

UNCLASSIFIED

AD 408 549

DEFENSE DOCUMENTATION CENTER

FOR

SCIENTIFIC AND TECHNICAL INFORMATION

CAMERON STATION, ALEXANDRIA, VIRGINIA



UNCLASSIFIED

NOTICE: When government or other drawings, specifications or other data are used for any purpose other than in connection with a definitely related government procurement operation, the U. S. Government thereby incurs no responsibility, nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use or sell any patented invention that may in any way be related thereto.

CATALOGED BY DDC

AS AD No. 408549

408 549

63-4-2

TM-1161 000 00

Sequential Designs of Experiments
for Alternative Objective Functions
in Automated Teaching Programs

9 May 1963

TECHNICAL MEMORANDUM

(TM Series)

This document was produced in connection with a research project sponsored by SDC's independent research program.

Sequential Designs of Experiments
for Alternative Objective Functions
in Automated Teaching Programs

Robert E. Dear

April 9, 1963

SYSTEM

DEVELOPMENT

CORPORATION

2500 COLORADO A

SANTA MONICA

CALIFORNIA



TABLE OF CONTENTS

	<u>Page</u>
<u>Introduction</u>	1
<u>Outline of the Theory of the Sequential Design of Experiments</u>	5
Strategies in Statistical Games	7
Normal Form of a Statistical Game	12
Sample Space of a Statistical Game	13
Sampling Plans and Terminal Decision Functions	16
Cost Functions and Loss Functions	17
Risk Function	18
Representation of a Sequentially Designed Statistical Game	19
Bayes Principle and Bayes Risk	19
<u>A Stimulus Sampling Teaching Model</u>	22
Mathematical Foundations of Stimulus Sampling Models	22
Single-Element Stimulus Sampling Models	24
A Two-Concept Teaching Model	25
Marginal Distributions of Responses at Trial t	28
The Set of All Possible Experiments, E_n	31
Conditional Sample Space of the Truncated Teaching Experiment $E_{k,n}$	35
<u>Objective, Loss, and Risk Functions in Teaching Experiments</u>	52
Objective Functions	52
Loss Functions	61
Risk Functions	65

	<u>Page</u>
<u>Solutions for Best Designs in Some Miniature Teaching Experiments</u> .	68
Pure Strategies in a 3-Trial Teaching Experiment	68
Normal and Extensive Forms of a Game	71
Solution for Best Designs by Backward Induction	77
Examples of Best Designs for 3-Trial Teaching Experiments	84
<u>Remarks</u>	104
Solution by Backwards Induction Yields Bayes Strategy.	104
Solutions for Best Designs in Terms of Posterior Probabilities .	108
Role of Sufficient Statistics	109
Relating Structures of Best Designs to Characteristics of . the Parameter Space	111
Relating Structures of Best Designs to Form of Loss and Cost Functions	112
Approximations to Solutions When the Numbers of Trials Are Large	113
<u>REFERENCES</u>	115

1. Introduction

A goal that many researchers in automated teaching would like to achieve is the development of branching procedures or designs of teaching programs that would in some sense be best tailored to the needs of the individual student. This paper is concerned essentially with giving theoretical foundations in terms of statistical decision theory to the pursuit of that goal which for the most part has enjoyed the status of a nice but vague idea. Several alternative design problems are formulated in this paper and a general technique for solution for best designs is outlined and illustrated.

The design problem for automated teaching programs or experiments is attacked in this paper from the standpoint of the theory of the sequential design of experiments. This seemed to be an especially interesting way to look at the design of automated teaching experiments since the use of a high-speed computer in conducting these experiments provides at least the possibility of making rapid calculations about whether a teaching experiment should be continued, and if so, what type of item should be presented at the next trial.

Part 2 of this report is devoted to outlining the general theory of the sequential design of experiments and the use of Bayesian procedures for determining best designs. Sequential experimentation in general is distinguished from fixed sample-size experimentation in that explicit consideration is given in sequential procedures to the cost of administering each subexperiment. In a sequentially designed experiment, the strategies available to the statistician or experimenter are made up of three components: the choice of an experiment,

the choice of a sampling plan, and the choice of a terminal decision function. In designing an automated teaching program these three components of a strategy respectively may be interpreted as: the choice of a rule for deciding what sequence of items shall be administered based on earlier item administrations and responses, the determination of the conditions under which the teaching experiment should be stopped for a student based on the set of possible outcomes of an experiment, and the determination of just what conclusion should be reached about the student's mastery of concepts at the termination of the program.

An effort was begun by Dear and Atkinson [11] to examine the branching problem in automated teaching by formulating a mathematical model of a teaching situation and then, within the framework of this model, mathematically searching for the best branching rule in a certain sense in a broad class of branching rules. A rather simple two-concept teaching situation was considered in this study in order to allow us to get some insight into the structure of these branching problems without being overwhelmed by the details that a mathematical representation of many current automated teaching programs would necessitate.

This two-concept teaching model was formulated in terms of the stimulus-sampling mathematical learning theory which was originally developed by Estes and is currently being widely applied and extended by many researchers. Two concepts labeled A and B were considered in our previous study and sets of equivalent items which embodied these concepts were assumed available for presentation at each trial. This two-concept model will be examined further in this paper since it is a sufficiently rich basis for study of the sequential

design problems which are to be considered here and it will quite well account for the results of a number of interesting learning studies that have been carried out.

Part 3 of this paper is devoted to a fairly detailed review of the foundations of stimulus sampling learning theory for the purpose of introducing the features of that theory which are needed to characterize a sequentially designed stimulus sampling teaching experiment and to identify certain parameters of these models which could lead to a number of different statistical decision problems. In Part 4, a number of alternative sequential design problems are developed in terms of several different identifications of states of nature or parameters of the relevant probability distributions, several different objective functions, and several alternative ways of expressing losses incurred by terminal decisions. The point is emphasized in Part 4 that solutions for best sequential designs of teaching experiments have the difficult complication over many current sequential design problems that the probability distributions on the sample spaces of these teaching experiments cannot be simply broken down over trials into independently distributed marginal components.

The technique of solution for best sequential designs of experiments called "backward induction" is outlined in Part 5. Solutions for best designs in several miniature 3-trial teaching experiments using this technique are then illustrated. The paper concludes, Part 6, with a discussion of characteristics that models of teaching processes need to have in order to be accessible to computation for best designs in full-scale teaching programs even when the

9 April 1963

4

TM-1161/000/00

backward induction technique is applied. The critical importance of coarse, sufficient partitions of the sample space of the teaching models is emphasized.

2. Outline of the Theory of the Sequential Design of Experiments

Sequential experimentation is a mode of carrying out statistical experiments by performing a sequence of subexperiments with the possibility that experimentation may be stopped at any point in the sequence and a terminal decision is then made at that stopping point. The theory of sequential analysis of experiments, many of whose early developments are due to Wald [22], is usually considered to refer to the situation where the identical subexperiment is repeated independently at the various points in the sample sequence. The extension of sequential theory to allow the possibility of selecting from a set of different subexperiments at each point in the experimental sequence is the distinguishing feature of the theory of the sequential design of experiments.

In addition to Wald's pioneering book [22], two general references are now available which give detailed consideration to various aspects of the theory of sequential experimentation. Blackwell and Girshick [4] present the foundations of sequential experimentation and they consider methods of solution for best sequential strategies in a number of specific situations; however, for the most part, these authors do not deal with the sequential design problem as they deliberately avoid the complication of sequential theory that the order in which subexperiments are performed may be important. In sequential games involving teaching or learning experiments one will, of course, expect the order of experimentation to be important. Raiffa and Schlaifer [18] consider a more general representation of sequential theory which does incorporate the possibility of performing different subexperiments at the various points in an experimental sequence. In the applications of this theory that these authors

then consider, they, too, take up only cases of independently distributed outcomes of the subexperiments.

A number of specialized papers on the sequential design of experiments have appeared in the literature. The so-called "two-armed bandit" problem and generalizations of this problem are the subject of several papers; Robbins [19], Bradt and Karlin [6] and Bradt, Johnson, and Karlin [7]. Chernoff [9] and Albert [1] have considered some hypothesis-testing problems from the standpoint of sequential design of experiments. DeGroot [12] has examined sequential design problems in terms of various measures of information in an experiment.

The design problems which Raiffa [17] has considered in his study of item selection procedures are, in terms of the subject of study, perhaps the most similar to the design problems in teaching experiments of any sequential design problems which have appeared. Raiffa considers both sequential and non-sequential experiments in this study which deals with the selection of items to develop psychometric tests and medical diagnosis procedures; however, in this study too, the outcomes of the subexperiments are assumed to be independently distributed.

An outline of the general structure of sequential statistical games will be sketched in the remaining sections of Part 2. For further detailed information about various aspects of the theory of statistical games, one may wish to consult in addition to the general references which have been mentioned, books such as the following: re game theory, Luce and Raiffa [15] and McKinsey [16]; and concerning statistical decision theory at a more elementary level, Chernoff and Moses [10] and Schlaifer [20].

Strategies in Statistical Games

The principal objective involved in the solution of any game is the identification of a rule of play or strategy which is in some sense a "good" way to play the game. Statistical game theory is primarily concerned with adapting a special class of games called two-person, zero-sum games to statistical decision problems. Under this adaptation, nature is represented as being one of the players and the statistician or experimenter is viewed as the second player.

To introduce some of the special design problems that arise in teaching experiments, a simple urn game will be considered. Let U_1 represent an urn containing m_1 white marbles and n_1 black marbles and U_2 be a second urn containing m_2 white marbles and n_2 black. Two players agree to play the following game using these two urns. Player 1 knows the identity of the two urns while the contents of the urns are not visible to player 2. Player 1 presents the two urns to player 2 in either the order (U_1, U_2) or (U_2, U_1) and player 2 is required to guess which presentation order is used (one may interpret the first position in this pair as the "urn on the left", U_L , and the second position as the "urn on the right," U_R). If player 2 guesses correctly, player 1 will pay him k dollars while if player 2 guesses incorrectly he is to pay player 1 k dollars.

Consequently, each play of the game results in an exchange of k dollars. These exchanges may be represented by payoff functions. For example, the payoffs that player 1 will receive from player 2 in this urn game are given in the matrix shown below:

		Player 2's Choice	
Player 1's Choice		(U_1, U_2)	(U_2, U_1)
(U_1, U_2)		-k	k
(U_2, U_1)		k	-k

This matrix gives all of the values of player 1's payoff function, say U_1 , for each possible play of the game. This game is a zero-sum, two person game since the payoff to player 2 is the negative of the payoff to player 1. Thus letting U_2 be player 2's payoff function, the values of this function are given by the following matrix:

		Player 2's Choice	
Player 1's Choice		(U_1, U_2)	(U_2, U_1)
(U_1, U_2)		k	-k
(U_2, U_1)		-k	k

The two choices of orderings of the urns (U_1, U_2) and (U_2, U_1) constitute the pure strategies for playing this game for both player 1 and player 2. It would not appear that there is any choice of a pure strategy for player 1 in this game which in conjunction with some choice of a pure strategy for player 2 results in a payoff which is a good compromise for both. Frequently, it is necessary for one or both of the players to resort to using more complicated

strategies called mixed strategies in order to obtain certain "good" ways to play the game or "good solutions" of the game. Mixed strategies are formed by defining a probability distribution over the set of pure strategies and hence selecting a pure strategy by first operating a random device which appropriately reflects the desired selection probabilities. The urn game illustrated here and all of the games considered in the paper have the characteristic that each of the two players has only a finite number of pure strategies. A fundamental theorem for such finite games establishes that, when mixed strategies are introduced, each player has at least one good strategy.

It can be shown for this urn game that if each of the two players independently employs the mixture of selecting the configuration (U_1, U_2) with probability $p = 1/2$ that this mixed strategy is a good strategy for each player. The expected payoff for either player using this strategy is then 0 dollars.

Consider next a modification of this game in which player 2 is allowed to pay one dollar to player 1 and, in return for this fee, player 2 is permitted to take a random draw of one marble from either of the two urns that player 1 presents. If player 2 decides not to pay the entry fee, he is still allowed to guess which configuration obtains, as in the original game. The pure strategies for this game and the payoffs to player 2 for each pair of choices of pure strategies are shown in the matrix which follows:

Player 2's Strategy	Player 1's Strategy	
	(U_1, U_2)	(U_2, U_1)
$[U_L, (W, (U_1, U_2)), (B, (U_1, U_2))]$	$k-1$	$-(k+1)$
$[U_L, (W, (U_1, U_2)), (B, (U_2, U_1))]$	$(k-1)p(W U_1) - (k+1)p(B U_1)$	$-(k+1)p(W U_2) + (k-1)p(B U_2)$
$[U_L, (W, (U_2, U_1)), (B, (U_1, U_2))]$	$-(k+1)p(W U_1) + (k-1)p(B U_1)$	$(k-1)p(W U_2) - (k+1)p(B U_2)$
$[U_L, (W, (U_2, U_1)), (B, (U_2, U_1))]$	$-(k+1)$	$k-1$
$[U_R, (W, (U_1, U_2)), (B, (U_1, U_2))]$	$k-1$	$-(k+1)$
$[U_R, (W, (U_1, U_2)), (B, (U_2, U_1))]$	$(k-1)p(W U_2) - (k+1)p(B U_2)$	$-(k+1)p(W U_1) + (k-1)p(B U_1)$
$[U_R, (W, (U_2, U_1)), (B, (U_1, U_2))]$	$-(k+1)p(W U_2) + (k-1)p(B U_2)$	$(k-1)p(W U_1) - (k+1)p(B U_1)$
$[U_R, (W, (U_2, U_1)), (B, (U_2, U_1))]$	$-(k+1)$	$k-1$
(U_1, U_2)	k	$-k$
(U_2, U_1)	$-k$	k

The notation for strategies involving experimentation, for example, consider the strategy $[U_L, (W, (U_1, U_2)), (B, (U_2, U_1))]$, is interpreted--the lefthand urn U_L is to be selected for a random draw; then if a white marble is drawn, claim that the correct configuration is (U_1, U_2) ; while if a black marble is drawn, claim the configuration to be (U_2, U_1) . Since random moves have now been introduced into the game, one sees that the payoff for certain of the pairs of the pure strategies for the two players must be expressed as expected values of the payoffs over probability distributions on the random moves.

This second simple example of an urn game has served to introduce the main characteristics of a sequential-design-of-experiments problem. The pure strategies for player 2 are seen to involve three important features: (1) a decision concerning how much experimentation should be done, (2) a choice of which experiment should be performed and (3) the final decision concerning what configuration of the urns that player 1 has presented. In more general sequential design problems these three components of a pure strategy for player 2 may be identified respectively as the choice of a sampling plan, the choice of a sequence of subexperiments, and the choice of a terminal decision function.

This urn game could be elaborated by allowing additional random draws by player 2. Most of the sequential design problems that have been considered in the literature have considered the situation where successive outcomes of subexperiments are independently distributed. In these urn games, the independence case would be effected by making the random draws with replacement. The designs of teaching experiments which are considered in this paper,

on the other hand, are more similar to strategies for urn games involving sampling without replacement.

Normal Form of a Statistical Game

The descriptions of the two urn games have included all the necessary ingredients to characterize a game in a mode called the normal form of a statistical game. Let U be the set of all pure strategies for player 1, and V be the set of pure strategies for player 2, and μ_2 be player 2's payoff function defined on the product set $U \times V$. The normal form of this game is the triple, say, $G = (U, V, \mu_2)$. Since statistical games are zero-sum, two-person games it is sufficient to specify either μ_2 or μ_1 in the triple to uniquely define the game G .

When mixed strategies are employed by each player, the sets of strategies U and V may be expanded to include all possible mixtures of the elements of each of these two sets. Let π represent a mixed strategy for player 1 and Π be the set of all his mixed strategies. Similarly, let η be a mixed strategy for player 2 and H be the set of all of player 2's mixed strategies. One defines the mixed extension of the game G to be the triple, say, $\Gamma = (\Pi, H, \mu_2)$.

Statistical games are often usefully represented in normal form to study various conditions under which good solutions to the games exist, to examine various relationships between certain classes of strategies, and to examine other fundamental problems in statistical game theory. Frequently, another equivalent representation of a statistical game called the extensive form is more suitable for the purpose of actually finding specific solutions to

statistical games. The characterization of the extensive form of a game will be deferred to Part 5 of this report where it is used to obtain best designs for several illustrative teaching experiments.

Sample Space of a Statistical Game

The features of a sequential-design-of-experiments problem which were introduced in a rough, intuitive way through the description of the two simple urn games will now be formalized in order to allow a general representation of sequential-design-of-experiments problems. Since statistical games typically will involve the use of experiments by the statistician, it is desirable to have a representation of all the possible outcomes of a statistical experiment. Conventionally, all possible outcomes of a statistical experiment are represented as a set, say Y , which is called the outcome space (or frequently the sample space) of the experiment. The outcome space will be defined in sequential design problems to be rich enough to include all possible experiments of interest and all conceivable outcomes of each experiment. Although the phrase "sample space of a statistical experiment" is often used synonymously with "outcome space" it is also used to represent a triple, say, $Z = (Y, \Omega, p(\cdot | \omega, e))$. The components of this triple are the outcome space, Y , a set Ω of parameters or indices of probability distributions which are defined on the outcomes of a particular experiment $e \subset Y$, and a probability distribution on the outcomes of a statistical experiment, $p(\cdot | \omega, e)$, which is defined when a parameter point $\omega \in \Omega$ and an experiment $e \subset Y$ are specified.

Although this representation of the sample space will be suitable for the development of most sequential design problems, a number of alternative modes

of representation of the sample space could be used. For example, one could let $Y_{\underline{e}}$ be the restriction of Y or the subset of Y consisting of all outcomes to the experiment $\underline{e}$. One can then define the conditional sample space given the experiment $\underline{e}$ to be the triple $Z_{\underline{e}} = (Y_{\underline{e}}, \Omega, p(\cdot | \underline{\omega}))$. These conditional sample spaces given a particular experiment $\underline{e}$ are the sample spaces considered in a typical sequential analysis problem where it is necessary only to determine a good sampling plan and a good terminal decision function. The important point to note about either representation of the sample spaces Z or $Z_{\underline{e}}$, is that it is necessary to specify both a parameter point $\underline{\omega}$ and an experiment $\underline{e}$ to define a probability distribution on the outcome space. In situations where it is well understood that a particular experiment is being employed, it is conventional to delete the subscript $\underline{e}$ from the definition of the conditional sample space.

