

FACULTY WORKING PAPERS

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

March 21, 1972

Statistical Learning Theory and Consumer Learning
An Experimental Investigation

Jagdish N. Sheth, University of Illinois
M. Venkatesan, University of Massachusetts
Eugene Kaczka, University of Massachusetts

#43

Jagdish N. Sheth is Professor, University of Illinois, M. Venkatesan and E. Kaczka are Associate Professors at the University of Massachusetts (Amherst). Financial support, in part, was received from the Research Council, University of Massachusetts, for which the authors are grateful.

Several researchers in consumer behavior have attempted to adapt and validate Bush-Mosteller (1) models of learning theories as explanations for the development of consumers' brand loyalty [see Sheth (9, 10) for review]. Some researchers [Kuehn (6), Sheth (10), Tucker (13)] have been able to support the position that systematic repetitive buying behavior can be described and explained in terms of statistical learning models. Others, however, have found this not to be true [Frank (3), Montgomery (8)].

In this paper the Bush-Mosteller statistical learning models are reviewed, the research literature in brand loyalty and learning theory is surveyed briefly and finally a study which was designed to test the role of statistical learning theory in consumer behavior is presented.

Description of Statistical Learning Models

Although there are several theories of learning [Hilgard (4)], such as classical conditioning vs. respondent conditioning, reinforcement vs. contiguity and the like, Bush and Mosteller (1) provided a single mathematical expression to measure the learning of systematic behavior. Essentially, it is a linear learning operator model in which it is assumed that following a response to an alternative, some event occurs (e.g., reinforcement or stimulus change) which has an effect on the probability of response to the same alternative next time the occasion arises.

This can be expressed as:

$$P_{t+1} = a_i + \alpha_i P_t \quad (1)$$

where p_{t+1} is the revised probability due to the consequences of the event, and it is summarized as a linear function of the probability of responding. This event and consequential effect may be positive in that it enhances the probability of responding or it may be negative in that it diminishes that probability. This is easily achieved if we define

$$\alpha_i = 1 - a_i - b_i$$

Then equation (1) can be rewritten as

$$p_{t+1} = p_t + a_i (1 - p_t) - b_i p_t \quad (2)$$

Equation (2) now is stated in a manner that if the event has positive effect, it is proportional to the largest possible gain in probability (namely $1 - p_t$ because it cannot exceed unity). On the other hand, if the event has a negative effect, it is proportional to the largest possible loss in probability (namely $-p_t$ because it cannot go below zero).

If there was complete learning after one event, the coefficients a_i and b_i would be unity. Hence if the event had a positive effect, the initial probability of responding would become unity, and if it had negative effect, it would be reduced to zero. However, most of the learning appears to be gradual over several trials and also it seldom is complete; it usually fluctuates between an upper limit or a lower limit. This concept of limit can be easily brought in if we define:

$$a_i = (1 - \alpha_i) \lambda_i$$

where λ_i is the limit. If we substitute this expression in equation (1), we get

$$p_{t+1} = \alpha_i p_t + (1 - \alpha_i) \lambda_i \quad (3)$$

If the probability of response p_t is equal to λ_i , we see that there is no further gain. If p_t is less than λ_i , then there is a proportional gain, and if p_t is greater than λ_i , there is a proportional decrease in the probability of responding next time. In fact, we can see that α_i and $1 - \alpha_i$ sum to unity and they are weights for p_t and λ_i .

Finally, we can with the use of equation (3), represent a learning or growth curve over several consecutive trials:

$$\begin{aligned} p_{t+1} &= \alpha_i p_t + (1 - \alpha_i) \lambda_i \\ p_{t+2} &= \alpha_i p_{t+1} + (1 - \alpha_i) \lambda_i \\ &= \alpha_i [\alpha_i p_t + 1 - \alpha_i \lambda_i] + (1 - \alpha_i) \lambda_i \\ &= \alpha_i^2 p_t + (1 - \alpha_i^2) \lambda_i \end{aligned}$$

and

$$p_{t+n} = \alpha_i^n p_t + (1 - \alpha_i^n) \lambda_i \quad (4)$$

Hence probability of responding after n trials is now a weighted average of initial probability (p_t) and the limit (λ_i). However, since α_i ranges between zero and unity, the greater the sequence of consecutive trials, the smaller it becomes such that it tends to become zero. And thus, the probability of responding after learning reaches the limit λ_i .

