

SOME DECISION PROCEDURES FOR USE WITH TAILORED TESTING

MARK D. RECKASE
UNIVERSITY OF MISSOURI

There are many applications of testing technology that require decisions to be made as to whether a person is above or below a criterion score. Criterion-referenced testing and its special case, mastery testing, are examples of such a decision. In the criterion-referenced testing application, it would be especially useful if decisions could be made quickly and conveniently for each student in an individualized instruction program. The recently developed technology of tailored/adaptive testing (Lord, 1970) has the potential to fulfill the requirements of such a testing system. However, there is no generally accepted procedure for making classification decisions using tailored testing, probably because these testing techniques are still relatively new. The few procedures that do exist are either based on randomly sampling items (Epstein, 1978; Sixtl, 1974), which does not take advantage of the power of tailored testing, or on heuristic techniques (Weiss, 1978), which do not have a sound theoretical base. The purpose of this paper is to present some decision procedures that operate sequentially and can easily be applied to tailored testing without loss of any of the elegance and mathematical sophistication of the examination procedures.

Tailored Testing Procedures

Numerous tailored (i.e., adaptive, response contingent, sequential) testing procedures now exist in the research literature, ranging from simple two-stage procedures (Betz & Weiss, 1973) to complex Bayesian procedures (Owen, 1969; see Weiss, 1974, for a good review of the tailored testing procedures that were developed prior to 1974.) Although many procedures exist, for the purposes of this paper only tailored testing procedures using item characteristic curve (ICC) theory and maximum likelihood ability estimation will be considered. It will also be assumed that the tests are administered to the examinees on a computer terminal and that the items are selected to maximize the value of the information function at the previous ability estimate. Despite the narrow definition of tailored testing used for this paper, the results should generalize to any procedure based upon ICC theory.

In applying the decision procedures discussed in this paper, two specific ICC models will be used: the 1- and 3-parameter logistic models. Although any other ICC model could just as easily have been used, these models were selected because of their frequent appearance in the research literature and because of the existence of readily available calibration programs (LOGIST, CALFIT) and tailored testing programs (Reckase, 1974).

Sequential Decision Procedures

A cursory review of the statistical literature indicates that much has been written about sequential estimation and classification procedures. Although somewhat more obscure than ANOVA and regression procedures, most intermediate level mathematical statistics books include at least one chapter on sequential analysis (for example, see Brunk, 1965, chap. 16). In an ongoing review of the extensive literature on this topic, it has been found that most procedures fall into one of three categories: 1) sequential probability ratio tests (SPRT; Wald, 1947), (2) Bayesian sequential procedures (e.g., DeGroot, 1970), and (3) curtailed single sampling plans (Dodge & Romig, 1929). Of these procedures, only the SPRT is narrowly specified--the other two refer to families of procedures rather than a single technique.

Although these statistical procedures are widely applied for quality control, little use has been made of them in the area of mental testing, probably because operable sequential testing procedures did not exist until recently. To date all references in the testing literature to sequential decisions have used the SPRT (Epstein, 1978; Reckase, 1978; Sixtl, 1974). The SPRT will therefore be described first, followed by the Bayesian procedures, since the curtailed sampling plans cannot readily be applied to the commonly used tailored testing procedures, they will not be discussed in this paper.

The Sequential Probability Ratio Test

The sequential probability ratio test (SPRT) was initially developed by Wald (1947) as a quality control device for use by the Armed Forces during World War II. In addition to Wald's (1947) excellent book on the subject, this procedure has been clearly described by Epstein (1978). It will, therefore, be only briefly described here in order to generalize the procedure so that it will more directly apply to tailored testing.

Application to Mastery Decisions

Wald originally developed the SPRT as a statistical test to decide which of two simple hypotheses is more correct. For example, it might be interesting to determine whether a student can answer correctly 60% or 80% of the items in an item pool. The basic philosophy behind the procedure used to decide between these two alternatives was to determine the likelihood of an observed response to an item under the two alternative hypotheses. If the likelihood were sufficiently larger for one hypothesis than the other, that hypothesis would be accepted. If the two likelihoods were similar, another observation would be taken. Wald (1947) has shown that one hypothesis will always be selected over another using a finite set of items.

To demonstrate this procedure, suppose an item is randomly selected from an item pool and administered to a student. If a correct response were obtained, the likelihood under H_1 (80% knowledge) would be .80, and the likelihood under H_0 (60% knowledge) would be .60. To evaluate these likelihoods, Wald takes the ratio of the two.

$$\frac{L(x = 1 | H_1)}{L(x = 1 | H_0)} = \frac{.80}{.60} = 1.67 \quad [1]$$

If the ratio is sufficiently large, H_1 is accepted; if it is sufficiently small, H_0 is accepted; and if it is near 1.0, another observation is taken. The values of this ratio that are considered sufficiently large or small depend upon what is considered acceptable for the two possible decision errors: (1) accepting H_1 when H_0 is true (α error) and (2) accepting H_0 when H_1 is true (β error).

Although Wald (1947) developed a procedure for determining the exact values of these decision points, the procedure is very complex and is seldom used. Instead, good approximations can be determined using the following formulas:

$$\text{lower decision point} = B = \frac{\beta}{1 - \alpha} \quad [2]$$

$$\text{upper decision point} = A = \frac{1 - \beta}{\alpha} \quad [3]$$

Thus, if the likelihood ratio is less than or equal to B , H_0 is accepted with error probability approximately β . If the likelihood ratio is greater than or equal to A , H_1 is accepted with error probability approximately α . If the ratio is between B and A , another item should be randomly sampled and administered and the decision rule implemented again. If $\alpha = .05$ and $\beta = .10$, for example, the decision points would be at $B = .105$ and $A = 18$. Since the likelihood ratio (1.67) is between these two values, no decision would be made, and another item would be selected and administered.