In order to simplify the description of a sequential game and the set of possible experiments that a statistician could choose from in this game, attention will be restricted to games which will continue for, at most, n steps or trials. Such sequential games are called truncated sequential games.

The outcome space Y of a truncated teaching game will be a set of n -dimensional sequences whose order is determined by the trial numbers. A notation which will be used generally in this paper to represent sequences and vectors is the employment of underlined lower case letters. Thus, for example, a representative element of the outcome space Y will be indicated as $\underline{y} = (y_1, y_2, \dots, y_n)$. The values y_j represent the coordinates or components of $\underline{y}$ at trial j . The parameter spaces Ω , which will be considered, will typically be multi-dimensional sets; consequently the elements of these sets or parameter points will be similarly indicated, i.e., $\underline{\omega} = (\omega_1, \omega_2, \dots, \omega_k)$.

The same type of notation will be used to designate experiments, components of experiments or subexperiments, and the set of all experiments of interest. Thus E shall be the set of all experiments of interest and the elements of E will be denoted by e . The description of an experiment in terms of its component subexperiments requires a somewhat more elaborate notational apparatus. One of the most convenient ways to describe an experiment is through the geometrical concept of a particular type of connected graph called a tree. For example, consider a sequential game truncated at 2 trials developed in terms of Bernoulli or binomial subexperiments. In a situation involving 2 different binomial subexperiments e_1 and e_2 (in the urn games considered earlier e_1 could be the selection of the lefthand urn, U_L , and e_2 the selection of the righthand urn, U_R), experiments such as the two shown in Figure 1 are possible.




Figure 1

The experiment e_1 may be characterized by the rule that subexperiment e_1 is to be used at both trials hence the rule does not depend on any outcomes. Experiment e_2 , on the other hand, does represent a rule which makes use of the information about outcomes at the first trial of the experiment.

Sampling Plans and Terminal Decision Functions

A sequential statistical game is distinguished from a fixed sample size or non-sequential game by the fact that sampling may terminate at any step (perhaps without even starting experimentation). The rules which specify when to continue and when to terminate sampling are usually called sampling plans. Each sampling plan may be represented as a partition of the outcome space Y_e into subsets called stopping regions.

The stopping regions which comprise a sampling plan are required to be cylinder sets of a particular type. Letting b represent a sampling plan, B the set of all sampling plans and b_i the stopping region for the i th step in a sequential game, it shall be required that if y_j and y_k are elements of Y_e and the values of the first i coordinates of y_k are equal to the corresponding values of y_j (one says in this case that y_j agrees with y_k in the first i coordinates), then $y_j \in b_i$ if and only if $y_k \in b_i$. The index i may range over the set of integers $0, 1, 2, \dots, n$; if $i = 0$, one uses the definition that all $y \in Y_e$ agree in the value of their 0th coordinate. Consequently, b_0 is either the entire set Y_e or the null set. A sampling plan b is then a sequence of cylinder sets $b = (b_0, b_1, b_2, \dots, b_n)$ which partitions the outcome space Y_e ; thus, given a sampling plan, one can tell for each $y \in Y_e$ the stopping region in which y is an element.

In every statistical decision problem, the ultimate goal is to make a choice of one, among a set of alternative actions. In sequential games, this set, say A , is referred to as the set of terminal actions. The choice of a good experiment and a good sampling plan is made only to improve one's basis for choosing a terminal action. Part of the statistician's strategy in a sequential game is typically formulated as the choice of a terminal decision function. For a sequential game truncated at n steps, one may define the set $I_n = \{i: i = 0, 1, 2, \dots, n\}$ as the set of possible stopping points. A terminal decision function is then defined as a function d whose domain is the product set $I_n \times Y_e$ and whose range is the set of terminal actions A .

One also constrains the terminal decision functions by a requirement that if two sample sequences y_j and y_k are elements of a stopping region b_i and these two sample sequences agree in the values of their first i coordinates, then the terminal decision reached for the sequence y_j must be the same as the decision reached for y_k ; that is, $d(i, y_j) = d(i, y_k) = a$.

Cost Functions and Loss Functions

Explicit recognition is given in sequential game theory to the fact that each additional subexperiment which is performed must in some sense be paid for. The cost of experimentation in truncated sequential games is represented by a non-negative, bounded function, say c , whose domain again is the product set $I_n \times Y_e$. A restriction similar to that imposed on the sampling plans and the terminal decision functions is levied on the cost functions. Thus, if y_j and y_k are each elements of a stopping region b_i , and y_j and y_k agree in

the values of their first i coordinates, then it shall be required that

$$c(i, \underline{y}_j) = c(i, \underline{y}_k).$$

When the statistician finally stops sampling and decides upon a terminal action a , then the return to him for making that final choice will also be dependent on the values of the parameter point that nature has chosen. It is customary in statistical games to express the consequences of the statistician's terminal actions against each of the possible choices of a parameter point or probability distribution as loss functions. A loss function in a statistical game, say L , is defined to be a non-negative, bounded function with domain, the product set $\Omega \times A$.

In general game theory, the consequences of a player's choices are usually expressed in terms of the values of his utility functions. In statistical games, the player's loss functions are defined to have as values, the negative of the values of his utility functions.

Risk Function

In statistical games, the payoff to player 2 (the statistician or experimenter) is expressed in terms of the value of a function called the risk function. For sequentially designed experiments, the risk function, say ρ , when only pure strategies are used by both players is defined as follows:

$$\rho(\underline{\omega}, (\underline{e}, \underline{b}, d)) = \sum_{i=0}^n \sum_{\underline{y} \in b_i} [c(i, \underline{y}) + L(\underline{\omega}, d(i, \underline{y}))] p(\underline{y} | \underline{\omega}).$$

Representation of a Sequentially Designed Statistical Game

One sees that the arguments of the payoff function ρ for the statistician in a sequentially designed game are the pairs of strategies, ω for nature and $(\underline{e}, \underline{b}, \underline{d})$ for the statistician. Thus, formally, these statistical games can be described as the triple $G = (\Omega, E \times B \times D, \rho)$ where the parameter set Ω represents the set of pure strategies for nature; the set of pure strategies for the statistician is the product set $E \times B \times D$ whose components are the set of all experiments E , the set of all sampling plans B , and the set of all terminal decision functions D ; and the risk function ρ is the statistician's payoff function.

Bayes Principle and Bayes Risk

One of the principles that has enjoyed considerable acceptance among statisticians as a means for determining a preference ordering on the set of strategies available to the statistician is Bayes principle. This principle asserts that the experimenter or statistician can designate a particular probability distribution that nature is using over the set of parameter points Ω or pure strategies for nature on the basis of his previous experience and background information available to him prior to the performance of any experiments. An alternative payoff function applies for the statistician when the Bayes principle is employed. This payoff function is called the risk function against the probability distribution, say, π over the parameter set Ω . The risk function evaluated for the strategy $(\underline{e}, \underline{b}, \underline{d})$ against π in a sequential design problem is defined as follows:

$$\rho(\pi, (\underline{e}, \underline{b}, \underline{d})) = \sum_{i=0}^n \sum_{\underline{y} \in \underline{b}_i} \sum_{\underline{\omega} \in \Omega} \left[c(i, \underline{y}) + L(\underline{\omega}, d(i, \underline{y})) \right] \pi(\underline{\omega}) p(\underline{y} | \underline{\omega}).$$

A Bayes solution to the design problem against π is a strategy which minimizes the risk function $\rho(\pi, (\underline{e}, \underline{b}, \underline{d}))$. The value of the risk function for a strategy which represents such a Bayes solution is called the Bayes risk. (However, frequently in the literature one will see any of the risk functions against a distribution π referred to as Bayes risk functions.) The probability distributions π are called a priori or prior distributions on Ω .

An important computational feature of Bayes procedure for finite statistical games of the sort which characterize the teaching experiments considered in the remainder of this paper is that it is sufficient to consider only pure strategies for the statistician. It is easy to show that for these kinds of problems that some pure strategy is at least as good as any mixture of pure strategies for every prior distribution π . However, even though the statistician has only a finite number of the pure strategies $(\underline{e}, \underline{b}, \underline{d})$, the total number of these strategies rapidly becomes so large with increasing truncation trial number n , that solution for best strategies by sheer enumeration of the values of $\rho(\pi, (\underline{e}, \underline{b}, \underline{d}))$ is not feasible.

The use of Bayes principle to define "bestness" for designs of teaching programs would appear to be a particularly appropriate mechanism to give substance to the notion of tailoring a teaching program to the needs of an individual student. The definition of the Bayes risk incorporates the concept that the design of a teaching program which is to be best for a game involving the responses of individual students will be dependent not only on the parameter

9 April 1963

21

TM-1161/000/00

values ω which govern the rate of the student's learning but will also be dependent on how well the experimenter can identify the student's learning capacity at the outset. The experimenter's probabilistic classification of students into the various possible populations, preliminary to the teaching experiment, may be represented in the prior distributions.

3. A Stimulus Sampling Teaching Model

It has been shown that in statistical games the moves and outcomes of these games are described in terms of the sample space of a statistical experiment. In Part 3, the sample space of a two-concept teaching model that is essentially the same model which was developed by Dear and Atkinson [11] will be described. However, it now seems apparent that more general experiments must be considered to determine best sequential designs for teaching experiments involving the sequence of responses of individual students than we considered in that earlier paper. One pays the price of having substantially more difficult probability distributions of responses to work with when the more extensive set of possible experiments is considered.

It will be useful to review first the mathematical foundations of the general stimulus sampling theory of learning in order to see how probability distributions are built up in the sample spaces of these models. The special assumptions that are involved in the single-element, two-concept teaching model which provides the setting for these sequential design studies will then be identified. Finally, the manner of constructing probability distributions on the outcome sequences of these two-concept models from certain elementary conditional probabilities and parameter values will be shown.

Mathematical Foundations of Stimulus Sampling Models

Estes and Suppes [14] have given a formal representation of the general stimulus sampling theory for simple learning situations as an axiom system. Since the model for teaching two related concepts that is being utilized here

is a special and, in many respects, a very simple version of stimulus sampling models, it will not be particularly useful to review the nature of their general axioms. However, it will be instructive to review their description of the sample space of these models in order to indicate the special restrictions that are imposed on the sample space in this two-concept teaching model.

The elements of the sample space of stimulus sampling models in general are sample sequences defined on a discrete time parameter set which consists of trial numbers. It is customary in mathematical learning theory to start counting with trial 1 so that the time parameter set of these processes is usually the set of all positive integers or some subset of the positive integers.

Estes and Suppes show that the coordinate at trial t of a sample sequence in these models consists of an ordered 6-tuple of values, (C, T, s, i, j, k) . The first term, C , in this expression denotes a conditioning function which partitions an abstract set of stimulus elements, say, S into subsets such that each element in any one of the subsets is conditioned to, or connected to, a particular response, and the various cells or subsets of the partition are each connected to a different one of the available responses. T denotes a subset of S that is chosen for presentation to the subject at a given trial and hence these authors call T a presentation set of stimuli. The third term, s , refers to the subset of T which the subject samples. The component i refers to the response that the subject makes on sampling the presented stimuli; while j denotes the outcome of the trial (receipt of food, avoidance of shock, being

informed of correctness of response, etc.) that the experimenter arranges to follow the various responses. The last value, k , in this 6-tuple designates an unobservable reinforcing event that may occur and alter the conditioning function at the next trial.

Single-Element Stimulus Sampling Models

A number of special cases of the general stimulus sampling model have been developed for application to specific experimental situations. One of the simpler versions of the model that has been applied extensively is the single-element stimulus sampling model. In this model, the set of stimulus elements is considered to consist of but a single element. The sampling axioms of the general stimulus sampling theory are usually modified in the single-element version to assert that the subject samples this element with probability 1 at each trial. The modification of the conditioning function from trial to trial is then assumed to be governed by a probabilistic process involving conditioning rate parameters.

The sample space of a typical single-element model consists of sequences whose coordinates at trial t can be reduced to a 4-tuple of values. Since the presentation set and the sampled subset in the usual versions of single-element models are at each trial the single stimulus element, these components can be deleted from the description of the sample sequences. Consequently, the description of the coordinate of a sample sequence at trial t in a single-element model can be reduced to $(C, i, j, k)_t$ — where these four components which are retained are defined as in the general theory.

Frequently, the first component of the 4-tuple is given in terms of the values of the conditioning function. For example, in a simple two-response situation ($i = 1$ or 2) one might let C_i represent the value of the conditioning function — the single element is conditioned to response i . These values of the conditioning function are often referred to as the states of the conditioning function or as the states of conditioning.

The sample space for the particular single-element stimulus sampling model which will provide the setting for the present study of sequential design of teaching experiments will be described in detail in the following section. For further information about the structure of single element models, a number of general references are available; see, for example — Suppes and Atkinson [21], Estes [13], Bower [5], and Atkinson and Estes [2].

A Two-Concept Teaching Model

A very simple model of a teaching situation, which is at the same time complex enough to reflect some of the branching problems that occur in automated teaching experiments, can be defined in terms of two related concepts. The two concepts which are considered will be labeled Concept A and Concept B. In the language of stimulus sampling theory each of these two concepts can be represented as an abstract stimulus element say respectively A and B. Two types of items are considered to be used in this experiment; they are called items of type A and of type B. The various items in the set of type A items are viewed as equivalent reproductions of the stimulus element A and a similar interpretation of items of type B as equivalent reproductions of the stimulus element B is

made. Consequently, the two types of items may be thought of as two presentation sets A and B and at each trial either the element A alone is presented or the element B alone is presented; no other presentations of stimuli occur. On the presentation of either element, it is assumed that the presented element is sampled with probability 1.

There are assumed to be only two possible responses to an A item or to a B item; consequently, the response index $i = 1, 2, 3$, or 4. However, it is further assumed that the responses are separated into two disjoint pairs such that, say, 1 and 2 are the only responses available to the subject on item A trials and 3 and 4 are the only responses available on item B trials. Further, it will be assumed that the experimenter wishes to have response 1 conditioned or connected to element A and response 3 conditioned to element B (responses 1 and 3 are respectively the correct answers to Concept A items and Concept B items).

The outcomes of each trial are limited to two values. The subject is told either that he has made the correct response, say $j = 1$ in this case, or he is told that he has made the incorrect response, $j = 2$. For simplicity, it is assumed these two outcomes have symmetric effects on the reinforcement of the correct response and that reinforcement occurs with probability 1 at each trial. Since reinforcement is a deterministic process by these assumptions, it will be possible for this model to delete the reinforcement component from the description of sample sequences.

In defining the sample space for this two-concept teaching model one could use the notation of the general theory and designate the coordinate of a sample sequence at trial t as $(C_i, T, i, j)_t$. In this expression the conditioning function

C can take four possible values C_1, C_2, C_3 , or C_4 , the presentation sets are either $T = A$ or $T = B$, there are four possible responses R_1, R_2, R_3, R_4 , and the experimenter's outcomes j take two values, say, $j = 1$ being the outcome, "you gave the correct answer" and $j = 2$ being the outcome, "you gave the incorrect answer."

The use of this notation would suggest the interpretation that this two-concept teaching model is a two-element stimulus sampling model. However, since the four responses in this model are grouped so that only a fixed pair are available on A trials and the remaining pair are available on B trials, it seems appropriate to interpret this model as consisting of two "linked" single-element processes. This interpretation is emphasized by using the following notation to define the sample sequences. For example, let $(C_A, C_B, A, R_A)_t$ represent the outcome at trial t that the Concept A conditioning function is in the state C_A (the stimulus element A is connected to the correct response), the Concept B conditioning function is in the state C_B , a Concept A item was presented at trial t , and the correct response, R_A , was made to the A item.

The Concept A and Concept B conditioning functions are each assumed to take two values. The two sets of values or states of these two conditioning functions are denoted respectively $C_A, \tilde{C}_A$ and $C_B, \tilde{C}_B$. The states C_A and C_B can be interpreted as states in which mastery of the concepts has occurred--that is, when a subject is in these states the stimulus element A or B is connected to the appropriate responses. The states $\tilde{C}_A$ and $\tilde{C}_B$ are interpreted as guessing states. When a subject is in these states he may guess the correct answers R_A or R_B with probabilities, say, g_A and g_B or may guess incorrect answers, say,

$\tilde{R}_A, \tilde{R}_B$. The complete set of possible combinations of these components that may occur at coordinate t of a sample sequence is the following:

$$\begin{aligned} & \left\{ (\tilde{C}_A, \tilde{C}_B, A, \tilde{R}_A)_t, (\tilde{C}_A, \tilde{C}_B, A, R_A)_t, (\tilde{C}_A, \tilde{C}_B, B, \tilde{R}_B)_t, (\tilde{C}_A, \tilde{C}_B, B, R_B)_t \right. \\ & (\tilde{C}_A, \tilde{C}_B, A, \tilde{R}_A)_t, (\tilde{C}_A, \tilde{C}_B, A, R_A)_t, (\tilde{C}_A, \tilde{C}_B, B, \tilde{R}_B)_t, (\tilde{C}_A, \tilde{C}_B, B, R_B)_t \\ & (\tilde{C}_A, C_B, A, \tilde{R}_A)_t, (\tilde{C}_A, C_B, A, R_A)_t, (\tilde{C}_A, C_B, B, \tilde{R}_B)_t, (\tilde{C}_A, C_B, B, R_B)_t \\ & \left. (C_A, C_B, A, \tilde{R}_A)_t, (C_A, C_B, A, R_A)_t, (C_A, C_B, B, \tilde{R}_B)_t, (C_A, C_B, B, R_B)_t \right\}. \end{aligned}$$

(The notational practice of grouping the components of the t^{th} coordinate of a sample sequence within parentheses with the trial number as a subscript will be used generally. Events or sets of the elementary sequences will usually have their trial numbers indicated in the same way. Departures from this notation will be defined as needed.)

Marginal Distributions of Responses at Trial t

The states of the conditioning function are typically not observable aspects of stimulus sampling learning experiments. On the other hand, the distributions of the item responses which are observable characteristics in these experiments, when the states of the conditioning function are given, constitute a set of time-independent Bernoulli distributions (in settings like the present problem which involves two possible responses). This results from the response axioms of stimulus sampling theory.

To clarify and emphasize this point about the response distributions, the marginal distributions of responses at a particular trial t will be described.

