Towards a Behavioral Foundation of Mathematical Proofs

by

Patrick Suppes

TECHNICAL REPORT NO. 44

January 2, 1962

PSYCHOLOGY SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

Applied Mathematics and Statistics Laboratories

STANFORD UNIVERSITY

Stanford, California

Towards a Behavioral Foundation of Mathematical Proofs[1/]

Patrick Suppes

Stanford University

Stanford, California

1. Introduction.

The logical theory of mathematical proofs has been developed intensively and with great success in this century. I do not need to review for a colloquium audience in Warsaw the main outlines of this development. What is surprising is that so little has been written about the psychological theory of mathematical proofs. However interesting they may be as literary documents I am not willing to count as scientific psychology mathematicians' testimonials of how they made discoveries. It is not that this material is not interesting. It is just that it holds very little promise scientifically. The psychological phenomena which lie at the base of any genuinely new mathematical discovery are surely as complicated and intellectually involved as any in the whole range of human behavior. Introspective accounts of these phenomena are as difficult to work with as basic data as are the descriptions of nature lovers of sunsets and storms in developing the science of meteorology.

1/ This research was performed pursuant to a contract with the United States Office of Education, Department of Health, Education and Welfare. The paper was originally presented at a colloquium on methodology in Warsaw, Poland on September 22, 1961.

Perhaps to the disappointment of some of you I shall approach the problem of providing a psychological analysis of mathematical proofs by considering examples of an almost ridiculous simplicity. The analysis shall proceed on the assumption that it is possible for sufficiently simple contexts to analyze written and spoken speech acts within the general framework of behavioral psychology (what I shall mean by "behavioral psychology" will become clearer in the sequel). As an act of faith I would also express the conviction that a still further reduction of behavioral psychology to neurophysiology will ultimately be possible, but I am not hopeful that this reduction will occur in the near future, particularly with reference to complicated intellectual processes.

As W. K. Estes wisely pointed out in the original discussion of this paper in Warsaw, I also do not believe that a detailed explanation in behavioral terms of the genuine discoveries of mathematicians can be given in the framework I am describing here. The aims of the psychological theory I shall set forth are schematic in the same way that physical theories are schematic. There are only a few phenomena caught in the raw, so to speak, in nature that are subject to any exact explanation and prediction of behavior in terms of existing physical theories, and we all have a good rough estimate of the relative power of physical and psychological theories. On the other hand, I do believe that the kind of theory and analysis I shall be giving of the simplest elements of the learning of mathematical proofs do provide the right sort of framework for the analysis of the most complicated mathematical activity. A phrase I used in this last sentence also provides an important indication of what I wish to mean by a behavioral foundation of mathematical proofs. I do not, it should be clear, mean a behavioral foundation for the written or inscribed proofs themselves but rather for the act of learning to give such proofs on the part of students of mathematics. In this sense, a behavioral foundation emphasizes the learning or discovery of proofs.

I will return to these general comments after giving a brief sketch of the relevant psychological theory and an indication of how it may be applied to a simple mathematical system.

2. Brief Sketch of Stimulus Sampling Theory.

Stimulus sampling learning theory was first given a quantitative formulation in 1950 by W. K. Estes but its basic concepts were developed by a number of psychologists running back to the beginning of the century; particularly important have been the general contributions of such figures as Pavlov, Watson, Hull and Guthrie. The great merit of Estes is to have shown how these ideas may be cast in a quantitative formulation subject to genuine mathematical analysis. 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 which identifies that one of the possible responses which 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.[2]

To illustrate how a quantitative theory may be developed with these ideas we 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.

---

[2] For an explicit axiomatic formulation of these ideas, see Suppes and Atkinson (1960. p. 5); for a more complete discussion of the technical aspects of axiomatization, see Estes and Suppes (1959).

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). In the simplest applications we wish to consider, which are also the ones pertinent to our subsequent discussion of mathematical proofs, we also assume that of the two responses available exactly one is always correct. (The restriction to two responses is not essential. What is needed is that the correct response in a given situation is well defined. We may classify the other responses simply as incorrect.) 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 non-trivial theoretical component.