Since the responses to the items follow a binomial distribution in this example, a general expression for the likelihood ratio can be developed for the administration of n items:

$$\begin{aligned} \frac{L(x_1, x_2, \dots, x_n | H_1)}{L(x_1, x_2, \dots, x_n | H_0)} &= \frac{p_1^{\sum x_i} (1 - p_1)^{n - \sum x_i}}{p_0^{\sum x_i} (1 - p_0)^{n - \sum x_i}} \\ &= \left(\frac{p_1}{p_0} \right)^{\sum x_i} \left(\frac{1 - p_1}{1 - p_0} \right)^{n - \sum x_i}, \end{aligned} \quad [4]$$

where

- x_i is the score on item i (0 or 1),
- p_1 is the proportion of i items known by the student in the item pool under H_1 , and
- p_0 is the proportion known in the item pool under H_0 .

If

$$\frac{L(x_1, \dots, x_n | H_1)}{L(x_1, \dots, x_n | H_0)} \geq A, \text{ accept } H_1. \quad [5]$$

If

$$\frac{L(x_1, \dots, x_n | H_1)}{L(x_1, \dots, x_n | H_0)} \leq B, \text{ accept } H_0. \quad [6]$$

Otherwise, continue administering items.

This procedure was originally developed to test simple hypotheses, but Wald (1947) has shown that the procedure operates in the same way for composite hypotheses. For example, suppose it is desirable to know whether a student knew more than some proportion, \underline{p}_1 , of the items in an item pool. In order to use the SPRT to make this decision, a region for which it does not matter which decision is made must first be selected around \underline{p} , say, $\underline{p}_0 < \underline{p} < \underline{p}_1$. If \underline{p}_0 is close to \underline{p}_1 , a very precise decision is required. If \underline{p}_0 and \underline{p}_1 define a wide indifference region around \underline{p} , a rather gross decision rule is all that is needed. The SPRT is then carried out in exactly the same fashion as above, using \underline{p}_0 and \underline{p}_1 as the values for hypotheses H_0 and H_1 , respectively. When the decision points A and B are computed as above, the error rates, α and β , hold for true values of \underline{p} at \underline{p}_0 and \underline{p}_1 . For true values of \underline{p} more extreme than \underline{p}_0 or \underline{p}_1 , the error rates are lower.

Evaluating Outcomes

In order to evaluate the properties of the SPRT, two functions have been derived: the operating characteristic (OC) function and the average sample number (ASN) function. The OC function is defined as the probability of accepting hypothesis H_0 as a function of the true proportion of the item pool known by the student. Although the derivation of the OC function is somewhat complex, the function can be approximated by the following two formulas:

$$p = \frac{1 - \left(\frac{1 - p_1}{1 - p_0}\right)^h}{\left(\frac{p_1}{p_0}\right)^h - \left(\frac{1 - p_1}{1 - p_0}\right)^h} \quad [7]$$

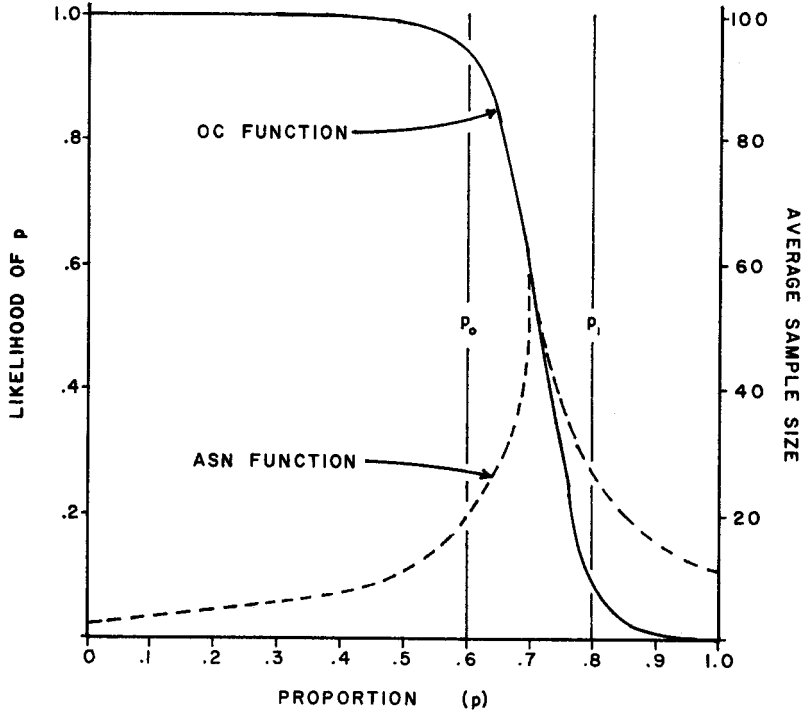
and

$$L(p) \approx \frac{\left(\frac{1 - \beta}{\alpha}\right)^h - 1}{\left(\frac{1 - \beta}{\alpha}\right)^h - \left(\frac{\beta}{1 - \alpha}\right)^h} \quad [8]$$

These equations are used by substituting various arbitrary values of h and solving for \underline{p} and $L(\underline{p})$. $L(\underline{p})$, the probability of accepting H_0 , is then plotted

against p to describe the OC function. Figure 1 shows an OC function for $\alpha = .05$, $\beta = .10$, $p_0 = .6$, and $p_1 = .8$. Note that at $p = p_0$ the height of the curve is equal to $1 - \alpha$, and at $p = p_1$, the height of the curve is equal to β . Note that the OC function is only dependent upon α , β , p_0 , and p_1 . Also, the steeper the curve, the more accurate the SPRT decision rule.