Letting R_t be the set of possible item administrations and responses at trial t , i.e.,

$$R_t = \{(A, R_A)_t, (A, \tilde{R}_A)_t, (B, R_B)_t, (B, \tilde{R}_B)_t\}$$

and letting C_t be the set of values of the two conditioning functions at trial t , i.e.,

$$C_t = \{(\mathcal{C}_A, \mathcal{C}_B)_t, (C_A, \tilde{\mathcal{C}}_B)_t, (\tilde{\mathcal{C}}_A, C_B)_t, (C_A, C_B)_t\}$$

one may then define the marginal sample space say Y_t as the triple $Y_t = (R_t, C_t, p\{. | c_t\})$. The conditional distributions on the elements r_t of the outcome space R_t given the values c_t of the conditioning function at trial t , are defined for this stimulus sampling model as follows:

$$p\{r_t = (A, R_A)_t | (A)_t, c_t = (C_A, C_B)_t\} = 1$$

$$p\{r_t = (A, R_A)_t | (A)_t, c_t = (C_A, \tilde{\mathcal{C}}_B)_t\} = 1$$

$$p\{r_t = (A, R_A)_t | (A)_t, c_t = (\tilde{\mathcal{C}}_A, C_B)_t\} = g_A, \quad 0 \leq g_A \leq 1$$

$$p\{r_t = (A, R_A)_t | (A)_t, c_t = (\tilde{\mathcal{C}}_A, \tilde{\mathcal{C}}_B)_t\} = g_A$$

$$p\{r_t = (B, R_B)_t | (B)_t, c_t = (C_A, C_B)_t\} = 1$$

$$p\{r_t = (B, R_B)_t | (B)_t, c_t = (\tilde{\mathcal{C}}_A, C_B)_t\} = 1$$

$$P\left\{r_t = (B, R_B)_t \mid (B)_t, c_t = (C_A, \tilde{C}_B)_t\right\} = g_B, \quad 0 \leq g_B \leq 1$$

$$P\left\{r_t = (B, R_B)_t \mid (B)_t, c_t = (\tilde{C}_A, \tilde{C}_B)_t\right\} = g_B$$

where the probabilities g_A and g_B represent the probabilities of guessing correct responses to Concept A and Concept B items when these concepts have not yet been mastered.

For the marginal sample space Y_t , it is evident that the values of the conditioning function C_t have been interpreted as parameters of the response distribution at trial t . However, in stimulus sampling models the time-dependency properties of these stochastic processes are characterized principally through the probability distributions defined on sequences of values of the conditioning functions. Estes and Suppes [14] have given conditions under which sequences of certain random variables will be finite-state Markov chains. Frequently, it turns out that the values of the conditioning functions can be taken as states of a 1st-order Markov chain. The manner in which it seems necessary to represent the set of possible pure experiments in studying the design of these two-concept teaching programs does not allow the sequence of values of the conditioning functions to be a 1st-order Markov chain. This results because the rules governing the choice of presentation sets or items to be administered must allow for all configurations of past histories of items presented and responses obtained.

The Set of all Possible Experiments, E_n

It will be convenient for the purpose of defining the set of all possible teaching experiments, say E_n , which the experimenter could perform first to restrict the numbers of trials to be allowed in a teaching program to be at most a finite number, say n . (This is not a practical restriction at all but it does impose some mathematical restrictions, principally on types of limiting processes that can be performed.) The teaching program under this restriction can thus be regarded as a sequential statistical game truncated at n trials.

It is clear that the rules for defining a complete experiment in a sequential teaching program can be based only on the observables, the types of items that have been administered and the responses that occurred to these items. For a teaching experiment involving two types of items and dichotomous responses to these items over n trials, one will find that E_n consists of $2^{(2n-1)}$ experiments. All of the component experiments in E_n may be described by enumerating all the trees or branching patterns that can be generated from consideration of the types of items that may be administered at each trial (A or B) and the responses to these items (R_A or $\tilde{R}_A$ and R_B or $\tilde{R}_B$).

It may help, to clarify further the concept of an experiment, to list several trees of experiments in this teaching model. For simplicity's sake, attention will again be restricted to small experiments--consisting in this illustration of 3 trials. Let $R = R_1 \times R_2 \times R_3$ be the outcome space for such a 3-trial experiment. It is readily shown that R consists of $4^3 = 64$ outcome sequences. An experiment e_3 will be a subset of R in which the responses at trial 3 are ignored.

A teaching experiment can be viewed then as a rule which determines a sequence of subexperiments or item administrations over all trials. Such rules can be illustrated by branching patterns or trees. Two trees that define two experiments for a three-trial learning situation are shown in Figures 2 and 3. The first tree, Figure 2, illustrates a rule which is conditioned only on the trial numbers. This rule ($e_{1,3}$) is, "Administer an A item at each odd-numbered trial and a B item at the even-numbered trial." The second tree, Figure 3 illustrates the rule ($e_{2,3}$), "Administer an A item at trial 1, and administer a B item following a correct response to any type of item, but administer an A item following an incorrect response to any type of item."




Figure 2

Item Administered at Trial 3


Tree of the Experiment $e_{2,3}$

Figure 3

These two trees illustrate the complete prescription of what types of items are to be administered at each trial under the conditions of the rules or experiments $e_{1,3}$ and $e_{2,3}$. When the two trees are extended to include the responses that could occur at Trial 3, there are then 8 branches to each tree. The 8 branches represent the possible sequences that can occur as elements of the outcome spaces, $R_{e_{1,3}}$ and $R_{e_{2,3}}$, which are determined as restrictions of the outcome space R respectively by the experiments $e_{1,3}$ and $e_{2,3}$. The elements of the outcome space $R_{e_{1,3}}$ are the eight sequences

$$\begin{array}{ll}
\left[(A, \tilde{R}_A)_1, (B, \tilde{R}_B)_2, (A, \tilde{R}_A)_3 \right] & \left[(A, R_A)_1, (B, \tilde{R}_B)_2, (A, \tilde{R}_A)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (B, \tilde{R}_B)_2, (A, R_A)_3 \right] & \left[(A, R_A)_1, (B, \tilde{R}_B)_2, (A, R_A)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (B, R_B)_2, (A, \tilde{R}_A)_3 \right] & \left[(A, R_A)_1, (B, R_B)_2, (A, \tilde{R}_A)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (B, R_B)_2, (A, R_A)_3 \right] & \left[(A, R_A)_1, (B, R_B)_2, (A, R_A)_3 \right]
\end{array}$$

while the eight sequences which comprise the outcome space of the experiment $e_{2,3}$ are

$$\begin{array}{ll}
\left[(A, \tilde{R}_A)_1, (A, \tilde{R}_A)_2, (A, \tilde{R}_A)_3 \right] & \left[(A, R_A)_1, (B, \tilde{R}_B)_2, (A, \tilde{R}_A)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (A, \tilde{R}_A)_2, (A, R_A)_3 \right] & \left[(A, R_A)_1, (B, \tilde{R}_B)_2, (A, \tilde{R}_A)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, \tilde{R}_B)_3 \right] & \left[(A, R_A)_1, (B, R_B)_2, (B, \tilde{R}_B)_3 \right] \\
\left[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, R_B)_3 \right] & \left[(A, R_A)_1, (B, R_B)_2, (B, R_B)_3 \right]
\end{array}$$

It is evident that these rules for determination of experimental sequences are deterministic or non-randomized rules concerning what type of item to administer next, given the history of item administrations and associated responses that has occurred prior to the current trial. It is clear that in the 3 trial situation illustrated here there are $2^5 = 32$ distinct trees

and hence 32 possible experiments. These experiments would represent all the pure strategies that the statistician could employ with respect to the allocation of types of items to the various trials. Randomized item allocation strategies can be developed by taking mixtures of the pure strategies.

Conditional Sample Space of the Truncated Teaching Experiment, $e_{k,n}$

For each experiment $e_{k,n} \in E_n$, one may describe a sample space for a truncated teaching program that will continue for at most n trials. Letting X be the outcome space for such a truncated teaching program, one can represent this outcome space as the product space

$$X = \left(\prod C_t \times R_T \right) \times C_{n+1}.$$

In this definition, one sees that although the observable item administrations and responses are truncated at trial n , the effects of the items administered at trial n and responses to these items may be carried along to modify the distribution on the set of conditioning states at trial $n + 1$, C_{n+1} .

Let the sample space for such an experiment be called $Z_{e_{k,n}}$ where

$$Z_{e_{k,n}} = \left(X, \Omega, P \left\{ \cdot \mid \underline{\omega}, e_{k,n} \right\} \right).$$

This triple consists of: (1) the outcome space X ; (2) the parameter set Ω whose elements $\underline{\omega}$ are vector-valued parameters which govern the changes in the distributions on the conditioning states and the response distributions; and

(3) the probability distribution $p\left\{ \cdot \mid \underline{\omega}, \underline{e}_{k,n} \right\}$ on the sequences $\underline{x}$ given a parameter point $\underline{\omega}$ and the restriction of the outcome space X by the experiment $\underline{e}_{k,n}$.

It will be convenient to introduce at this point a notation to distinguish the conditioning state components and the observable components of the outcome sequences $\underline{x}$. The symbols that will be used to represent generically these two sub-sequences of components in any outcome sequence $\underline{x}$ will be $\underline{c}$ to represent the sub-sequence of conditioning state components and $\underline{r}$ to represent the sub-sequence of item administrations and responses, thus: $\underline{x} = (\underline{c}, \underline{r})$.

Frequently, the values of the sequence of conditioning functions in stimulus sampling models will form a Markov chain; but it has been noted that the manner of definition of the set of experiments E_n generally will not permit a simple definition of a 1st-order Markov chain on the conditioning states in the present problem. However, it is possible to utilize some of the matrix theory associated with the theory of finite Markov chains to simplify the representation of the distribution of response sequences. For this reason, the vector of initial probabilities of being in the various states of the conditioning function will be defined and two matrices of probabilities of transition from state to state will be defined.

Let the vector of initial state probabilities be called $\underline{P}'_1$, that is,

$$\underline{P}'_1 = \left[p\left\{ (\mathcal{C}_A, \mathcal{C}_B)_1 \right\}, p\left\{ (C_A, \mathcal{C}_B)_1 \right\}, p\left\{ (\mathcal{C}_A, C_B)_1 \right\}, p\left\{ (C_A, C_B)_1 \right\} \right].$$

The transitions from state to state will be governed by four condition-
ing rate parameters, say, $\theta_A, \theta_{AB}, \theta_B, \theta_{BA}$. These parameters constitute the non-
zero entries in the two transition matrices which will be called P_A and P_B .
The P_A matrix applies to those trials where an A item is used and conversely
the P_B matrix applies to B item trials. The structures of these two matrices
are shown below:

State at Trial $t + 1$

$(\tilde{C}_A, \tilde{C}_B) (C_A, \tilde{C}_B) (\tilde{C}_A, C_B) (C_A, C_B)$

State at Trial t

$$P_A = \begin{matrix} & \begin{matrix} (\tilde{C}_A, \tilde{C}_B) \\ (C_A, \tilde{C}_B) \\ (\tilde{C}_A, C_B) \\ (C_A, C_B) \end{matrix} & \begin{bmatrix} 1-\theta_A & \theta_A & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1-\theta_{AB} & \theta_{AB} \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix}$$

and

State at Trial $t + 1$

$(\tilde{C}_A, \tilde{C}_B) (C_A, \tilde{C}_B) (\tilde{C}_A, C_B) (C_A, C_B)$

State at Trial t

$$P_B = \begin{matrix} & \begin{matrix} (\tilde{C}_A, \tilde{C}_B) \\ (C_A, \tilde{C}_B) \\ (\tilde{C}_A, C_B) \\ (C_A, C_B) \end{matrix} & \begin{bmatrix} 1-\theta_B & 0 & \theta_B & 0 \\ 0 & 1-\theta_{BA} & 0 & \theta_{BA} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix}$$

The four conditioning rate parameters may be interpreted as follows: the parameter θ_A represents the rate of transition from $\tilde{C}_A$ to C_A when Concept B has not been mastered and a similar interpretation is given to θ_B ; the parameter θ_{AB} represents the rate of transition from $\tilde{C}_A$ to C_A when Concept B has been mastered; and the parameter θ_{BA} represents the rate of transition from $\tilde{C}_B$ to C_B when Concept A has been mastered. Minimal restrictions on these parameters are that $0 < \theta_A \leq \theta_{AB} \leq 1$ and $0 < \theta_B \leq \theta_{BA} \leq 1$.

The parameter space Ω will include coordinates that serve as parameters of either the response distributions or of the total stochastic process defined on the outcome sequences $\underline{x}$. A representative parameter point $\underline{\omega}$ will be defined to be a 10-tuple of the following form:

$$\underline{\omega} = \left[\xi_A, \xi_B, \theta_A, \theta_{AB}, \theta_B, \theta_{BA}, P\left\{(\tilde{C}_A, \tilde{C}_B)_1\right\}, P\left\{(C_A, \tilde{C}_B)_1\right\}, P\left\{(\tilde{C}_A, C_B)_1\right\}, P\left\{(C_A, C_B)_1\right\} \right].$$

One can develop the joint probability distribution on the outcome space of a stochastic process such as this teaching experiment in many ways. Since the teaching experiment is a truncated experiment, this joint distribution is defined on only a finite-dimensional domain and the representation of the joint distribution is straightforward. In that part of stochastic process theory which deals with discrete state and time-parameter sets one usually regards the individual random sequences as elementary events. In principle, probabilities may then be assigned to each elementary event in the sample space and probabilities of more general events of interest are derived from the probabilities of the individual outcome sequences.

Typically, it is conceptually too difficult to assign probabilities directly to the individual outcome sequences of a process such as this teaching experiment. The joint distribution will often be constructed from certain marginal and conditional probabilities. To illustrate the development of the joint probability distribution on the outcome space of this teaching experiment a representative sequence will be considered. For example, the following sequence is an elementary event or point in the outcome space of the experiment $e_{2,3}$:

$$\left[(\tilde{c}_A, \tilde{c}_B, A, R_A)_1, (\tilde{c}_A, \tilde{c}_B, B, R_B)_2, (\tilde{c}_A, \tilde{c}_B, B, R_B)_3, (\tilde{c}_A, \tilde{c}_B)_4 \right].$$

Consider then the probability of this sequence given the experiment $e_{2,3}$ and a parameter point ω_0 ; i.e., the conditional probability

$$P\left\{\left[(\tilde{c}_A, \tilde{c}_B, A, R_A)_1, (\tilde{c}_A, \tilde{c}_B, B, R_B)_2, (\tilde{c}_A, \tilde{c}_B, B, R_B)_3, (\tilde{c}_A, \tilde{c}_B)_4 \right] \mid \omega_0, e_{2,3} \right\}.$$

The computation of the probability of this sequence will be sketched below and some brief remarks will be made in justification of various steps of the computation. From the multiplication law for the probability of joint events, one can write that

$$\begin{aligned}
& P\left\{(\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B, B, R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&= P\left\{(R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&\quad \cdot P\left\{(B)_3 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B)_3, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&\quad \cdot P\left\{(R_B)_2, (\tilde{C}_A, \tilde{C}_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B)_2, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&\quad \cdot P\left\{(B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B)_2, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&\quad \cdot P\left\{(R_A)_1, (\tilde{C}_A, \tilde{C}_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&\quad \cdot P\left\{(A)_1 \mid (\tilde{C}_A, \tilde{C}_B)_1, \underline{\omega}_0, \underline{e}_2, 3\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_1 \mid \underline{\omega}_0, \underline{e}_2, 3\right\}.
\end{aligned}$$

Consider, in order, the evaluation of the terms on the right-hand side of this expression:

$$\begin{aligned}
& P\left\{(R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_2, 3\right\} \\
&= P\left\{(R_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_2, 3\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_2, 3\right\}
\end{aligned}$$

(by the assumptions, for these stimulus sampling models, that probabilities of responses and next states of the conditioning function are dependent only on the current state of the conditioning function, and that conditional probabilities given the current state of the condition function of current responses and next states of the conditioning function are independent).

$$\begin{aligned} P\{(B)_3 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B)_3, \omega_0, e_{2,3}\} \\ = P\{(B)_3 \mid (A, R_A)_1, (B, R_B)_2, \omega_0, e_{2,3}\} \end{aligned}$$

(by the definition of the experiment $e_{2,3}$).

By similar arguments one has that

$$\begin{aligned} P\{(R_B)_2, (\tilde{C}_A, \tilde{C}_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B)_2, \omega_0, e_{2,3}\} \\ = P\{(R_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, B)_2, \omega_0, e_{2,3}\} P\{(\tilde{C}_A, \tilde{C}_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, B)_2, \omega_0, e_{2,3}\} \\ P\{(R_A)_1, (\tilde{C}_A, \tilde{C}_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \omega_0, e_{2,3}\} \\ = P\{(R_A)_1 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \omega_0, e_{2,3}\} P\{(\tilde{C}_A, \tilde{C}_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \omega_0, e_{2,3}\} \end{aligned}$$

and

$$P\{(B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B)_2, \omega_0, e_{2,3}\} = P\{(B)_2 \mid (A, R_A)_1, \omega_0, e_{2,3}\}.$$

In summary, when these simplifying values are substituted into the expression for the conditional probability of the complete sequence, one obtains the result that

$$\begin{aligned}
 & P\left\{(\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B, B, R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &= P\left\{(R_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_{2,3}\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &\quad \cdot P\left\{(B)_3 \mid (A, R_A)_1, (B, R_B)_2, \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &\quad \cdot P\left\{(R_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, B)_2, \underline{\omega}_0, \underline{e}_{2,3}\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, B)_2, \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &\quad \cdot P\left\{(B)_2 \mid (A, R_A)_1, \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &\quad \cdot P\left\{(R_A)_1 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \underline{\omega}_0, \underline{e}_{2,3}\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_2 \mid (\tilde{C}_A, \tilde{C}_B, A)_1, \underline{\omega}_0, \underline{e}_{2,3}\right\} \\
 &\quad \cdot P\left\{(A)_1 \mid \underline{\omega}_0, \underline{e}_{2,3}\right\} P\left\{(\tilde{C}_A, \tilde{C}_B)_1 \mid \underline{\omega}_0, \underline{e}_{2,3}\right\}.
 \end{aligned}$$

To evaluate the probability of this sequence, one uses the values given explicitly in the parameter point $\underline{\omega}_0$ and other values implied by the definition of the experiment $\underline{e}_{2,3}$; thus,

$$P\left\{(\tilde{C}_A, \tilde{C}_B, A, R_A)_1, (\tilde{C}_A, \tilde{C}_B, B, R_B)_2, (\tilde{C}_A, \tilde{C}_B, B, R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid \underline{\omega}_0, \underline{e}_{2,3}\right\}$$

$$\begin{aligned}
&= g_{B,o}(1-\theta_{B,o}) \cdot 1 \cdot g_{B,o}(1-\theta_{B,o}) \cdot 1 \cdot g_{A,o}(1-\theta_{A,o}) \cdot 1 \cdot p_o\{\tilde{c}_A, \tilde{c}_B\}_1\} \\
&= g_{B,o}^2 g_{A,o} (1-\theta_{B,o})^2 (1-\theta_{A,o}) p_o\{\tilde{c}_A, \tilde{c}_B\}_1\}.
\end{aligned}$$

This example shows how the probability distribution on the outcome sequences may be evaluated in terms of certain marginal and conditional distributions. The probabilities which will be of chief interest in problems dealing with the design of teaching experiments are probabilities of response sequences and certain marginal probability distributions of the states of the conditioning functions. Matrix operators may be defined which will provide a convenient way to compute these two types of probabilities in a manner very similar to the matrix operator calculation of response probabilities that is employed in linear models of learning [e.g. see Bush and Mosteller, 8].

