_1-\pi_1\theta_2 \end{bmatrix} \quad [1]$$

Similarly on S_2 trials identifying states 1 and 2 referring to A_1 and A_2, respectively, being conditioned to S_2 the transition matrix is

$$\begin{bmatrix} 1-\theta_1-\pi_2\theta_2 & \theta_1+\pi_2\theta_2 \\ (1-\pi_2)\theta_2 & 1-(1-\pi_2)\theta_2 \end{bmatrix} \quad [2]$$

Using the fact that the sequence of S_1's and S_2's occur in accordance with a binomial distribution with parameter γ, we may combine the above processes to obtain $P_n(A_1|S_1)$ and $P_n(A_1|S_2)$ representing the probabilities of A_1 responses on the nth trial of the full experiment given S_1 and S_2 stimulus events, respectively;

$$P_n(A_1|S_1) = \frac{\frac{\theta_1}{\theta_2}+\pi_1}{\frac{\theta_1}{\theta_2}+1} + \left[\delta - \frac{\frac{\theta_1}{\theta_2}+\pi_1}{\frac{\theta_1}{\theta_2}+1}\right] \times [1-\gamma(\theta_1+\theta_2)]^{n-1} \quad [3]$$

$$P_n(A_1|S_2) = \frac{1-\pi_2}{\frac{\theta_1}{\theta_1}+1} + \left[1-\delta - \frac{1-\pi_2}{\frac{\theta_1}{\theta_2}+1}\right] \times [1-(1-\gamma)(\theta_1+\theta_2)]^{n-1} \quad [4]$$

$$n = 1, 2, \cdots$$

The psychophysical constants δ and $(1-\delta)$ have been taken as the initial probabilities of A_1 responses on S_1 and S_2 type trials, respectively. It is clear that since $0 \leq \theta_1 + \theta_2 \leq 1$, the right hand terms of each expression vanish for large n, and hence, the asymptotic results are given by the first terms to the right of the equality signs. We have graphed the asymptotic quantities as a function of the ratio θ_1/θ_2 for relevant values of π_1 and π_2 in Figure 1.

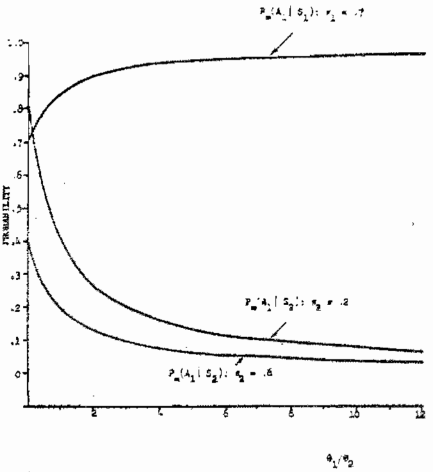

FIG. 1. Asymptotic behavior of $P(A_1|S_1)$ and $P(A_1|S_2)$ as a function of the ratio θ_1/θ_2.

METHOD

Stimuli

The stimuli for these experiments were pairs of parallel lines projected on a field 46 inches wide by 32 inches high; the lines were oriented so that a line passing through their midpoints would be perpendicular to each. Subjects were run under two different stimulus conditions, which we shall designate the "hard" and "easy" conditions to refer to their involving hard and easy discriminations, respectively. Within each difficulty condition (hard or easy) there were two types of slides used (A and B in Table 1), and the orientation of the line pairs was permuted so as to produce $2 \times 2 \times 2 = 8$ different line pairs (e.g., lines vertical, lines horizontal, long line on left, long line on right, etc.). The lines were labeled 1 and 2, respectively, for identification by the subjects. In order to avoid the effect of position habits on the data, for half the subjects the upper horizontal and left vertical lines were always labeled 1 and for the other half the bottom and right lines were called 1.

In Table 1, which gives a description of the line pairs, a is the distance between the lines

TABLE 1

DESCRIPTION OF STIMULI

Slide Type	Hard				Easy			
	a	b	c	δ	a	b	c	δ
A	15.5	5.12	.240	.72	15.5	2.56	.310	.92
B	15.5	7.75	.250	.72	15.5	7.75	.340	.92

Note.—a, b, and c are measured in inches.