Figure 1
Example of the OC and ASN Functions



The ASN function is defined as the expected number of items required to make a decision at the various values of the true proportion of known items, $E(n|p)$. The formula for the ASN function for the binomial case described above is

$$E(n|p) = \frac{L(p) \ln B + (1 - L(p)) \ln A}{p \ln\left(\frac{p_1}{p_0}\right) + (1 - p) \ln\left(\frac{1 - p_1}{1 - p_0}\right)}, \quad [9]$$

where all of the symbols are as described above and the logarithms are to the base e . Figure 1 also shows the ASN function for the example presented above. Note that the ASN function is highest between the points p_0 and p_1 and that the closer together the values of p_0 and p_1 are, the higher the curve in that region. In general, the lower the ASN curve, the more efficient the decision rule.

Application to Tailored Testing

Although the SPRT as defined above is a valuable procedure for decision-

making in many situations, it makes an implicit assumption that limits its usefulness for tailored testing. The model as presented assumes that the probability of a correct response is the same for all items in the pool. This assumption is reasonable if items are randomly selected and p is the proportion of the items that a student can answer correctly, but it is not reasonable if items are selected to maximize information at an ability level. Under the tailored testing model assumed by this paper, the probability of a correct response changes with each item, requiring a modification of the model.

Fortunately, a detailed analysis of Wald's (1947) work indicates that the sequential random sample assumption is not necessary for the application of the SPRT but is needed only for the derivation of the OC and ASN functions. The SPRT can then be directly applied to tailored testing, but the OC and ASN functions must be determined in a different manner. One approach to determining these functions will be presented later.

To demonstrate the application of the SPRT to tailored testing as defined by this paper, suppose that a tailored test is being used to determine whether a student has exceeded the criterion specified for a criterion-referenced test. Although the method for selecting this criterion is currently not well specified, assume that a value, θ_c , has been determined and that students above this value on the latent achievement scale pass the unit, while those below θ_c are given more instruction.

In order to use the SPRT, a region must be specified around θ_c for which it does not matter whether a pass or a fail decision is made. If high accuracy is desired for the decision rule, a narrow indifference region must be specified, but more items will be required to make the decision. As the region gets wider, the decision accuracy declines, but fewer items are required. Values of θ , θ_0 , and θ_1 mark the boundaries of this indifference region ($\theta_0 < \theta_c < \theta_1$). Once these values have been selected, the likelihood ratio can be defined as

$$\frac{L(x_1, \dots, x_n | \theta_1)}{L(x_1, \dots, x_n | \theta_0)} = \frac{\prod_{i=1}^n P_i(\theta_1)^{x_i} Q_i(\theta_1)^{1-x_i}}{\prod_{i=1}^n P_i(\theta_0)^{x_i} Q_i(\theta_0)^{1-x_i}} \quad [10]$$

where

$L(\underline{x}_1, \dots, \underline{x}_n | \theta_k)$, $k = 0, 1$, is the likelihood of the student's response string of n items administered so far;

x_i is the 0, 1 score on item i ;

$P_i(\theta_k)$ is the probability of a correct response to item i assuming ability θ_k determined from the appropriate ICC model; and

$Q_i(\theta_k) = 1 - P_i(\theta_k)$.

If the 1-parameter logistic model is used as a basis for the tailored test-

ing procedure, Equation 10 becomes

$$\frac{L(x_1, \dots, x_n | \theta_1)}{L(x_1, \dots, x_n | \theta_0)} = \frac{\prod_{i=1}^n \frac{e^{x_i(\theta_1 - b_i)}}{1 + e^{(\theta_1 - b_i)}}}{\prod_{i=1}^n \frac{e^{x_i(\theta_0 - b_i)}}{1 + e^{(\theta_0 - b_i)}}}, \quad [11]$$

where b_i is the difficulty parameter for item i . Equation 11 can be simplified to

$$\frac{L(x_1, \dots, x_n | \theta_1)}{L(x_1, \dots, x_n | \theta_0)} = e^{\sum_{i=1}^n x_i(\theta_1 - \theta_0)} \prod_{i=1}^n \frac{1 + e^{(\theta_0 - b_i)}}{1 + e^{(\theta_1 - b_i)}} \quad [12]$$

The values of this likelihood ratio can then be used to test whether the student is above or below θ_c using the same method presented earlier. If the ratio is greater than $A = \frac{(1 - \beta)}{\alpha}$, the student is classified as being above θ_c ; if it is below $B = \frac{\beta}{(1 - \alpha)}$, the student is classified below the criterion; otherwise, another item is administered. If the 3-parameter logistic model is the basis for the tailored testing procedure, the SPRT procedure is applied in exactly the same manner as above, except that

$$P_i(\theta_k) = c_i + (1 - c_i) \frac{e^{D\alpha_i(\theta_k - b_i)}}{1 + e^{D\alpha_i(\theta_k - b_i)}} \quad [13]$$

is used in Equation 10 instead of the simple logistic form.

The evaluation of the OC and ASN functions cannot be performed as easily as for the simple binomial model due to the presence of the item parameters in the formula for computing the probability of a correct response. Since the item parameters for the next item to be administered are dependent on the item pool used and on the responses to the previous items, the derivation of these functions depends on a complex string of conditional expectations. The conditional probabilities involved make the derivation of these functions, for all practical purposes, impossible. Therefore, the OC and ASN functions can only be approximated using simulation techniques, but these approximations should be adequate for most purposes. Some OC and ASN functions for tailored tests based on the 1- and 3-parameter logistic models will be presented later in this paper. Note, however, that although the full OC function cannot be derived, the value of the function is equal to $1 - \alpha$ at θ_0 and to β at θ_1 , assuming that the item parameters are known. In reality, these two points are not known either, since in all cases except simulations the item parameters are only estimated.

Bayesian Sequential Decision Procedure

The Bayesian decision procedure is an alternative to the SPRT for deciding whether or not a student has exceeded the criterion, θ_c . Although this procedure is much more complicated than the SPRT, it has the capability of using additional information in making the decision. This added information may improve the decision process.