Conditional probabilities of the various responses, given the current state of the conditioning function, were presented in the illustration of the computation of the probability of an outcome sequence, for example,

$p\{(R_B)_3 | (C_A, C_B)_3, \omega_o, e_{2,3}\}$. Since these probabilities are assumed to be constant over trial numbers, they will be collected into four matrices. Let $D_A, \tilde{D}_A$ and $D_B, \tilde{D}_B$ be the following four diagonal matrices:

$$D_A = \begin{matrix} & (R_A)_t & (R_A)_t & (R_A)_t & (R_A)_t \\ \begin{matrix} (\tilde{C}_A, \tilde{C}_B, A)_t \\ (C_A, \tilde{C}_B, A)_t \\ (\tilde{C}_A, C_B, A)_t \\ (C_A, C_B, A)_t \end{matrix} & \begin{bmatrix} g_A & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & g_A & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix}$$

$$\tilde{D}_A = \begin{matrix} & (\tilde{R}_A)_t & (\tilde{R}_A)_t & (\tilde{R}_A)_t & (\tilde{R}_A)_t \\ \begin{matrix} (\tilde{C}_A, \tilde{C}_B, A)_t \\ (C_A, \tilde{C}_B, A)_t \\ (\tilde{C}_A, C_B, A)_t \\ (C_A, C_B, A)_t \end{matrix} & \begin{bmatrix} 1-g_A & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1-g_A & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix}$$

$$D_B = \begin{matrix} & (R_B)_t & (R_B)_t & (R_B)_t & (R_B)_t \\ \begin{matrix} (\tilde{C}_A, \tilde{C}_B, B)_t \\ (C_A, \tilde{C}_B, B)_t \\ (\tilde{C}_A, C_B, B)_t \\ (C_A, C_B, B)_t \end{matrix} & \begin{bmatrix} g_B & 0 & 0 & 0 \\ 0 & 1-g_B & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix}$$

$$\tilde{D}_B = \begin{matrix} & (\tilde{R}_B)_t & (\tilde{R}_B)_t & (\tilde{R}_B)_t & (\tilde{R}_B)_t \\ \begin{matrix} (\tilde{C}_A, \tilde{C}_B, B)_t \\ (C_A, \tilde{C}_B, B)_t \\ (\tilde{C}_A, C_B, B)_t \\ (C_A, C_B, B)_t \end{matrix} & \begin{bmatrix} 1-g_B & 0 & 0 & 0 \\ 0 & 1-g_B & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix}$$

The principal diagonal entries in these matrices are the probabilities of various responses given the current state of the conditioning function and the type of item administered as specified by the row label. Off-diagonal entries are defined to be zero. The D_A matrix gives conditional probabilities of correct response to A items; the $\tilde{D}_A$ matrix gives conditional probabilities of incorrect responses to A items. The matrices D_B and $\tilde{D}_B$ are similarly defined for trials involving B items.

It has been shown that the conditional probability of a state of the conditioning function at the next trial and a particular response at the current trial given the conditioning state at the current trial and the type of item that was administered may be broken down into a certain product of two conditional probabilities. For example,

$$P\{(R_B)_3, (\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \omega_0, e_{2,3}\} \\ = P\{(R_B)_3 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \omega_0, e_{2,3}\} P\{(\tilde{C}_A, \tilde{C}_B)_4 \mid (\tilde{C}_A, \tilde{C}_B, B)_3, \omega_0, e_{2,3}\}.$$

One will recognize that given the experimental rule $e_{2,3}$, the remaining parts of the specification of the two conditional probabilities on the right-hand side of this expression are contained in the matrices D_B and P_B . In general, the various sets of conditional probabilities of the type given on the left-hand side of the above expression may be computed as the product of one of the diagonal matrices with either the P_A or P_B matrix. The resulting conditional probabilities along with the specification of the experimental rule

and the parameter values ω provide the basic quantities needed to compute the relevant probabilities.

Several more matrices and vectors will be defined to complete a basis for convenient representation of the computation of important probabilities.

Let P_{R_A} , $P_{\gamma_{R_A}}$ be matrix operators or transformation matrices where

$$P_{R_A} = D_A P_A = \begin{bmatrix} g_A(1-\theta_A) & g_A\theta_A & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & g_A(1-\theta_{AB}) & g_A\theta_{AB} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$P_{\gamma_{R_A}} = \tilde{D}_A P_A = \begin{bmatrix} (1-g_A)(1-\theta_A) & (1-g_A)\theta_A & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & (1-g_A)(1-\theta_{AB}) & (1-g_A)\theta_{AB} \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$P_{R_B} = D_B P_B = \begin{bmatrix} g_B(1-\theta_B) & g_B\theta_B & 0 & 0 \\ 0 & g_B(1-\theta_{BA}) & 0 & g_B\theta_{BA} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$P_{R_B} = \tilde{D}_B P_B = \begin{bmatrix} (1-g_B)(1-\theta_B) & (1-g_B)\theta_B & 0 & 0 \\ 0 & (1-g_B)(1-\theta_{BA}) & 0 & (1-g_B)\theta_{BA} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

These four matrix operators permit one to calculate the probabilities of the states of the conditioning function which correspond to the various branches of the trees of the different experiments. For example, experiment $e_{2,3}$ uses an A item at trial 1. The joint probabilities of the states of the conditioning function at trial 2, in conjunction with a correct response, R_A , to the A item at trial 1, are computed by applying the transformation matrix P_{R_A} to the vector of initial state probabilities P_1' ; thus,

$$P\left\{\left[\tilde{C}_A, \tilde{C}_B\right)_2, (C_A, \tilde{C}_B)_2, (\tilde{C}_A, C_B)_2, (C_A, C_B)_2\right], (A, R_A)_1\right\} = P_1' P_{R_A}.$$

If an incorrect response had occurred at trial 1, the joint probabilities of the conditioning states at trial 2, in conjunction with an incorrect response $\tilde{R}_A$, at trial 1, would be computed by applying the matrix $P_{R_A}^v$; i.e.,

$$P\left\{ \left[(\tilde{C}_A, \tilde{C}_B)_2, (C_A, \tilde{C}_B)_2, (\tilde{C}_A, C_B)_2, (C_A, C_B)_2 \right], (A, \tilde{R}_A)_1 \right\} = \underline{P}_1' P_{R_A}^v.$$

Note that P_{R_A} and $P_{R_A}^v$ are not stochastic matrices (P_{R_B} and $P_{R_B}^v$ also are not stochastic matrices); consequently the vectors $\underline{P}_1' P_{R_A}$ and $\underline{P}_1' P_{R_A}^v$ do not represent probability distributions. However, the sum of these two vectors does yield the conditional probability distribution of the states of the conditioning function at trial 2, given that an A item was administered at trial 1, since

$$\underline{P}_1' P_{R_A} + \underline{P}_1' P_{R_A}^v = \underline{P}_1' (D_A + \tilde{D}_A) P_A = \underline{P}_1' I P_A = \underline{P}_1' P_A$$

and P_A is a stochastic matrix.

Definitions of four additional vectors will be introduced here in order to simplify the computation of various response probabilities. Let

$$\underline{G}_A = \begin{bmatrix} \underline{g}_A \\ 1 \\ \underline{g}_A \\ 1 \end{bmatrix} = \begin{bmatrix} P\left\{ (R_A)_t | (\tilde{C}_A, \tilde{C}_B, A)_t \right\} \\ P\left\{ (R_A)_t | (C_A, \tilde{C}_B, A)_t \right\} \\ P\left\{ (R_A)_t | (\tilde{C}_A, C_B, A)_t \right\} \\ P\left\{ (R_A)_t | (C_A, C_B, A)_t \right\} \end{bmatrix}$$

$$\tilde{G}_A = \begin{bmatrix} 1-g_A \\ 0 \\ 1-g_A \\ 0 \end{bmatrix} = \begin{bmatrix} P\left\{(\tilde{R}_A)_t | (\tilde{C}_A, \tilde{C}_B, A)_t\right\} \\ P\left\{(\tilde{R}_A)_t | (C_A, \tilde{C}_B, A)_t\right\} \\ P\left\{(\tilde{R}_A)_t | (\tilde{C}_A, C_B, A)_t\right\} \\ P\left\{(\tilde{R}_A)_t | (C_A, C_B, A)_t\right\} \end{bmatrix}$$

$$\underline{G}_B = \begin{bmatrix} g_B \\ g_B \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} P\left\{(R_B)_t | (\tilde{C}_A, \tilde{C}_B, B)_t\right\} \\ P\left\{(R_B)_t | (C_A, \tilde{C}_B, B)_t\right\} \\ P\left\{(R_B)_t | (\tilde{C}_A, C_B, B)_t\right\} \\ P\left\{(R_B)_t | (C_A, C_B, B)_t\right\} \end{bmatrix}$$

$$\tilde{G}_B = \begin{bmatrix} 1-g_B \\ 1-g_B \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} P\left\{(\tilde{R}_B)_t | (\tilde{C}_A, \tilde{C}_B, B)_t\right\} \\ P\left\{(\tilde{R}_B)_t | (C_A, \tilde{C}_B, B)_t\right\} \\ P\left\{(\tilde{R}_B)_t | (\tilde{C}_A, C_B, B)_t\right\} \\ P\left\{(\tilde{R}_B)_t | (C_A, C_B, B)_t\right\} \end{bmatrix}.$$

Probabilities of sub-sequences of responses along the various branches of an experimental tree may readily be computed using these vectors. For example, in the experiment $e_{2,3}$ the probability of a correct response to the A item at trial 1 is $\underline{P}'_1 \underline{G}_A$ while the corresponding probability of an incorrect response is $\underline{P}'_1 \underline{\tilde{G}}_A$.

Conditional probability distributions on the states of the conditioning function given a response sequence may now be calculated. The conditional distribution of these states at trial 2 given a correct response to the A item at trial 1, for example, is given by

$$P\left\{[(\tilde{C}_A, \tilde{C}_B)_2, (C_A, \tilde{C}_B)_2, (\tilde{C}_A, C_B)_2, (C_A, C_B)_2] \mid (A, R_A)_1\right\} = \underline{P}'_1 P_{R_A} / \underline{P}'_1 \underline{G}_A.$$

Joint probabilities of the states of the conditioning function with higher level sub-sequences of responses are computed by applying the transformation matrix appropriate to the item administration and response that occurred at the current trial. To illustrate, the response sequence for experiment $e_{2,3}$ which involves only correct responses will be considered. One obtains the joint probabilities

$$P\left\{[(\tilde{C}_A, \tilde{C}_B)_3, (C_A, \tilde{C}_B)_3, (\tilde{C}_A, C_B)_3, (C_A, C_B)_3], (A, R_A)_1, (B, R_B)_2\right\} = \underline{P}'_1 P_{R_A} P_{R_B}$$

and

$$P\left\{[(\tilde{C}_A, \tilde{C}_B)_4, (C_A, \tilde{C}_B)_4, (\tilde{C}_A, C_B)_4, (C_A, C_B)_4], (A, R_A)_1, (B, R_B)_2, (B, R_B)_3\right\} \\ = \underline{P}'_1 P_{R_A} P_{R_B} P_{R_B}.$$

The corresponding conditional distributions of the states of conditioning given these response sequences are respectively

$$\begin{aligned}
 & P\left\{ \left[(\tilde{C}_A, \tilde{C}_B)_3, (C_A, \tilde{C}_B)_3, (\tilde{C}_A, C_B)_3, (C_A, C_B)_3 \right] \mid (A, R_A)_1, (B, R_B)_2 \right\} \\
 &= \underline{P}'_1 P_{R_A} P_{R_B} / \underline{P}'_1 P_{R_A} \underline{G}_B
 \end{aligned}$$

and

$$\begin{aligned}
 & P\left\{ \left[(\tilde{C}_A, \tilde{C}_B)_4, (C_A, \tilde{C}_B)_4, (\tilde{C}_A, C_B)_4, (C_A, C_B)_4 \right] \mid (A, R_A)_1, (B, R_B)_2, (B, R_B)_3 \right\} \\
 &= \underline{P}'_1 P_{R_A} P_{R_B} P_{R_B} / \underline{P}'_1 P_{R_A} P_{R_B} \underline{G}_B.
 \end{aligned}$$

These matrix calculations provided a convenient scheme for computing all of the probabilities that were necessary to determine best designs in the illustrations that are taken up in Part 5 of this report. A few additional types of probabilities are required for the determination of the sequential designs in that section; however, the examples of computations that have been given here should be sufficient to indicate the general scheme for computing probabilities of events in these truncated statistical games.

4. Objective, Loss, and Risk Functions in Teaching Experiments

The set of states of nature Ω which was described in Part 3 for the two-concept, stimulus-sampling teaching model consists of multi-dimensional parameter points whose components are initial probabilities of the various states of the conditioning functions, learning rate parameters, and guessing probabilities. If the statistical game that one wished to consider for these teaching experiments were to be concerned with estimating the values of the components of these learning parameters, the set Ω would be the appropriate set of states of nature or set of pure strategies for nature. In a pure estimation problem, Ω would no doubt be represented as a continuous set--probably a hypercube embedded in a multidimensional real space.

Objective Functions

More modest goals seem appropriate when considering the design of many teaching experiments. Frequently, one will be more concerned with the occurrence of events such as the mastery of the more difficult of the two concepts, or the mastery of both concepts, or the occurrence of correct responses rather than with obtaining values for an estimator, say, $\hat{\omega}$ of the parameters of the original game. Hence from the original two-concept teaching game $G = (\Omega, E \times B \times D, \rho)$ several alternative games involving simpler objectives will be "extracted" or perhaps it is better to consider that G will be "restricted" to simpler objectives. The expression of these alternative objectives can be done in several ways; conveniently, one can express these objectives either in terms of certain sets of events in the outcome space X or in terms of "objective functions" defined on these sets of events.

Letting Λ represent generally such sets or classes of objective events in X , one can define the new game restricted by Λ in some instances as $G = (\Lambda \times \Omega, E \times B \times D, \rho)$ and in other instances simply as $G = (\Lambda, E \times B \times D, \rho)$. On the other hand, letting ϕ be an objective function defined on Λ , one can characterize the restricted games in terms of ϕ as either $G = (\phi(\Lambda) \times \Omega, E \times B \times D, \rho)$ or in some cases $G = (\phi(\Lambda), E \times B \times D, \rho)$ where $\phi(\Lambda)$ represents the range space of ϕ . Specific examples of plausible objective functions are given in the following illustrations.

Mastery of both concept A and concept B

In teaching situations involving the presentation of two concepts (or more generally k concepts) one might wish to design a teaching program which in some sense was best for the students mastering both concepts. It appears useful to distinguish two classes of events relating to this objective:

$$\text{let } \Lambda_{(C_A, C_B)}^1 = \left\{ (C_A, C_B)_1^1, (C_A, C_B)_2^1, \dots, (C_A, C_B)_{n+1}^1, (\widetilde{C_A, C_B})_{n+1}^1 \right\}$$

where $(C_A, C_B)_t^1$ represents the event that both concept A and concept B were mastered for the first time at trial t ($t = 1, 2, \dots, n+1$) and $(\widetilde{C_A, C_B})_{n+1}^1$ represents the event that not both concept A and concept B have been mastered by the beginning of trial $n+1$ (the experiment having been truncated at trial n),

$$\text{also let } \Lambda_{(C_A, C_B)} = \left\{ (C_A, C_B)_1, (C_A, C_B)_2, \dots, (C_A, C_B)_{n+1}, (\widetilde{C_A, C_B})_{n+1} \right\}$$

where $(C_A, C_B)_k = \bigcup_{t=1}^k (C_A, C_B)_t^1$ represents the event that both concepts have been mastered by trial k ($k = 1, 2, \dots, n+1$).

These two classes of events seem to be the basic classes related to the mastery of both concepts that are worthy of consideration. In the language of stimulus-sampling theory, mastery of concept A and concept B means transition into the state C_A for the concept A conditioning function and transition into the state C_B for the concept B conditioning function. The class of events $\Lambda_{(C_A, C_B)}^1$ is seen to be a partition of the outcome space X while the class $\Lambda_{(C_A, C_B)}$ consists of the event $(\widetilde{C_A, C_B})_{n+1}$ and a monotone-increasing class of subsets of the complement of $(\widetilde{C_A, C_B})_{n+1}$.

Instead of considering these classes of events, it may be more useful in certain circumstances to consider numerical-valued functions defined on these classes. For example, with respect to the class $\Lambda_{(C_A, C_B)}^1$ it may be desirable to define an objective function such as:

$$\varphi((C_A, C_B)_t^1) = t \quad \text{for } t = 1, 2, \dots, n+1$$

and

$$\varphi((\widetilde{C_A, C_B})_{n+1}) = 0.$$

For the most part, the effects of the restriction of the teaching games by the classes of objective events or objective functions are reflected only in the loss functions for these games. That is, the loss functions L in the restricted games are defined on $\Lambda \times A$ or $\varphi(\Lambda) \times A$ instead of on the product set formed between the parameter set Ω of the original game and the action set A .