TABLE 2

GROUP DESCRIPTIONS

Group Number	π_1	π_2	δ	Description
I	.7	.2	.92	Easy, low symmetry
II	.7	.2	.72	Hard, low symmetry
III	.7	.6	.72	Hard, medium symmetry

of the pair; b is the half-length of the shorter line; c is the difference in half-lengths between the longer and shorter lines; δ is the probability of the subject making a correct discrimination when no experimenter controlled reinforcing events are influencing his behavior.

Room and Apparatus

The line pairs were projected by an Argus 300 Projector with a Sylvania 300-watt projector lamp on to a beaded screen 146 inches away. The subjects (from one to five in number) were seated at a distance of about 96 inches from the front of the screen. The room illumination was about .3 Weston II units.

The subjects responded on answer sheets prepared in essentially the same way as standard IBM score sheets. For half the subjects the left column was used to indicate an A_1 response and for half the right.

Experimental Groups

Three groups of subjects whose conditions are described in Table 2 were run. The degree of symmetry under "Description" refers to relative similarity of reinforcing situations on S_1 and S_2 type trials. For all groups $\gamma = .5$.

Subjects

The subjects, who were students at Stanford University, were obtained from the student employment service, an introductory psychology class, and a university dormitory. With the exception of a few of those from the psychology class all subjects were male. There were 26, 25, and 18 subjects in Groups I, II, and III, respectively. Because of certain technical difficulties in preparing the stimuli, subjects were not placed randomly in groups.

Procedure

There were from one to five subjects per experimental session. The subjects, having been seated and given answer sheets, were given the following instructions. (A sample slide was projected on the screen throughout the instructions.)

> This is an experiment on judgment of length. I am going to flash a number of slides on the screen each of which has two lines on it; your job in each case will be to decide which of the two lines is the longer and to mark your judgment on the answer sheet in the appropriate column opposite the number of the slide we are on. For example: if on the first slide I showed, the line next to the figure one on the screen was longer, you would mark the answer sheet like this [show marked answer sheet]. Sometimes the pair of lines will be vertical rather than horizontal; then the line here would be line one and the one here line two [point to screen to illustrate what I am saying]. Please mark the answer sheets heavily and completely as though you were filling in an IBM score sheet. I will announce the number of the slide we are on before each judgment and will tell you the correct answer after each judgment. You must make your decisions quickly and record them BEFORE I tell you the answer. Each slide will appear for about 2 seconds. One of the two lines will always be longer than the other; however, some of the judgments may be difficult. If you are not sure which line is longer, then guess. Since I want you to work completely independently of each other, I must request that you remain absolutely silent during the experiment. Are there any questions?

The trials then proceeded at a rate of about seven per minute. Each stimulus presentation lasted for approximately 2 seconds. Occasionally the experimenter found it necessary to repeat early in the sessions "this is a test of sorts; you must remain completely silent." Each subject received a total of 250 trials.

After the experiment the subjects were instructed: "Will you please write a brief statement of your reactions to this experiment?" The responses to this question indicated that the subjects understood and believed the instructions.

ESTIMATION METHODS

Two methods will be presented for estimating the parameters θ_1 and θ_2 from the data of our experiments. First we shall make maximum likelihood estimates based on the separate S_1 and S_2 response subsequences

(Anderson & Goodman, 1957); then, taking advantage of the mutual independence of the S_1 and S_2 learning processes, we shall maximize the joint likelihood function to obtain single estimates of θ_1 and θ_2.

It can be shown that the first method is essentially equivalent to setting the expressions for the theoretical transition probabilities equal to their estimates and solving for θ_1 and θ_2.

The estimates based on maximizing the joint likelihood function are not quite so simple. The derivation, which is reasonably straightforward but tedious, leads to a tenth degree equation for the unknown, θ_2; solutions were obtained by approximation methods.

RESULTS AND DISCUSSION

Tables 3 and 4 present learning parameter estimates obtained by the two methods described above.

The most important observation to be made is that for all estimates θ_1/θ_2 is less on hard than on easy discriminations. One can go even further and point out that with the exception of the S_2 trial estimate in Group III this learning parameter ratio is greater than one for the easy discriminations and less than one for the hard ones. This says essentially (for our experimental situation) that if the physical stimulus situation is fairly unambiguous, subjects will "prefer" to derive their reinforcement from it via Ω_1 and Ω_2 while if it is not, they will depend primarily on social support. This makes good intuitive sense and in fact has been adopted as a postulate by Festinger (1954) in his theory of social comparison processes.