Basic Concepts

Initially, it is assumed that a population of students exists such that each student has some definable achievement level, θ . Individual achievement levels are labeled θ_1 . Each person is to be tested and a decision is to be made concerning placement above or below the criterion. The decision to place above the criterion score is labeled d_1 ; and the decision to place below the criterion score, d_2 .

In order to decide upon a decision rule using Bayesian methodology, three pieces of information are required in advance. These are (1) a prior distribution of θ , (2) a loss function relating the achievement levels to the decisions, and (3) the cost of each observation. Using these three types of information, a decision rule (technique for selecting a decision) and a stopping rule (technique for deciding when a decision should be made) can be determined.

The basic concept used in choosing a decision rule is the concept of risk. Risk is defined as the expected loss, given a decision. Obviously, the decision that minimizes the risk is the desired one. When a Bayesian prior is used, this minimum risk is called the Bayes risk.

The stopping rule used with the Bayesian sequential decision procedure is also based upon the Bayes risk concept. If the expected risk after taking another observation plus the cost of the observation is less than the risk before the observation is taken, the sampling should go on. However, if the expected risk plus the cost of a new observation is greater than the risk without the observation, then sampling should cease. In some cases, it is best not to take any observations at all, because the expected risk plus the cost of an observation is greater than the initial risk of a guess based on the prior distribution of achievement.

Based on this framework, theorems have been proven showing that an optimal procedure exists and that the optimal procedure will reach a decision after some finite number of observations (DeGroot, 1977). If the risk decreases with each observation, the procedure is called a regular sequential decision procedure. Only regular procedures will be considered here, since it is assumed that each item administered yields some positive information rather than providing some misinformation.

Simplified Example

Although this example is not realistic, it demonstrates the basic concepts without requiring complicated mathematical expressions. The extension of the

procedure to realistic situations is direct, but the mathematics is cumbersome. Suppose that two types of individuals exist in the population of interest, those with $\theta_1 = -.8$ and those with $\theta_1 = +.8$ on a latent achievement dimension. A tailored test is to be used to classify the individuals into two groups--those above the criterion score 0.0 and those below. Thus, two decisions are possible: (1) classify as \underline{d}_1 those above the criterion and (2) classify as \underline{d}_2 those below the criterion.

If persons with ability $-.8$ are classified above the criterion, a loss of 25 is incurred in each case. If they are classified below the criterion, there is no loss. If persons with ability $+.8$ are classified above the criterion, there is no loss, whereas a loss of 15 is incurred for each person classified below the criterion. This loss function is summarized in Table 1; it should be noted that these loss function values are totally arbitrary.

Table 1
Loss Function

| Ability (θ_1) | Decision | |
|------------------------|-------------------|-------------------|
| | \underline{d}_1 | \underline{d}_2 |
| $+.8$ | 0 | 15 |
| $-.8$ | 25 | 0 |

Suppose that the prior belief that a randomly selected person has ability $+.8$ is $.6$ and the prior belief that he/she has ability $-.8$ is $.4$. Then, the first step in using a Bayesian sequential decision process is to determine the risk associated with \underline{d}_1 and \underline{d}_2 when no observations are taken. The expected loss (risk) if decision \underline{d}_1 is made is

$$\begin{aligned}
 E(\text{loss} | \underline{d}_1) &= P(\theta_1) \ell(\underline{d}_1 | \theta_1) + P(\theta_2) \ell(\underline{d}_1 | \theta_2) & [14] \\
 &= .4 \times 25 + .6 \times 0 \\
 &= 10,
 \end{aligned}$$

where $P(\theta_1)$ is the prior probability of θ_1 and $\ell(\underline{d}_j | \theta_1)$ is the loss from making decision \underline{d}_j when θ_1 is true. The expected loss (risk) if \underline{d}_2 is made is

$$\begin{aligned}
 E(\text{loss} | \underline{d}_2) &= P(\theta_1) \ell(\underline{d}_2 | \theta_1) + P(\theta_2) \ell(\underline{d}_2 | \theta_2) & [15] \\
 &= .4 \times 0 + .6 \times 15 \\
 &= 9.
 \end{aligned}$$

Thus, the Bayes decision when no observation is taken is \underline{d}_2 , and the Bayes risk is 9. The decision \underline{d}_2 is obviously chosen because it has the lower risk.

Although the proper decision has been determined for the case when no observations have been taken, it has not been determined whether or not an obser-

vation should be taken. To do that, the expected risk after one observation plus cost must be compared to the Bayes risk without an observation. Determining the expected risk after an observation requires several steps, the first of which is determining the posterior distribution of ability after an observation.

Suppose that an item of 0.0 difficulty is administered to a person with ability +.8 or -.8. Depending upon whether the response is correct or incorrect, a Bayesian posterior can be determined using Bayes theorem

$$P(\theta_i | x) = \frac{P(x | \theta_i) P(\theta_i)}{\sum_{i=1}^2 P(x | \theta_i) P(\theta_i)} \quad [16]$$

If a correct response to the item is obtained, the posterior probability of a +.8 ability is given by

$$P(.8 | x = 1) = \frac{P(1 | .8)P(.8)}{P(1 | .8)P(.8) + P(1 | -.8)P(-.8)} \quad [17]$$

The probabilities of an ability of +.8 or -.8 were given in the prior distribution as .6 and .4, respectively. The probability of a correct response, given the known ability, can be determined from the appropriate ICC model. For example, using the 1-parameter logistic model,

$$P(1 | .8) = \frac{e^{(.8 - 0)}}{1 + e^{(.8 - 0)}} = .69 \quad [18]$$

where $P(1 | -.8) = .31$. The posterior probability of +.8 is then $P(.8 | 1) = .77$. Similarly, the posterior probability of -.8 is $P(-.8 | 1) = .23$. The posterior probability of the +.8 and -.8 abilities, given an incorrect response, can likewise be determined using Equation 16. The posterior probabilities, given an incorrect response, are $P(.8 | 0) = .37$ and $P(-.8 | 0) = .63$.