Bush and Mosteller (1) proposed three specific types of statistical learning models which encompass all varieties of learning situations. The first type is referred to as experimenter-controlled situation in which the consequence of events (reward and punishment or stimulus configuration change) following a choice among alternatives is non-contingent upon the

specific response that an individual chooses to make. Instead, the consequences following from occurrence of specific events are fixed and determined by the experimenter. Most of the experiments with rats in T-maze or Y-maze in which conditions and proportions of rewards and punishments are controlled by the experimenter are representative of this situation. The statistical model in this situation predicts that in the long run (equilibrium state) the proportion of responses to various alternatives equals the proportion of times those alternatives are reinforced. Hence if in a two-choice situation, alternative A is rewarded 65 percent of the time, the level of systematic behavior toward them is predicted to be 0.65 and 0.35 respectively.

The second type of learning is called subject-controlled situation in which events following responses to specific alternatives are directly a function of the specific responses. Hence consequences are contingent upon the choice among a set of alternatives; each alternative is presumed to have entailing consequences of various magnitudes. A good example of a subject-controlled learning situation is the Solomon and Wynne (11) experiment in which dogs learned to jump a barrier to avoid an intensive electric shock; the latter is completely predicated upon jumping by the dog within a prespecified time. Once again, the level of learning, to respond a specific alternative, in the long run is determined by the number of times the consequences of a response are found to be reinforcing.

The third type of learning is called experimenter-subject controlled situation. As the name implies, the occurrence of an event with entailing consequences is partly contingent upon the choice of alternative by the subject and partly by the experimenter. The most common are the learning

experiments with rats using T-maze or Y-maze in which the rat chooses the left or the right turn, and the experimenter controls the rate of reinforcement at the end of each turn.

Learning Theories and Brand Loyalty

The work by Kuehn (6) was the first effort to attempt to describe consumer brand choice with a generalized form of the Bush-Mosteller stochastic learning model. "Factorial analysis" was performed on panel data to determine the effect of the four preceding purchases of frozen orange juice on the probability of selecting a particular brand of the fifth purchase.

In another study, Frank (3) analyzed consumer panel data on coffee purchases and suggested a model which involved constant response (purchase) probabilities which are different for different consumers. He then used simulation to demonstrate when aggregating such heterogeneous consumers an effect may be obtained which appears as if learning is occurring.

Carman (2) used consumer panel data on dentifrice purchases to test the linear learning model proposed by Kuehn and to test the hypothesis suggested by Frank's work that the learning effect within homogeneous groups of consumers is negligible. His results indicated purchase probability behavior which was consistent with the generalized linear learning model. Further, an analysis based upon the division of the panel into "brand loyal" and "brand switcher" groups indicated that the learning effect cannot be completely explained by the aggregation of data. Consumer panel data however pose problems as a source of

data for testing the learning model, for one lacks control over the environment in which the purchase decisions are made.

The model developed by Montgomery (8) is an extension of some of Coleman's work in mathematical sociology. It is a binary choice model which allows the response probabilities of different consumers to be different, and to change through time. He tested his model against much of the same dentifrice purchasers panel data that Carman had used. This study demonstrated that the model provided a very good fit to the data and as such it appears to have some empirical viability.

Unfortunately, most of the research in applying statistical learning theory to consumers' development of brand loyalty seems to have suffered from at least two limitations in construct validation methods. The first limitation is related to the inappropriateness of the empirical reality of consumer behavior in which statistical learning theory has been applied. For example, it has been tested on standard (commercial) purchase panel data in which product classes and brands such as coffee or toothpaste are all very well known. In such a case, one would expect the consumers to have already learned brand preferences prior to the time period chosen for analysis, and therefore they would manifest steady habit behavior in the analysis time period. In addition, the reinforcement aspect inherent in statistical learning theory has been missing in empirical situations so that validation of learning construct is at best incomplete.

The second limitation is related to problems of data analysis. One of the basic issues is the number of alternatives involved in the learning situation. Instead of working deductively from the theory, most analyses have grouped alternatives that are not even mutually exclusive,

much less being exhaustive. For example, the total set of brands involved in the choice situation is not known from the panel data, and often the alternative of "not buying" is included among alternatives of choosing a set of brands once the consumer has decided to buy. Consequently, the role of statistical learning theory in consumer behavior has still remained untested.

The objective of this study is to validate, to the extent possible, the major statistical learning models under simulated conditions of consumers' choice behavior. This research effort was conducted in a laboratory setting. The experimenter-controlled and the subject-controlled models were tested using a consumer product, viz., razor blades. With each model two and three choice situations were presented to groups of involved and non-involved subjects. The time interval between selections was identical for all groups eliminating any confounding which might be attributed to differentials in usage rates.

