In H. O. A. Wold (Scientific Organizer), Model Building in the Human Sciences, Entretiens de Monaco en Sciences Humaines, Session 1964. Monaco: Union Europeenne D'Editions. Pp. 39-49.

APPLICATIONS OF MATHEMATICAL MODELS OF LEARNING IN EDUCATION

by

Patrick SUPPES

I. — INTRODUCTION

Since 1950 there has been a vigorous and recently quite extensive development of mathematical models of learning. This is not the first time that mathematics or mathematical models have made an appearance in psychology, but there is little question that the developments of the past decade and a half are more extensive than those of any other previous period in psychology. One of the most encouraging signs for the scientific viability of this development has been the constant insistence that any new model be tested in a detailed fashion against appropriate experimental data. Probably one of the more significant effects of the recent developments has been the effect on the standards of data analysis. The point is that when a precisely formulated quantitative model is available for testing, inevitably a wealth of detailed questions arise that would not be raised in a framework of sheer empiricism.

Let me give just one example of this impact on data analysis. Probabilistic schedules of reinforcement or reward have been intensively studied with both human and animal subjects. In the simplest experimental situation with humans, the subject is asked on each trial to predict which one of two lights will flash. He makes this prediction by pressing one of two keys placed directly underneath a light. The subject is told that his objective is to make as many correct predictions as he can. In the typical experiment of this sort he is not told much about the schedule of reinforcement. A simple linear response-reinforcement model that can be used to analyze data is the following. It is postulated that on each trial the probability of a response is increased by a linear function when that response, that is, that prediction, is rewarded or proved correct. A corresponding linear decrement takes place in the probability of the other response. In the simplest formulation of the theory, the probability of an increase or decrement in probability according to whether or not a response is reinforced is independent of the response actually made. Let us designate the two responses A_1 and A_2, with the corresponding designation E_1 and E_2 for the two reinforcing events, in this case, the two flashing lights. Let p_n be the probability of an A_1 response

on trial n, and thus $1 - p_n$ is the probability of an A_2 response on trial n. When an E_1 reinforcement occurs on trial n

$$(1) \qquad p_{n+1} = (1 - \theta)\, p_n + \theta.$$

If an E_2 reinforcement event occurs on trial n

$$(2) \qquad p_{n+1} = (1 - \theta)\, p_n.$$

If we apply the simplest probabilistic reinforcement schedule, that is, one in which the probability of an E_1 event on each trial is π, independent of the response made or the preceding pattern of events, then it is easy to show that at asymptote, that is, as n approaches ∞, $p_\infty = \pi$, which result is termed the probability matching law. The striking thing about the probability matching law predicted by this linear model is that it is not at all the law of behavior that would follow from a rationalistic analysis of the prediction situation in which the subject is placed. Obviously for this reinforcement schedule at asymptote a subject should always predict an E_1 reinforcement if $\pi > \frac{1}{2}$ in order to maximize the number of correct predictions he makes. However, comparison of the rationalistic analysis with the learning-theory prediction is not my concern here. The point I want to emphasize, rather, is that we can move from the prediction of the asymptote $p_\infty = \pi$ to a great many additional predictions particularly in terms of conditional probabilities. As might be expected, these additional predictions ordinarily involve the learning parameter θ as well as π, and thus a test of the predictions requires an estimation of θ from the data. As is easily seen from the two equations given above, when an E_1 reinforcement occurs on a preceding trial the probability of an A_1 response on the next trial goes up; a corresponding decrement occurs when an A_1 response is preceded by an E_2 reinforcement. Preceding occurrences of A_1 and A_2 responses themselves have similar effects. I use the notation $P(A_1/E_1A_2)$, for example, to mean the conditional probability of an A_1 response at asymptote given an E_1 reinforcement and an A_2 response on the immediately preceding trial. Six typical asymptotic predictions in terms of π and θ are the following (Estes and Suppes 1959):