The next step is to determine the risk using the posterior distributions just computed. If a correct response is obtained, the expected loss for \underline{d}_1 is $.23 \times 25 + .77 \times 0 = 5.75$. The expected loss for \underline{d}_2 is $.77 \times 15 + .23 \times 0 = 11.55$. Thus, if a correct response is obtained, the Bayes decision is \underline{d}_1 with a Bayes risk of 5.75. If an incorrect response is obtained, the expected loss for \underline{d}_1 is $.63 \times 25 + .37 \times 0 = 15.75$, while the expected loss for \underline{d}_2 is $.37 \times 15 + .63 \times 0 = 5.55$. Thus, after an incorrect response, \underline{d}_2 is the Bayes decision with a Bayes risk of 5.55.

Since it is not known whether a correct or incorrect response will be given, the expected risk must be computed regardless of the response. To compute the overall expected risk, the probability of a correct and an incorrect response is needed. The probability can be obtained using the following formula:

$$P(1) = P(1 | .8)P(.8) + P(1 | -.8)P(-.8) \quad [19]$$

$$\begin{aligned} &= .69 \times .6 + .31 \times 4 \\ &= .538 \end{aligned}$$

$$P(0) = 1 - P(1) = .462 .$$

The expected risk after a response can now be determined from

$$\begin{aligned} E(\text{risk}|\text{response}) &= E(\text{loss}|1)P(1) + E(\text{loss}|0)P(0) & [20] \\ &= 5.75 \times .538 + 5.55 \times .462 \\ &= 5.66 . \end{aligned}$$

At this point, whether or not another observation should be taken can be determined. If the expected loss after an observation plus cost is greater than the risk before an observation, then administration of items should cease. If the risk before an observation is taken is greater, then another item should be administered. In the example given here, assume the cost of a response is 1 unit. The expected loss after a response plus cost is then $5.66 + 1 = 6.66$. Since the Bayes risk with no items administered was 9, another item should be administered. Depending on the response to the item, decision \underline{d}_1 or \underline{d}_2 could be selected. After the item is administered, the appropriate posterior becomes the new prior and the process continues as above. A flowchart of the entire decision process is presented in Figure 2.

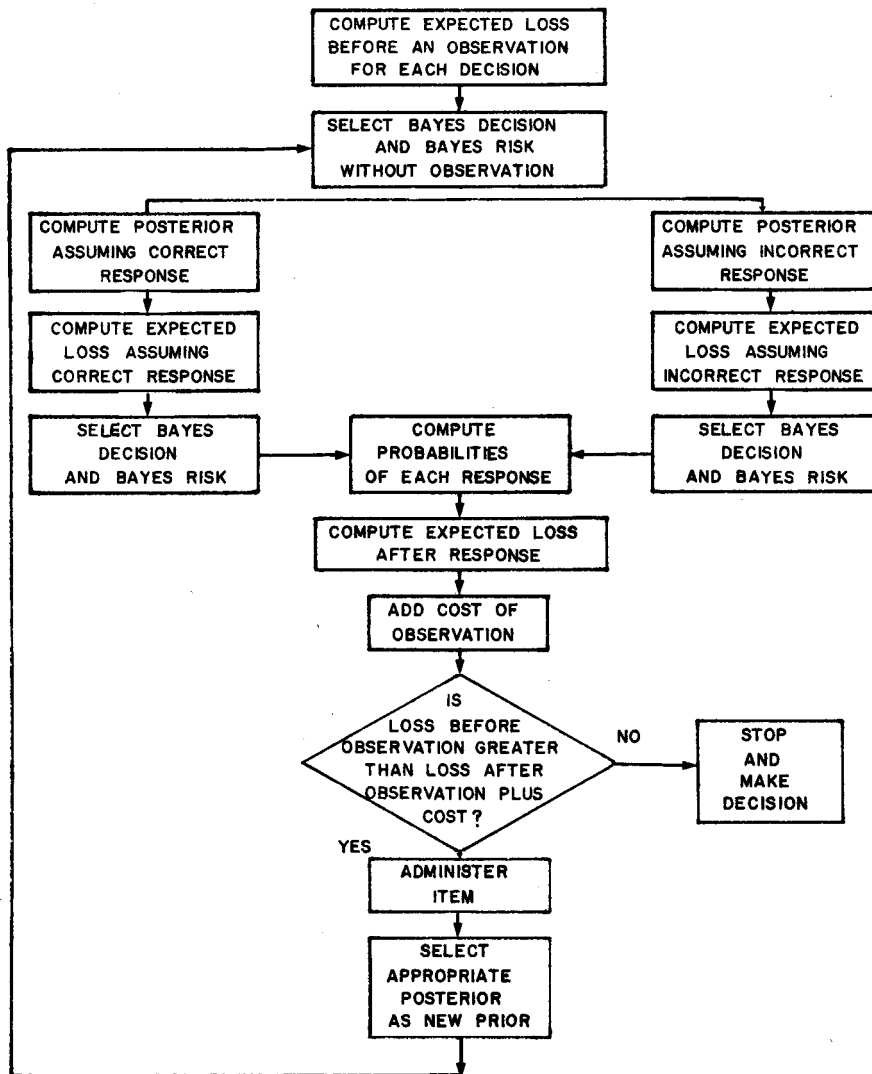
Limitations

Although there are many positive factors in the use of the Bayesian procedure, the very information that makes the control of the testing situation more precise also makes it difficult to implement initially. For example, specifying reasonable loss functions on the same metric as the cost of an observation is difficult for most educational applications. What is the cost of misclassifying persons below the criterion's score when they really should be classified above it? Some attempts have been made by this author to specify loss functions for tailored testing applications, but no satisfactory results have been obtained so far.