METHOD

Subjects

The subjects were 209 male and female college students, all undergraduates from the University of Massachusetts, School of Business Administration. They served in the experiment during the duration as part of a requirement for an introductory course in Business Administration. This total was comprised of 168 male and 41 female students. Based on their responses to a preliminary questionnaire administered to all of them, this pool of subjects was separated into "users" and "nonusers" of razor blades. One-hundred-nineteen males and 24 females from the

"user" group were randomly assigned to the four experimental conditions (involved) and 49 males and 17 females from the "nonuser" group were assigned to the four "uninvolved" (control) conditions.

Design and Procedure

Four experimental conditions were created: Two experimenter-controlled situations (Groups I and II) and two subject-controlled situations (Groups III and IV). Two choice alternatives were provided for Groups I and III and three choices were provided to Groups II and IV. The task involved a choice among two or three brands of double-edged razor blades over a period of time.

The procedure required the subjects to come to the laboratory three times a week (once every Monday, Wednesday and Friday) and to indicate a choice among the brands of blades indicated for his group. In the first meeting with the subject, he was told that the experiment would last several weeks and he (she) is required to come every Monday, Wednesday and Friday to make a choice. If the subject agreed to participate for several weeks, then, based on the group to which he was assigned, he was told to indicate a choice among two or three brands of blades and was told to continue making a choice each time from among these (two or three) alternatives only. For those who were assigned to the experimenter-controlled situations, the subject was told that while he is required to make a choice on each visit, among the alternatives indicated, he will be given a razor blade (free) as indicated for him by the computer for that trial, irrespective of his choice. The subjects in the subject-controlled situations were told that on each visit the subject has to indicate a

choice on the choice sheet and he would be given a blade (free) if his choice matched the choice that is indicated for him (her) by the computer. In other words, the subject was told, that if his choice matched the choice indicated for him by the computer, he would get a free blade, otherwise no blade would be given him on that trial.

The reward schedules were determined for each block of ten trials in advance. For each subject there was a folder, in which there was a "choice sheet" to indicate the subject's choice and there was a sheet for the experimenter in which just before each trial, the free blade choice (computer choice) was indicated. Thus the research assistant was not in a position to know the "computer choice" earlier. Each subject was run individually. Two separate rooms and two research assistants were used to separate the experimenter-controlled situations from the subject-controlled situations.

On entering the laboratory, the subject was asked to indicate his choice for that time period on his sheet. Then the research assistant made sure that the alternative indicated was among the applicable set of alternatives for the group to which that subject was assigned and then looked into his folder to see the computer choice. In the experimenter-controlled situations, the subject received a blade free, in accordance with the computer choice. For the subject-controlled situations, the subject was told whether his choice matched or did not match to that made by the computer, and the subject was given a blade if there was a "match." The blades had been individually packaged with the name written on top of the small envelope and thus the participating subjects were unaware of all the brands that were involved in the experiment. The subjects had been

told that there were several studies that were in progress using different brands of blades and was cautioned not to compare his situation with that of others.

The choices involved and the reward schedule that was used are indicated in Table 1. The tasks for Groups V, VI, and VII and VIII were similar in every way (choices, etc.) with the exception that these subjects were told that since they were "dry" shavers, we wanted them to play the "game," and no mention was made of any free blade being given away. The reason given to them was that we were interested in seeing how well they would be able to guess the computer's choices.

At the end of 30 trials the experiment was terminated and subjects were debriefed as to the nature and purpose of the experiment. A few subjects who had missed two or three trials were allowed to complete them at their last time period. There is no reason to suspect that there was any more than natural interaction among the subjects during the duration of the experiment.

Insert Table 1 about here

RESULTS AND DISCUSSION

The results from all the four experiments are summarized in Table 2. The basic statistic under consideration is the proportion of subjects at each trial who chose the brand of blades that had the greatest reward schedule: That is, Personna brand of blade in the two-choice and three-choice experimenter-controlled situations and Wilkinson brand of blade in the two-choice and three-choice subject-controlled situations.