Mastery of the more difficult concept B

Another type of interesting teaching objective which might be pursued concerns itself with achieving conditioning of the more difficult of the two concepts, concept B. Mastery of concept A may or may not be pursued as a subgoal insofar as it promotes faster learning of concept B. Two classes of objective events will be defined for this type of a goal. Letting $(\cdot, C_B)_t^1$ be the event that conditioning or mastery of concept B occurred for the first time at trial t ($t = 1, 2, \dots, n+1$) and $(\cdot, \tilde{C}_B)_{n+1}^1$ be the event that concept B has not been mastered after n trials, define the following class of events:

$$\Lambda(\cdot, C_B)^1 = \left\{ (\cdot, C_B)_1^1, (\cdot, C_B)_2^1, \dots, (\cdot, C_B)_{n+1}^1, (\cdot, \tilde{C}_B)_{n+1}^1 \right\}$$

or as an alternative definition of an objective of this general type one might wish to consider the class of subsets of the outcome space defined below:

$$\Lambda(\cdot, C_B) = \left\{ (\cdot, C_B)_1, (\cdot, C_B)_2, \dots, (\cdot, C_B)_{n+1}, (\cdot, \tilde{C}_B)_{n+1} \right\}$$

where $(\cdot, C_B)_t = \bigcup_{t=1}^k (\cdot, C_B)_t^1$.

Weighted response scores

An appealing and frequently used objective in the development of psychometric tests consists of assigning numerical scores to the correct responses and errors that can occur on administration of the items of the test. Applying this objective to the situation of teaching the two concepts A and B, one might define random variables or scoring functions, say, W^* as follows:

$$\text{let } W^*((R_A)_t) = w_A, \quad W^*((\tilde{R}_A)_t) = -w_A$$

$$\text{and } W*((R_B)_t) = w_B, \quad W*((\tilde{R}_B)_t) = -w_B$$

$$\text{where } 0 < w_A \leq w_B$$

The values w_A and w_B could be thought of as the utility to the experimenter of correct responses respectively to concept A items and concept B items. Consequently, following the statistical convention of using non-negative loss functions instead of utility functions, one could define an equivalent score function, say W , in terms of losses as follows:

$$\text{let } W*((R_A)_t) = w_B - w_A, \quad W*((\tilde{R}_A)_t) = w_B + w_A$$

$$\text{and } W*((R_B)_t) = 0, \quad W*((\tilde{R}_B)_t) = 2w_B$$

Although one could perhaps use such a system of response scores to develop an objective for a standard sequential design game involving teaching experiments, this type of setup provides a basis for introducing a related aspect of sequential game theory which has not been dealt with in the outline of sequential design of experiments presented earlier in this paper. The sequential games which were outlined allow the experimenter to make use of a sequence of subexperiments to gain partial information about nature's choices; finally, the experimenter makes a terminal decision and conceptually a payoff is then made in loss units at the end of each single play of the sequential game. Some investigations have been made into the theory for playing sequences of games (see Luce and Raiffa [15] for a survey of various types of situations involving plays-in-sequence of games or sequential compounding of games).

Although there are many similar features found in the theory of sequential games and the theory of sequences of games, it should be evident that the characterization of terminal decision functions especially will vary considerably between the two theories. On the other hand, questions such as whether the sequences of subexperiments or sequences of games will terminate with probability 1 arise in both of the two theories.

The theory of sequences of games will not be elaborated here but it is important to emphasize that certain objectives and their associated payoff functions in teaching experiments are best represented as a sequence of games rather than as a sequential game. This distinction can be illustrated by consideration of the simple urn game that was discussed in Part 2. An analogy between that urn game and the teaching game involving the teaching of the two concepts could be made by representing, say, concept B as the urn U_1 and concept A as the urn U_2 . Further, let the draw of a white marble be interpreted as equivalent to a correct response and the draw of a black marble be equated with an incorrect response. Let a payoff in loss units be defined by the scoring function W .

An example of a sequential game involving the order of presentation of the two urns has been given. Recall that player 1 had two pure strategies--present the urns in the order (U_1, U_2) or in the order (U_2, U_1) . Player 2's terminal actions consisted of these same two identifications of the order of presentation of the urns. Payoffs were defined in terms of whether player 2's choice of a presentation order agreed or disagreed with the choice that player 1 had actually used. Player 2 was also allowed to perform experiments for a

fee in order to gain partial information about the composition of the two urns before making his terminal decision.

As an example of a sequence of these urn games consider the situation described as follows: Let player 1 have four urns. Let two of the urns be identified as concept-B urns with respect to the scoring function W and two as concept-A urns with respect to W . Let the concept-B urns be relabeled as U_{B1} and U_{B2} and suppose that the proportion of white marbles in U_{B1} exceeds the proportion of white marbles in U_{B2} . Let the concept-A urns be similarly relabeled as U_{A1} and U_{A2} and the proportion of white marbles in U_{A1} be greater than the proportion of white in U_{A2} .

The first game in the sequence proceeds in the following way: Player 2 is allowed to request a draw from either a concept-B urn or a concept-A urn but can only specify a choice from one of these two pairs and not a specific one of the four urns. Suppose that player 2 selected the U_B urns, then player 1 is allowed to present either U_{B1} or U_{B2} for sampling. Player 2 draws a marble at random from the presented urn and a payoff is made for the first game as prescribed by the scoring function W . A second game is now played following the same rules, only the composition of the urn presented in the first game is changed by allowing sampling without replacement. A number of variations of rules governing the total sequence of these games could be introduced, e.g., a fixed number of games could be played, player 2 could be charged an entry fee for each game and given an initial purse to gamble with, etc. Roughly, a good strategy for player 2 in these sequences of games would be the selection of type of urn in each new game, given his

history of selections of the types of urns and resulting draws which would minimize his losses for the total sequence.

Similar sequences of games could be developed for teaching situations involving two concepts and a scoring or loss function such as W . In these situations the probabilities of obtaining correct or incorrect response to items of the two types would be characterized, however, by more complicated relationships involving the learning rate parameters θ_A , θ_{AB} , θ_B , and θ_{BA} rather than simply by the transitions determined by sampling without replacement. Good strategies for such sequences of teaching games essentially should minimize the sum over the sequences of the expected values of weighted responses.

Expected trial of first reaching the state, $(\cdot, C_B)$

In our earlier paper [11] dealing with the optimal design of teaching experiments, we defined an optimal strategy as one which minimized the expected trial of first mastering the more difficult concept B or first reaching the state $(\cdot, C_B)$ by appropriate choice of certain probabilities of allocating type A or type B items at each trial, given the item and response at the immediately preceding trial. In terms of sequential game theory, this principle, of designing the teaching experiment to minimize the expected trial of first mastery of concept B, is perhaps best represented as a sequence of games. Clearly, we did not incorporate a sampling plan with its component stopping regions into the strategies considered in that earlier study, and consequently we did not incorporate a terminal decision function of the standard type used in sequential games, either.

The item allocation rules which constituted our basic strategies in that problem were concerned only with the choice of which subexperiment one should continue with. One could think of that design situation as consisting of a sequence of games where player 2 may select the type of game he wishes to play at each trial (choice of an A item or a B item) and he is given the information concerning which type of game he played at the preceding trial and its outcome. The payoff to player 2 could be thought of as t units of loss at trial t if the student reaches the state $(\cdot, C_B)$ at trial t and the payoff to player 2 is zero at trial t otherwise. Consequently, the function

$\sum_{t=1}^{\infty} t p\{(\cdot, C_B)_t\}$ could be interpreted as the experimenter's expected loss in an infinite sequence of these related games.

More general objectives

Several examples of objectives in these two-concept teaching games have been presented. There are a number of other more general objectives which are much more difficult to formulate precisely.

In educational settings it is often desired that the consequences of insufficient training on a set of stimulus materials be expressible in terms of the subsequent rates of learning and performance on similar stimulus materials used in related courses. To make precise such evaluations of losses in terms of degree of transfer of training to other stimulus sets, one would need to expand the mathematical model of the teaching process to include the appropriate parameters that govern the transfer.

One often hears it said that a major reason for attempting to regulate branching in a teaching program in accordance with the general abilities and history of performance of individual students is that such branching procedures promote and sustain the student's motivation to do well in the teaching program. While there probably is a good deal of informal evidence to support this reason for incorporating flexible branching rules in teaching programs, it would, at this stage of our abilities to measure motivational aspects of behavior, be the vaguest sort of speculation to discuss seriously the "optimal design" of teaching experiments to promote and maintain students' motivation levels.

Loss Functions

In the preliminary discussion of objectives and objective functions, it was noted that the effects of the choice of alternative objectives enter the representation of a sequential game through the loss function of the game. That is, the set of pure strategies for nature is restricted by the statement of an objective to either the set of objective events implied by the statement of the objective or to the set of values of an objective function defined on these events.

Again, several illustrations will be given to show, in this case, various plausible loss functions that one might choose to use. For this purpose, it will be sufficient to take up only one of the illustrative objectives: the objective of achieving conditioning for both concept A and concept B.

Losses evaluated in terms of occurrence of mistakes

A simple, plausible loss function that is often used when the set of terminal actions for the experimenter is identical to the set of states of nature consists essentially of assigning a loss of 0 when the terminal action selected agrees with the state of nature and a constant positive loss is assigned to all other pairs of states of nature and terminal actions (mistakes or incorrect actions). For example, let the set of states of nature be given by the class of events $\Lambda_{(C_A, C_B)}^1$ and suppose that the set of terminal actions A is given by $\Lambda_{(C_A, C_B)}$. Losses evaluated in terms of mistakes might be defined in a teaching experiment truncated at n trials as the following loss matrix indicates:

Terminal Actions

	$(C_A, C_B)_1$	$(C_A, C_B)_2$	$(C_A, C_B)_3$	...	$(C_A, C_B)_{n+1}$	$(\widetilde{C_A}, \widetilde{C_B})_{n+1}$
States of Nature						
$(C_A, C_B)_1^1$	0	0	0	...	0	1
$(C_A, C_B)_2^1$	1	0	0	...	0	1
$(C_A, C_B)_3^1$	1	1	0	...	0	1
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$(C_A, C_B)_{n+1}^1$	1	1	1	...	0	1
$(\widetilde{C_A}, \widetilde{C_B})_{n+1}$	1	1	1	...	1	k

The losses in this matrix are considered determined up to a scalar multiplier. The value k in the lower right-hand corner of this matrix will be allowed to range over the interval $[0,1)$.

This loss function might express the payoffs in a no-data game; however, it should be evident that the terminal action $(C_A, C_B)_n$ should always be at least as preferable as any of the terminal actions $(C_A, C_B)_t$ for all $t < n$. Rather than attempting to modify the loss functions to remove such strong inherent determination of the preference ordering on the terminal actions, frequently it will be most suitable to impose restrictions on what values the terminal decision functions may take at various stages of experimentation. Illustrations of such restrictions are given in the design problems which are solved in Part 5.

Absolute error and quadratic loss functions

In some teaching situations it might be desired to evaluate losses associated with terminal decisions by more stringent standards which consider not only the occurrence of errors but also the magnitude of the errors. For example, let the class of events $\Lambda_{(C_A, C_B)}^1$ define an objective and again let the set of terminal actions A be identified with $\Lambda_{(C_A, C_B)}^1$. Suppose that an experimenter expressed his losses in terms of the absolute value of the difference between the trial number which he claimed was the trial when a student first mastered both concepts, say $(C_A, C_B)_t^1$, and the correct trial number when first mastery occurred, say $(C_A, C_B)_t^1$. That is, the absolute errors generally would be defined as $|t-t'|$. The complete loss function defined in terms of absolute errors might have the structure shown in the following loss matrix (perhaps defined up to a scalar multiplier):

Terminal Actions

	$(c_A, c_B)_1^1$	$(c_A, c_B)_2^1$	$(c_A, c_B)_3^1$	...	$(c_A, c_B)_{n+1}^1$	$(\widetilde{c_A, c_B})_{n+1}$
States of Nature						
$(c_A, c_B)_1^1$	0	1	2	...	n	n
$(c_A, c_B)_2^1$	1	0	1	...	n-1	n
$(c_A, c_B)_3^1$	2	1	0	...	n-2	n
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$(c_A, c_B)_{n+1}^1$	n	n-1	n-2	...	0	n
$(\widetilde{c_A, c_B})_{n+1}$	n	n	n	...	n	m

The loss m , which occurs when a correct decision has been made that the student had not learned after n trials, would perhaps best be restricted by a condition such as $0 \leq m < n$.

In some situations it might be more appropriate to treat the errors whose losses are shown, below the principal diagonal of this matrix, differently from corresponding values above this diagonal. For example, the loss matrix shown below might be more suitable than the matrix defined symmetrically in absolute errors.

Terminal Actions

	$(C_A, C_B)_1^1$	$(C_A, C_B)_2^1$	$(C_A, C_B)_3^1$	...	$(C_A, C_B)_{n+1}^1$	$(C_A, C_B)_{n+1}$
States of Nature	$\left[\begin{array}{cccccc} 0 & 0 & 0 & \dots & 0 & n \\ 1 & 0 & 0 & \dots & 0 & n \\ 2 & 1 & 0 & \dots & 0 & n \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ (C_A, C_B)_{n+1}^1 & n & n-1 & n-2 & \dots & 0 & n \\ (C_A, C_B)_{n+1} & n & n & n & \dots & n & m \end{array} \right]$					
$(C_A, C_B)_1^1$						
$(C_A, C_B)_2^1$						
$(C_A, C_B)_3^1$						
$\vdots$						
$(C_A, C_B)_{n+1}^1$						
$(C_A, C_B)_{n+1}$						

Frequently, it may be more appropriate and especially may be mathematically more convenient to define losses in terms of the square of the magnitude of the error rather than in terms of absolute error. Loss functions whose values are defined in this manner are usually called quadratic loss functions. Obviously, if the individual elements in the first of the two illustrative loss matrices shown above were squared, the resulting matrices would constitute listings of all the values of two quadratic loss functions.

Risk Functions

The payoff functions in statistical games are commonly called risk functions. Given a prior distribution $\underline{\pi}$ and an objective Λ for a game involving the sequential design of a teaching program, and letting λ represent

an objective event in a class Λ which defines an objective, a risk function ρ will be defined as follows:

$$\rho(\pi, (\underline{e}, \underline{b}, d)) = \sum_{t=1}^n \sum_{\underline{r} \in b_t} \sum_{\underline{\omega} \in \Omega} \left[c_t(\underline{r}) + L(\lambda, d(t, \underline{r})) \right] p\{\underline{r} | \underline{\omega}, \underline{e}\} \pi\{\underline{\omega}\}$$

where $\underline{e}$ is an experiment (truncated at n trials)

$\underline{b}$ is a sampling plan consisting of stopping regions

b_t ($t = 1, 2, \dots, n$)

d is a terminal decision function

$\underline{r}$ is a response sequence, the observable components of the sequence $\underline{x}$ which comprise the outcome space X

$\underline{\omega}$ is a parameter point

λ is an objective event

c is a cost function

L is a loss function

and $p\{\underline{r} | \underline{\omega}, \underline{e}\}$ is a probability distribution of the outcome sequences

$\underline{r} \in X$ when a parameter point $\underline{\omega}$ and an experiment $\underline{e}$ are specified.

More general risk functions than the above might be required in some decision situations (see Raiffa and Schlaifer [18] for descriptions of such generalizations) but the definition of risk in terms of additive effects of sampling costs and terminal losses is predominant in sequential game theory. In fact, a further specialization of the risk function is commonly employed in many studies of sequential games; the cost function c is considered to depend only on the number of subexperiments used, t , and not on the outcomes

of the subexperiments performed (i.e. not on the values in the sub-sequences $(r_1, r_2, \dots, r_t)$). The restriction of the risk function to the case of constant costs will be used in the development of optimal designs in Part 5.

A Bayes solution for an optimal design of a teaching program with risk function $\rho(\underline{\pi}, (\underline{e}, \underline{b}, \underline{d}))$ can be obtained by finding a pure strategy $(\underline{e}, \underline{b}, \underline{d})$ which will minimize ρ when the prior distribution $\underline{\pi}$ is given. A mathematical programming technique for finding an optimal design in accordance with the Bayes principle will be elaborated in the next part of this paper.

5. Solutions for Best Designs in Some Miniature Teaching Experiments

A technique will be developed in this section of the report to solve for best designs (or best strategies) in teaching experiments. Since the trees of these games rapidly grow a large number of branches when the last trial n is even a modest-sized number, it will be desirable for purposes of exposition to keep n as small as possible. For this reason, 3-trial truncated experiments were chosen for examples as they are just large enough to allow some interesting design features to be revealed and small enough to permit detailed graphing of the overall structure of these experiments.

Other severe simplifications in these examples are made too, for the purpose of promoting clarity of exposition of the results. These examples obviously are not intended as serious efforts to design optimally any experiments for specific teaching situations but hopefully they should illustrate a general technique for solving for best designs. Among the more prominent of the further simplifications that are made are (1) the restriction of consideration to only two types or populations of students and (2) the restriction of the response distributions to the case of no guessing. Again, these restrictions were made for simplicity of exposition of the optimization technique; the removal of these restrictions, for the most part, has little effect on the procedure for solutions for best designs.

Pure Strategies in a 3-Trial Teaching Experiment

Recall from the earlier description of a sequential statistical game that a pure strategy for the statistician is a triple $(\underline{e}, \underline{b}, \underline{d})$ where, in the

case of a teaching experiment, e is an item allocation rule, b is a sampling plan, and d is a terminal decision function. It would appear that in most truncated learning experiments appropriate terminal decision functions should take only the value which indicates a conclusion that conditioning has occurred at any trial prior to the last trial of the truncated experiment. In many non-truncated learning experiments, it will be appropriate to define the set of terminal actions A to consist of but the single conclusion that conditioning has occurred.

In the several examples of solutions for best teaching strategies that are taken up in this section, the objective which is adopted is the teaching of both concepts A and B . Thus, the set of terminal actions will be defined to consist of two values a_c and $a_{\sim c}$ where


a_c represents the conclusion that a subject is in the state (C_A, C_B) and

$a_{\sim c}$ represents the conclusion that a subject is in the complementary state

$$(\widetilde{C_A, C_B}) = (\widetilde{C_A}, \widetilde{C_B}) \cup (C_A, \widetilde{C_B}) \cup (\widetilde{C_A}, C_B)$$

To clarify the restriction imposed on the terminal decision functions in these problems--since these d 's are defined on the product set $I_{n+1} \times X$ (and here I_{n+1} is the set of trial numbers $t \in I_4 = \{1, 2, 3, 4\}$), the restriction on the terminal decision functions is that $d(t, \underline{x}) = a_c$ when $t \leq n$. For $t = n+1$, $d(t, \underline{x})$ may in the present examples take either the value a_c or $a_{\sim c}$.

