The Behavior Analyst 1990, 13, 67-77 No. 1 (Spring)

PSI: Two Unorthodox Studies

Donald A. Cook
JHM Corporation

The Personalized System of Instruction (PSI) as developed by Fred S. Keller, J. Gilmour Sherman, and their colleagues, is an application of reinforcement principles to the design of a complex human learning environment, which is both well-defined and has been well-studied. Questions remain, however, concerning the variables that contribute to its effectiveness, the possible limits of that effectiveness, and the measures most useful in studying its operation. An ongoing PSI setting is not merely an application of already stated principles; it is a laboratory for their refinement and the development of new principles—which no longer need to be exclusively "imported" from some other realm of discourse, such as experiments with nonhuman animals or the remote laboratory of the pure researcher. These sources may never run dry, but to them is now added the laboratory of the PSI environment, when it is sufficiently stable yet rich.

In this paper one such PSI course, which evolved over a number of years and reached a fairly stable level of operation, will be described. Two somewhat independent studies, employing the same ongoing system, will be presented, one of which addresses an issue of system effectiveness, and the other of which examines the utility of some possible measures available to students of PSI. The reader is invited to discern, while reading, wherein these studies might be called "unorthodox."

The PSI course was an introductory psychology course organized into eleven self-paced units taken over a 10-week quarter. This course evolved over a ten-year period at Northeastern University (Terman, 1978). Each unit represented from 4–8 hours of study in a textbook, parts of which were programmed; most units included sessions with interactive videotapes shown on a continuous schedule at a designated campus location. Unit mastery quizzes consisted of ten multiple-choice questions, with nine questions answered correctly as the criterion for passing and advancement. Coaching sessions took place immediately after each quiz, during the same hour and at the same location, and provided feedback on quiz performance and tutorials on missed questions and their objectives. (Tutorials were conducted by undergraduate upperclass students enrolled in courses in tutoring principles.) Any one student was assigned three quiz-opportunity hours each week, and thus had a total of 30 opportunities to take tests over the 11 units of the course. Alternate forms of mastery quizzes for each unit were employed to preclude repeated encounters with identical test items.

STUDY 1: THE SURVIVAL OF INDIVIDUAL DIFFERENCES

Personalized systems of instruction typically require mastery at a high level of each successive unit of the course, but permit repeated attempts on the student's own timetable, in the achievement of that mastery. The self-pacing feature of the course leads to characteristic high-variance distributions of progress rates which have been widely reported. These variations are explicitly provided for in order to minimize the differential impact upon course achievement of fixed-pace instructional presentation, which selectively penalizes students for whom the fixed pace is rapid. It has been hoped by many, and assumed by some, that PSI's

The construction and maintenance of a PSI environment of the complexity described here requires the coordinated efforts of a staff with considerable behavioral competence. Among the many such who contributed to the work at Northeastern University upon which this paper is based, I especially want to acknowledge the contributions of these colleagues: David R. Barkmeier, Susan Ott, Marilyn L. Rumph, and Robin R. Rumph.

TABLE 1

Mean score on the first attempt on each unit

	Number of units completed				
	11	10	9	5–8	1–4
Number of students	307	133	123	311	326
Unit no.					
1	9.54	9.34	9.36	9.03	8.56
2	9.19	8.80	8.48	7.95	7.29
3	9.09	8.97	8.58	7.93	7.28
4	8.99	8.39	8.11	7.25	6.52
5	8.65	8.47	7.67	6.65	
6	9.22	9.13	8.83	8.08	
7	9.06	8.64	8.33	7.45	
8	8.84	8.47	8.40	7.16	
9	8.94	8.54	8.29		
10	9.34	9.13			
11	9.02				

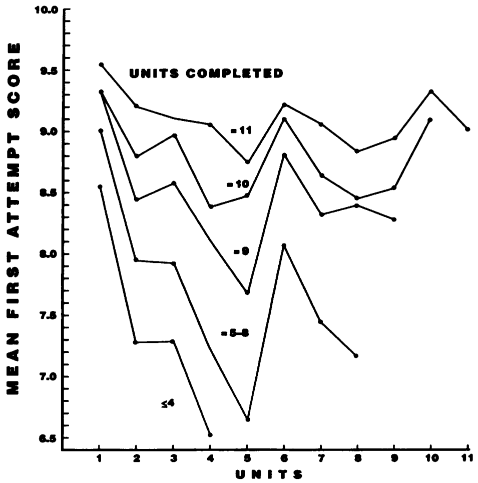

Figure 1. Mean scores on first attempts on mastery quizzes for each unit of the course. Students are grouped on the basis of the total number of units finally completed, and data for each group are plotted separately. Mastery quizzes consist of ten questions. For number of students in each group, see Table 1.

accommodation to variations in rates of progress would serve to narrow the gap between "rapid" and "slow" learners, at least regarding measures of performance other than progress rate itself.