$$P(A_1/E_1) = (1 - \theta)\,\pi + \theta$$

$$P(A_1/E_2) = (1 - \theta)\,\pi$$

$$P(A_1/E_1A_1) = \frac{1 - \theta}{2 - \theta}\,[2\,\pi(1 - \theta) + 0] + \theta$$

$$P(A_1/E_1A_2) = \frac{(1 - \theta)}{2 - \theta}\,[2\,\pi(1 - \theta) + \theta]$$

$$P(A_1/E_2A_1) = 1 - \frac{(1 - \theta)}{2 - \theta}\,[2(1 - \pi)(1 - \theta) + \theta]$$

$$P(A_1/E_2A_2) = (1 - \theta)\left[1 - \frac{2(1 - \pi)(1 - \theta) + \theta}{2 - \theta}\right].$$

The important point to note about the six predictions is that none of them simply repeats the asymptotic means prediction p_∞. In every case the probability of an A_1 response is sensitive to the immediately preceding events. To compare these predictions to one experiment, we may use the pseudomaximum-likelihood estimate of θ described in Suppes and Atkinson

(1960) and compare the resulting predictions with observed data. Table 1 shows the results for such a set of data reported in CHAPTER 10 of SUPPES and ATKINSON (1960). In this experiment $\pi = .6$ and the estimate of θ is .19. A fit to the linear model of the data.

TABLE 1.

Asymptotic predictions of the linear model in a noncontingent probability-learning experiment.

	Observed	Predicted
$P(A_1/E_1)$	.667	.676
$P(A_1/E_2)$	.489	.486
$P(A_1/E_1A_1)$	.715	.710
$P(A_1/E_1A_2)$	.602	.625
$P(A_1/E_2A_1)$	.535	.520
$P(A_1/E_2A_2)$	.413	.435

shown in Table 1 is reasonably satisfactory. On the other hand, I do not mean to suggest that the fit of the linear model is the best that may be obtained. Estes' pattern model of stimulus sampling (ESTES, 1959), also discussed in CHAPTER 10 of SUPPES and ATKINSON (1960), gives a better fit than the linear models to the same data, but it will not be necessary to consider the pattern model in order to make the point about data analysis I am presently concerned with. Table 1, simple as it is, gives a good sense of the kind of detailed analysis one expects to see in connection with experiments when any sort of quantitative model is available for the analysis of data.

In the area of probability-learning experiments one expects as a natural thing to find a detailed analysis of data. It is perhaps more true to say that once such an analysis is presented one has a feeling of incompleteness and uneasiness about the exact character of a probability matching experiment in which these details are not reported. A good example of a report which generates such uneasiness is a recent article by GRAF et al (1964.) These authors discuss various experiments with animals (rats, fishes, and pigeons for instance) in which probability matching is obtained for some experimental procedures and conditions, but not in others. Their gross remarks about these experiments are undoubtedly correct. What is missing is the kind of detailed analysis I have just been discussing. From the standpoint of what is now known about probability matching experiments, statements like the following are certainly unsatisfactory without detailed presentation of evidence. "For example, a correspondence found in the rat by HICKSON (1961) turned out on further analysis to be due to reinforcement-following—the rats displayed a strong tendency to choose on each trial the correct alternative of the immediately preceding trial." (p. 151) Moreover, nowhere in this article is there discussion of conditional probability predictions of the sort shown in Table 1.

It is one of the theses of this paper that the kind of impact that mathematical learning has had on a wide class of learning experiments is similar to the kind of impact mathematical models will have in education as adequate models of subject-matter learning are developed and tested.

The kind of data analysis described for the simple probability-learning experiment is the kind we may eventually hope to have for significant educational material, but there are a number of roadblocks of both a scientific and technological sort that must be overcome before any significant part of the curriculum is analyzed in this kind of detail. From a scientific standpoint it is probably fair to say that current mathematical models of learning

are as yet too simple in structure adequately to analyse the learning of any parts of the curriculum. It is not simply that the structure of the models is too simple, however; it is also the case that we probably do not yet have the appropriate set of scientific concepts to embody in a model in order to make the right sort of analysis.