The tree of a pure strategy for a 3-trial teaching experiment adopting the objective of teaching both concepts is shown in Figure 4 below. In this example the item allocation rule or experiment $e_{2,3}$ is used in conjunction with



Tree of the Pure Strategy, (e_2, y, b_{jo}, d_{ko})

Figure 4

the specific sampling plan b_{j_0} and decision function d_{k_0} . The sampling plan b_{j_0} partitions the outcome space $X_{e_2,3}$ into the following sequence of cylinder sets

$$\begin{aligned} b_{0,j_0} &= \emptyset \text{ (null set)} \\ b_{1,j_0} &= \emptyset \\ b_{2,j_0} &= \{x: [r_1, r_2] = [(A, R_A)_1, (B, R_B)_2]\} \\ b_{3,j_0} &= \{x: [r_1] = [(A, \tilde{R}_A)_1] \cup [r_1, r_2] = [(A, R_A)_1, (B, \tilde{R}_B)_2]\}. \end{aligned}$$

The terminal decision function d_{k_0} must take the value a_c for all outcome sequences x which are elements of b_{2,j_0} . The terminal decision function d_{k_0} further partitions b_{3,j_0} into two subsets. This terminal decision function takes the value a_c for that subset of b_{3,j_0} for which $[r_2, r_3] = [(A, R_A)_2, (B, R_B)_3]$ and the function takes the value $a_{\tilde{c}}$ when its argument is any of the remaining sequences of b_{3,j_0} .

Normal and Extensive Forms of a Game

An obvious way to solve for the best strategy in this type of sequential statistical game would be for each pure strategy to compute its risk against a given prior distribution π , and then to pick out the strategy (or strategies) having the smallest risk. Unfortunately, although there are only a finite number of pure strategies in the games considered here, the number of pure strategies becomes so large so fast with increase in the truncation number n that even with the aid of large high-speed computers, solution by sheer enumeration of risks is almost never feasible.

A considerable reduction in computational effort can be gained by representing these statistical games in an alternative equivalent manner. The characterization of a game as a triple $G = (I, J, \mu_1)$ which was given in Section 2 is referred to as the normal form of a game. Essentially, the normal form of a two-person, zero-sum game consists of listing every pure strategy for player 1 (a pure strategy for player 1 being a complete prescription of what choices he will make at each of his decision points in the game given all the available information of personal choices and random moves made prior to these decision points) in the set I and every pure strategy for player 2 in the set J . The specification of the utility function for player 1, μ_1 , defined on $I \times J$ completes the description of the normal form of such a game.

The representation of a game in extensive form roughly consists of:

- (1) specifying in order each move available at the various stages of the game,
- (2) identifying whether the move is a personal move to be made by one of the players or a random move and identifying which player is to make the choice for personal moves, (3) specifying the set of alternatives available at each move,
- (4) specifying the probability distribution on the sets of alternatives for each random move, and (5) finally giving the numerical-valued payoff to one of the two players for each realizable play of the game.


The extensive form of a game may be diagrammed as a tree also. It is evident that the tree which describes the complete extensive form of a 3-trial teaching experiment will have many more nodes or vertices than does, for example, the tree which illustrates the representative pure strategy $(e_{2,3}, b_{j_0}, d_{k_0})$. However, the feature of the extensive form which makes it the more convenient

representation for solving for best strategies in many problems is that the number of branches to the tree which represents the extensive form is typically a considerably smaller number than the number of pure strategies for the statistician which appear in the normal form.

Space limitations prohibit the diagramming of the complete tree of the extensive form of this teaching game except for a very small number of trials. In fact, it will be sufficient here to illustrate the tree of a single-trial teaching experiment. All pure strategies for the experimenter will be derived from the overall tree.


Several additional notations will be required to graph the extensive form of this teaching experiment. Following common practice, nature will be designated as player 1, the experimenter or statistician will be designated as player 2, and random moves will be assigned to an umpire or player 0. If nature is in the state (C_A, C_B) at any trial, this state will be indicated by the abbreviation C^* . Conversely, if nature is in the state $(\tilde{C}_A, \tilde{C}_B) \cup (C_A, \tilde{C}_B) \cup (\tilde{C}_A, C_B)$, this state will be abbreviated as $\tilde{C}^*$.

The tree of the extensive form of a single-trial teaching experiment of the two-concept type under consideration here is shown in Figure 5. From the tree of the extensive form of this single-trial teaching experiment, 9 trees may be derived which represent all of the pure strategies for the statistician in this case. The trees of these 9 pure strategies are diagrammed in Figure 6 below:



Tree of the Extensive Form of a Single-Trial Teaching Experiment

Figure 5



Trees of the 9 Pure Experimenter's Strategies Derived from Figure 5

Figure 6

Consequently, for this exceedingly simple single-trial teaching experiment there are only 9 pure strategies for the experimenter while a count of the number of terminal vertices of the graph shown in Figure 5 reveals that the tree of the extensive form of the single-trial game has 18 branches. In the single-trial case only, enumeration of the risks associated with each of the experimenter's pure strategies would be the simpler way to determine the best pure strategy. One can readily derive the following two formulas which give respectively the number of branches of the tree of the extensive form of this teaching experiment and the number of pure strategies for the experimenter in this game as a function of the truncation trial number, n :

Number of Branches to Tree of Extensive Form, $\tau_1(n)$

$$\tau_1(n) = \sum_{t=1}^n 2^{2(t-1)+1} + 2^{2(n+1)} \text{ for } n = 1, 2, 3, \dots$$

Number of Experimenter's Pure Strategies, $\tau_2(n)$

$$\tau_2(1) = 9$$

$$\tau_2(n) = 1 + 2(\tau_2(n-1))^2 \text{ for } n = 2, 3, 4, \dots$$

The number of pure strategies for the experimenter has been derived here under the assumption that the experimenter has perfect recall of all of his earlier moves whenever he reaches another choice point. Frequently, the number of pure strategies for a player in an abstract game is computed in a redundant fashion yielding much larger numbers of strategies than the values of $\tau_2(n)$ given here (e.g., see McKinsey [16] for a discussion of ways of counting all

strategies). However, for a rough comparison of the computational efforts required by sheer enumeration of the risks of each pure strategy to find the minimum-risk strategy versus the computational technique which will be presented that considers the extensive form of the game, the values of $\tau_1(n)$ and $\tau_2(n)$ are appropriate.

In Table 1 the values of $\tau_1(n)$ and $\tau_2(n)$ are listed for $n = 1, 2, 3, 4, 5$. It is evident that even $\tau_1(n)$ increases quite rapidly with n but at a rate considerably less than $\tau_2(n)$.

Table 1

n	$\tau_1(n)$	$\tau_2(n)$
1	18	9
2	74	163
3	298	53,139
4	1,194	5,647,506,643
5	4,778	63,788,662,565,458,258,899

Solution for Best Designs By Backward Induction

A general technique is known for finding best designs for sequential experimentation when the set of states of nature Ω and the set of terminal experimenter's actions A are finite. In this case, algorithms can be set up which with the computational assistance of a large high-speed computer permit one to solve for best designs in a number of circumstances. Blackwell and Girshick [4] refer to this technique as "backwards induction." Raiffa and Schlaifer [18] suggest that the procedure might better be called "averaging out and folding back." The folding back stages of the solution are done in accordance with the "principle of optimality" of dynamic programming

[Bellman,3] and indeed the backward induction procedure can be viewed as an application of dynamic programming.

To complete the representation of the extensive form of the single-trial teaching experiment which has been partially depicted by the tree in Figure 5, one must assign the numerical-valued payoff to one of the players of each complete play or branch of the game. Furthermore, a probability distribution may be defined over the branches of the tree by assuming, at the outset of the backwards induction process, that the conditional probability of the experimenter's selection of each alternative at any of his choice points is 1, given the entire history along the path leading to these choice points. As the folding-back process proceeds, these conditional probabilities of selection of an alternative become 1, for the alternative chosen to be best at a particular choice point, and 0, for the remaining alternatives at that point.

The backwards induction technique will be illustrated by considering the tree of the single-trial experiment shown in Figure 5. The objective in this procedure is to find which of the 9 pure strategies depicted in Figure 6 is the best strategy for a particular game. The essentials of the computational process in this simple case can be shown by establishing two tables of values.

In Table 2.1, all of the branches of the tree which represent the utilization of either an A item or a B item experiment are listed (each of these branches may be identified as including 5 nodes or vertices). Along with each play or game sequence, the probability of the sequence is listed. The payoff of each sequence or play of the game to player 2 is listed in column 3 of Table 2.1 as the loss to player 2. For the moment, the values of these

probabilities and losses may be considered to be given numbers. The basis for derivation of these particular values will be indicated later.

The 16 sequences in Table 2.1 have been grouped there into 8 pairs. Each pair represents the choice of one of the 2 alternative terminal actions a_c or $a_{\bar{c}}$ when one of the 4 experimental outcomes $(A, \tilde{R}_A)$, (A, R_A) , $(B, \tilde{R}_B)$, or (B, R_B) has occurred. Within each of these 8 pairs, player 1 or nature has chosen the state $\tilde{C}^*$; however, player 2 does not have precise knowledge of which state player 1 has chosen but only knows the probabilities of each sequence. The losses to player 2, if he chooses a_c or $a_{\bar{c}}$ at "Trial 2" when nature chooses $\tilde{C}^*$ or C , are the values given in column 3. Since player 2 does not know the precise loss he will incur from taking either of his terminal actions in view of any of the four experimental outcomes, he computes his expected loss over the probabilities given in column 2. The resulting 8 values of expected losses are given in column 4. This phase of the process represents an "averaging out" stage.

On the basis of the expected losses, the best actions at the "Trial 2" level of the game are now determined for each of the 4 possible experimental outcomes that could occur. For example, if a type A item had been administered and incorrect response $\tilde{R}_A$ occurred, then the expected loss of taking the action $a_{\bar{c}}$ is .26875 while the expected loss of taking action a_c in this circumstance is .05. Consequently, the best action, or the action which minimizes the expected loss in this case, is a_c . The 4 best actions at "Trial 2" in this game are listed in column 5 of Table 2.1.

The sets of sequences representing the best actions at "Trial 2" are now "folded back" and considered along with terminal actions permitted at Trial 1

in order to complete the determination of the best pure strategy for player 2. In Table 2.2, the events or sets of sequences which correspond to the best actions at "Trial 2" are listed in column 1 along with the only additional terminal action permitted at Trial 1, a_c . The probabilities of these events are given in column 2 of Table 2.2 (the first four values of this column are carried along from Table 2.1 while the last two values may be considered to be given numbers).

In the third column of Table 2.2., losses to player 2 associated with each of the 6 events are listed. The first four values of this column are the expected losses of these best four actions at "Trial 2." The fifth and sixth values in this column are the given values of taking the terminal action a_c when respectively nature chooses the state $\tilde{C}^*$ or C^* .

The 6 events in Table 2.2 have also been grouped into pairs. Within each of the three pairs, it is again necessary to "average out" the losses associated with the choices of player 2. In the case of the pairs $(A, \tilde{R}, a_c)$, (A, R_A, a_c) and $(B, \tilde{R}_B, a_c)$, (B, R_B, a_c) , the expectations are computed over the probability distributions of the responses $\tilde{R}_A, R_A$ and $\tilde{R}_B, R_B$. The expected loss of the two expected losses from the "Trial 2" level associated with the choice of the item A(or B) is computed by merely summing the two higher level expected losses. The expected loss of the action a_c at Trial 1, is computed by averaging the losses of that action given in column 3 over the corresponding probabilities given in column 2.

The best pure strategy for this particular single-trial teaching game may now be identified as the one of the three strategies which has the minimum

9 April 1963

81

TM-1161/000/00

expected loss or risk. The three expected losses in question are given in column 4 of Table 2.2. The strategy with minimum risk is seen to be the strategy which includes the two sets of branches $(A, \tilde{R}_A, a_c) \cup (A, R_A, a_c)$. The tree of this best strategy is tree (4) in Figure 6.

Table 2.1

Branch or Sequence	Probability of Sequence	Loss to Player 2 Associated with Sequence	Expected Loss of Terminal Action at "Trial 2" to Player 2	Best Action at "Trial 2"
$(A, \tilde{R}_A, \tilde{a}_C, \tilde{C}^*)$	.025	.85	.26875	$(A, \tilde{R}_A, a_C)$
$(A, \tilde{R}_A, \tilde{a}_C, C^*)$	.225	1.10		
$(A, \tilde{R}_A, a_C, \tilde{C}^*)$	.025	1.10	.05	
$(A, \tilde{R}_A, a_C, C^*)$	.225	.10		
$(A, R_A, \tilde{a}_C, \tilde{C}^*)$	.25	.85	.7625	(A, R_A, a_C)
$(A, R_A, \tilde{a}_C, C^*)$	.50	1.10		
$(A, R_A, a_C, \tilde{C}^*)$	.25	1.10	.325	
(A, R_A, a_C, C^*)	.50	.10		
$(B, \tilde{R}_B, \tilde{a}_C, \tilde{C}^*)$	.05	.85	.2625	$(B, \tilde{R}_B, a_C)$
$(B, \tilde{R}_B, \tilde{a}_C, C^*)$	.20	1.10		
$(B, \tilde{R}_B, a_C, \tilde{C}^*)$	.05	1.10	.075	
$(B, \tilde{R}_B, a_C, C^*)$	.20	.10		
$(B, R_B, \tilde{a}_C, \tilde{C}^*)$	.25	.85	.7625	(B, R_B, a_C)
$(B, R_B, \tilde{a}_C, C^*)$	.50	1.10		
$(B, R_B, a_C, \tilde{C}^*)$	.25	1.10	.325	
(B, R_B, a_C, C^*)	.50	.10		

Table 2.2

Event	Probability of Event	Loss to Player 2 Associated with Event	Expected Loss of Pure Strategy to Player 2	Best Pure Strategy
$(A, \tilde{R}_A, a_c)$	.25	.05	.375	$(A, \tilde{R}_A, a_c)$
(A, R_A, a_c)	.75	.325		$\cup$ (A, R_A, a_c)
$(B, \tilde{R}_B, a_c)$	.25	.075	.40	
(B, R_B, a_c)	.75	.325		
$(a_c, \tilde{C}^*)$	.50	.1	.50	
(a_c, C^*)	.50	0		

Examples of Best Designs for 3-Trial Teaching Experiments

The single-trial example of a teaching experiment illustrated the general form of the backward induction computational process but it was such an abbreviated example that it did not permit a very interesting look at item allocation problems. Furthermore, it will be desirable to explore the effects that alternative loss functions and prior probability distributions on the parameters have on the structure of best sequential designs. For these reasons, five examples of best designs for 3-trial, two-concept teaching experiments were computed. The results of these computations are described in the remainder of Section 5.

Description of the alternative populations of students

The role of prior probability distributions on the parameter set Ω may be introduced by reference to several classes or populations of students. Thus each point $\underline{\omega} \in \Omega$ will be viewed as determining a population of students. It will be sufficient for the objectives of these examples to restrict Ω to consist of two values, say, $\underline{\omega}_1$ and $\underline{\omega}_2$. The components of $\underline{\omega}_1$ were given values which might be considered representative of a relatively slow-learning population of students, while the components of $\underline{\omega}_2$ were given values that might represent a relatively fast-learning population of students. Recall that each parameter point $\underline{\omega}$ in the two-concept teaching process under consideration has 10 components, i.e.,

$$\underline{\omega} = \left[g_A, g_B, \theta_A, \theta_{AB}, \theta_B, \theta_{BA}, P\left\{(\tilde{c}_A, \tilde{c}_B)_1\right\}, P\left\{(c_A, \tilde{c}_B)_1\right\}, P\left\{(\tilde{c}_A, c_B)_1\right\}, P\left\{(c_A, c_B)_1\right\} \right]$$

In each of the five design problems which follow, these values were chosen to represent the population of slow-learning students, ω_1 , and the population of rapid-learning students, ω_2 :

$$\omega_1 = \begin{bmatrix} 0, 0, .40, .50, .10, .20, .50, .25, .25, 0 \end{bmatrix}$$

and

$$\omega_2 = \begin{bmatrix} 0, 0, .80, .90, .60, .80, 0, .25, .25, .50 \end{bmatrix}$$

The choice of the value 0 for the probabilities of guessing correct responses, g_A and g_B , was made in order to make the resulting best strategies somewhat more intuitively clear. This choice has relatively minor effects on the nature of the computations. Values of the other components of these two parameter points were picked so that the probability distributions on the sequences of these two processes would be well separated. Furthermore, the values of these components were also deliberately chosen at levels which would allow the effects of sequential experimentation to show up even though the experiments are truncated at the small number, 3 trials.

Computation of probabilities of sequences in the extensive form

The matrix operator computational apparatus which was outlined in Section 3 offers a relatively simple method for computing the probabilities of the sequences required for the backward induction solution. The two parameter points ω_1 and ω_2 each provide the values necessary to define a vector of initial state probabilities, two matrices of transition probabilities (one for item A trials

and one for item B trials), and four diagonal matrices of conditional response probabilities. That is, given an ω_h ($h = 1, 2$) the following matrices are defined:

$$P_{1,h}, \quad P_{R_{A,h}} = D_{A,h} P_{A,h}, \quad P_{\tilde{R}_{A,h}} = \tilde{D}_{A,h} P_{A,h},$$

$$P_{R_{B,h}} = D_{B,h} P_{B,h}, \quad P_{\tilde{R}_{B,h}} = \tilde{D}_{B,h} P_{B,h}.$$

In addition to these matrices, two vectors which will determine the probabilities of nature's being respectively in the states $C^* = (C_A, C_B)$ and $\tilde{C}^* = (\tilde{C}_A, \tilde{C}_B) \cup (C_A, \tilde{C}_B) \cup (\tilde{C}_A, C_B)$ will be defined; let

$$\underline{U}' = [0, 0, 0, 1]$$

and

$$\underline{\tilde{U}}' = [1, 1, 1, 0].$$

Finally, let $\underline{\pi}' = [\pi_1, \pi_2]$ be the vector of prior probabilities that nature is in the state ω_1 or ω_2 .