The present study was designed to examine this supposition by an analysis of data concerning several performance measures taken from a PSI course whose initial enrollment was over 800 students.

A null-hypothesis formulation was that we would find no significant correlation between measures of rate of progress on the one hand and time-independent achievement measures on the other. Put in this extreme form, probably the hypothesis would have few defenders. The question remains, though: how far does PSI go in bringing us near to the fulfillment of the hope?

Of particular interest in this study is the first unit of the course, which consisted of the course pamphlet explaining the rules and procedures of the course, and contains no subject-matter material of any kind. The quiz on this pamphlet—the first unit quiz of the course—is taken with the pamphlet available for consultation during the quiz. Performance on this first unit and associated quiz may thus be regarded as reflecting individual differences in entering skills—from reading level to information processing abilities, however defined.

The Data

Categories were constructed to index the rate-of-progress variable by dividing all students into five groups on the basis of number of course units completed by the end of the academic quarter, as follows: all 11 units completed, 10 units completed, 9 units completed, between 4–8 units completed, and fewer than 4 units completed. The performance variables examined were (a) average score on the first attempt on a given unit; and (b) average number of attempts required to pass each quiz with a criterion score of nine or better.

Table 1 shows the average score attained on the first quiz attempt for each course unit, for the five groups of students. (Empty cells occur when subjects in a group do not reach the unit corresponding to that cell.) The same data are plotted in Figure 1.

These results establish that for any unit, the average score attained on the first attempt is higher for those groups which complete more units of the course. This departure from a random pattern is con-

TABLE 2

Mean no. of attempts required to achieve mastery

	Number of units completed				
	11	10	9	5–8	1–4
Number of students	307	133	123	311	326
Unit no.					
1	1.11	1.14	1.34	1.25	1.42
2	1.29	1.46	1.64	2.04	2.04
3	1.31	1.41	1.62	1.92	1.93
4	1.34	1.78	2.07	2.61	2.33
5	1.52	1.68	2.11	2.39	
6	1.24	1.22	1.46	1.85	
7	1.30	1.53	1.76	1.87	
8	1.45	1.65	1.82	2.18	
9	1.38	1.59	1.59		
10	1.21	1.21			
11	1.25				

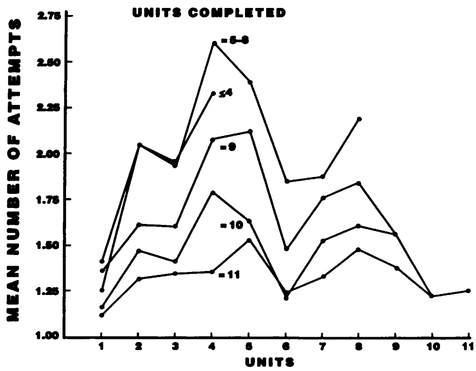

Figure 2. Mean number of attempts required to reach mastery for each unit of the course. Students are grouped on the basis of the total number of units finally completed, and data for each group are plotted separately. Mastery criterion is 9 or better on each ten-item quiz. Data from Table 2.

sistent, with but one reversal in the orderly pattern in the entire set of data. Of particular interest is the fact that the ordering of groups on the basis of eventual fate in the course can be discerned in the earliest units, even unto the very first unit, which teaches the course procedures and is quizzed upon with an "open pamphlet."

Table 2 shows the average number of attempts to achieve mastery for each course unit, for the five groups in the experiment. The data are plotted in Figure 2. It is apparent that those subjects who complete more of the course require fewer attempts to pass each unit—even on the early units. Although there are a few small reversals in the overall pattern, the departure from randomness is mostly consistent. And as with the score on the first attempt, the ordering of results can be discerned in the results for the very first unit.