In spite of the fact that we get these intergroup differences for all estimates, it is interesting and rather odd that, as will be seen by studying

TABLE 3

SEPARATE MAXIMUM LIKELIHOOD ESTIMATES OF θ_1, θ_2, AND DERIVED QUANTITIES

Type Trials on Which Estimates Based	Quantity	Group I	Group II	Group III
S_1	$\hat{\theta}_1$	.60	.38	.30
	$\hat{\theta}_2$	.20	.54	.62
	$\hat{\theta}_1/\hat{\theta}_2$	3.0	.71	.48
	$\hat{\theta}_1+\hat{\theta}_2$	.78	.93	.91
S_2	$\hat{\theta}_1$	.42	.36	.48
	$\hat{\theta}_2$	.27	.36	.45
	$\hat{\theta}_1/\hat{\theta}_2$	1.6	1.0	1.1
	$\hat{\theta}_1+\hat{\theta}_2$	.69	.72	.92

Table 3, the effects described above show up much more markedly on S_1 than on S_2 type trials. Since the actual stimulus and response labeling was randomized over subjects, and since the stimuli and responses were physically symmetrical, the relation between π_1 and π_2 must be responsible for this rather surprising phenomenon.

Reference to the tables under discussion will show that the over-all effectiveness of reinforcement, $\theta_1+\theta_2$, increased from the easy to the hard conditions. One might argue that the harder a discrimination, the more attentive the subject will be to all events surrounding him; thus, reinforcements are generally more effective when δ is small. Although it is true that via Equations 3 and 4 one can estimate the ratio θ_1/θ_2 directly from asymptotic data, we have preferred to make much more efficient

TABLE 4

JOINT MAXIMUM LIKELIHOOD ESTIMATES OF θ_1, θ_2, AND DERIVED QUANTITIES

Quantity	Group I	Group II	Group III
$\hat{\theta}_1$	.44	.39	.40
$\hat{\theta}_2$	.25	.42	.52
$\hat{\theta}_1/\hat{\theta}_2$	1.8	.92	.77
$\hat{\theta}_1+\hat{\theta}_2$	.69	.81	.92

estimates of this ratio by making strong use of the Markov properties of response subsequences. Specifically, the ratio $\hat{\theta}_1/\hat{\theta}_2$, where $\hat{\theta}_1$ and $\hat{\theta}_2$ are maximum likelihood estimates, yields the maximum likelihood estimate of θ_1/θ_2. Consequently, we wish to emphasize that the extent of agreement between observed asymptotic behavior and the asymptotes *predicted* using the maximum likelihood estimates of θ_1 and θ_2 is an empirical question representing one sort of test of the model. The predicted asymptotic conditional response probabilities have been included in Table 5, which also presents for purposes of comparison the relevant relative frequencies over the last 50 and 100 trials for each group. Although the fit is not perfect, the observed and predicted rank orders across groups are in agreement for the predictions based on the joint maximum likelihood estimates.

Observed mean learning curves on which the predicted asymptotes (of the last line of Table 5) are indicated as dashed lines are plotted in Figure 2. In order to determine whether a significant amount of learning occurred, the conditional response proportions from the first and last 50 trials of each experiment were compared by two-tailed sign tests (Siegel, 1956); the significance levels α are shown next to the appropriate curves.

The sign test, which was used here because the distribution assumptions necessary for it are almost negligible, does not produce acceptable significance levels. Higher level ($\alpha < .01$) tests may be obtained if we are willing to make the rather strong assumption that the individual responses of a particular subject are independent Bernoulli events. However, this assumption is too grossly violated to allow one to place much confidence in results based upon it.