It will be instructive to consider two simple applications of this one-element model to experiments. I first consider a typical paired-associate experiment. The subject is shown a succession of nonsense syllables on printed cards; each nonsense syllable constitutes a stimulus, and on each trial he sees exactly one stimulus. What the subject must learn to do is to make an appropriate response when a stimulus is shown. Typically, he might be asked to respond with one of the numerals 1, 2, 3, or 4. Given a list of twenty nonsense syllables, the experimenter would arbitrarily assign five of them to each of the four numerals. On the first trial the subject is in the unconditioned state; if the experiment had been well designed, the probability p of guessing the correct answer should be .25. There is considerable evidence (Bower [1961], Estes, [1960]) to show that when the subject does make a firm association between a nonsense syllable and the correct response, he makes this association on an all-or-none basis and retains it throughout the experiment. The probability c of moving from the unconditioned state may be estimated from the experimental data.

The paired-associate experiment provides a paradigm of the stimulus-response conditioning connection, but it is not necessary for application of this model that the conditioning connection be conceived as holding between a particular stimulus display and a given response. We have also applied this model extensively to concept formation in children (Suppes and Ginsberg [1961], [1962]). Here I shall describe briefly an experiment concerned with the concept of identity of sets. The subjects were 48 children of first-grade age (6 or 7 years). On each trial the child's task was to indicate whether two sets were identical or not. There were a total of 56 trials on 28 of which the stimulus display showed identical sets and the remaining 28 non-identical sets. The subjects were instructed to press one of two buttons when the stimulus pairs presented were "the same" and the other button when they were "not the same". In this experiment no stimulus display on any trial was repeated for individual subjects. In this case the conditioning connection may be postulated as holding at different levels of abstraction. To begin with, we may assume there is a single concept of identity of sets and the child is learning to establish the appropriate connection between this concept and the two responses. Until this connection is established he is guessing the correct answer with probability p, and after it is made he makes the correct response with probability 1.

The next natural level of analysis is in terms of two concepts, one for the pairs of sets that are identical and one for the pairs of sets that are not identical. The one-element model may then be applied to the subsequences of trials on which identical sets are displayed and again to the complementary subsequences of trials on which non-identical sets are displayed. As the model has been formulated above and as applied to paired-associate data, it is assumed that the probabilities of conditioning are statistically independent for the two subsequences. For the analysis of any concept formation experiment in terms of more than one concept it is necessary directly to test this assumption of statistical independence against the data.

At a still more refined level, we may analyze the stimulus displays in terms of pairs of sets that are identical in the sense of ordered sets, pairs of sets that are identical but not in the sense of ordered sets, pairs of sets that are not identical but equipollent, and pairs of sets that are not equipollent. In this case, we consider four concepts rather than one or two.

It is not the purpose of this paper to evaluate the empirical adequacy of any of the alternative ways of analyzing an experiment on identity of sets. It is worth emphasizing, however, that there is no direct way of building from the individual stimulus displays to these various concepts by simple stimulus connections when no stimulus display is repeated for an individual subject. Admittedly, it is not fully satisfactory intellectually to analyze the learning of concepts simply in terms of a conditioning connection between the concept and the correct response. Theories which postulate more details about the learning process in concept formation are needed to go beyond the present analysis. There are two things to be said about such theories at the present time. In the first place, there seems little doubt but that a good first approximation to the data may be obtained in terms of theories formulated in terms of notions of hypothesis and strategies (compare the discussion in Suppes and Atkinson [1960, Sec. 1.7], Restle [1961]). On the other hand, at the moment, these theories have little more to offer than the simple one-element model in terms of detailed analysis of actual experimental data.

It is also not appropriate here to consider in detail statistical methods of analyzing the goodness of fit of the simple one-element model to experimental data. However in order to sketch briefly some results for a pilot experiment at the end of this paper, I recapitulate briefly the ideas set forth in Suppes and Ginsberg [1961]. There are just two basic ideas needed for essentially complete statistical analysis of the one-element all-or-none conditioning model. In the first place, the assumption that there is a constant guessing probability p that the subject responds correctly before he is conditioned implies that the sequence of responses prior to the last error of the subject is a sequence of Bernoulli trials with binomial distribution parameter p. Classical statistical tests for stationarity of the response probability, independence from trial to trial, and the actual binomial distribution of responses in blocks of fixed size may be applied to the data considered in terms of responses made prior to the last error. The conditioning parameter c on the other hand, enters only in terms of the distribution of the last error. Across a group of subjects this last error may be estimated from the essentially geometric distribution of the last error as derived from the theory formulated above. A standard goodness of fit test may then be performed to see if the assumption of a homogeneous conditioning parameter for all subjects in a given experiment is acceptable.