Discussion

These results cast strong doubt upon a null hypothesis which states that no performance differences other than time are to be expected where certain key features of PSI are implemented—features such as well-defined units, immediate feedback and tutorials, and a self-pacing system which permits repeated attempts. A tautological defense might be constructed by extending the view of tutorials to include the correction of any problem presented by the student, no matter how unusual or deep in nature. Rather than argue at such an abstract level, we suggest that PSI courses be regarded as evolving laboratories for the specification of these problems and the development of instructional and tutoring strategies which might address them. Our own data suggest that at least some of the differences which account for variations in achievement late in the course can be identified at the very outset with materials which have nothing to do with the formal subject-matter of the course. The specification and, where possible, remediations of these problems, in an environment designed to promote effective instruction, would add a new feature to the world of PSI—remediation of prerequisites on the basis of initial or entering assessment. This might seem a dangerous flirtation with the world of "testing." But it might also result in PSI learning outcomes which are closer to those initially hoped for.

STUDY 2: A GROUP MEASURE OF TUTORING EFFECTIVENESS

A key feature of the PSI approach is that it permits students to make several attempts—often as many as needed to attain a specified level of mastery—on each unit of a course, with progress to the next unit contingent upon the attainment of mastery on the prior unit. Typically, a PSI course will provide for some kind of feedback or tutoring after each attempt—especially if the attempt did not succeed in attaining the stipulated criterion level of mastery. The precise nature of this feedback, offered between successive attempts on the same unit, varies widely in practice, and constitutes an important independent variable in the field of instructional psychology. Farmer, Lachter, Blaustein and Cole (1972) have studied the effect of varying the relative frequency of this feedback; Robin and Heselton (1977) have studied variations in the nature of the feedback itself. The present study is concerned with the appropriate dependent variable, and its formulation in a manner sensitive to variations in important properties of the tutoring system.

Score Distributions

When an individual student takes a test, a numerical score—such as number of items correct or percent correct—may be available. When such measures are available and preserved for a large number of students, a distribution of such scores can serve to characterize the performance of the system as a whole on any given quiz or attempt.

Figure 3 shows the distribution of scores (number of items correct out of ten multiple choice questions) on the first attempt of each of eleven successive units of an introductory psychology course enrolling about 800 students at Northeastern University. (The number declines slightly on later units thanks to dropouts.) These data are taken from a mature stage of the course—many difficulties had been overcome over several years. This situation is reflected in the positive skew seen in all quiz functions; these curves depart widely from the normal distribution, as is appropriate—but not easy to achieve—for "mastery-oriented" instruction. (As a check on the possibility that the test items are too easy, or over-prompted, similar data have been collected on the same items in a pre-test mode, prior to exposure to instruction. The resulting distributions were strongly skewed in the opposite direction.)

Such data are extremely useful in telling us how well the primary instructional components of the PSI course are doing. But if our interest is in the characterization of a tutoring system, then a single set of such distributions describing first attempts will not suffice. We must examine, for each unit of interest, a number of such distributions which describe scoring patterns on successive attempts; and we must look for changes which can be ascribed to properties of the tutoring system. The number of functions required becomes large, and it is tempting to pass to some summary measure, such as a mean value (e.g., mean score on each successive attempt). But here a second problem must be noted, that is, that scores of this kind (such as items correct) may not always be available. The observation which *is* available in any such system is whether any given student passed or failed a given attempt. It *must* be available in any instructional system in which advancement is made contingent upon passing; furthermore, it is likely to remain near the surface—that is, to be *easily* available—in the record-keeping processes of such courses. We propose, then, to examine further some properties of the measure *percent passing* as a measure of the effectiveness of a tutoring system. Specifically, we want to observe changes in percent passing over a series of successive attempts on the same unit, on the part of a group of students each of whom continues these attempts only until a passing performance is attained, and then stops (typically turning his or her attention to the next unit).

Percent Passing

The measure percent passing is neither new to psychology generally nor to PSI specifically. With regard to PSI, Gallup (1969) pointed out some time ago the uses of the percent passing datum in the evaluation of course materials and of tu-