If we draw learning curves (Equations 3 and 4) using the maximum likelihood estimates of θ_1 and θ_2, we find that the curves rise to their asymptotes much more rapidly than do the observed curves of Figure 2. We wish to point out that this does not represent a fundamental problem for the present theory although it limits the usefulness of the specific form of the model which we have presented. The problem arises from our having assumed that only one stimulus "element" or "pattern" (see Estes, 1959) was associated with each stimulus event, S_i. By positing a larger number of such stimulus patterns and retaining the supposition that exactly one is sampled on each trial, the

TABLE 5

OBSERVED AND PREDICTED ASYMPTOTIC CONDITIONAL RESPONSE PROBABILITIES

	Group I		Group II		Group III	
	$P(A_1\|S_1)$	$P(A_1\|S_2)$	$P(A_1\|S_1)$	$P(A_1\|S_2)$	$P(A_1\|S_1)$	$P(A_1\|S_2)$
Observed: Over last 100 trials	.92	.31	.83	.40	.79	.19
Observed: Over last 50 trials	.94	.28	.84	.39	.74	.18
Theoretical: Based on S_1 trial estimates	.92	.20	.83	.48	.79	.27
Theoretical: Based on S_2 trial estimates	.88	.31	.85	.40	.85	.19
Theoretical: Based on joint MLE	.89	.29	.84	.42	.83	.23

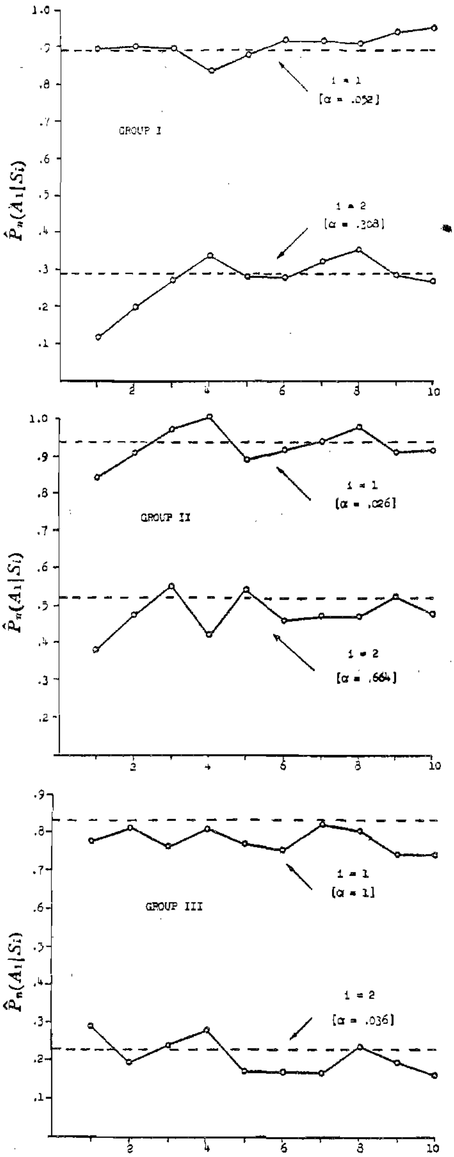

FIG. 2. Mean learning curves. (Blocks of 25 trials.)

theoretical learning rate may be decreased at the expense of considerably more tedious computations.

We may take a compromise approach, however. If we suppose there to be N stimulus elements associated with each stimulus situation, then although most formulae become considerably complicated, the learning curves 3 and 4 are changed only in the last term; let $\theta_1/\theta_2 = \epsilon$ then we may write

$$P_n(A_1|S_1) = \frac{\epsilon+\pi_1}{\epsilon+1} + \left(\delta - \frac{\epsilon+\pi_1}{\epsilon+1}\right) \times \left[1 - \gamma\frac{\theta_2}{N}(\epsilon+1)\right]^{n-1}$$

$$P_n(A_1|S_2) = \frac{1-\pi_2}{\epsilon+1} + \left(1 - \delta - \frac{1-\pi_2}{\epsilon+1}\right) \times \left[1 - (1-\gamma)\frac{\theta_2}{N}(\epsilon+1)\right]^{n-1}$$

We may estimate ϵ directly from asymptotic data and choose θ_2/N to make a least squares fit of the theoretical to the empirical learning curves.

CONCLUSIONS

In conclusion it is our opinion that considerable analytic advantage has derived from application of a quantitative theory of individual behavior in the present context. Through its use we have been able to make highly specific remarks about the extent of operation of certain events which are not directly observable but which may reasonably be postulated to exist, namely secondary self-reinforcement of responses well established in everyday experience. Also we have been able to compare the magnitude of the effects of these secondary reinforcements to that of the more directly controlled social support. As far as we can see, it would be considerably more difficult to obtain comparable information from experimentation not oriented about a quantitative theory.