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. It is possible to think of a paired-associate experiment as such a discrimination experiment. On the other hand, because the discrimination itself is not difficult, it is not ordinarily described as such. However, in many cases, the problem of discrimination is one of discriminating between two stimuli that are highly similar in their 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 the same way that we have applied it above to a paired-associate experiment. For instance, suppose the subject is a rat in a T-maze. At the choice point of the T-maze a white card is placed on some trials, and on other trials, a black card. The animal must learn to turn left in order to receive food when the card is white, and to turn right when it is black. We may analyze such an experiment exactly in the manner indicated for paired-associate situations. The ideas of discrimination to be mentioned below will implicitly assume the simple context in which there is no problem of stimulus overlap.

3. An Utterly Trivial Mathematical System.

The simple mathematical system we shall analyze in terms of the behavioral ideas just discussed is concerned with production of finite strings of 1's and 0's. Any finite string of 1's and 0's is a well formed formula of the system. The single axiom is the single symbol 1. The four rules of inference are:

R1. $S \rightarrow S11$

R2. $S \rightarrow S00$

R3. $S1 \rightarrow S$

R4. $S0 \rightarrow S$,

where S is a non-empty string. A theorem of the system is, of course, either the axiom or a finite string that may be obtained from the axiom by a finite number of applications of the rules of inference. A general characterization of all theorems is immediate: any finite string is a theorem if and only if it begins with 1 . A typical theorem in the system is the following one, which I have chosen because it uses all four rules of inference.

Theorem. 101

(1)	1	Axiom
(2)	100	R2.
(3)	10	R4.
(4)	1011	R1.
(5)	101	R3.

It is apparent that a shorter proof of this theorem could not be given, and this is generally true of this system. A proof of minimal length of any theorem is easily found, but it is not the case that there is exactly one proof of minimal length. For instance, if we want to prove the theorem 111 we may apply rule R1 twice to obtain 1111 and then remove the last 1 by applying R3; or we may interchange the position of the application of R3 and prove the theorem by first using R1 and then R3 and then R1 again. Counting the introduction of the axiom as one line, the two proofs are both of length four in terms of number of steps. (I have not fully formalized the system here by giving a recursive definition of proof, etc. because it is completely obvious how these matters go for a system of this kind; I want to give only enough formal detail to make the mathematical system definite.)

For simple reference in the behavioral analysis to follow let us call this mathematical system, the system $\mathcal{U}$.

4. Behavioral Analysis of Proofs in the System $\mathcal{U}$.

Initially, it will be simplest to ignore the possibility of more than one proof of minimal length and consider only an analysis that will always 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 parenthesis. 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. On the assumption of statistical independence made earlier, we may then analyze separately the four subsequences of trials on which the four stimulus discrimination categories appear. It is to be emphasized again that the four guessing parameters p_i and the four conditioning parameters c_i are to be estimated from the experimental data.

The example given in the preceding section of two proofs of minimal length shows that the analysis just stated represents a slight oversimplification. For example, if exactly one 1 is needed and two 1's are at the end of the string standing as the last line of proof, it will be just as efficient first to apply R3 and then R1 as to apply R1 and then R3. It does not require serious modification of the behavioral analysis of the system $\mathcal{U}$ to take account of this fact. Reinforcement can be given randomly of either of the rules that are correct in terms of rendering a minimal proof and either one of these responses can be counted as correct when made. This leaves the analysis in terms of the two states of conditioning untouched. It does change the relation between the state of conditioning and the probability of a response. The weaker requirement for this discrimination is that the probability of making one of the two correct responses is 1. We could if we so desired analyze the system $\mathcal{U}$ in such a way that there were more than four discriminating situations in order to take account of the cases in which more than one response was correct. For an extensive experiment this would be desirable. In terms of the pilot study to which I would now like to turn, it is not necessary.

The pilot study was conducted with a group of first-grade children (ages 6 and 7) in an elementary school near Stanford University. Initially, we considered doing the experiment with fourth-grade children (ages 9 and 10), but preliminary testing with a few children of this age indicated that the experimental problem was far too easy for them; most of the fourth-grade children made no errors at all. There is a variety of evidence to indicate that the reasoning abilities of young children are far superior to their ability to put an argument in written form (see e.g., Hill, 1961). For this reason we attempted to avoid entirely a written context for the exhibiting of proofs in the system $\mathcal{U}$. Our procedure was the following.[3/]

We considered only theorems which are of length greater than one and less than seven. There are 62 members of this class of theorems. We ordered their proofs according to the following four criteria of simplicity.