Probably of an even more pressing character are the technological problems confronting any systematic behavior analysis of subject-matter learning. To illustrate the difficulties, let us compare the total number of observations in the probability-learning experiment cited above with typical subject-matter learning data. In the case of the probability-learning experiment there were 30 subjects each of whom made 240 responses and each response was followed by 240 reinforcements. Thus the observational data for the experiment consisted of 7 200 responses and 7 200 reinforcements. These 14 400 discrete observations are manageable from the standpoint of data collection and data analysis without facing any problems of instrumentation. In an experiment of a long-term sort we are currently running with gifted first graders, we are finding that the collection of data on their mathematics learning is scarcely really practical from a technical standpoint if the data are collected and analyzed by hand. Let me give you some typical figures. In the first 14 weeks of our work with this class the average six-to-seven-year-old child in the group worked approximately 4 000 problems. This means that in the short period of 14 weeks we accumulated data on 160 000 problems. Many of these problems have more than one response; in terms of responses we probably have something like 250 000 to 300 000 observations. I emphasize that this figure is just for the first 14 weeks. By the time these children have been in the experiment for 50 or 100 weeks the amount of accumulated data will far exceed what we are able to analyze unless we use, as we plan to do, fairly sophisticated computer techniques.

In spite of the formidable scientific and technological difficulties we face, it is my own feeling that we are now on the edge of a very rapid increase in the seriousness and depth of research in educational learning. The two examples I discuss in the remainder of this paper are a modest contribution to this development. The first example deals with the optimal arrangement of a list of vocabulary items in learning a second language, and in the corresponding arrangement problems that must be faced in curriculum organization with respect to many lists of a simple character. The second example illustrates the kind of detailed analysis of learning familiar from applications of mathematical models in psychology but not yet customary in pedagogical settings. This example can be viewed as a psychological experiment, but one directed toward a significant pedagogical problem and therefore one that lies halfway between psychology and education. In particular, the subject of the experiment is the analysis of the learning of simple systems of abstract mathematics and the learning of simple mathematical proofs in exceedingly simple mathematical systems by six- and seven-year-old children.

II. — OPTIMAL ORGANIZATION OF A LIST OF SIMPLE ITEMS

Let us suppose that we want to think of acquiring in the early phases of learning a second language, a vocabulary of 200 or 300 words, and let us also think of the learning of this vocabulary. We separate this problem from the problem of learning the grammatical structure of the language. Let us also suppose for the sake of simplicity at the moment that all the vocabulary items are of equal difficulty. The problem we want to consider is

how should we organize this list in order to learn it. Suppose we have time enough to look or listen to each vocabulary item ten times. Should we group the items in blocks of 30, 40 or 50, going through each block and then repeating it until we have had the ten exposures to each item and thereby hopefully have learned it, or would it be more efficient to choose a larger block size like 100 or a much smaller block size like 5 or 6.

Some very simple theoretical assumptions about learning can provide quite useful guidelines as answers to these questions. We may, in the customary way, think of each presentation of a vocabulary item as a trial On trials when an item is presented, we can assume that the mean rate of learning satisfies a simple linear equation of the sort described earlier. In this case, let p_n be the probability of a correct response and $q_n = 1 - p_n$ the probability of an incorrect response on trial n. Replacing $1 - \theta$ by α and writing the equation in terms of q_n rather than p_n we obtain the simple equation

$$q_{n+1} = \alpha q_n.$$

Making this assumption about the *mean* rate of learning has a good deal of empirical support; it is important to realize that it does not commit us to a detailed assumption about the process of learning by an individual subject but only to the mean rate. Expressed in terms of α and the initial probability of an error, q_1, we have

$$q_{n+1} = \alpha^n q_1.$$

We may make a similar sort of assumption about forgetting. We postulate that on all those trials on which the item is not presented the same kind of mean phenomena is occurring with respect to forgetting. Let now β be the parameter representing the rate of forgetting corresponding to the above equation. We may express the equation for forgetting by

$$p_{n+1} = \beta p_n.$$

Note that in the case of forgetting p_n plays the role that q_n plays for learning. In order to evaluate the proper sublist or block size it is necessary to state explicitly the criterion with respect to which optimization is sought. One standard condition is to fix the total number of trials, with a definite number of the terminal trials allocated as test items. Mean performance on these test items is then used as the criterion in terms of which performance is evaluated. The optimization problem is to decide how the remaining trials, apart from the test, should be allocated for learning purposes.