We wish to emphasize again that the central ideas behind the present analysis are ones of considerable generality; similar techniques of analysis can be applied to many social psychological situations by a fairly straightforward extension of the ideas presented here.

REFERENCES

ANDERSON, T. W., & GOODMAN, L. A. Statistical inference about Markov chains. *Ann. math. Statist.*, 1957, **28**, 89–109.

ASCH, S. E. Studies of independence and submission to group pressure: 1. A minority of one against a unanimous majority. *Psychol. Monogr.*, 1956, **70**(9, Whole No. 416).

ATKINSON, R. C., & SUPPES, P. An analysis of two-person game situations in terms of statistical learning theory. *J. exp. Psychol.*, 1958, **55**, 369–378.

ATKINSON, R. C., & SUPPES, P. Application of a Markov model to two-person noncooperative games. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory.* Stanford: Stanford Univer. Press, 1959. Ch. 3.

BURKE, C. J. Applications of a linear model to two-person interactions. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory.* Stanford: Stanford Univer. Press, 1959. Ch. 9.

BUSH, R. R., & ESTES, W. K. (Eds.), *Studies in mathematical learning theory.* Stanford: Stanford Univer. Press, 1959.

ESTES, W. K. Component and pattern models with Markovian interpretations. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory.* Stanford: Stanford Univer. Press, 1959. Ch. 1.

FESTINGER, L. A theory of social comparison processes. *Hum. Relat.*, 1954, **7**, 117–140.

HAYS, D. G., & BUSH, R. R. A study of group action. *Amer. sociol. Rev.*, 1954, **19**, 693–701.

SHERIF, M. A study in some social factors in perception. *Arch. Psychol.*, 1935, No. 187.

SIEGEL, S. *Nonparametric statistics for the behavioral sciences.* New York: McGraw-Hill, 1956.

SUPPES, P., & ATKINSON, R. C. *Markov learning models for multiperson interactions.* Stanford: Stanford Univer. Press, 1960.

(Received October 5, 1959)

APPENDIX A

The model (Model I) originally proposed for the experimental situation studied in the body of this paper differed in its reinforcement mechanisms from the one (Model II) actually discussed there. Since rather extensive work was done on it and since we consider some of its failings to be instructive, we present here some detailed comments on Model I.

Let us call those trials on which the subscripts of E_i and S_j do not agree "conflict trials." Rather than introducing secondary reinforcers Ω_1 and Ω_2, we assume that the reinforcing event E_k is effective with probability θ_A on nonconflict trials and with probability $\theta_B < \theta_A$ on conflict trials. While we make no explicit psychological assumptions regarding the reasons for attenuating the learning parameter on conflict trials, the intuitive reasons are clear.

Markov process state identifications are exactly the same as in Model II, and we here present those results for Model I which correspond to those indicated for Model II in the body of the paper. On S_1 trials the transition matrix is

$$\begin{bmatrix} 1-(1-\pi_1)\theta_B & (1-\pi_1)\theta_B \\ \pi_1\theta_A & 1-\pi_1\theta_A \end{bmatrix} \quad [A1]$$

and on S_2 trials it is

$$\begin{bmatrix} 1-\pi_2\theta_A & \pi_2\theta_A \\ (1-\pi_2)\theta_B & 1-(1-\pi_2)\theta_B \end{bmatrix} \quad [A2]$$

The learning curves are

$$P_n(A_1|S_1)=\frac{\pi_1}{\pi_1+\frac{\theta_B}{\theta_A}(1-\pi_1)}+\left(\delta-\frac{\pi_1}{\pi_1+\frac{\theta_B}{\theta_A}(1-\pi_1)}\right)\times\{1-\gamma[\theta_B(1-\pi_1)+\theta_A\pi_1]\}^{n-1} \quad [A3]$$

$$P_n(A_1|S_2)=\frac{1-\pi_2}{\frac{\theta_A}{\theta_B}\pi_2+(1-\pi_2)}+\left((1-\delta)-\frac{1-\pi_2}{\frac{\theta_A}{\theta_B}\pi_2+(1-\pi_2)}\right)\times\{1-(1-\gamma)[\theta_B(1-\pi_2)+\theta_A\pi_2]\}^{n-1} \quad [A4]$$

And the graphs of the asymptotes as a function of θ_A/θ_B, which are presented in Figure A1, should be compared with the corresponding ones for Model II.