A second difficulty in the application of this procedure is in specifying the prior distribution of achievement for a group. This is not as serious a problem as determining loss functions, since performance data are usually available from previous groups. Of course, the more accurate the prior distribution, the more accurate the decision based on the procedure.

It should be realized that the procedure presented here is a simplification of a procedure that would be used for actual tailored testing applications. Achievement levels are usually continuous rather than discrete, as presented here; and the loss due to an incorrect decision is a function of the person's distance from the criterion score rather than a constant value. The procedure can also be modified by changing the cost of observations with increasing test length to allow for fatigue effects. Unfortunately, the Bayesian decision pro-

Figure 2
Flowchart of Bayesian Decision Process



cedure as described here has not yet been implemented in conjunction with an operational tailored testing procedure. Plans are being developed, however, to evaluate an operational version at the Tailored Testing Research Laboratory at the University of Missouri.

Research Design

The purposes of this research were (1) to obtain information on how the SPRT procedure functioned when items were not randomly sampled from the item pool; (2) to gain experience in selecting the bounds of the indifference region, θ_0 and θ_1 ; and (3) to obtain information on the effects of guessing on the accuracy of classification when the 1-parameter logistic model was used.

Tailored Testing Procedure

To determine the effects of these variables, the computation of the SPRT was programmed into both the 1- and 3-parameter logistic tailored testing procedures that were operational at the University of Missouri-Columbia. Since these procedures have been described in detail previously (Koch & Reckase, 1978), they will be merely summarized here. The programs implementing both models used a fixed stepsize method for branching through an item pool until both a correct and an incorrect response had been given. After that point, all ability estimates were obtained using an empirical maximum likelihood estimation procedure. Items were selected for both models to maximize the item information at the previous ability estimate.

To evaluate the decision-making power of the SPRT, subjects with known ability were needed. Therefore, a simulation routine was built into the tailored testing program in place of the responding live examinee. At the beginning of each simulation run, the true ability of the simulated examinee was input into the program. This value was used to determine the true probability of a correct response to the administered items based on the model used (1- or 3-parameter logistic) and the estimated item parameters. A number was then randomly selected from a uniform distribution in the range from 0 to 1. If the randomly selected number was less than or equal to the probability of a correct response, the item was scored as correct. If the randomly selected number was greater than the probability of a correct response, the item was scored as incorrect. This procedure continued for each item in the tailored test.

Tailored tests were simulated 25 times at each true ability using different seed numbers for the random number generator. True abilities from -3 to +3 at .25 intervals were used for both the 1- and 3-parameter models to evaluate the performance of the SPRT. In addition, simulations were run on a composite procedure in which tailored test procedure and the probability ratio calculations (Equation 11) were based on the 1-parameter model, but the item responses were determined by using the 3-parameter model. This was done to determine the effects of guessing on correct classification using the 1-parameter logistic model.

Criterion Values

In computing the probability ratios, three sets of limits of the indifference regions were used: $\underline{+}.3$, $\underline{+}.8$, $\underline{+}1$. A criterion of $\theta_c = 0$ was assumed in all cases. The ratios were computed after each item was administered, and the results were compared to an A value of 45 and a B value of .102. These were determined based on $\alpha = .02$ and $\beta = .10$. A classification was made the first time these limits were exceeded. If the limits were not exceeded before 20 items had been administered (an arbitrary upper limit on test length), the values above 1.0 were classified as above θ_c and the values below 1.0 were classified as below θ_c . This is called a truncated SPRT. At each true ability used for the simulation, the proportion of the 25 administrations classified below θ_c and the average number of items administered were computed. Plots of these values against the true abilities approximate the OC and ASN functions, respectively.

These plots were made for each combination of indifference region and tailored testing method, yielding nine plots of the OC and ASN functions.

Item Pools

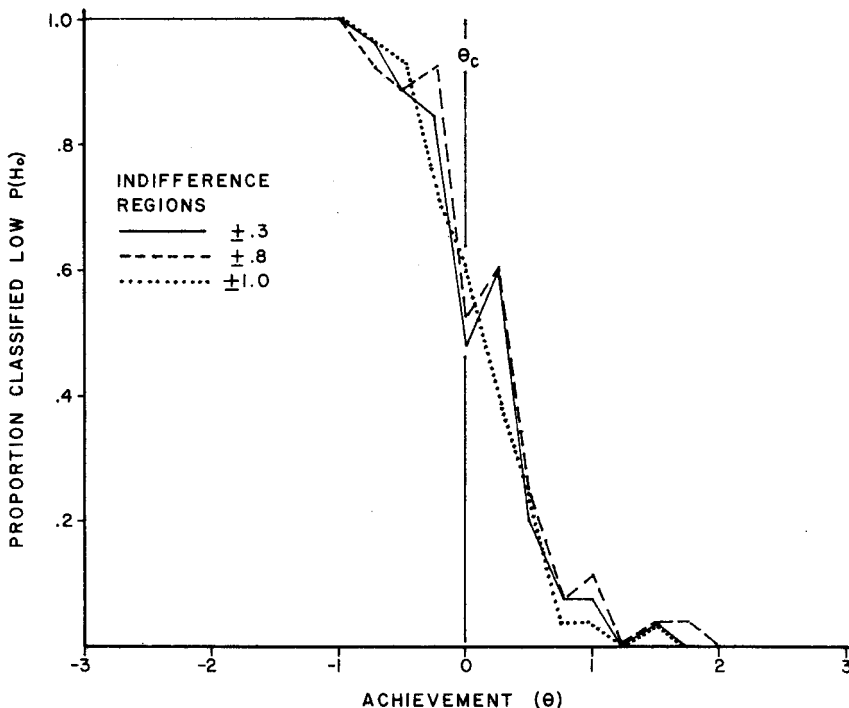
Two different item pools were used for this study. For the analyses using just the 1-parameter or the 3-parameter model, an existing pool of 72 vocabulary items were used. This item pool had an approximately normal distribution of difficulty parameters. For the 1-parameter tailored test using 3-parameter responses, an item pool with 181 items, rectangularly distributed between -3 and +3 in difficulty was used. These simulated items had constant discrimination parameters of .588 (this value yields a 1.0 when multiplied by $D = 1.7$) and a pseudo-guessing parameter of .12. This simulated item pool was selected over the real vocabulary pool to have better control over the guessing parameters. The 1-parameter procedure used only the b -values from the pool.