It is shown in SUPPES (1964) that it follows from our basic assumptions about learning and forgetting that the block size should be chosen as large as possible when learning is faster than forgetting, i.e., when $\alpha < \beta$. When the learning process is slower than forgetting, i.e., $\beta < \alpha$, the block size should be chosen as small as possible, which means that in this case a block size of one should be used.

Let E(q) be the expected error, summed across items, on the terminal test, and for simplicity of notation, let $q = q_1$, $\varepsilon = \alpha/\beta$, let N be the total number of trials on all items prior to the terminal test, m the number of items, n the number of trials on which each item appears (thus $N = mn$), and k the block size. Then it may be shown that

$$E(q) = \left\{ 1 - \frac{(1 - \beta^N)(1 - \alpha)}{(1 - \beta)\, m} \left(\frac{1 - \varepsilon^n x^n}{1 - \varepsilon x} \right) \left(\frac{1 - x}{1 - x^n} \right) \right\} q$$

where $x = \beta^{\varepsilon}$. If we look at the derivative with respect to x, the whole matter hinges on the sign of the derivative of

$$\left(\frac{1 - \varepsilon^n x^n}{1 - \varepsilon x}\right)\left(\frac{1 - x}{1 - x^n}\right),$$

because

$$\frac{(1 - \beta^N)(1 - \alpha)}{(1 - \beta)\,m}$$

is positive. However, it is easily shown that the derivative is positive if $\varepsilon > 1$, negative if $\varepsilon < 1$, and of course, zero if $\varepsilon = 1$, for the range of x we are considering, namely, $0 < x < 1$. Clearly these results imply those about block size described in the preceding paragraph.

Extensive study of the empirical correctness of these predictions is now under way in our laboratory at STANFORD, and I hope soon to be able to report in detail on them. Our investigations are mainly concerned with the learning of Russian as a second language by Americans.

III. — THE ERROR ANALYSIS OF THE LEARNING OF MATHEMATICAL PROOFS

To indicate how fairly simple theoretical ideas of learning may be applied in an analysis of what on the surface is a rather complicated learning phenomenon and in fact is a phenomenon that for many people represents the highest sort of educational learning, I turn now to a discussion of some recent work I have been doing on the analysis of how students learn mathematical proofs. Before beginning this discussion, I hasten to add that what I shall have to say does not constitute anything like a realistic analysis of what goes on at an advanced level when a mathematician is creating new mathematics or even when students at a university are learning classical mathematics. There is a sense in which the kind of analysis I am going to present does relate to these more complicated cases, but the detailed stimulus material to be presented and the theory to be applied to explain the behavioral data are fartoo simple to admit to any direct extension to mathematics beyond an elementary-school grade level.

The discussion of this example naturally falls into four parts. First, I present the quite simple mathematical theory itself that is involved. Second, I describe the actual experiment with six=year=old children and the way in which the theory itself was interpreted in the experimental study for subjects of this age. Third, I sketch the mathematical learning models applied to the analysis of the data. And finally, I briefly comment on some of the wider implications of this analysis.