Several methods of learning parameter estimation were used for this model, and each led to the same rather disquieting result: for all methods the estimates of θ_A were greater than one. Before considering the reason for this anomaly, let us consider the various methods of estimation from which we obtained estimates of θ_A and θ_B.

Method 1. This method is precisely the same in principle as the first method presented in the body of the paper for θ_1 and θ_2. It is a true maximum likelihood estimate on each of the two response subsequences. As in Model II the method is for a fairly large number of counted transitions per subject equivalent to setting the observed transition matrices equal to the theoretical ones and solving. A comparison of the transition matrices for the two models shows that using this estimation procedure $\hat{\theta}_2 = \hat{\theta}_B$.

Method 2. The joint maximum likelihood estimate here, which is considerably simpler than in Model II, may be obtained explicitly by solving a second degree equation.

Method 3. The next procedure, which is rather complicated in the present experiments, we consider especially interesting, because it seems to yield many of the advantages of maximum likelihood estimation in situations where true maximum likelihood estimation is unfeasible.

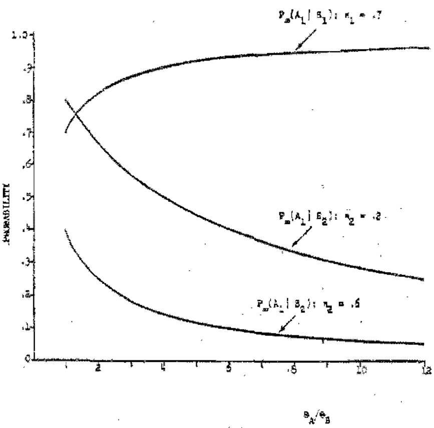

FIG. A1. Asymptotic behavior of $P(A_1|S_1)$ and $P(A_1|S_2)$ as a function of the ratio θ_A/θ_B in Model I.

Consider the full sequence of responses (S_1 and S_2 type trials in the same sequence now) for a particular subject. It is possible to compute for large n

$$\lambda_{ij} = P_\infty\ (A_j \text{ on Trial } n \mid A_i \text{ on Trial } n-1)$$

independently of the trial number, n. Thus, while it is true that the sequence of responses for a subject is a chain of infinite order, we may as a first approximation treat it as a stationary Markov chain with transition probabilities, λ_{ij}, and compute under this assumption the maximum likelihood estimate of the parameter pair (θ_A, θ_B); such a procedure is an example of *pseudo maximum likelihood estimation.* The pseudo likelihood function is

$$L(\theta_A, \theta_B) = \prod_{i,j} \lambda_{ij}^{n_{ij}} P_\infty^{n_1}(A_1) P_\infty^{n_2}(A_2)$$

where $P_\infty(A_i)$ is taken asymptotically, n_i is the number of occurrences of response A_i on the first trial of the sequence over which the pseudo maximum likelihood estimate is being taken, and n_{ij} is the observed number of transitions from response A_i to A_j over the sequence in question.

Method 4. One simple estimate of θ_A/θ_B was used in order to provide a check on the above procedures. Let $\hat{P}_\infty(A_1|S_1)$ and $\hat{P}_\infty(A_1|S_2)$ be the estimated (over the last 50 trials) asymptotic conditional probabilities of response. Then the value of θ_A/θ_B which minimizes

$$\left| \hat{P}_\infty(A_1|S_1) - \frac{\pi_1}{\pi_1 + \dfrac{\theta_B}{\theta_A}(1-\pi_1)} \right| + \left| \hat{P}_\infty(A_1|S_2) - \frac{1-\pi_2}{\dfrac{\theta_A}{\theta_B}\pi_2 + (1-\pi_2)} \right|$$

is a rather natural estimate of θ_A/θ_B although it uses very little of the information in the data.

The various estimates of θ_A and θ_B are given in Table A1.