Results

1-Parameter Model

Figure 3 shows the OC functions for the 1-parameter logistic model based on the vocabulary item pool. The figure shows three graphs, one for each of the $\pm.3$, $\pm.8$, and ± 1 indifference regions. Note that the curves are similar regard-

Figure 3
One-Parameter OC Functions
for Three Indifference Regions



less of the indifference region. The data indicate that in all three cases the classification accuracy was nearly the same.

The values of the curves at the limits of the indifference region give further evaluative information. At the lower point the OC function should pass through $1 - \alpha$. At the -0.3 value the curve is in fact $.85$ when it should be $.98$, showing the degrading effects of restrictive stopping rules used by the tailored testing procedure. At the -0.8 and -1 points for the corresponding curves, the results are about as expected, being $.94$ and 1.00 rather than $.98$.

At the upper limit of the indifference region, the OC function should have a value of $.1$. For the $+0.3$ case it is in fact $.5$ rather than $.1$, again showing the effects of truncating the procedure. At the values of $+0.8$ and $+1$ the values of the OC function were near or better than what they should have been, based on the theoretically expected results.

The ASN functions for the 1-parameter model are given in Figure 4. The curves plotted correspond to the ASN functions, using indifference regions for ± 0.3 , ± 0.8 , and ± 1 . It can immediately be seen that there was a substantial difference in the average number of items needed to reach a decision, with the greatest number required when the indifference region was narrowest. It can also be seen that the largest expected number of items was near the criterion score 0.0 and that the average number dropped off at the extreme abilities. The slight lack of symmetry in the curves is due to the fact that α was not equal to β . For abilities beyond $+1$, an average of only about 3 to 5 items was needed for classification for the wider regions, but 6 to 11 items were needed for the ± 0.3 indifference region. Note that the ± 0.3 curve approached the arbitrary 20-item limit for the tailored tests.

Figure 4
One-Parameter ASN Functions
for Three Indifference Regions

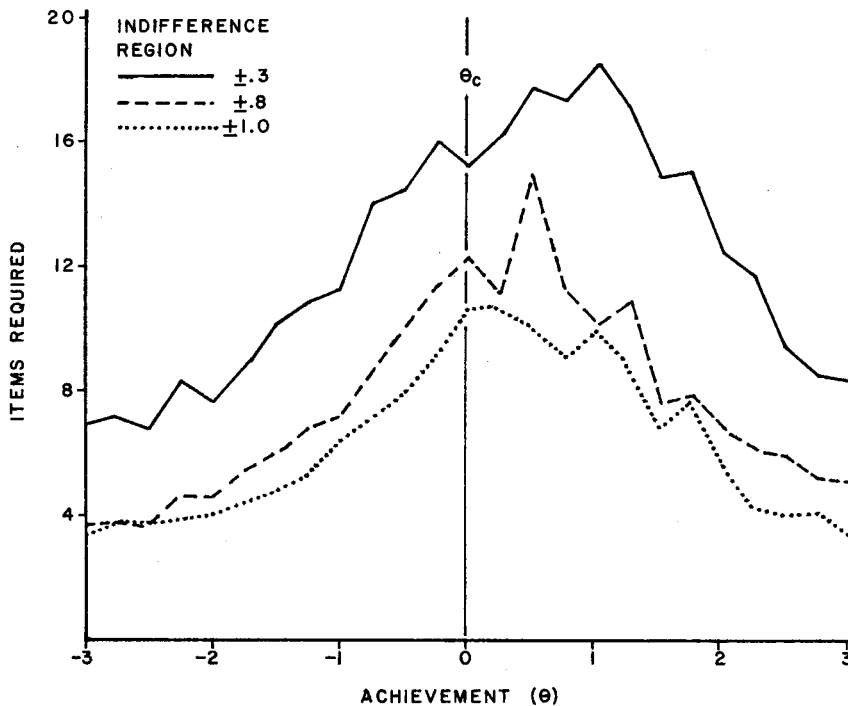



Figure 5 shows, for comparison purposes, the theoretical curves for the ASN and OC functions based on the $\pm .3$ indifference region. An infinite number of items with difficulty 0.0 was assumed for the theoretical functions, and the tests were assumed to have no upper limit on the number of items administered. A comparison of Figures 3 and 4 with Figure 5 shows that the OC curve for the theoretical function is steeper at the cutting point than the simulated curves, and that the ASN function is substantially higher. The difference in the theoretical and simulated OC curves shows the effect of the 20-item stopping rule and the selection of items of differing difficulty.

Figure 5
Theoretical OC and ASN Functions




3-Parameter Model

The results of the simulation of the 3-parameter logistic tailored test are given in Figures 6 and 7. Figure 6 presents the OC functions for the 3-parameter model, again using the indifference regions of $\pm .3$, $\pm .8$, and ± 1 . Notice that as with the 1-parameter model, the OC curves are fairly similar for the three indifference regions throughout most of the range of ability. However, there are discrepancies for the ± 1 indifference range curve near the ± 1 and -1 points, indicating a decline in decision precision for that region. At the $-.3$ value for the $\pm .3$ indifference range, the value of the curve is $.96$, fairly close to the $.98$ theoretical value. At the upper end ($+.3$), however, the value is $.2$ instead of the $.1$ value that it should be. This may show the effects of guessing on the decision process. The $\pm .8$ and ± 1 indifference regions again yield better error probabilities than would be expected from the theory.