The mathematical systems considered in the experiment are the simplest possible sort and therefore I have baptized them U_4 and U_5 with the intention that the letter U stands for " utterly trivial ". Let me briefly describe system U_4 and then state how we modify it to obtain U_5. The formulas or grammatical strings of the theory are finite strings of 1's and 0's. Put another way, the alphabet of the language of the theory consists of the two symbols 1 and 0 and the grammatical expressions are finite sequences made up from these two symbols. The single axiom of U_4 is the symbol 1. A theorem of the system

is defined as any string that may be obtained from the axiom by application of the following four rules of inference, or in this context, as they are sometimes called, rules of production :

$$\begin{array}{ll} R_1 & S \rightarrow S11 \\ R_2 & S \rightarrow S00 \\ R_3 & S1 \rightarrow S \\ R_4 & S0 \rightarrow S \end{array}$$

The arrow in the statement of the rules indicates that from the string S shown on the left of the arrow we may obtain the new string shown on the right. It is understood that the letter S always stands for non-empty strings. The shortest theorem in system U_4 that uses all four rules of inference is the following.

Theorem. 101

Proof :

(1)	1	Axiom
(2)	100	R_2
(3)	10	R_4
(4)	1011	R_1
(5)	101	R_3

It should be clear that any string of 1's and 0's is a theorem of system U_4 if and only if it begins with a 1.

System U_5 is obtained from U_4 by adding the additional rule of inference R_5 as follows :

$$R_5.\ S \rightarrow 1S.$$

To those who are appalled by the simplicity and triviality of these two systems, I emphasize that they were designed for work with six=year=old children, who have no experience with mathematical proofs whatsoever and very little experience with any sort of mathematics including arithmetic. Even the way the theory has been expressed here is too complicated for work with children of this age. The isomorphic representation we used was the following. Our young subjects entered the experimental room and were placed in front of a medium-sized table with a vertical panel near the leading edge. On this vertical panel occurred two rows of squares one directly above the other. Each of the squares could be lighted with red or green lights. In the experiment the top row served for expressing the theorems with 1 corresponding to red and 0 to green. Thus the theorem proved above would be expressed on the apparatus by a red square followed by a green square followed by a red square. A child selected the successive steps of his proof by pushing one of four buttons (in the case of U_4 and 5 in the case of U_5), the buttons corresponding to the rules of production. In the case of U_4, the left light in the bottom row was always lighted red corresponding to the single axiom. The initial situation was more variable in the case of U_5. In order to simplify the presentation of data, from this point on I shall mainly restrict myself to U_4.

I turn now to the mathematical models of learning used to make a behavioral analysis of the experimental data from the fifty six-year-old children used as subjects.

In a highly simplified form the basic ideas run as follows. The organism is presented with a sequence of trials, on each of which he makes a response that is one of several possible choices. In any particular experiment it is assumed that there is a set of stimuli from which the organism draws a sample at the beginning of each trial. It is also assumed that on each trial each stimulus is conditioned to at most one response. The probability of making a given response on any trial is postulated to be simply the proportion of sampled stimuli conditioned to that response, unless there are no conditioned stimuli in the sample, in which case, there is a "guessing" probability for each response. Learning takes place in the following way. At the end of a trial a reinforcing event occurs that identifies which one of the possible responses was correct. With some fixed probability the sampled stimuli become conditioned to this response, if they are not already, and the organism begins another trial in a new state of conditioning. The sequence of events occurring on a given trial may be illustrated by the following diagram.

State of Conditioning → Stimuli Sampled → Response Made → Reinforcement Occurs →

Reconditioning of Sampled Stimuli → New State of Conditioning

Note that the trial begins with a certain kind of conditioning and ends with a new state of conditioning. This change of conditioning is the kernel of the learning process.

To illustrate how a quantitative theory may be developed I shall consider what is perhaps the simplest possible version. We assume that there is exactly one stimulus element and that this element is sampled on every trial by the subject. In the scheme sketched above this reduces to triviality the sampling process. Secondly, once the single element is conditioned the response is known whenever the conditioning of that single element is known. A mathematical model that arises from this simple one-element assumption can be described in the following way. On every trial the subject—the learner in the experiment—is in one of two states. Either the single element is conditioned (state C) to the correct response or it is unconditioned (state U). We then formulate the mathematical background of the model in such a way that the subject's behavior forms a Markov process in these two states with the transition matrix indicated below.