Insert Table 2 about here

In accordance with the statistical learning theory prediction, at the end of thirty trials the response rate should be equal to the asymptotic level of learning. The latter turns out to be 0.70 in all the four situations. It is obvious from Table 2 that both choice situations in experimenter-controlled conditions failed to reach the level of learning predicted by the model. In fact, in the three-choice situation the proportions are not far better than what one would expect by chance, and in the case of the two-choice situation, the proportions hardly reached the 0.50 level one would expect by chance if the choices were random.

The results for the two subject-controlled conditions provide a different picture. In the three-choice situation, the proportions are significantly different from chance proportions, although the asymptotic level of learning is not attained. In the two-choice situation, not only are the proportions systematically different from what we would expect by chance, but the rate of learning has reached or even surpassed the asymptotic level predicted by the model.

Statistically, it would be more appropriate to compare the observed proportions over 30 trials with what the model theoretically predicts. However, in order to obtain the theoretical learning curves, three parameters are needed: the initial probability of response to the alternative (p_0), the asymptotic level of learning (λ_1) and the rate of learning ($\alpha_1 = 1 - a_1 - b_1$). The first two parameters are given by the dictates of the model: If there is no prior learning and if there are no individual differences, the initial probability p_0 is equal to chance probability.

In a two-choice situation, this would be 0.50 and in a three-choice situation, it will be 0.34. Similarly, the asymptotic level of learning would equal the proportion of times a response is reinforced. In all the four experimental conditions, the value of λ_1 is 0.70. However, the parameter values of α_1 need to be estimated from the data.

Bush and Mosteller (1) provided a variety of estimation procedures primarily to permit assumptions related to the inequality of consequences following the reward as opposed to punishment events. However, not knowing whether matching the brand of blade that a subject had chosen (reward) is different from not matching the brand of blade (nonreward), we have assumed that their respective effects on the probability of choosing an alternative are about the same although inversely related. The rate of learning (α_1) in all the four experiments is accordingly estimated with the method suggested by Bush and Mosteller (1, p. 281).¹

With the use of estimated α_1 , the theoretical proportions for subject-controlled conditions were calculated. In the case of experimenter-controlled sequence experiments, it was clear from the data that observed values were consistently lower than theoretical (fitted) values. In fact, there was not even a single trial when the observed proportions were equal to or greater than the theoretical proportions. On the other hand, both the subject-controlled experiments approximate the theoretical proportions better, as shown in Figures 1 and 2. Since the estimated values are high, it may be indicative of slow rate of learning ($\alpha_1 - 1 - a_1 - b_1$). Thus, the rate of learning is greater generally in the subject-controlled conditions, and in particular, for the two-choice situation.

¹See Appendix for calculations.

Insert Figures 1 and 2 about here

Two types of tests of goodness-of-fit were performed on the observed and theoretical proportions. The first is a runs test proposed by Swed and Eisenhart (12). At each trial, if the observed value is greater than the theoretical value, a plus (+) sign is given to that trial, and if it is less, a minus (-) sign is given. Then the number of consecutive pluses or minuses (runs) is calculated. This number is compared to what would be expected by chance alone. If the runs are too many or too few as compared to expected number of runs, it indicates that there are significant differences between the observed and theoretical proportions of the choices over 30 trials.

A normal deviate is computed using the following formula in cases when the trials are large in number. It is as follows:

$$Z = \frac{d - E(d)}{\sigma_d}$$

where d = number of runs of consecutive pluses or minuses in the data,

$$E(d) = \frac{2n_1n_2}{n_1+n_2} + 1 \quad \text{where } n_1 = \text{number of pluses}$$

$$n_2 = \text{number of minuses,}$$

$$\text{and } \sigma_d^2 = \frac{2n_1n_2(2n_1n_2 - n_1 - n_2)}{(n_1+n_2)^2(n_1+n_2-1)}$$

Only for the two-choice subject-controlled condition, the actual runs were significantly more than the expected number of runs (number of runs = 19, $Z = 2.15$, significant at .05 level). For all the other conditions, the runs turned out to be fewer than would be expected.

We can conclude from these tests that in most of the cases experimental data do not match the behavior predicted by learning models. However, the theoretical models were based on certain assumptions which may not be true in the real-life situations. For example, the models presumed that there is no prior learning or that there are no differences among subjects when participating in the experiments. Our examination of the data revealed that there were set preferences for certain brands of blades that were used in the experiments: Wilkinson had been found to be generally more preferred and used by the subjects prior to their participation in the experiments, and Personna was found to be less preferred and used.