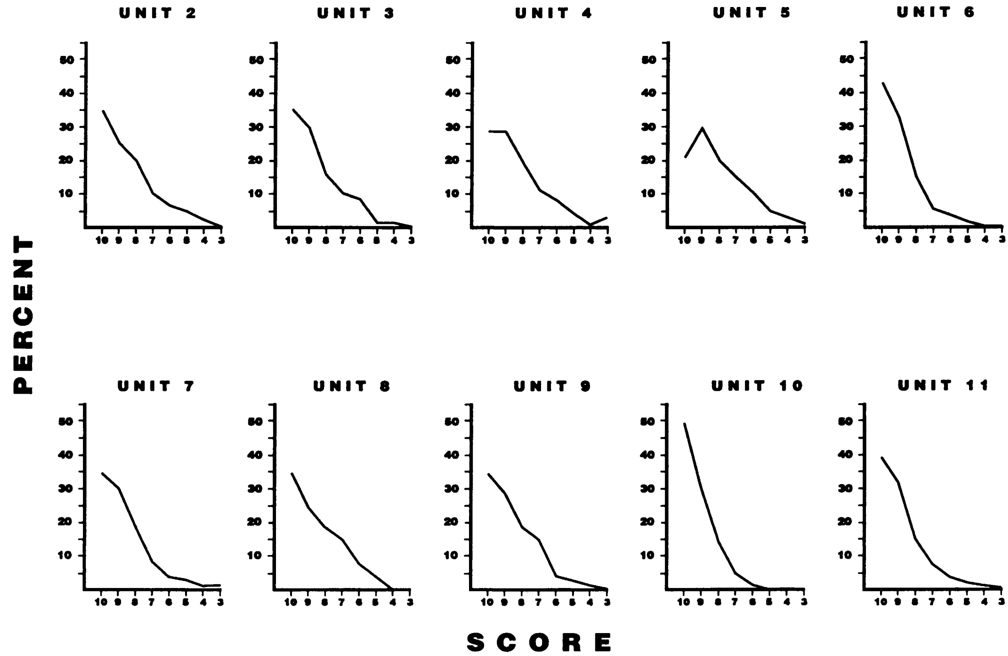

Figure 3. Distribution (by percent) of quiz scores on the first attempt for each unit of the course for all students remaining in the course at each unit. Unit 1 (the course procedures handbook) is omitted. Quizzes are 10 items in length, and the mastery criterion is 9 or better.

toring effectiveness. The argument is elaborated further by Cook, Cole, Gallup, Pennypacker, and Schiller (1976). Neither of these sources discusses, however, which of several variant forms this measure can take is best suited to reflect important properties of the tutoring system. That is our concern here.

The percentage value which is perhaps simplest to obtain routinely is the percent of the entire group passing on each successive attempt for each unit in the course. Figure 4 shows such a function for the same eleven successive units of Northeastern's introductory psychology course. This function is related to that presented in Figure 3 in that once the passing criterion is stated—nine or more items correct in our case—then the first point for each unit in Figure 4 (percent passing on the first try) is seen to be the sum of the first two points (percent obtaining scores of ten or nine) from the corresponding graphs in Figure 3. The remaining data in Figure 4—percent passing on later attempts—represent additional information which is not derivable from Figure 3. Thus, although these two sets of curves resemble each other superficially, they tell us very different things.

From Figure 4 much can be seen: for most units, about two-thirds of the students pass on the first attempt; a steadily declining number of passes occur on each successive attempt; units four and five are the most difficult and perhaps in need of repair (a fact which can also be surmised from Figure 3); and so on. But other questions cannot be answered; in particular, it is difficult to go beyond the gross impression that tutoring "seems to be working."

Cumulative Percent Passing

To ascertain whether the tutoring-restudy process eventually reaches everyone, the same data can be plotted cumulatively to see if each curve reaches an asymptote of 100%. This is done in Figure 5. This form of such data is the most advantageous to show to deans and to use in raising money. Those familiar with cumulative curves will be familiar