	C	U
C	1	0
U	c	1 — c

The meaning of this matrix is exceedingly simple. When the subject is in the unconditioned state, on each trial there is a probability c that he will move to the conditioned state. Once he becomes conditioned he remains so, as indicated by the probability 1. Secondly, we postulate that the subject *guesses* the correct response with probability p when he is in the unconditioned state, and responds correctly with probability 1 when he is in the conditioned state. In spite of the simplicity of this model, it is interesting to note that the two states of conditioning, that is, being conditioned or being unconditioned, are not themselves directly observable. In this sense even this simplest formulation of the theory already has a nontrivial theoretical component.

Before turning to the specific context of mathematical proofs, there is one further remark

about applications of the one-element all-or-none conditioning model which is needed. This is the application to simple discrimination experiments Essentially a discrimination experiment is one in which the subject needs to learn to discriminate between two or more stimuli and make the appropriate response to each. In many cases, the problem of discrimination is one of discriminating between two stimuli that are highly similar in ther perceptual characteristics. When it is assumed that the similarity is negligible, or in more technical terms, that there are no common stimuli between the two stimulus displays, the one-element model may be applied to the discrimination experiment in a simple and direct fashion.

Subjects were always reinforced or corrected in terms of a proof of minimal length. For illustrative purposes here it will be simplest to ignore the possibility of more than one proof of minimal length and consider only an analysis that will alway yield exactly one proof of minimal length. The stimulus discrimination facing the subject on each trial is simply described. He must compare the last line of proof in front of him with the theorem to be proved. This comparison immediately leads to a classification of each last line of a proof into one of four categories : additional 1's need to be added to match the theorem (R1); additional 0's need to be added to match the theorem (R2); a 1 must be deleted to continue to match the theorem (R3); or, finally, a 0 must be deleted in order to match the theorem (R4). The rule that should be applied once the stimulus comparison has been made is indicated in parentheses. When the subject is completely conditioned to all four stimulus discrimination situations, he will make the response corresponding to applying that rule. For each of the four discriminations with respect to which he is not yet conditioned, there is a guessing probability p_i, $i = 1, 2, 3,$ or 4, that he will guess the correct rule and thus a probability $1 - p_i$ that he will guess incorrectly. Also, when the subject is unconditioned, for any one of the discrimination comparisons there is a probability c_i that he will become conditioned on the next trial.

Let us now see how well these theoretical ideas are borne out by the data on the sixteen theorems of U_4 used in the experiment. We may begin with the response data on Rule R2 (of system U_4) of the twenty-five subjects who began with U_4. The one-element model predicts a stationary or horizontal learning curve prior to last error. The data in quartiles before last error, summed across subjects are shown in Table 2. The x^2 test of stationarity with 2df has a value of 1.30, which is certainly not significant. (The criterion used in determining last error was the last error prior to five successively correct R2 responses.)

TABLE 2.

Number of successes and errors, and probability of a success for Rule R2, 25 subjects beginning with U_4.

Quartile	Successes	Errors	Prob. of a Success
1	44	70	.386
2	40	74	.351
3	47	67	.412
4	40	74	.351

The one-element model also predicts a binomial distribution of successes and errors prior to last error. The observed and predicted distributions of error frequencies in blocks of four responses are given in Table 2. In this case, the x^2 test of goodness of fit with 3df is highly significant, having a value of 18.5. The main contribution to the discrepancy

TABLE 3.

Distribution of error frequencies in blocks of four trials prior to last error, Rule R2, 25 subjects beginning with U4.

No. of errors	Observed	Predicted
0	5	2.36
1	20	15.46
2	33	37.92
3	27	41.35
4	29	16.91

between the observed and predicted distributions is to be found in the observed and predicted number of blocks consisting entirely of errors, the predicted number falling considerably below the observed value.

