5

Probability Learning and Sequence Learning

Jerome L. Myers

University of Massachusetts, Amherst

I. INTRODUCTION

In 1954, Estes and Straughan published a paper on binary prediction that proved to be the forerunner of a stream of theoretical and experimental articles dealing with such behavior. The task they employed was not new; it had been developed 15 years earlier by Humphreys (1939) as an analog to classical conditioning. However, not until publication of the Estes and Straughan (1954) paper was the binary prediction task the focus of major efforts by investigators of learning. The impact of this article undoubtedly resulted from the fact that it was one of the first published tests of Estes' statistical learning theory (SLT) which had appeared a few years earlier (Estes, 1950). "Toward a statistical theory of learning" (Estes, 1950) had held forth the promise of a learning theory capable of precise, quantitative prediction, something beyond the curve fitting which had typified previous mathematical theories. The Estes and Straughan paper began to realize that promise, not only providing support for the theory, but also presenting a simple paradigm in which it could be tested. Thus, it signaled an era in which both proponents and opponents of SLT would derive and test its consequences in a variety of prediction tasks.

In this early research on "probability learning," random event sequences were typically employed, and marginal event probability was a major independent variable. In more recent experiments, structured sequences have been employed and learning as a function of structure has been of major interest. This later work was originally motivated in large part by a desire to test alternatives to simple conditioning models of choice behavior; it has been maintained, however, by an interest in sequential information processing. In contrast to earlier research in probability learning, which was primarily concerned with the validity of a general theory of learning for which choice behavior was a convenient testing ground, current studies are focused on such issues as the acquisition and representation of sequential information in memory and the way in which such information influences choice behavior. In this chapter, these two phases of work on sequential choice behavior are reviewed. The review is not exhaustive but attempts to provide a sense of the major theoretical developments, the basic issues, and the relevant findings in probability learning and sequence learning.

II. PROBABILITY LEARNING

The basic paradigm employed in studying probability learning is quite simple. Upon presentation of the ready signal that initiates a trial, the subject is required to predict which of several events (E_i ; $i = 1, 2, \ldots, r$) will occur next; prediction of E_i is designated as A_i . The prediction is usually followed by feedback, presentation of one of the alternative events. In the simplest case, that of noncontingent events, E_i occurs with constant probability, π_i , and is independent of previous responses and events. Thus, the task is analogous to that of predicting the fall of a die or the outcome of a spin of a roulette wheel. Typically, an experimental session consists of several hundred trials, spaced at intervals of roughly 3–5 sec.

Estes and his co-workers have assumed that event prediction reflects an underlying representation of event probability which develops by means of a simple conditioning process. Several alternative models have been derived within this conditioning framework. Two of these, the linear and pattern models, have been extensively investigated. The linear model assumes that some proportion of a large population of stimulus elements is randomly sampled on each trial, that probability of response A_i on Trial n [$P(A_{i,n})$] is equal to the proportion of the sample conditioned to that response, and that, after presentation of the reinforcing event, the entire sample becomes conditioned to the response of predicting that event. In contrast, the pattern model assumes a small set of N elements, or patterns, only one of which is sampled on each trial. The conditioning state of the sampled N pattern completely determines the response for that trial. With some probability, c, the sampled pattern becomes conditioned to the reinforced response for that trial and with probability, $1 - c$, the conditioning state is unchanged. Atkinson and Estes (1963) have provided a more formal treatment of stimulus sampling models. Despite apparent differences, the two models yield numerous predictions that are either identical or very similar. In view of the fact that much of the probability learning literature has implications for these models, it will be helpful to note certain common predictions.

A. Some Predictions from Statistical Learning Theory

One fundamental consequence of both the linear and pattern models is that $P(A_{i,n})$ should approach π_i as n increases. This prediction of asymptotic *probability matching* is of interest not only because it is a strong, parameter-free prediction, but also because it is somewhat surprising from the viewpoint of decision theory; maximization of the expected number of correct predictions requires that the subject always predict the most frequent event. A second strong prediction of both models, which follows directly from their conditioning axioms, is positive recency; if the occurrence of an event increases the probability of predicting that event, then $P(A_i)$ should monotonically approach 1.0 as the length of a run of consecutive E_i's increases.

In the following Section II.B, the status of these predictions—asymptotic probability matching and positive recency—is considered for experiments in which noncontingent event schedules have been employed. Concern here is with delineating the conditions under which these predictions are, and are not, verified, and with attempting to gain some understanding of the implications of this pattern of results for probability learning in general, and statistical learning theory (SLT) in particular.

B. Some Basic Results

1. *Asymptotic Response Probability, $P_\infty(A_i)$*

a. Extended practice. It is generally believed that probability matching, the asymptotic approach of $P(A_i)$ to π_i , is a robust phenomenon, readily demonstrated in probability learning studies. The facts, in contrast to the impression, are somewhat more complicated. It is true that in some studies extremely stable probability matching has been demonstrated for several terminal trial blocks; Neimark and Shuford (1959) provide an excellent example. However, it is also true that $P(A_1)$ consistently overshoots π with extended practice. An experiment by Friedman, Burke, Cole, Keller, Millward, and Estes (1964), often cited as strong evidence for probability matching, is a case in point. In each of the last 7 12-trial blocks of a series of 288 trials with $\pi = .8$, $P(A_1)$ exceeded .8; the average deviation was only .03, small but nonetheless troublesome. Similar departures from matching have been observed by other investigators, in fact, by almost anyone who has run subjects for more than 300 trials. For example, with π at .6, .7, and .8, Myers, Fort, Katz, and Suydam (1963) obtained values of $P(A_1)$ of .616, .753, and .871 for Trials 301–400.