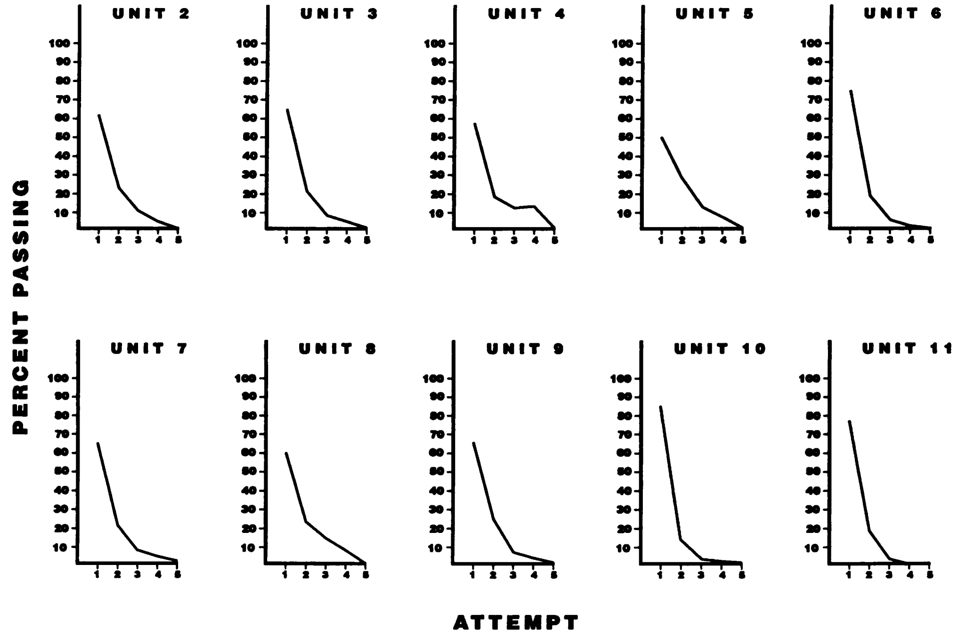

Figure 4. Percent of students passing unit quizzes on each consecutive attempt for each unit quiz. Successive attempts on the same unit employed alternate forms of equivalent 10-item quizzes on the same objectives. The initial points for each unit in this figure can be obtained by adding the percent obtaining 10 and 9 correct in Figure 3.

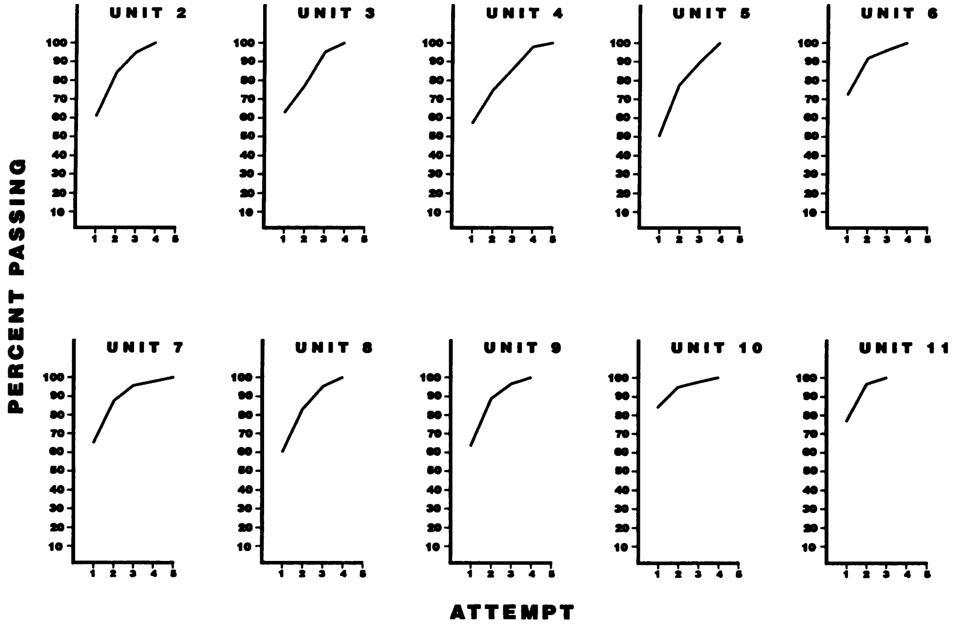

Figure 5. Cumulative percent of students passing unit quizzes on each consecutive attempt for each unit quiz. The points for each unit in this figure are obtained by cumulating the corresponding values from Figure 4.

TABLE 3

Numerical relations in "cybernetic" scoring

			Passing		
Attempt no.	Total attempting	Pass/fail	Cumulative	%	Cumulative % of total
1	800	480/320	480	60	60
Residual % passing		60%			
2	320	224/96	704	28	88
Residual % passing		70%			
3	96	77/79	781	9.6	97.6
Residual % passing		80%			
4	19	17/2	798	2.1	99.7
Residual % passing		90%			

with their "positive" flavor, reminiscent of the sundial which "tells only the sunny hours."

But if we want to ask more refined questions concerning tutoring effectiveness, yet another function, derived from the same data, is more useful.

Residual Percent Passing

That function, which might be called the "residual percent passing," is obtained by employing the same numerator as in "percent of total passing,"—namely, the number of students who pass on a given attempt; but as a denominator we use the *total number of students who had not yet passed and who therefore made the attempt again.* This yields the percent passing from among those who made this specific attempt. We have thus a measure of the relative success of each attempt. Such information is implicit in the functions discussed so far. It could be obtained from the "percent of total passing" by subtractions which adjust the denominator of each percentage point; or it could be regarded as the slope or derivative of the cumulative percent passing function. Its particular significance for an instructional system will now be explained.