1. If proof P_1 may be obtained from proof P_2 by interchanging 1's and 0's in all lines, then P_1 and P_2 are of equal simplicity.

2. If $m < n$ then a proof of length m is simpler than a proof of length n .

3/ I am much indebted to John M. Vickers for his contributions to the detailed design and actual execution of this pilot study. Mrs. Susan Matheson assisted Mr. Vickers in running the experiment.

3. If P_1 and P_2 are minimal proofs of the same length, P_1 is one of several alternative minimal proofs and P_2 is a unique minimal proof, then P_1 is simpler than P_2.

4. If P_2 is not simpler than P_1 by virtue of Criteria 2 or 3 and P_1 uses a smaller number of different rules of inference than does P_2, then P_1 is simpler than P_2.

These four criteria arrange the 62 members into 17 equivalence classes of increasing complexity. Within an equivalence class the theorems were randomized and each subject was presented with a sequence of 17 theorems, as determined by random selection from each of the classes.

The four criteria of simplicity are not necessarily the only ones or even the best ones with which to begin. They did provide jointly a workable basis for arranging the theorems of length not greater than six.

The apparatus consisted of a plywood board, approximately 20 in. by 6 in., placed horizontally. The top half of the board was slotted for insertion of cards 20 in. by 2 1/2 in. Theorems were written on these large cards. The bottom half of the board had hooks at 2 3/4 in. intervals to permit hanging 2 in. by 2 1/2 in. cards on which were printed either a 0 or a 1. Each theorem was written on one of the large cards in such a way that it could be matched directly below it by hanging the appropriate small cards on the hooks.

Subjects were instructed as follows:

"This is a string of zeros and ones. Whenever I put a string of zeros and ones up here you are to put a string just like it down here.

Every string that I put up begins with a one, and so we'll leave these ones in the first places.

"Now of course if you could hang zeros and ones on the board any way you wanted, the game would be terribly easy. Why don't you try that now: Make a string of zeros and ones just like this one underneath it.

"That's right. Now that was too easy to be much fun, wasn't it? In the game we are going to play now there are some rules about the ways in which you may make strings. These are the rules:

"First, you may put two zeros on the end of a string you have made, but you may not put one zero by itself on the end. With ones it's the same way: you may put two ones on the end, but you may not put one one by itself on the end.

"If you wish though, you may take the last card off the end of a string, whether it's a one or a zero, and put it back in its box, but you may only remove the last card, none of the others.

"So the rules are: You may add two zeros to the end of a string, add two ones to the end of a string, take a zero off the end of a string, or take a one off the end of a string, except that you may not take away this first one. (practice). Now would you tell me the rules, just to make sure you understand them?

"Now each time you put cards on or take a card away, I want you to tell me what you're doing; what rule it is that you're following. That is, say, "I'm adding two zeros", or "I'm adding two ones", or "I'm taking off the last zero", or "I'm taking off the last one"."

Subjects were divided into two groups. Subjects in the correction group were corrected for each wrong step in each proof. Subjects in the other group (the discovery group) were stopped only when a valid proof was not completed in three times the length of a minimal proof. At the end of each proof, subjects in the discovery group were shown a minimal proof or--in the event that the subject constructed a minimal proof--told that the proof constructed was correct. For theorems which have several alternative proofs, we considered either proof correct and when demonstrating a minimal proof selected randomly.

Subjects made correction responses overtly, that is, in the correction group they removed the cards themselves and put up correct cards (or took down the correct cards) as instructed. In the discovery group they executed a minimal proof following instructions.

Actually not all subjects were required to prove 17 theorems. The following criterion rule was used. We considered that a subject had learned how to give minimal proofs in the system when four correct theorems were proved in succession. However, it was required that each subject prove at least ten theorems. All subjects, except for two in the discovery group, satisfied this criterion by the time the seventeenth theorem was reached. (It should be noted, however, that the nine subjects in each group represent a net figure; approximately this same number of subjects were discarded because of various kinds of problems that arose: seeming failure to comprehend the instructions at all, lack of attention after the first few theorems, etc.)