A representative pair of plays from the extensive form of this 3-trial teaching experiment will be considered. For example, the following two plays or sequences are representative of the class of longest plays in this game:

$$[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, R_B)_3, (a_c)_4, (\tilde{C}^*)_4]$$

$$[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, R_B)_3, (a_c)_4, (C^*)_4]$$

If it were certain that a student whose responses were described by the above sequences was from the population ω_1 , then the probability of the first sequence would be given by

$$\underline{P}_{1,1} \tilde{P}_{R_{A,1}} P_{R_{A,1}} P_{R_{B,1}} \underline{\tilde{U}},$$

while the probability of the second sequence would be

$$\underline{P}_{1,1}' \tilde{P}_{R_{A,1}} P_{R_{A,1}} P_{R_{B,1}} \underline{U}.$$

On the other hand, if it were certain that the population ω_2 obtained, the probabilities of the first and second sequences would be given respectively by

$$\underline{P}_{1,2}' \tilde{P}_{R_{A,2}} P_{R_{A,2}} P_{R_{B,2}} \underline{U}$$

and

$$\underline{P}_{1,2}' P_{R_{A,2}} P_{R_{A,2}} P_{R_{B,2}} \underline{U}.$$

These computations represent the situations where the prior distributions are:

$$\pi_{1,1}' = [1,0] \text{ and } \pi_{1,2}' = [1,0].$$

For an arbitrary prior distribution vector $\pi_u' = [\pi_{1,i}, \pi_{2,i}]$, the marginal probability of the first play is obtained by computing the weighted sum

$$\pi_{1,i} \left(\underline{P}_{1,1}' \tilde{P}_{R_{A,1}} P_{R_{A,1}} P_{R_{B,1}} \underline{U} \right) + \pi_{2,i} \left(\underline{P}_{1,2}' \tilde{P}_{R_{A,2}} P_{R_{A,2}} P_{R_{B,2}} \underline{U} \right)$$

and the probability of the second play is obtained by computing the weighted sum

$$\pi_{1,i} \left(\underline{P}_{1,1}' \tilde{P}_{R_{A,1}} P_{R_{A,1}} P_{R_{B,1}} \underline{U} \right) + \pi_{2,i} \left(\underline{P}_{1,2}' \tilde{P}_{R_{A,2}} P_{R_{A,2}} P_{R_{B,2}} \underline{U} \right).$$

These illustrative computations should be sufficient to indicate the general scheme for computing the probability of each play in the extensive form of this game. It should also be evident from these illustrations that the extension of this game from the case of 2 populations to m populations

should not increase the computational effort beyond feasible bounds if m is not too large (perhaps in the range 10 to 25).

Description of the cost function

In these examples of solution for best designs, the simple case of constant cost of experimentation for each trial will be considered. It was found that a cost of .1 unit per trial was suitable to reveal the effects of sequential experimentation in these examples. Since the costs and losses in these problems need to be measured along some common scale, for the purposes of these illustrations, a scale will be used whose unit will arbitrarily be called a utile.

Thus, the complete cost function used in these examples, is defined by the following table:

t	1	2	3	4	.
c(t)	0	.1	.2	.3	

Best designs when losses are expressed symmetrically in terms of errors

The basic loss function which will be considered throughout the five examples to follow is the loss function common to two-action decision problems. In the statistical literature the most frequent case of such a two-action problem is the testing of a simple hypothesis against a simple alternative. The restrictions on the terminal decision functions which have been imposed on these teaching experiments lead to the determination of two loss matrices. At the level of the terminal decisions for "trial $n + 1$ "

(in these examples, "Trial 4") the two actions a_c and $a_{\tilde{c}}$ are permitted. Consequently, at this level the following loss matrix obtains:

State of Nature	Terminal Action	
	a_c	$a_{\tilde{c}}$
C^*	0	1
$\tilde{C}^*$	1	0

On the other hand, only the terminal action a_c has been chosen to be permissible at trials $1, 2, \dots, n$ in this teaching experiment; thus, the loss matrix which is applicable at these trials is the first column of the above matrix.

Two examples of best teaching strategies were computed using this familiar loss function. The principal motivation in solving for these two best strategies was to highlight the fact that one must consider the characterization of the experimenter's losses very carefully otherwise one will arrive at strategies which may be best in terms of minimizing the specified risk but hardly could be considered best in terms of representing an acceptable teaching strategy.

Example 1: Best Design for a Rapid Learners Population

In this first example, the prior distribution $\pi'_{i_2} = [0, 1]$ was assumed. The parameter points ω_1 and ω_2 and the constant cost function $c(\cdot)$ with increment .1 utile are used throughout all five examples. This example and the next example assume the loss function which has been described above.

To evaluate the payoff of each play in the extensive form of this teaching experiment to player 2, one must consider both the sampling cost and the terminal loss associated with the play. These components of the payoff can be collected

into four matrices, say, $L_1(1), L_1(2), L_1(3)$ and $L_1(4)$ which correspond to the sets of plays which respectively use 0 trials, 1 trial, 2 trials, and 3 trials.

These matrices are defined below:

$$L_1(4) = \begin{array}{cc} & \begin{array}{c} a_c \\ a_{\sim c} \end{array} \\ \begin{array}{c} C^* \\ \tilde{C}^* \end{array} & \begin{bmatrix} .3 & 1.3 \\ 1.3 & .3 \end{bmatrix} \end{array}$$

while

$$L_1(3) = \begin{array}{cc} & \begin{array}{c} a_c \\ a_{\sim c} \end{array} \\ \begin{array}{c} C^* \\ \tilde{C}^* \end{array} & \begin{bmatrix} .2 & 1.2 \\ 1.2 & .2 \end{bmatrix} \end{array} \quad L_1(2) = \begin{array}{cc} & \begin{array}{c} a_c \\ a_{\sim c} \end{array} \\ \begin{array}{c} C^* \\ \tilde{C}^* \end{array} & \begin{bmatrix} .1 & 1.1 \\ 1.1 & .1 \end{bmatrix} \end{array} \quad L_1(1) = \begin{array}{cc} & \begin{array}{c} a_c \\ a_{\sim c} \end{array} \\ \begin{array}{c} C^* \\ \tilde{C}^* \end{array} & \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \end{array}$$

The expected payoff to the experimenter of the representative pair of plays, $[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, R_B)_3, (a_c)_4, (C^*)_4]$ and $[(A, \tilde{R}_A)_1, (A, R_A)_2, (B, R_B)_3, (a_{\sim c})_4, (C^*)_4]$ is obtained by computing:

$$1.3 \left(\begin{array}{c} P'_{1,1} \\ P_{R_{A,2}} \\ P_{R_{A,2}} \\ P_{R_{B,2}} \end{array} \begin{array}{c} \tilde{U} \\ U \end{array} \right) + .3 \left(\begin{array}{c} P'_{1,1} \\ P_{R_{A,2}} \\ P_{R_{A,2}} \\ P_{R_{B,2}} \end{array} \begin{array}{c} \tilde{U} \\ U \end{array} \right).$$

This computation represents a typical "averaging out" calculation required in the course of the backwards induction process. It should be evident that these computations would be only slightly more complicated if an arbitrary prior distribution vector had been assumed. The expected payoff would then be computed as:


$$1.3 \left[\pi_{1,1} \left(\begin{array}{c} P'_{1,1} \\ P_{R_{A,1}} \\ P_{R_{A,1}} \\ P_{R_{B,1}} \end{array} \begin{array}{c} \tilde{U} \\ U \end{array} \right) + \pi_{2,1} \left(\begin{array}{c} P'_{1,2} \\ P_{R_{A,2}} \\ P_{R_{A,2}} \\ P_{R_{B,2}} \end{array} \begin{array}{c} \tilde{U} \\ U \end{array} \right) \right]$$

$$+ .3 \left[\pi_{1,1} \left(\begin{matrix} P'_{1,1} & P_{R_{A,1}} & P_{R_{A,1}} & P_{R_{B,1}} & U \end{matrix} \right) + \pi_{2,1} \left(\begin{matrix} P'_{1,2} & P_{R_{A,2}} & P_{R_{A,2}} & P_{R_{B,2}} & U \end{matrix} \right) \right].$$

A computational procedure for solving for the best strategy in these examples was set up and carried out on a desk calculator since the computational effort required did not merit the development of a computer program. The resulting best design for this population of students and this familiar loss function is graphed in Figure 7. Two features of this design require some discussion although in general the design seems to represent a "reasonable" teaching strategy.

The branches $[(B, \tilde{R}_B)_1, (B, R_B)_2, (a_c)_3]$ and $[(B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, R_B)_3, (a_c)_4]$ of the tree of this pure strategy may appear to represent unusual terminal actions at first glance, but inspection of the set of particular values of initial state probabilities in this example confirms the appropriateness of these actions. Since $g_B = 0$ in these examples, the response $\tilde{R}_B$ at Trial 1 indicates that this student must have initially been in the subpopulation whose initial state of conditioning is $(C_A, \tilde{C}_B)$ (the initial probability of the only other applicable state $(\tilde{C}_A, \tilde{C}_B)$ having been assumed to be zero). Therefore, when the correct response R_B occurs at trial 2 or 3, this is errorless evidence that the student is now in the desired state (C_A, C_B) .

The branch $[(B, R_B)_1, (A, \tilde{R}_A)_2, (a_c)_3]$ of this tree will be found to be a situation where there is some probability of being in error by taking this terminal action in the face of the evidence from the responses $(R_B)_1$ and $(R_A)_2$, but the probability of making an error is so small that it does not pay to use another trial. The branch $[(B, R_B)_1, (B, R_B)_2, (B, R_B)_3, (a_c)_4]$ represents



Example 1: Tree of the Best Design

Figure 7


a circumstance where it paid to use all 3 trials but where it is now best to take the terminal action a_c at "Trial 4" since the probability of the students being in the state (C_A, C_B) at "Trial 4" is sufficiently high.

Example 2: Best Design for a Slow Learners Population

The only difference between Example 2 and Example 1 is that the prior distribution vector $\pi_{i_1} = [1, 0]$ is used instead of the vector π_{i_2} . That is, in the present example one assumes at the outset the student is definitely from the slow-learning population instead of from the rapid-learning population.

The tree of the best design for this example is given in Figure 8. Although this is the best strategy for this particular population under the conventional loss function being employed, most educators would probably question the "bestness" of this strategy from the standpoint of appropriate teaching objectives. One might describe this strategy as a "defensive teacher's strategy"--the apparent objective being to give the hardest possible sequence of items in order to conclude with minimum risk that most individuals from this population did not learn the required concepts.

This anomalous result, it can be shown, occurred primarily because of the inappropriateness of the considered loss function for decision problems such as this teaching situation. One must look particularly at the preferences for taking the two terminal actions a_c and $a_{\bar{c}}$ implied by the loss matrix $L_1(4)$. The backward induction solution is started by pairing plays of the game with common histories up to the last vertex of the play, but one member of the pair has the choice $\bar{C}^*$, and the other member has the choice C^* as the value of its



Example 2: Tree of the Best Design

Figure 8

terminal vertex. Let $\underline{h}$ be the generic symbol to stand for a representative history along a play of the game up to the last vertex and then consider the conditional probabilities $p\{(C^*)_4 | \underline{h}\}$ and $p\{(\tilde{C}^*)_4 | \underline{h}\}$. The set of all these conditional probabilities for the representative pair of plays $[\underline{h}, (C^*)_4]$ and $[\underline{h}, (\tilde{C}^*)_4]$ is shown in Figure 9. The graph of the set of these representative pairs of conditional probabilities is shown to be the line connecting the points (0,1) and (1,0) in the plane.

The best terminal action among the pairs of these longest plays in the game is determined by computing the expected loss over the distributions defined on the states C^* and $\tilde{C}^*$ for the actions a_c and $a_{\tilde{c}}$. The best action at the level of the outermost vertices of the class of longest plays is then the action which minimizes this expected loss. Consequently, it can be seen that the loss matrix $L_1(4)$ implies a partition of the sets of conditional probabilities $p\{(C^*)_4 | \underline{h}\}$ and $p\{(\tilde{C}^*)_4 | \underline{h}\}$ into three subsets:

- (1) Region of preference for the action, $a_{\tilde{c}}$

--the terminal action $a_{\tilde{c}}$ will be best when $p\{(C^*)_4 | \underline{h}\} < .5$

- (2) Indifference point

--the terminal actions a_c and $a_{\tilde{c}}$ are equally preferred when

$$p\{(C^*)_4 | \underline{h}\} = .5$$

- (3) Region of preference for the action, a_c

--the terminal action a_c will be best when $p\{(C^*)_4 | \underline{h}\} > .5$




Figure 9

In Example 2, the initial-state distribution and the learning-rate parameters are such that, for practically all of these pairs of plays of the game, the conditional probabilities $p\{(C^*)_4 | \underline{h}\}$ and $p\{(\tilde{C}^*)_4 | \underline{k}\}$ fall in the region of preference for the action a_c . Consequently, the best strategy, so to speak, seeks out the item allocation scheme which will maximize the conditional probabilities $p\{(\tilde{C}^*)_4 | \underline{h}\}$ in the various pairs of longest plays of the game in order to minimize the overall Bayes risk of the design.

Best designs when losses are expressed differentially in terms of errors

It would appear that the preference regions for the terminal actions implied by the definition of the loss matrix $L_1(3)$ do not reflect the preferences for these terminal actions that sound teaching objectives would dictate. It seems safe to say that most educators would prefer to expand the region of preference for the terminal action a_c ; that is, they would prefer to conclude that a student had learned the required concepts unless the probability that the student had mastered the materials was quite small.

A simple modification of the loss matrix $L_1(3)$ will be considered in the next three examples which leads to more appropriate strategies from the standpoint of teaching objectives. Let the preference partition of the sets of conditional probabilities $p\{(C^*)_4 | \underline{h}\}$, $p\{(\tilde{C}^*)_4 | \underline{h}\}$ be prescribed in the following alternative way:

(1') Region of preference for the action a_c

--the terminal action a_c shall be best when $p\{(C^*)_4 | \underline{h}\} < .2$

(2') Indifference point

--the terminal actions a_c and $a_{\tilde{c}}$ are equally preferred when

$$P\{(C^*)_4 | \underline{h}\} = .2$$

(3') Region of preference for the action, a_c

--the terminal action a_c shall be best when $P\{(C^*)_4 | \underline{h}\} > .2$

The choice of the particular alternative indifference point (.2, .8) was made only to illustrate the general direction of change in the preference partition which should be followed.

There are, of course, many ways to alter the loss matrix $L_1(4)$ which will satisfy this alternative preference partition. However, it seems especially suitable to alter only the second column of the matrix to preserve the common definition of losses in taking the action a_c over the trial numbers 1, 2, 3, and 4. Thus, to satisfy this preference partition the loss plus cost matrix $L_2(4)$ will be used, where

$$L_2(4) = \begin{array}{cc|c} & a_c & a_{\tilde{c}} & \\ \hline & .3 & 1.3 & C^* \\ & 1.3 & 1.05 & \tilde{C}^* \end{array}$$

The vectors $L_1(1)$, $L_1(2)$, and $L_1(3)$ in the following three examples are not changed.


Example 3: Best Design for a Rapid Learners Population

The conditions in this problem are identical to Example 1 with the single exception that the $L_2(4)$ matrix is used instead of $L_1(4)$. The best design for

this example turned out to be the same design found to be best for Example 1. Since most of the conditional probabilities $p\{(C^*)_4 | \underline{h}\}$ for the various histories $\underline{h}$ exceeded .5 at the end of the third trial, it should not be too surprising that the best design for Example 3 turns out to be identical to the best design in Example 1. (The detailed illustration of the solution for the best single-trial strategy which was shown in Table 2.1 and Table 2.2 used the parameters, cost function, and loss function of Example 3. One may derive the values in those tables which were presented as given numbers by performing the required computations using the input values defined for Example 3.)

Example 4: Best Design for a Slow Learners Population

The conditions in this problem, other than the substitution of the $L_2(4)$ matrix for the $L_1(4)$ matrix, are identical to the conditions of Example 2. The tree of the best design which resulted for Example 4 is shown in Figure 10. Comparison of the trees in Figure 8 and Figure 10 will reveal that the modification of the loss function carried out in the $L_2(4)$ matrix has had a marked effect on the structure of the best strategy. Foremost among the effects one observes is that the new best strategy employs item allocations which lead, with a single exception, to the terminal conclusion that the student has mastered the concepts, the conclusion a_c . Also the general rule is detected in this strategy that, after starting with a B item, one switches to A items following a correct response, and, for both A and B items, one stays with the same type of item if an incorrect response occurs. This procedure intuitively would seem sound in view of the fact that the guessing probabilities, g_A and g_B , are assumed to have the value 0 in these examples.



Example 4: Tree of the Best Design

Figure 10

Example 5: An Equal Mixture of the Populations ω_1 and ω_2


To further highlight the role of the prior distribution in determining the structure of the best strategy, the vector of prior probabilities, say, $\pi'_1 = [.5, .5]$, was used in this example. The other input values for this example were identical to the conditions of Examples 3 and 4.

The tree of the best design for this example is shown in Figure 11. The item allocation configuration which occurs in this tree does not differ radically from the corresponding configurations of the trees of the best designs for Examples 3 and 4. Roughly speaking, one could also say that the terminal actions for Example 5 agree fairly closely with those of Examples 3 and 4. The sampling plan component of these best pure sequential strategies seems to be the component which has changed the most from example to example.

Incidental remarks about the best strategies in these five examples

It is interesting to note that throughout all five examples the best strategies were initiated by the administration of the more difficult type B item. Surely, this cannot represent a universal rule in these two-concept problems, as one could easily alter this result by appropriate choice of the initial distribution on the four states of the conditioning function. However, in these examples, where in each case the initial probabilities of being in the states $C = (C_A, \bar{C}_B) \cup (C_A, C_B)$ and $C_B = (\bar{C}_A, C_B) \cup (C_A, C_B)$ are equal, the general rule that it is best to allow more trials for exposure to Concept B than to Concept A may obtain.