Cybernetic System

A tutoring system could be regarded as "perfect" if all problems remaining after the first attempt were detected and corrected by the second attempt. In such a case, the second point of a percent-of-total curve would be 100% minus the value of the first point; and the cumulative function would reach 100% on the second attempt. Also, the measure "residual percent passing" would be 100%. Such

TABLE 4

Numerical relations in "stochastic" scoring

			Passing		
Attempt no.	Total attempting	Pass/fail	Cumulative	%	Cumulative % of total
1	800	480/320	480	60	60
Residual % passing		60%			
2	320	192/128	672	24	84
Residual % passing		60%			
3	128	77/51	749	9.6	93.6
Residual % passing		60%			
4	51	31/20	780	3.8	97.4
Residual % passing		60%			

TABLE 5

Numerical relations in "differential" scoring

Attempt no.	Total attempting	Pass/fail	Passing: Cumulative	Passing: %	Passing: Cumulative % of total
1	800	480/320	480	60	60
Residual % passing		60%			
2	320	160/160	640	20	80
Residual % passing		50%			
3	160	64/96	704	8	88
Residual % passing		40%			
4	96	28.8/67	733	3.6	92.6
Residual % passing		30%			

an ideal may rarely be achieved, but may usefully serve to define the end-point of a continuum, indicating the degree to which a tutoring system responds to residual problems by detecting and eliminating them. We may define as a "cybernetic" tutoring system, one in which an increasing proportion of the residual students pass on each successive attempt. Table 3 shows the numerical relations in a hypothetical case. The residual proportions passing increase over successive tries, leading to a very rapid approach to the asymptote. Such a system contains powerful self-corrective elements, which permit an increasingly efficient attack upon whatever problems remain after each attempt.

Stochastic System

In contrast, we can define a second kind of system as "stochastic." Table 4 illustrates such a system, in which the proportion of the residual which passes on each successive attempt remains constant. The relative efficiency in addressing its unsolved problems is unchanging from one trial to the next; the cumulative curves will climb more slowly to their asymptotes. The system as a whole might be regarded as neither "learning" nor "deteriorating."

Differential System

A third possible case—called "differential"—is illustrated in Table 5. Here the residual percent passing become smaller on each successive attempt. Such a case could arise if the "best" students pass on early attempts, so that those students encountered on later attempts are progressively harder to tutor, lacking in needed prerequisites, weak in motivation, etc. Such a deteriorating situation could arise as a result of a shift in relatively fixed individual differences from try to try, or it could arise from variables introduced by the operation of the system itself: tutors punish those who fail on a given attempt, with the result that students study less effectively for later attempts, or approach such attempts in a disturbed state. A mixture of individual differences introduced by a heterogeneous student body, together with reactions on the part of an inexperienced tutoring staff to these differences, may be a sure-fire formula for deterioration in a differential system.

Figure 6 presents graphically the three functions we have discussed, for the three types of system, characterized on the basis of whether the residual percent passing increases with successive attempts ("cybernetic"), remains constant ("stochastic"), or declines ("differential"). On the basis of the "percent passing" or the "cumulative percent passing" data alone, these differences cannot be sorted out; the three cases all look satisfactory—all are doing their job in the sense that most students eventually pass each unit. In contrast, inspection of the residual func-

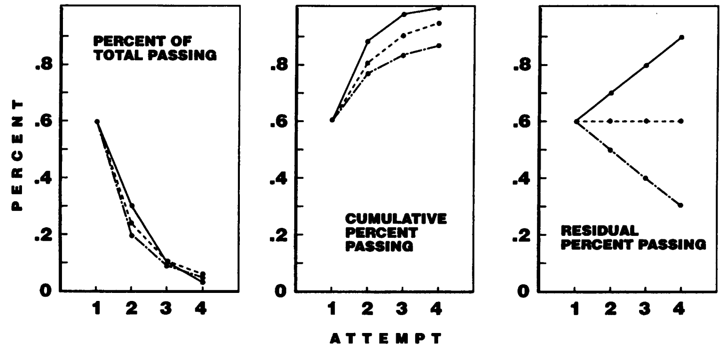

Figure 6. Three graphs of PSI system performance. The hypothetical data of Tables 3, 4, and 5, represented in three possible ways. The "cybernetic" (solid line), "stochastic" (dashed line) and "differential" (dot-and-dash) cases are most readily distinguished in the plot of residual percent passing.

tions permits the three cases to be easily distinguished.

If we think back to the empirical data we have examined in this paper, plotted in the first two of the three possible ways we have discussed, it is not clear whether the system we are studying is cybernetic, stochastic, or differential. If we plot the same data in this third way—what will we see?