These preferences clearly state that initial probability (p_0) is not likely to be equal to chance probability and hence our estimations of initial probabilities should be other than the equal chance probabilities that had been used in the calculations. Secondly, the reinforcement schedules are likely to be more or less effective depending upon prior preferences or prejudices toward the brands. Hence the asymptotic levels which were presumed to be equal to the levels of reinforcement schedules should be revised.

The initial probabilities were re-estimated from the data: the first five trials were examined in their proportions and the mean level of these proportions was chosen as the estimate of initial probability. The asymptotic levels (w_1) were reduced from 0.70 to 0.60 in both of the experimenter-controlled sequences because the alternative under consideration, namely Personna blade, was less preferred. On the other hand, the asymptotic levels were raised to 0.80 in both of the subject-controlled experiments because Wilkinson blade was more preferred by the subjects.

The new estimates of rates of learning (α_1) based on the new estimated values of initial probabilities and asymptotic levels of learning turned out to be not substantially different from the previous estimates indicating that the rates of learning are not affected by bringing in the prior experiences.

Comparisons between the experimental data and the new theoretical proportions revealed that the new estimates are considerably closer to the experimental data particularly in the initial stages of learning. However, the runs test over all the 30 trials did not show any improvement in the goodness-of-fit between experimental and theoretical values.

One of the important cognitive aspects relevant to consumer learning is "involvement." Krugman (7) has suggested that learning may take place even without involvement.

In order to examine the effects of non-involvement on the learning process, four control conditions had been created, Groups VI, VII, VIII, and IX. In Table 3 the proportions of choice of the most rewarded blade are given for the non-involved group. Comparison of the non-involved and involved groups reveal some very interesting similarities.

Insert Table 3 about here

First, in the case of experimenter-controlled conditions, the proportions are relatively very similar. In fact, the proportions for the non-involved group seem to match better than the involved group with the predicted proportions based on the theoretical models. In view of the fact that the experimenter-controlled conditions are more like game playing for both the groups, and hence their levels of involvement may in fact be the same as that of the control groups.

Secondly, in both of the subject-controlled conditions proportions are less for the non-involved groups than the involved group. This can be explained by two factors: (1) the subject-controlled conditions are more realistic and simulate consumer choice behavior, inasmuch as the consequences are directly a function of the choice. Hence the involved group would be expected to learn more rapidly and manifest greater systematic behavior; and (2) the involved group had prior preferences for the Wilkinson blade and the choices and reward schedule involved this brand of blades.

CONCLUSIONS

Based on our experimentation with a model based on statistical learning theories, it appears reasonable to conclude that even when these models are modified to make them realistic to consumer learning situations, they do not fully predict brand choice behavior. On the positive side, the experiment indicates that learning (systematic behavior) does take place, but the particular form of learning or model that would satisfactorily explain brand loyalty phenomenon is yet to be found. In examining the data, however, it appears that the subjects at first seem to manifest systematic behavior (as measured by the size of proportions) to a brand and then switch to the other alternatives and again come back to the first alternative. This cycling is occurring more than once in each of the experimental conditions. This may be indicative that learning may be fast enough for individuals in consumer learning situations as simple as this experimentation attempted to simulate, so that the subject may have

been switching possibly for exploratory purposes. Such post hoc explanations seem to support the cyclical phenomenon which Howard and Sheth (5) have called the "psychology of simplification and complication," need to be systematically investigated in the future.

REFERENCES

1. Robert R. Bush and Frederick Mosteller. Stochastic Models for Learning
New York: John Wiley & Sons, 1955.
2. James M. Carmen, "Brand Switching and Linear Models," Journal of Advertising Research, 6, (June, 1966), 23-31.
3. Ronald E. Frank, "Brand Choice as a Probability Process," Journal of Business, 35, (1962), 43-56.
4. E. R. Hilgard. Theories of Learning. New York: Appleton-Century-Crofts, 1956.
5. John A. Howard and Jagdish N. Sheth. The Theory of Buyer Behavior.
New York: John Wiley & Sons, 1969.
6. Alfred A. Kuehn, "Consumer Brand Choice as a Learning Process,"
Journal of Advertising Research, 2, (1962), 10-17.
7. Herbert E. Krugman, "The Impact of Television Advertising: Learning without Involvement," Public Opinion Quarterly, 29, (Fall, 1965), 349-356.
8. David B. Montgomery. "A Probability Diffusion Model of Dynamic Market Behavior." Working Paper No. 205-66, Sloan School of Management, MIT, 1966.
9. Jagdish N. Sheth, "A Review of Buyer Behavior," Management Science, 13, (1967), B718-B756.
10. Jagdish N. Sheth, "How Adults Learn Brand Preference," Journal of Advertising Research, 8, (1968), 25-38.
11. R. L. Solomon and L. C. Wynne, "Traumatic Avoidance Learning: The Principles of Anxiety Conservation and Partial Irreversibility," Psychological Review, 68, (1954), 353-385.