For the purpose of providing some basis of expressing the reduction in the risk that is gained by following the optimal strategies, the Bayes risks



Example 5: Tree of the Best Design

Figure 11

of the best strategies and the risks that would occur if no experimentation were performed have been listed for each example in Table 3 below:

	Risks	
	No Experimentation	Best Strategy
Example		
1	.50	.232
2	1.00	.3775
3	.50	.232
4	1.00	.86425
5	.75	.560625

Table 3

6. Remarks

Solutions for best designs of teaching experiments by backward induction down the trees of the extensive form of such games were shown in Part 5. These solutions reduce the computational effort required to obtain the Bayes solution or best strategy for these games by an appreciable amount over the effort demanded for sheer enumeration of the risk of each pure strategy. In this final part of the report an outline will be given of the proof that the backward induction technique does yield the strategy which is Bayes against the prior distribution π . The role of sufficient statistics for further reducing the computational effort to solve for Bayes strategies will then be examined briefly. Finally some methods for determining "reasonably good" strategies will be discussed.

Solution by Backwards Induction Yields Bayes Strategy

A formal proof that the backward induction solution technique yields a pure strategy $(\underline{e}, \underline{b}, d)$ which is Bayes against the given prior distribution π can be carried out by straightforward generalization of well-known proofs. These are proofs that the backward induction technique yields the Bayes strategy $(\underline{b}, d)$ in the sequential game involving a fixed experiment $\underline{e}$ (See Blackwell and Girschick [4, Chapter 9] for such a proof). The extension of the proof that the backwards induction technique which yields the sampling plan $\underline{b}$ and terminal decision function d that is Bayes against π for the fixed experiment sequential design problem is outlined below.

It should be noted at the outset that the sequential teaching games under consideration here involve finite sets of terminal actions; consequently, at

each folding-back step in the backward induction process it is possible to select the terminal action $a \in A$ which actually achieves the minimum expected terminal loss. It is the finiteness of the set of pure strategies for nature and of the set of terminal actions which results in a Bayes solution that is a pure strategy.

Since there are only a finite number of pure experiments $\underline{e}$ in the set of all possible pure strategies for these truncated teaching games, one can in principle find the Bayes solution by backward induction for each fixed experiment $\underline{e} \in E$ and then select, as the Bayes solution for the sequential design problem, the strategy $(\underline{e}, \underline{b}, d)$ which has minimum risk against π over the set of all Bayes solutions for each fixed experiment. That is, let $\underline{e}_1, \underline{e}_2, \dots, \underline{e}_m$ be the set of all pure experiments possible in these truncated sequential games, m in number, and for a fixed experiment $\underline{e}_k$ let $(\underline{b}^*, d^*)$ be the sampling plan and terminal decision function which minimizes the risk

$$\rho_{\underline{e}_k}(\pi, (\underline{b}, d)) = \sum_{t=1}^n \sum_{\underline{r} \in \underline{b}_t} \sum_{\underline{\omega} \in \Omega} \left[c_t(\underline{r}) + L(\lambda, d(t, \underline{r})) \right] p_{\underline{e}_k}(\underline{r} | \underline{\omega}) \pi(\underline{\omega}).$$

Thus let the set $\left\{ (\underline{e}_1, \underline{b}^*, d^*), (\underline{e}_2, \underline{b}^*, d^*), \dots, (\underline{e}_m, \underline{b}^*, d^*) \right\}$ represent the set of Bayes solutions for each of the sequential games involving fixed pure experiments $\underline{e}_1, \underline{e}_2, \dots, \underline{e}_m$. Further, let $(\underline{e}^*, \underline{b}^*, d^*)$ be that strategy in this set of strategies which has minimum risk among the members of the set. It is clear that $(\underline{e}^*, \underline{b}^*, d^*)$ is the Bayes solution to the sequential design problem, i.e., this strategy minimizes

$$p(\pi, (\underline{e}, \underline{b}, d)) = \sum_{t=1}^n \sum_{\underline{r} \in \underline{b}_t} \sum_{\underline{\omega} \in \Omega} \left[c_t(\underline{r}) + L(\lambda, d(t, \underline{r})) \right] p(\underline{r} | \underline{\omega}, \underline{e})$$

The extension of the proof that the backward induction solution for fixed experiment sequential games is Bayes against π to proving that the backwards induction solution in the sequential design situation also yields a strategy which is Bayes against the prior distribution π could perhaps be made clearer had we represented the extensive form of the sequential design game by a tree such as the one illustrated in Figure 12. In that figure, a tree of a 2-trial truncated teaching game is graphed. For the two-concept model there are 8 pure experiments possible and each of these includes the possibility of no experimentation. In the tree of Figure 12, the 8 possible pure experiments have been segregated and the experimenter's initial choice set at Trial 1 requires a choice of one of the 8 pure experiments or the choice of the dummy experiment $\underline{e}_0$ which involves no experimentation. The backwards induction solution process can thus be carried out separately for each of the 8 pure experiments, and finally at the level of Trial 1, the choice of the best sequential design is made by seeking among the best strategies in each of the 9 experiments for the strategy with minimum risk.

The trees of the type shown in Figure 5 seemed a more efficient representation of these sequential design problems for computational purposes. The equivalence of the games represented by the two types of trees shown in Figures 5 and 12 is easily shown by reduction of each game to normal form.




Figure 12

Solutions for Best Designs in Terms of Posterior Probabilities

The best designs of teaching experiments which were obtained in Part 5 tell the experimenter what he should do for each sample sequence that may occur in the sample space of the best experiment. In certain problems, it is possible to obtain closed-form solutions for best sequential designs by considering the space of posterior distributions of the parameters, given the various sub-sequences of sample values. For example, when there are only two states of nature and two terminal actions, then the Bayes strategy is equivalent to the sequential probability ratio procedure. Particularly when the components of the sample sequences are independently distributed, the sequential probability ratio test can be readily determined by finding two critical values which partition the unit interval into three regions. These critical values can frequently be solved for by closed-form procedures.

Other sequential games with finite and equal numbers of states of nature and terminal actions plus special types of loss functions can be solved by fairly straightforward methods. In the examples considered in this paper, there did not appear to be any special advantages to seeking for best design in terms of the space of posterior distributions rather than in the sample space of the observable outcomes, because of the loss functions employed and because of the numbers of states of nature distinguished ($n+2$ states). We are currently examining some special cases of sequential strategies in teaching programs, for which it appears that closed-form solutions may be obtainable in terms of decision regions in the space of posterior distributions on the parameters. Results of these studies will be published in subsequent reports.

Role of Sufficient Statistics

In many statistical games it is possible to partition the sample space into sets of outcomes such that all the elements in a given cell of the partition are, in a certain sense, equally informative about the parameters or states of nature. Such partitions of the sample space are called sufficient partitions of the sample space. A standard definition of a sufficient partition, say S , of a sample space X is to call a partition S of X sufficient if for every subset A of X and for every cell or subset s of the partition S for which $p(s) > 0$, the conditional probability $p(A|\underline{\omega}, s)$ is independent of the parameter points $\underline{\omega} \in \Omega$. An alternative, equivalent definition stated in the language of Bayesian theory is that a partition S of the sample space X is sufficient if the posterior distribution on Ω , given that an outcome $\underline{x}$ is an element of a cell s of S , is the same as the posterior distribution, given the actual value of the outcome $\underline{x}$ for every prior distribution π on Ω and for every outcome $\underline{x}$ such that $\sum_{\underline{\omega} \in \Omega} p(\underline{x}|\underline{\omega})\pi(\underline{\omega}) > 0$.

A function or random variable T whose domain is a sufficient partition S of a sample space X is called a sufficient statistic. Since the elements of many sets can be put in one-to-one correspondence with the cells s of S , a sufficient partition determines many sufficient statistics. The range of a sufficient statistic T is a new outcome space generated by T . In many situations, the range space of T has a considerably smaller number of elements than the original outcome space X . When this kind of reduction of the number of relevant outcomes to be considered in the representation of the

tree of the extensive form of a sequential design game is possible through the identification of sufficient statistics, the computational effort to carry out the backward induction solution will be correspondingly reduced.

It is evident that even for models of teaching processes of the simple type that have been considered in this paper, the computation required even by the backward induction process will rapidly become excessive with increase in the truncation trial number n . For some models of teaching process which admit of coarse sufficient partitions of the sample spaces, it may be possible to extend the truncation trial number to a level approaching common practice in certain types of teaching experiments. Many models of learning or teaching processes take infinite-dimensional sets for their outcome spaces. Both the backward induction solution concept and the concept of reduction of the sample space to the range space of a sufficient statistic must be modified in the case of infinite-dimensional outcome sequences.

Raiffa and Schlaifer [18] have very comprehensively examined a large number of statistical decision problems involving some familiar probability distributions by representing these problems as games in extensive form. They do not consider problems involving time-dependent processes, such as the learning process being considered here, but they do make several observations about solving statistical games in extensive form that are pertinent also to these stochastic process problems. These authors emphasize two features of a statistical game which markedly simplify the computations required to determine best strategies by the backward-induction-solution technique:

- (1) The existence of sufficient statistics of fixed, small dimensionality (essentially this condition on a sufficient statistic T means that its range space should have a small number of elements).
- (2) The existence of prior distributions on Ω which have a property that they call "being a natural-conjugate" of the conditional distributions of the sufficient statistics given an experiment e and parameter ω .

Their second condition results in simplification of the computation of the expected losses associated with each individual play of the game or each value of the sufficient statistic. Raiffa and Schlaifer refer to this phase of the solution as terminal analysis. It would appear that the probability distributions on the outcome sequences, or sufficient contractions thereof, of many current learning models are not apt to admit of conjugate prior distributions at least of reasonably simple forms.

Relating Structures of Best Designs to Characteristics of the Parameter Space

The backward induction technique is a very generally applicable method of solving for best designs in teaching experiments; however, without some further analysis of the relationship of the structures of best designs to characteristics of the parameter space of these experiments one would have little basis for predicting general directions of change in the designs as changes are made in the parameter space. For example, some general theory should be developed concerning the relationship of changes in the item-allocation portion of the structure of a strategy as one changes the values of the initial distribution on the states of the conditioning function,

or the values of the various learning-rate parameters, or the values of the probabilities of guessing correct answers.

In the illustrations of designs of teaching experiments given in Part 5 the parameter space Ω was restricted to contain only two elements. Thus in these illustrations the design problem was viewed in part as a simple classification problem (classifying each individual student as a member of one of two populations) along with other objectives concerned with deciding when a student had mastered the concepts in the teaching program. For many practical applications of these design techniques it may be adequate to define the parameter space Ω to consist of a small, finite number of "well-separated" points. In such applications, one would anticipate that an experimenter could assign prior probabilities directly to each element of Ω for every student who is brought into the automated teaching situation.

It may frequently turn out, on the other hand, that the representation of Ω as a small, finite set is not practically adequate nor analytically convenient. One may find, for example, that it is much more appropriate to expand Ω so that the experimenter can specify his prior distributions by simply specifying a few parameters of distributions over Ω . Raiffa and Schlaifer [18] discuss this mode of assignment of prior distributions at length and specify a number of desiderata that prior distributions should satisfy in order to make the solutions for best strategies tractable.

Relating Structures of Best Designs to Form of Loss and Cost Functions

It is unfortunate that the structure of best strategies in statistical games often **vary** quite sensitively with changes in the specification

of the loss and cost functions. Again, it would be desirable to examine these learning or teaching experiment situations to supplement solution techniques with some general analysis of the relationship of directions of changes in the structures of best designs to changes in the loss and cost functions.

Approximations to Solutions When the Numbers of Trials Are Large

If one is faced with solving for a best design in a problem involving a large number of trials where it is not possible to reduce the many branches of the tree to a more feasible number, one might hope to get a good approximation to the best design by considering solutions for best strategies for a number of subproblems consisting of "looking forward" as many trials as computational time and costs will allow.

For example, in problems like those considered in Part 5, if costs of experimentation allowed for continuation through, say, 30 trials but the computation of best strategies could only be afforded for the 3-trial subproblems one could put together an approximation to the best strategy with (at most) 10 pieces or substrategies. In such situations, one would terminate experimentation at any point where the terminal decision was made that the subject had mastered the concepts, a_c , but a continuation for another three trials would be made at any terminal vertex where the decision a_c was made.

In Example 5 of Part 5, it is found that only a single terminal vertex would lead to such a continuation. The sequence associated with that vertex is $[(B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3, (a_c)_4]$. In this circumstance, one could deter-

mine a continuation that would be the best strategy for the next three trials by computing the posterior distribution on Ω given this outcome sequence; e.g.,

$$\pi_1 p\{\omega_1 | (B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3\} = \frac{\pi_1 p\{(B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3 | \omega_1\}}{\pi_1 p\{(B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3 | \omega_1\} + \pi_2 p\{(B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3 | \omega_2\}}.$$

In Example 5 the prior distribution was assumed to be $\pi' = (.5, .5)$; thus when these and the other known values are substituted into the right-hand side of this expression, one obtains

$$p\{\omega_1 | (B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3\} = \frac{.5(.565)}{.5(.565) + .5(.010)} = .983;$$

consequently, $p\{\omega_2 | (B, \tilde{R}_B)_1, (B, \tilde{R}_B)_2, (B, \tilde{R}_B)_3\} = .017$. Using this posterior distribution of the parameters, one could then seek the best strategy for the next three trials by the techniques of Part 5.

In some circumstances, one might also reduce the computations for finding best strategies in truncated teaching experiments involving a large number of trials by redefining a trial to include a block of items. The success of such reductions of the computational difficulties would of course hinge largely on what effects such aggregations had on the relative complexity of the probability distributions of the sequences of item blocks. Other reductions of the complexity of these problems suggest themselves too; prominently, the various objectives which one may try to achieve can be modified in a number of ways to simplify the task of determining best terminal actions.

REFERENCES

1. Albert, A. E., "The sequential design of experiments for infinitely many states of nature," Ann. Math. Stat., Vol. 32 (1961), pp. 774-799.
2. Atkinson, R. C. and Estes, W. K., "Stimulus sampling theory," in Handbook of Mathematical Psychology, Bush, R. R., Galanter, E., and Luce, R. D., Eds., John Wiley and Sons, New York, in preparation.
3. Bellman, R., Dynamic Programming, Princeton University Press, Princeton, 1957.
4. Blackwell, D. and Girshick, M. A., Theory of Games and Statistical Decisions, John Wiley and Sons, New York, 1954.
5. Bower, G. H., "Application of a model to paired-associate learning," Psychometrika, Vol. 26 (1961), pp. 255-280.
6. Bradt, R. N. and Karlin, S., "On the design and comparison of certain dichotomous experiments," Ann. Math. Stat., Vol. 27 (1956), pp. 390-409.
7. Bradt, R. N., Johnson, S. M., and Karlin, S., "On sequential designs for maximizing the sum of n observations," Ann. Math. Stat., Vol. 27 (1956), pp. 1060-1074.
8. Bush, R. R. and Mosteller, F., Stochastic Models for Learning, John Wiley and Sons, New York, 1955.
9. Chernoff, H., "Sequential design of experiments," Ann. Math. Stat., Vol. 30 (1959), pp. 755-770.
10. Chernoff, H. and Moses, L. E., Elementary Decision Theory, John Wiley and Sons, New York, 1959.
11. Dear, R. E. and Atkinson, R. C., "Optimal allocation of items in a simple, two-concept automated teaching model," in Programmed Learning and Computer-Based Instruction, Coulson J., Ed., John Wiley and Sons, New York, 1962.
12. Degroot, M. H. "Uncertainty, information and sequential experiments," Ann. Math. Stat., Vol. 33 (1962), pp. 404-419.
13. Estes, W. K. "Component and pattern models with Markovian interpretations," Chapter 1 in Studies in Mathematical Learning Theory, Bush, R. R. and Estes, W. K. Eds., Stanford University Press, Stanford, 1959.

14. Estes, W. K. and Suppes, P., "Foundations of statistical learning theory. II. The stimulus sampling model for simple learning," Stanford University: Institute for Mathematical Studies in the Social Sciences, Applied Mathematics and Statistics Laboratories, 1959 (Technical Report No. 26, Contract Nonr. 225 (17)).
15. Luce, R. D. and Raiffa, H. Games and Decisions, John Wiley and Sons, New York, 1958.
16. McKinsey, J. C. C., Introduction to the Theory of Games, McGraw-Hill Book Co., New York, 1952.
17. Raiffa, H., "Statistical decision theory approach to item selection for dichotomous test and criterion variables," in Studies in Item Analysis and Prediction, Solomon, H., Ed., Stanford University Press, Stanford, 1961.
18. Raiffa, H. and Schlaifer, R., Applied Statistical Decision Theory, Division of Research Graduate School of Business Administration, Harvard University, Boston, 1961.
19. Robbins, H. E., "Some aspects of the sequential design of experiments," Bull. Amer. Math. Soc., Vol. 58 (1952), pp. 527-535.
20. Schlaifer, R., Probability and Statistics in Business Decisions, McGraw-Hill Book Co., New York, 1959.
21. Suppes, P. and Atkinson, R. C. Markov Learning Models for Multiperson Interactions, Stanford University Press, Stanford, 1960.
22. Wald, A., Sequential Analysis, John Wiley and Sons, New York, 1947.

UNCLASSIFIED

System Development Corporation,
Santa Monica, California
SEQUENTIAL DESIGNS OF EXPERIMENTS FOR
ALTERNATIVE OBJECTIVE FUNCTIONS IN
AUTOMATED TEACHING PROGRAMS.

Scientific rept., TM-1161/000/00, by
R. E. Dear. 9 April 1963, 116p.,
22 refs., 12 figs., 3 tables

Unclassified report

DESCRIPTORS: Education. Decision Making.
Sampling (Mathematics).

Presents the theoretical foundation, in
terms of statistical decision theory,
for the development of branching

UNCLASSIFIED

procedures or designs of automated
teaching programs best tailored to
individual student needs. Attacks
the design problem for automated
teaching programs or experiments from
the standpoint of the theory of the
sequential design of experiments.
Outlines the general theory of the
sequential design of experiments and
the use of Bayesian procedures for
determining best designs. Outlines
the technique of solution for best
sequential designs of experiments called
"backward induction". Discusses the
characteristics that models of teaching
processes need to have in order to be
accessible to computation for best designs
in full-scale teaching programs even
when the backward induction technique is
applied. Emphasizes the critical importance
of coarse sufficient partitions of the
sample space of teaching models.

UNCLASSIFIED

UNCLASSIFIED