Figure 7 shows the empirical data first presented in Figure 4, now replotted in terms of the residual percent passing. Inspection of the individual cases—one for each unit—suggests all three types of system, with no apparent reliable trend. Probably the most sensible interpretation is that the system is too noisy to be characterized. This is the price of looking too deeply underneath smooth curves!

It is not out of the question that each unit should have its own characteristic function. If the units differed in some specifiable manner—say some being heavily mathematical while others were more verbal or graphic—this notion could be investigated. But in the absence of such an independent sorting, such a possibility cannot be pursued.

Mathematical Modeling

Curves such as these could be fit mathematically in many ways, but the distinctions here drawn do not incorporate sufficient assumptions to favor any specific approach. Several remarks may nonetheless be in order. The "stochastic"

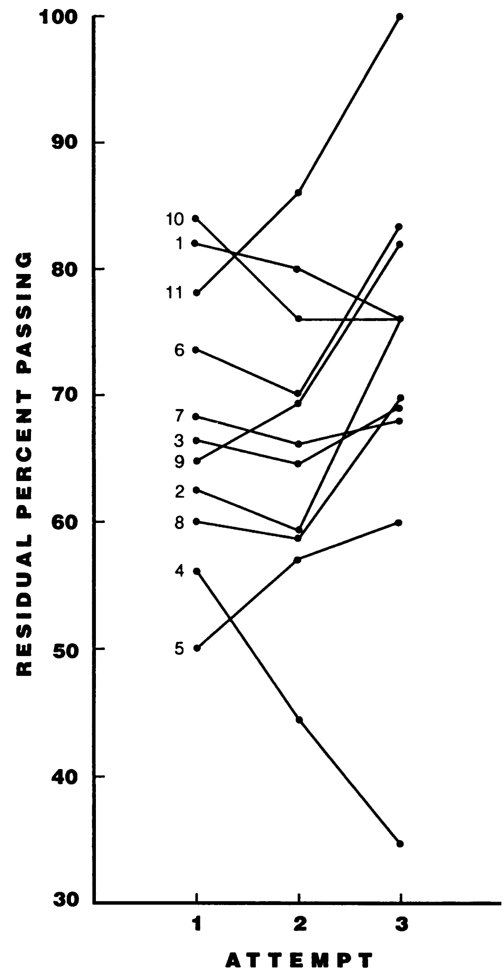

Figure 7. Residual percent passing over three successive attempts, for each of the 11 units of the psychology course. These empirical data should be compared with the third panel (residual percent passing) of Figure 6.

case, in which the gain per trial is proportional to the residual number of students yet to pass, is of special interest. It is parallel to the negatively accelerated positive growth functions of learning theory, which have been stated and employed in many forms to describe the growth of changes in the individual subject on successive learning trials. Perhaps the best-known instance is to be found in Clark Hull's formulation of the growth of habit strength (Hull, 1943, Ch. 8):

$$sHr = M(1 - e^{-kt}) \quad (1)$$

where

sHr = habit strength
M = upper limit of habit strength
t = trials
k = a constant expressing a learning rate
e = the base of natural logarithms

Equation (1) is built upon the assumption that the gain per trial is proportional to the amount remaining to be learned (before the final limit M is reached). That assumption can be expressed as a simple differential equation:

$$dH/dT = k(M - H) \quad (2)$$

where "H" is a simpler expression for habit strength, replacing the more cumbersome sHr. From this equation, expressing the assumption, the first equation is derived by integration.

In later work similar assumptions appear, albeit in other forms. Estes, for example, uses linear operators (1953) to express the probability that a given choice will be made on the nth trial of an experiment as a function of previous reinforcements for that choice:

$$P_{n+1} - P_n = \theta[1 - P_n] \quad (3)$$

The trials are discrete and we are dealing with difference equations, in which the increase in probability from one trial to the next is a constant fraction of the shift which has not yet occurred. The resulting curve will resemble Hull's, differing chiefly in that it is defined only for integral values of n.

Estes points out the formal similarity in the two statements:

$$H = k(M - H) \quad (4)$$

and

$$P = \theta(1 - P)$$

which express the assumptions embodied in (2) and (3), respectively, in a common notation applicable to both continuous and discrete cases (Estes, 1960).

A similar formulation appears in the mathematical models of Bush and Mosteller (1955). The ubiquity of the approach in theories of learning has been pointed out by Restle (1959) and by Sternberg (1963). And even in the more recent work of Atkinson, Boer, Suppes, and others, which builds upon rather different notions of the underlying psychological events, the same equation has been identified by Coombs, Dawes, and Tversky (1970, pp. 291 ff.).