Because there were only nine subjects in each group the empirical data are not to be taken seriously and no detailed statistical analysis shall be presented here. I have summarized in Table 1 below the mean proportion of errors prior to the last error in blocks of twelve trials

Insert Table 1 about here

for each group and for the two groups combined. A trial in this instance is defined as a step in a proof and not as an entire proof. There were more than sixty trials, that is, more than a total of sixty lines of proof in the seventeen theorems, but because very few subjects needed the entire seventeen theorems to reach criterion, it has been necessary to terminate the mean curves with the last block ending on trial 60. Several things are to be noted about the data in Table 1. In the first place, the correction group seems to have done better than the discovery group, which result is consistent with experiments of a similar character dealing with the effects of immediate reinforcement. Secondly, the discovery group is more or less stationary (i.e., the learning curve in terms of responses prior to the last error is approximately flat). If anything, there is a tendency for the proportion of errors to increase with trials, whereas the correction group is clearly not stationary, and there is a definite tendency for the proportion of errors to decrease with trials. When the two groups are combined, an approximately stationary learning curve is obtained. The problem for future investigation is to discover which of these effects will be observed in larger and more stable bodies of data. In interpreting Table 1 it should be emphasized again that these figures are based on responses prior to the last error. Naturally, if the full

Table 1

Observed Proportion of Errors Prior to Last Error for the Correction, Discovery and Combined Groups. Blocks of 12 trials.

Group	Block 1	2	3	4	5
Correction	.28	.23	.15	.00	.10
Discovery	.23	.20	.40	.30	.33
Combined	.25	.21	.30	.18	.24

set of data were considered the learning curves would approach 1 at trial 60. It is also important to emphasize that the data of Table 1 are based on considering only a single sequence of trials. There is no analysis of the data into the separate subsequences defined by the various stimulus comparisons. This lumping together of the four sets of stimulus discrimination situations could in itself account for the lack of stationarity for the correction group, because responses were not deleted from consideration after the appropriate rule became conditioned to the stimulus discrimination. The rule for considering responses prior to the last error was invoked only for the whole sequence and not for the subsequences.

There was also some evidence against the independence of responses as shown by the figures given in Table 2. Here are shown the conditional

Insert Table 2 about here

probabilities of a correct response following a correct response and following an incorrect response for the correction, discovery and combined groups. The biggest difference occurs for the discovery group, which has a mean probability of .81 that a correct response will follow a correct response in comparison to a mean probability of .58 that a correct response will follow an incorrect response.

From this preliminary evidence it is perhaps doubtful that the one-element all-or-none conditioning model will fit very well the fine structural details of experimental data derived from young children learning proofs in the mathematical system $\mathfrak{A}$. On the other hand, this model does seem to give a pretty good first approximation to actual behavior. To give some idea of the immediate range of alternative possibilities I shall briefly sketch three other models that might be applied to data similar to those obtained from our pilot study.

The simplest alternative model is the linear incremental model with a single operator. The intuitive idea of this model is precisely the opposite of the all-or-none conditioning model. The supposition is that learning proceeds on an incremental basis. Let q_n be the probability of an error on trial n. Then the model is formulated by the following recursive equation.

(1) $$q_{n+1} = (1 - \theta)q_n ,$$

where $0 < \theta \leq 1$. It is simple to show but somewhat surprising that this purely incremental model has precisely the same mean learning curve as the all-or-none model if we set $c = \theta$. (To obtain this identity of the learning curves we must consider all responses and not simply responses prior to the last error.) The incremental model does differ sharply from the all-or-none model in the kind of learning curve predicted for responses prior to the last error, as is evident from equation (1).

The second simple alternative is a model which represents a kind of compromise between the all-or-none model and the incremental model. It assumes that associated with each discrimination situation there are two elements. Each of these (unobserved) elements is conditioned on an all-or-none basis but the two parameters of conditioning may be adjusted to produce various incremental effects on the response

Table 2

Conditional Probability Prior to Last Error of a Correct Response Following a Correct Response and Following an Error.

Group	After Correct R.	After Error
Correction	.85	.77
Discovery	.81	.58
Combined	.83	.64

probabilities. A model of this kind, as is pointed out in Suppes and Ginsberg [1961], could account fairly well for the kind of data shown in Tables 1 and 2. Probably its main inadequacy for accounting for more extensive data obtained from a large number of subjects would be found in connection with the problem of the assumed independence of the subsequences defined by the four types of stimulus comparisons.