TABLE A1

MAXIMUM LIKELIHOOD ESTIMATES OF θ_A, θ_B, AND θ_A/θ_B

Estimation Method	Quantity	Group		
		I	II	III
1	θ_A on S_1 trials	1.0	1.1	1.0
	θ_B on S_1 trials	.20	.54	.62
	θ_A/θ_B on S_1 trials	5.3	2.0	1.7
	θ_A on S_2 trials	2.4	2.2	1.1
	θ_B on S_2 trials	.27	.36	.45
	θ_A/θ_B on S_2 trials	8.9	6.0	2.4
2	θ_A	1.3	1.2	1.1
	θ_B	.24	.42	.52
	θ_A/θ_B	5.2	2.9	2.2
3	θ_A	2.2	1.7	.56
	θ_B	.22	.16	.07
	θ_A/θ_B	10.0	11.0	8.0
4	θ_A/θ_B	10.0	6.0	2.4

The occurrence of learning parameter estimates greater than one makes interpretation of Model I very tenuous. It will be instructive to consider the problem in more detail. It will be seen from Equations A1 and A2 on which our estimation procedures are primarily based, that θ_A and θ_B appear in no equations which would restrict their values to the unit interval. It is only in interpretation and in the derivation of transition matrices that we have made use of the learning parameters' roles as probabilities.

Although we may obtain bounded estimates of θ_A and θ_B by looking at detailed sequential properties of the data, we find that it is certain inadequacies of the learning model at precisely this level of analysis which are the cause of our excessively large learning parameter estimates.

To be more specific let $A_{1,n}$ represent response A_1 on trial n, and let $E_{1,n}$ be defined analogously. Then, it easily follows from our learning assumptions that on S_1 type trials

$$P(A_{1,n+1}|E_{2,n}, A_{2,n}) = 0 \qquad [A5]$$

$$P(A_{1,n+1}|E_{1,n}, A_{2,n}) = \theta_A \qquad [A6]$$

Although it would now seem quite natural to use Equation A6 as the basis for a bounded estimate of θ_A, we find that it is precisely because Equation A5 is not empirically valid that our previous estimation procedures have yielded estimates of θ_A greater than one. It can be easily shown that using estimation Method 1, for example, θ_A is indeed restrained to the unit interval if the observed relative frequency associated with the left side of Equation A5 is near zero, which in the present study it unfortunately is not.

In spite of its other difficulties the present model gives a sufficient description of important qualitative aspects of the data. First, *θ_A is consistently greater than θ_B*; second, *comparing Group I (easy, low symmetry) and Group II (hard, low symmetry), for which π_1 and π_2 are the same, Group I, which involved the easier discrimination, has a greater learning parameter ratio θ_A/θ_B.* Although the pseudo maximum likelihood estimates of the learning parameters weakly contradict this latter statement, it is our feeling that we may place more confidence in the true maximum likelihood estimates, since their optimality properties are well known.

We should expect under this model no change in the θ values as a function of π_1 and π_2 when the same stimuli are used. Unfortunately there is such a change;

indeed, we notice that when the subject makes the less frequently reinforced response and *is* reinforced, this reinforcement tends to have a greater probability of being effective than it would have otherwise had. This can be seen by comparing the θ estimates made on S_1 and S_2 type subsequences; the θ_B estimates in Group I, however, are contrary to this generalization.

APPENDIX B

We consider it to be of some general interest to present the method of obtaining the learning curves 3, 4, 7, and 8, since essentially the same technique may be used to derive learning curves in an N element pattern model (Estes, 1959) with a probability distribution over the elements.

Let

$$\sigma_{i,n,k} = \begin{Bmatrix}\text{Event of precisely } k \text{ samplings} \\ \text{of stimulus element } i \text{ in first} \\ n \text{ trials, } n = 1, 2, \cdots\end{Bmatrix}$$

Let

$$\alpha_{m+1}^{(i)} = P(A_{1,n} | S_{i,n} \cap \sigma_{i,n-1,m}) \quad \text{[B1]}$$