12. F. S. Swed and C. Eisenhart, "Tables for Testing Randomness of Grouping in a Sequence of Alternatives," Annals. of Math. Statistics, 14, (1943), 66-87.
13. W. T. Tucker, "The Development of Brand Loyalty," Journal of Marketing Research, 1, (1964), 32-35.

APPENDIX

Estimation of (α_1) rate of learning

$$\alpha = 1 - \frac{\pi_1 - V_{1,0}}{N\pi_1 - \bar{T}}$$

where α = rate of learning parameter,

π_1 = asymptotic level of learning,

$V_{1,0}$ = proportion of responses to the alternative at the initial trial (p_0),

N = number of trials, and

$\bar{T} = \frac{1}{k} \sum_{i=1}^k T_i$ = average number of responses to the alternative over all trials.

With the a priori knowledge of π_1 and $V_{1,0} = p_0$ for all the four experimental conditions, it is easy to determine α_1 for various types of learning. The estimates are calculated below:

p_0 = average of first five trials

$\pi_1 = 0.60$ in experimenter-controlled situations

$\pi_1 = 0.80$ in subject-controlled situations

1. Experimenter-controlled situation, two-choice:

$$1 - \frac{\pi_1 - V_{1,0}}{N\pi_1 - \bar{T}} = 1 - \frac{.60 - .32}{30(.6) - 11.2} = 1 - \frac{.28}{6.8} = .959$$

2. Experimenter-controlled situation, three-choice:

$$1 - \frac{\pi_1 - V_{1,0}}{N\pi_1 - \bar{T}} = 1 - \frac{.60 - .18}{30(.6) - 9.2} = 1 - \frac{.42}{8.8} = .952$$

3. Subject-controlled situation, two-choice:

$$1 - \frac{\pi_i - V_{1,0}}{N\pi_i - \bar{T}} = 1 - \frac{.80 - .64}{30(.8) - 20.4} = 1 - \frac{.16}{3.6} = .956$$

4. Subject-controlled situation, three-choice:

$$1 - \frac{\pi_i - V_{1,0}}{N\pi_i - \bar{T}} = 1 - \frac{.80 - .41}{30(.8) - 14.2} = 1 - \frac{.39}{9.8} = .960$$

Table 1
EXPERIMENTAL CONDITIONS

| N | Group* | Choices and Reward Schedule |
|----|----------------------------|--|
| 32 | I Experimenter-controlled | Two-choice situation Personna 70%, Gillette 30% |
| 22 | II Experimenter-controlled | Three-choice situation Personna 70%, Wilkinson 20%, Gillette 10% |
| 31 | III Subject-controlled | Two-choice situation Wilkinson 70%, Personna 30% |
| 30 | IV Subject-controlled | Three-choice situation Wilkinson 70%, Personna 20%, Gillette 10% |

*Groups V (n = 18), VI (n = 15), VII (n = 17), and VIII (n = 16) are matching "uninvolved" (control) groups for these four experimental groups.

Table 2

PROPORTIONS OF CHOICE OF MOST REWARDED ALTERNATIVE

| Trial | Exp-Controlled Two-Choice | Exp-Controlled Three-Choice | Sub-Controlled Two-Choice | Sub-Controlled Three-Choice |
|-------|------------------------------|--------------------------------|------------------------------|--------------------------------|
| 1 | .32 | .06 | .50 | .27 |
| 2 | .18 | .13 | .64 | .37 |
| 3 | .32 | .19 | .68 | .40 |
| 4 | .32 | .22 | .61 | .43 |
| 5 | .27 | .31 | .75 | .57 |
| 6 | .32 | .19 | .32 | .23 |
| 7 | .50 | .19 | .75 | .43 |
| 8 | .32 | .38 | .46 | .47 |
| 9 | .18 | .34 | .61 | .47 |
| 10 | .27 | .31 | .64 | .50 |
| 11 | .32 | .31 | .75 | .40 |
| 12 | .50 | .28 | .68 | .37 |
| 13 | .41 | .38 | .82 | .40 |
| 14 | .32 | .38 | .82 | .37 |
| 15 | .23 | .38 | .71 | .47 |
| 16 | .50 | .44 | .61 | .47 |
| 17 | .41 | .22 | .86 | .57 |
| 18 | .50 | .44 | .39 | .60 |
| 19 | .36 | .28 | .75 | .40 |
| 20 | .46 | .38 | .61 | .57 |
| 21 | .50 | .38 | .75 | .53 |
| 22 | .36 | .34 | .75 | .53 |
| 23 | .46 | .41 | .79 | .57 |
| 24 | .27 | .28 | .68 | .60 |
| 25 | .41 | .34 | .71 | .47 |
| 26 | .41 | .34 | .82 | .57 |
| 27 | .40 | .34 | .68 | .47 |
| 28 | .46 | .28 | .75 | .57 |
| 29 | .50 | .38 | .79 | .50 |
| 30 | .36 | .31 | .71 | .63 |