The third alternative is to start with one of the three models already sketched and to introduce in a natural way dependencies among the subsequences. To introduce these dependencies we define a new process whose states are now ordered quadruples. The first coordinate of the quadruple indicates the state of conditioning of Rule R1, the second coordinate the state of conditioning of Rule R2, etc. One natural direction is then to define conditioning parameters c_{ij} for each Rule i, where j is the number of other rules already conditioned. By assuming further that the parameters depend only on j and not on i, we once again obtain a process with four conditioning parameters but with the parameters defined in an entirely different way. Without a large set of data to analyze and thereby to decide among these various alternatives it does not seem profitable to pursue them in any detail. Experiments are now underway in our laboratory and I hope to be able to report soon which models are most able to account for the fine structure of the data.

5. General Comments.

I would like to conclude with two general comments. In the discussion following the original presentation of this paper in Warsaw, Professor Kalmar appropriately raised the question of how the reinforcement schedule, i.e., the correction procedure, would be defined for more complicated mathematical systems than $\mathcal{U}$, in particular for systems which do not possess a decision procedure. It should be apparent that the behavioral theory outlined above is certainly not yet powerful enough to specify clear recipes for laying out the schedule of reinforcements. At the present time for more complicated systems, for example, the elementary algebra of integers and real numbers, the only practical procedure seems to be to proceed in a manner very similar to that used by Newell, Shaw and Simon (1956, 1957) in working out a program for proving theorems of elementary logic. Essentially their procedure is to abstract those heuristic principles that seem most useful in giving the set of proofs under consideration. My own conjecture is that in this area we shall find a substantial intersection between the work of mathematical psychologists interested in behavior theory and scientists like Simon who are interested in artificial intelligence and computer simulation of human behavior.

My second general comment is to emphasize that I am under no illusions about the fragmentary character of the behavioral foundations sketched in this paper. The next step forward it seems to me is to provide a theory at the following level of generality. Suppose we retained the problem of proving theorems in systems whose well formed formulas are strings of 1's and 0's. As rules of inference we have various rules of production of the kind given for the system $\mathcal{U}$. The problem is to formulate in sufficiently general terms the behavior theory that will lead to appropriate conditioning connections in at least a fairly wide class of systems similar to $\mathcal{U}$. The weakness of the present theory is easily brought out by considering the problem of using the theory to build a machine to prove theorems in such systems. It is clear that a machine could not be programmed, on the basis of the present theory, in a general way to prove theorems in systems similar to $\mathfrak{U}$. It is of course a trivial matter to program a machine to prove theorems in systems like $\mathfrak{U}$ if the programming is done for that particular system after the rules of inference of the system are specified. The much deeper problem of programming a machine to accommodate itself to proofs in a variety of systems similar to $\mathfrak{U}$ seems to me to be one of the most pressing problems to solve in order to provide a more adequate behavioral foundation of mathematical proofs. The solution of this problem will be of direct help in constructing a more general theory to predict the "proof-giving" behavior of our young subjects.

References

Bower, G. Application of a model to paired-associate learning. Psychometrika, 26 (1961), 255-280.

Estes, W.K. Learning theory and the new mental chemistry. Psychological Review, 67 (1960), 207-223.

Estes, W.K. and Suppes, P. Foundations of statistical learning theory, II. The stimulus sampling model. Tech. Report No. 26, October, 1959. Psychology Series, Institute for Mathematical Studies in the Social Sciences, Stanford University.

Hill, S.A. A study of the logical abilities of children. Ph.D. dissertation, unpublished, Stanford University, 1961.

Newell, A., Shaw, J.C., and Simon, H.A. Empirical explorations of the logic theory machine. Proceedings of the 1957 Western Joint Computer Conference, February, 1957.

Newell, A. and Simon, H.A. The logic theory machine. Transactions on Information Theory, IT-2, No. 3, September, 1956.

Restle, F. Statistical methods for a theory of cue learning. Psychometrika, 26 (1961), 291-306.

Suppes, P. and Atkinson, R.C. Markov Learning Models for Multiperson Interactions. Stanford University Press, 1960.

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. Tech. Report No. 39, September, 1961. Psychology Series, Institute for Mathematical Studies in the Social Sciences, Stanford University.

Suppes, P. and Ginsberg, R. Application of a stimulus sampling model to children's concept formation with and without overt correction responses. Journal of Experimental Psychology (in press, 1962)