Now,

$$P(A_{1,n}|S_{i,n}) = \sum_{\nu=0}^{n-1} \frac{P(A_{1,n} \cap S_{i,n} \cap \sigma_{i,n-1,\nu})}{P(S_{i,n})} \quad \text{[B2]}$$

$$= \frac{1}{P(S_{i,n})} \sum_{\nu=0}^{n-1} P(A_{1,n} | S_{i,n} \cap \sigma_{i,n-1,\nu}) \times P^{\nu}(S_{i,n} | \sigma_{i,n-1,\nu}) P(\sigma_{i,n-1,\nu}) \quad \text{[B3]}$$

$$= \sum_{\nu=0}^{n-1} \alpha_{\nu+1}^{(i)} \binom{n-1}{\nu} P^{\nu}(S_i) \times (1 - P(S_i))^{n-1-\nu} \quad \text{[B4]}$$

In a general 2×2 Markov process if x_n is the probability of being in state one on trial n and

$$\bar{P} = \begin{pmatrix} 1-b & b \\ a & 1-a \end{pmatrix} \quad \text{[B5]}$$

then

$$x_n = (1-a-b)^n \left[x_0 - \frac{a}{a+b} \right] + \frac{a}{a+b} \quad \text{[B6]}$$

here, $n = 0, 1, \cdots$.

But for the transition matrices for S_1 and S_2 type trials, respectively, we have

$$S_1 \text{ trials:} \quad \begin{aligned} b &= (1 - \pi_1)\theta_2 \\ a &= \theta_1 + \pi_1\theta_2 \end{aligned} \quad \text{[B7]}$$

so

$$a + b = \theta_1 + \theta_2$$

$$S_2 \text{ trials:} \quad \begin{aligned} b &= \theta_1 + \pi_2\theta_2 \\ a &= (1 - \pi_2)\theta_2 \end{aligned} \quad \text{[B8]}$$

so

$$a + b = \theta_1 + \theta_2$$

Thus,

$$\alpha_{\nu+1}^{(i)} = \alpha_\infty^{(i)} + (\alpha_0^{(i)} - \alpha_\infty^{(i)}) \times [1 - \theta_1 - \theta_2]^{\nu} \quad \text{[B9]}$$

$$\nu = 0, 1, \cdots,$$

which substituting in Equation B4 gives

$$P(A_{1,n}|S_{i,n}) = \sum_{\nu=0}^{n-1} [\alpha_\infty^{(i)} + (\alpha_0^{(i)} - \alpha_\infty^{(i)}) \times (1 - \theta_1 - \theta_2)^{\nu}] \binom{n-1}{\nu} \times P^{\nu}(S_i)(1 - P(S_i))^{n-1-\nu} \quad \text{[B10]}$$

$$= \alpha_\infty^{(i)} \sum_{\nu=0}^{n-1} \binom{n-1}{\nu} P^{\nu}(S_i) \times [1 - P(S_i)]^{n-1-\nu} + (\alpha_0^{(i)} - \alpha_\infty^{(i)}) \sum_{\nu=0}^{n-1} \binom{n-1}{\nu} \times [(1 - \theta_1 - \theta_2)P(S_i)]^{\nu} \times [1 - P(S_i)]^{n-1-\nu}$$

$$= \alpha_\infty^{(i)} + (\alpha_0^{(i)} - \alpha_\infty^{(i)}) \times [(1 - \theta_1 - \theta_2)P(S_i) + 1 - P(S_i)]^{n-1}$$

$$= \alpha_\infty^{(i)} + (\alpha_0^{(i)} - \alpha_\infty^{(i)}) \times [1 - P(S_i)(\theta_1 + \theta_2)]^{n-1} \quad \text{[B11]}$$

This becomes for S_1 trials, letting

$$\alpha_0^{(1)} = \delta$$

$$P(A_{1,n} | S_{1,n}) = \frac{\theta_1 + \pi_1\theta_2}{\theta_1 + \theta_2} + \left[\delta - \frac{\theta_1 + \pi_1\theta_2}{\theta_1 + \theta_2}\right] \times [1 - \gamma(\theta_1 + \theta_2)]^{n-1} \quad [B12]$$

and for S_2 trials becomes, letting

$$\alpha_0^{(2)} = 1 - \delta$$

$$P(A_{1,n} | S_{2,n}) = \frac{(1-\pi_2)\theta_2}{\theta_1+\theta_2} + \left[(1-\delta - \frac{(1-\pi_2)\theta_2}{\theta_1+\theta_2}\right] \times [1-(1-\gamma)(\theta_1+\theta_2)]^{n-1} \quad [B13]$$