Thus familiar methods are available which make it easy to identify the "stochastic" case, in which a constant proportion of the remaining quantity is transformed with each trial. And by the same token, deviations in either direction from this familiar case are easy to identify. A more sober statement might be: we are able to identify several kinds of possible order, if we find it.

Discussion of the Residual Measure Approach

For our concerns here the interest in these equations is not in the psychology of learning which they may embody, for we are not "psychologizing." We are describing a system in which individual people are merely the elements; they are the molecules in a gas, and it is the behavior of the gas we are trying to characterize. The interest lies in the articulation of a process which, on each attempt made, closes a constant proportion of the gap which separates its present position from its desired outcome. The "stochastic" model offers a baseline or touchstone against which to examine a multi-trial instructional system. It does not describe what it is—even hypothetically—or what it should be. But it does pose a challenge. Can we hope to surpass a stochastic system as we produce a powerful cybernetic system which detects and eliminates remaining problems ever more rapidly from one trial to the next? Or should we be happy to even achieve a stochastic model, pleased to have surmounted the lagging inertias of a differential system?

It is certainly too early to tell. And as we study instructional systems using this scheme, we may find that more terms are needed to sketch out an actual description of what is happening. It may be that an additional term will be needed to recognize the independence of the instructional component from the tutoring process. (Notice in Figure 7 the suspicion that the residual value deteriorates from the first attempt—based on instruction without tutoring—to the results of the first tutorial effort in the second attempt, and then improves between the second and third attempts.)

At this point, we have merely offered a dependent variable of possible use in detecting the fine grain of the corrective properties of an instructional system permitting several attempts on each quiz, with tutoring after attempts which fail. The residual percent passing, like its counterpart the first derivative, tells an important story about changes in the process taking place. It should be useful in the analysis and improvement of that process to a degree that conventional single-valued outcome measures cannot achieve. All that is needed, then, is the identification of the independent variables of which this measure is a function.

CONCLUDING UNSCIENTIFIC POSTSCRIPT

This paper has as its subtitle, "Two Unorthodox Studies." The reader was invited to identify, in the course of reading, those respects wherein lay the deviations from orthodoxy.

The author's view is that this paper is unorthodox in these respects:

In both studies, there is no independent variable. It might be worth discussing in what way these are studies.

In both studies, group data are employed. All behavioral data are in the form of averages or percents.

In one of the studies, some use is made of mathematics for modeling purposes.

So much for unorthodoxy. But what, then is the orthodoxy? The answer to this question can be left as an exercise for the reader. The facts have been given. Let him extrapolate who will.

REFERENCES

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning.* New York: John Wiley.

Cook, D. A., Cole, B. K., Gallup, H. F., Pennypacker, H. S., & Schiller, W. J. (1976). Contexts for instructional evaluation. In F. S. Keller & B. V. Koen (Eds.), *The personalized system of instruction: State of the art.* A report to the Alfred P. Sloan Foundation. Austin, TX: The Engineering Institute of the University of Texas.

Coombs, C. H., Dawes, R. M., & Tversky, A. (1970). *Mathematical psychology: An elementary introduction.* Englewood Cliffs, NJ: Prentice-Hall.

Estes, W. K. (1953). Models for learning theory. In *Symposium on psychology of learning basic to military training problems* (HR-HTD 210/1, 21–28). Washington, DC: Committee on Human Resources, Research and Development Board.

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

Farmer, J., Lachter, G. D., Blaustein, J. J., & Cole, B. K. (1972). The role of proctoring in personalized instruction. *Journal of Applied Behavior Analysis, 5,* 401–405.

Gallup, H. F. (1969). Individualized instruction in introductory psychology. In W. C. Shepphard (Ed.), *Proceedings of the Conference on Instructional Innovations in Undergraduate Education* (pp. 63–88). Eugene, OR: University of Oregon.

Hull, C. L. (1943). *Principles of behavior.* New York: Appleton Century Crofts.

Restle, F. (1959). A survey and classification of learning models. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory* (pp. 415–427). Palo Alto, CA: Stanford University Press.

Robin, A. L., & Heselton, P. A. (1977). Academic and attitudinal responses engendered by "enriched" versus "lean" feedback in a personalized instruction course. *Journal of Personalized Instruction, 2,* 136–142.

Sternberg, S. (1963). Stochastic learning theory. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 2) (pp. 1–120). New York: John Wiley.

Terman, M. (1978). Personalizing the large enrollment course. *Teaching of psychology, 5,* 72–75.