The ASN function for the 3-parameter model (Figure 7) also shows similar

Figure 7
Three Parameter ASN Functions
for Three Indifference Regions




eter model was used to determine the responses. The figure shows three graphs, one for each of the $\pm .3$, $\pm .8$, and ± 1 indifference regions. Note that the curves are fairly similar regardless of the indifference region but that they are shifted substantially to the left compared to the previous OC curves. This indicates that the probability of classifying a person below θ_c has dropped off substantially until an ability of about -2 has been reached. In other words, it is much easier to be classified above the criterion score with this procedure than when guessing does not enter into the decision. Instead of being at zero, the effective criterion has been shifted down to -1.5 . Clearly, the values of the OC function at the limits of the indifference region are entirely different from the theoretical values.

The ASN functions for the three indifference regions— $\pm .3$, $\pm .8$, and ± 1 —are shown in Figure 9. The difference between these graphs and those presented in Figure 4 are that the curves are higher (more items were required) and the highest point of the curve is shifted to the steepest part of the OC curve. The relationship between the height of the ASN function and the width of the indifference region still holds; however, as the region gets wider, the average number of items decreases.

Summary and Conclusions

The purpose of this paper has been to describe two procedures for making


Figure 8
Composite OC Functions
for Three Indifference Regions



binary classification decisions using tailored testing--the sequential probability ratio test (SPRT) and a Bayesian decision procedure--and to present some simulation data showing the characteristics of the operation of the SPRT for two ICC models. The first procedure described, the SPRT, was developed by Wald for quality control work. It has not been widely applied for testing applications because the assumption of an equal probability of a correct response was made to facilitate the derivation of the operating characteristic (OC) and average sample number (ASN) functions. Since this assumption can only be met for testing applications by randomly sampling items for administration, the procedure has not been used with tailored testing. In this paper the probability of a correct response was allowed to vary from item to item, although it made the derivation of the OC and ASN functions impossible. Simulation procedures were then used to estimate these functions.

The SPRT procedure described is operational at the Tailored Testing Research Laboratory of the University of Missouri-Columbia in two forms: a live tailored testing procedure and a simulated procedure. The results of the application of the simulation procedure to three studies were described in this paper. The first study estimated the OC and ASN functions for a 1-parameter logistic based tailored testing procedure in which the size of the indifference region around the criterion score was varied. The results of the study showed that the average number of items needed for classification was quite low when the true ability of a simulated person was not too close to the criterion score

Figure 9
Composite ASN Functions
for Three Indifference Regions



and that the width of the indifference region did not greatly affect the OC function. The width of the indifference region did have a substantial effect on the ASN function. The accuracy of classification of the simulated tailored test was not quite as good as administering a large number of items with difficulty values equal to the criterion score. This result was explained by the arbitrary 20-item limit imposed on the tailored test and by the variation in the difficulty parameters of the items administered.

The second study estimated the OC and ASN functions for a 3-parameter logistic tailored testing procedure, also varying the size of the indifference region. The results were similar to those for the 1-parameter model, but even fewer items were generally needed for classification. The results of these first two studies both indicated that the SPRT could be successfully applied to tailored testing.

The third simulation study estimated the OC and ASN functions for the 1-parameter model when guessing was allowed to enter into the responses to the items administered. The results showed that, in effect, guessing lowered the criterion score, making it easier to classify an examinee above the criterion and raising the average number of items needed for classification. This spurious shift in the criterion greatly increased the error rates in classification. The effect was strong enough to preclude the use of the 1-parameter model for classification decisions when guessing is a factor.

The second decision procedure described in this paper allows the use of a greater amount of information in making a decision than the SPRT. The Bayesian procedure includes a prior distribution of student achievement, a loss function for incorrect decisions, and the cost of observations in the development of the decision rule. The basic philosophy of this procedure is to administer items until the expected loss incurred in making a decision is less than the expected loss after the next item is administered plus the cost of administration. At that point a decision is made that minimizes the expected loss. The Bayesian procedure is described in detail, and a simple example is given of its use. The Bayesian procedure is not yet operational for making decisions under tailored testing because appropriate loss functions for educational decisions have not been determined. However, simulation studies of the procedure will commence in the near future.

Both of the decision procedures described in this paper show promise for use in tailored testing. Both also require substantial research effort before they can be applied with confidence. It is hoped that this paper will help to stimulate that research.

REFERENCES

- Betz, N. E., & Weiss, D. J. An empirical study of computer-administered two-stage ability testing (Research Report 73-4). Minneapolis: University of Minnesota, Department of Psychology, Psychometric Methods Program, October 1973. (NTIS No. AD 768993)
- Brunk, H. D. An introduction to mathematical statistics (2nd ed.). New York: Blaisdell, 1965.
- DeGroot, M. Optimal statistical decisions. New York: McGraw-Hill 1970.
- Dodge, H. F., & Romig, H. G. A method of sampling inspection. Bell System Technical Journal, 1929, 8, 613-631.
- Epstein, K. Applications of sequential testing procedures to performance testing. Proceedings of the 1977 Computerized Adaptive Testing Conference. Minneapolis: University of Minnesota, Department of Psychology, Psychometric Methods Program, 1978.
- Koch, W. R., & Reckase, M. D. A live tailored testing comparison study of the one- and three-parameter logistic models (Research Report 78-1). Columbia: University of Missouri, Tailored Testing Research Laboratory, June 1978.
- Lord, F. M. Some test theory for tailored testing. In W. H. Holtzman (Ed.), Computer-assisted instruction, testing, and guidance. New York: Harper & Row, 1970.
- Owen, R. J. A Bayesian approach to tailored testing (Research Bulletin RB-69-92). Princeton, NJ: Educational Testing Service, 1969.