Table 3

COMPARISON OF EXPERIMENTAL AND CONTROL GROUPS

| | Exp.-Controlled Two-Choice | | Exp.-Controlled Three-Choice | | Subject-Controlled Two-Choice | | Subject-Controlled Three-Choice | |
|----|-------------------------------|-------------|---------------------------------|-------------|----------------------------------|-------------|------------------------------------|-------------|
| | Involved | Noninvolved | Involved | Noninvolved | Involved | Noninvolved | Involved | Noninvolved |
| 1 | .32 | .47 | .06 | .24 | .50 | .35 | .27 | .31 |
| 2 | .18 | .47 | .13 | .29 | .64 | .53 | .37 | .31 |
| 3 | .32 | .53 | .19 | .24 | .68 | .35 | .40 | .38 |
| 4 | .32 | .67 | .22 | .24 | .61 | .24 | .43 | .38 |
| 5 | .27 | .40 | .31 | .41 | .75 | .35 | .57 | .38 |
| 6 | .32 | .40 | .19 | .41 | .32 | .29 | .23 | .19 |
| 7 | .50 | .60 | .19 | .35 | .75 | .47 | .43 | .44 |
| 8 | .32 | .53 | .38 | .23 | .46 | .24 | .47 | .31 |
| 9 | .18 | .33 | .34 | .29 | .61 | .35 | .47 | .38 |
| 10 | .27 | .40 | .31 | .35 | .64 | .47 | .50 | .31 |
| 11 | .32 | .33 | .31 | .47 | .75 | .59 | .40 | .31 |
| 12 | .50 | .53 | .28 | .47 | .66 | .41 | .37 | .56 |
| 13 | .41 | .47 | .38 | .53 | .82 | .59 | .40 | .25 |
| 14 | .32 | .53 | .38 | .47 | .82 | .53 | .37 | .31 |
| 15 | .23 | .53 | .38 | .47 | .71 | .53 | .47 | .44 |
| 16 | .50 | .40 | .44 | .59 | .61 | .59 | .47 | .38 |
| 17 | .41 | .33 | .22 | .41 | .86 | .59 | .57 | .38 |
| 18 | .50 | .47 | .44 | .35 | .39 | .41 | .60 | .43 |
| 19 | .36 | .47 | .28 | .41 | .75 | .53 | .40 | .38 |
| 20 | .46 | .53 | .38 | .41 | .61 | .65 | .57 | .44 |
| 21 | .50 | .53 | .38 | .35 | .75 | .53 | .53 | .44 |
| 22 | .36 | .33 | .34 | .47 | .75 | .35 | .53 | .38 |
| 23 | .46 | .40 | .41 | .65 | .79 | .47 | .57 | .50 |
| 24 | .27 | .47 | .28 | .29 | .68 | .53 | .60 | .38 |
| 25 | .41 | .40 | .34 | .24 | .71 | .59 | .47 | .44 |
| 26 | .41 | .40 | .34 | .41 | .82 | .41 | .57 | .44 |
| 27 | .46 | .53 | .34 | .35 | .68 | .53 | .47 | .25 |
| 28 | .46 | .53 | .28 | .35 | .75 | .53 | .57 | .44 |
| 29 | .50 | .53 | .38 | .65 | .79 | .47 | .50 | .38 |
| 30 | .36 | .53 | .31 | .47 | .71 | .47 | .63 | .56 |