The one-element model predicts independence of the guessing probability from trial to trial prior to last error, and here too there is a significant discrepancy for Rule R2.

$$P(\text{Success/Success}) = .491$$
$$P(\text{Success/Error}) = .302$$

The X^2 test of independence with 1 df and a value of 15.9 is highly significant.

Rather better results are obtained for Rule R4. The data in quartiles before the last error are shown in Table 4. There is a tendency for the probability of a success to increase, but the x^2 test of stationarity has a value of 5.7 and is not significant at the .05 level. The

TABLE 4.

Number of successes and errors, and probability of a success for Rule R4, 25 subjects beginning with U4.

Quartile	Successes	Errors	Prob. of a Success
1	26	52	.333
2	25	53	.321
3	38	40	.487
4	30	48	.381

observed and predicted distributions of error frequencies in blocks of four responses are given in Table 5. In this case, the goodness of fit is extremely good and the x^2 test is not significant. The test of independence of successive responses to a discrimination

TABLE 5.

Distribution of error frequencies in blocks of four trials prior to last error, Rule R4, 25 subjects beginning with U4.

No. of errors	Observed	Predicted
0	2	1.59
1	11	10.50
2	28	25.88
3	21	28.37
4	16	11.66

calling for an R_4 response is also not significant ($x^2 = 1.7$), and is reflected in the conditional probabilities :

$$P \text{ Success/Success}) = .457$$
$$P \text{ (Success/Error)} = .384,$$

even though there is the same tendency, as in the case of Rule R_2, for the probability of a success to be higher after a success than after an error.

The data for Rules R_2 and R_4 with respect to the twenty-five subjects beginning with U_5 are very similar to those just cited for U_4 and will not be given here.

It is a fair summary of the situation to say that the one-element model is able to give a surprisingly good account of the experimental data, although systematic discrepancies are found. A more extensive review of its adequacy to account for mathematical concept formation in childrenis to be found in SUPPES & GINSBERG (1963), SUPPES (1965), (1966). The natural extension of the one-element model to a two-element or two-stage model is given in SUPPES & GINSBERG (1963), but it is also fair to say that we are still a long way from having mathematical models of learning that are sufficiently rich in their conceptual structure to account for many important aspects of mathematical concept formation. Until theories of learning are able to deal with the complexity of subject matter taught in the classroom we cannot be satisfied with our scientific results.

REFERENCES

1. ESTES, W. K. Component and pattern models with Markovian interpretations. In R. R. Bush and W. K. Estes (Eds.) *Studies in Mathematical Learning Theory.* Stanford : Stanford University Press, 1959. pp. 9-52.
2. ESTES, W. K. and SUPPES, P. Foundations of linear models. In R. R. Bush and W. K. Estes (Eds.) *Studies in Mathematical Learning Theory.* Stanford : Stanford University Press, 1959, pp. 137-179.
3. GRAF, V., BULLOCK, D. H. and BITTERMAN, M. E. Further experiments on probability-matching in the pigeon. *J. Exptl. Analysis Behavior*, 1964, **7**, 151-157.
4. HICKSON, R. H. Response probability in a two-choice learning situation with varying probability of reinforcement. *J. exp. Psychol.*, 1961, **62**, 138-144.
5. Problems of optimization in learning a list of simple items. In Maynard W. Shelly, II, and Glenn L. Bryan (Eds.) *Human Judgments and Optimality.* New York : Wiley, 1964, 116-124.
6. On the behavioral foundations of mathematical conceps. *Monographs of the Society for Research in Child Development*, 1995, **30**, 60-96.
7. Mathematical concept formation in children. *Amer. Psychologist*, 1966, **21**, 139-150,
8. SUPPES, P. and ATKINSON, R. C. *Markov Learning Models for Multiperson Interactions.* Stanford : Stanford University Press, 1960.
9. SUPPES, P, and GINSBERG, R. A fundamental property of all-or-none models, binomial distribution of responses prior to conditioning, with application to concept formation in children. *Psych. Rev.*, 1963, **70**, 139-161.