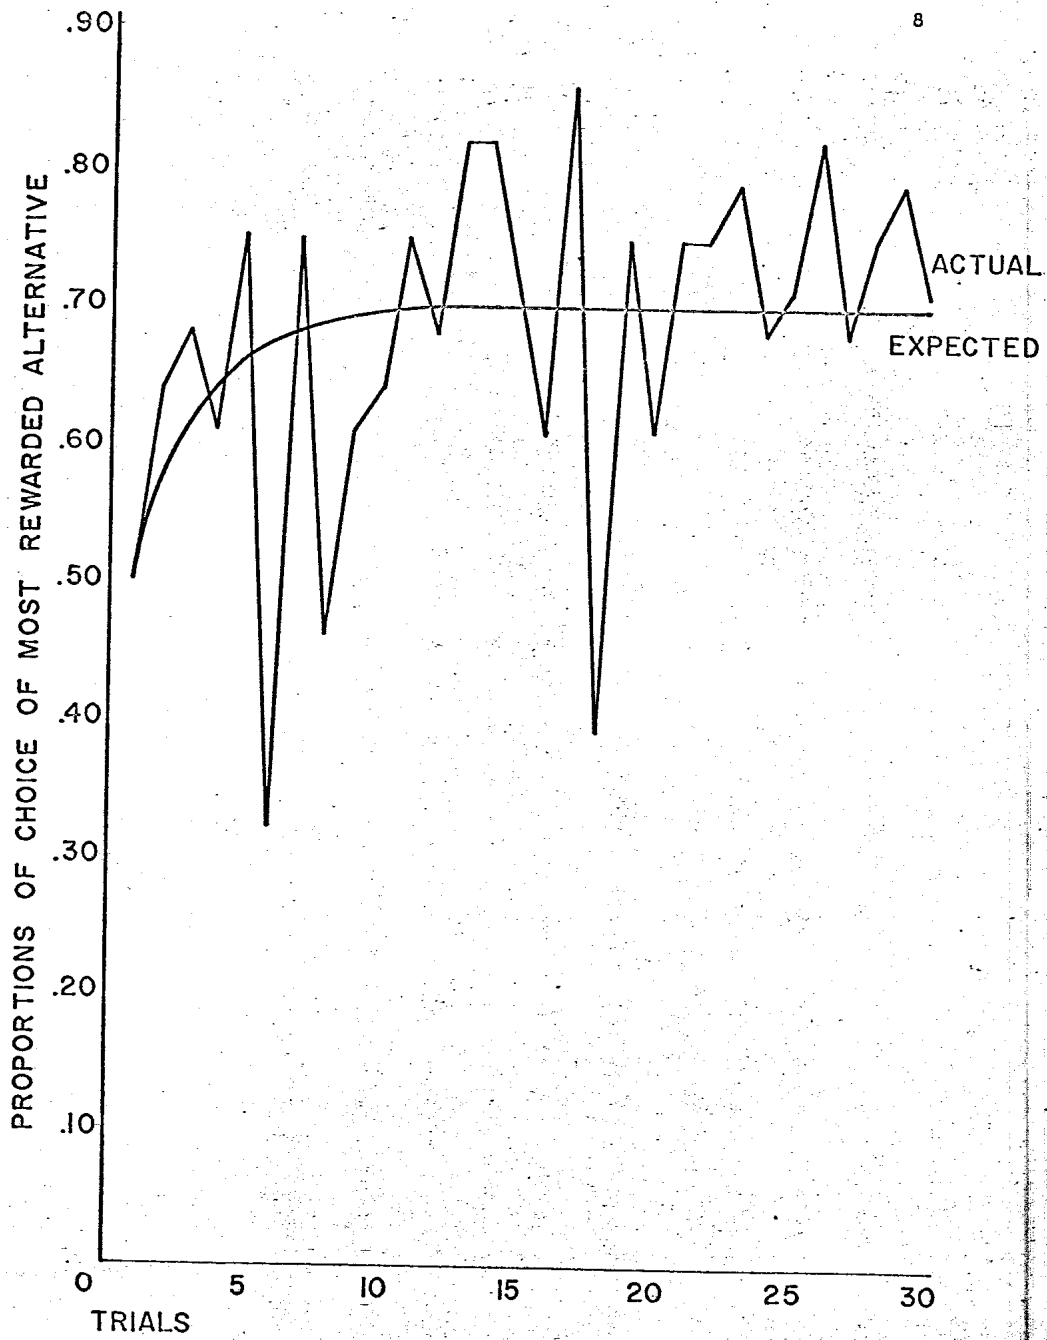


FIGURE 1. Subject-Controlled Condition--Two Choice Situation