Probability matching should hold in the noncontingent case for more than two choices. Early experiments employing three choices for 200 and 100 trials (Detambel, 1955; Neimark, 1956) did, in fact, support the probability matching theorem. However, as in binary prediction experiments, extended practice resulted in overshooting. Furthermore, experiments by Gardner (1957) and Cotton and Rechtschaffen (1958) have demonstrated that overshooting on the most frequent alternative is more pronounced in the three- than in the two-choice case. With three choices $P_\infty(A_1)$ on Trials 286–450 was about .67 and .80 for π of .6 and .7, respectively. A subsequent study by Gardner (1958), employing from 2 to 8 choices for 420 trials, reaffirmed the overshooting result and indicated that the amount of overshooting increased with number of choices. It also appears that several curves were still rising at the end of the session.

Although the basic finding that overshooting occurs with extended practice is well established, its implications for the validity of the linear and pattern models are not clear. Estes (1964, 1972), although conceding that overshooting occurs, has argued that SLT provides an essentially correct account of the course of probability learning. He ascribes the apparent failure of the probability-matching theorem under extended practice to extraneous factors that influence predictive behavior so that it is no longer an adequate reflection of the underlying state of learning. One possibility is that the theory describes predictive behavior quite accurately as long as the probability of a correct response continues to increase. When the basic learning process asymptotes at π and, consequently, there is no further increase in probability of a correct response, the subject may test various strategies designed to further improve his performance. According to this analysis, overshooting should be more likely to occur and should be of greater magnitude in conditions under which it represents a large improvement in probability correct, relative to probability matching. This seems to be the case. The difference in probability correct between matching and optimization (100% prediction of the most frequent event) is a quadratic function of π, and also increases with number of choices when there is a uniform distribution over the less frequent events; the degree of overshooting follows the same pattern.

Measures other than prediction probabilities may be helpful in assessing Estes' interpretation of overshooting. In particular, reasons advanced for overshooting—boredom, fatigue, experiments in which subjects make attempts to increase the percentage correct—do not obviously apply to direct estimates of π. If SLT provides a valid account of probability learning, as opposed to predictive behavior, such estimates of π might well prove more stable over prolonged series of trials than choice proportions. Whether this will, in fact, occur remains to be seen. Neimark and Shuford (1959) and Beach, Rose, Sayeki, Wise, and Carter (1970) have found that estimates of π match π in the last few blocks of a 300-trial sequence. Unfortunately, more trials are required to resolve the issue of asymptotic stability.

As one might expect, degree of practice is not the only variable that constrains the applicability of the probability matching theorem. Investigators who are determined to do so can produce overshooting in a variety of ways; undershooting is somewhat more difficult to achieve, but possible. Most of these departures from probability matching have been obtained under experimental conditions outside of the intended scope of SLT and are thus not relevant to evaluating the validity of the theory. Nevertheless, there are implications for alternative models, for an understanding of decision making capabilities, and for methodology in probability learning experiments. Therefore, two other factors are now considered that produce departures from matching.

b. Instructions. The premise underlying this research appears to be that subjects are capable of more intelligent decisions than SLT, and early findings of probability matching, imply. In particular, probability matching has been assumed to be a result of the subject's failure to detect the randomness inherent in the event sequence and his consequent attempts to find a perfectly predictable pattern (Flood, 1954). This assumption gains support from the finding of 5–10% overshooting when events are displayed in such a way as to appear randomly sequenced (Nies, 1962; Peterson & Ulehla, 1965). Under these conditions of apparent randomness, explanations which emphasize that the odds are constant over trials and that event runs are irrelevant cues have no additional effect. Nies found no difference between two groups differing only with respect to the presentation of such an explanation, and Beach and Swensson (1967), who employed such an explanation in addition to random shuffling of an event deck, obtained the same 8% overshooting observed by Peterson and Ulehla (1965) without such instructions. In the absence of a clear appearance of randomness, explanations must be very strongly worded to have an effect. Studies by McCracken, Osterhout, and Voss (1962), and Braveman and Fischer (1968) have demonstrated that merely telling subjects that the sequence is random or that there is no fixed pattern has little effect, nor are subjects unduly influenced by being instructed to avoid a trial-by-trial approach or by being told that it is impossible to be correct on every trial. Subjects appear to understand what is expected of them only when both knowledge of randomness and the desirability of maximizing correct responses over blocks of trials is communicated; then, terminal values of $P(A_1)$ are obtained that exceed, by 10–20%, those for subjects merely instructed to attempt to predict correctly on each trial.

Although carefully worded instructions, or displays of randomness, can elicit overshooting, the results cited above hardly stand as a testimonial to man's decision making capabilities. Under conditions in which the subject is all but instructed to predict the more frequent event, overshooting by only 10% is obtained. Individual subject protocols are not more impressive. For example, only 29 and 27% optimalize (always predict the more frequent event) in the last trial block of the Peterson and Ulehla (1965) and Beach and Swensson (1967) studies, respectively; this result occurs despite the fact that the former experiment involved monetary incentives, which by themselves elicit overshooting, and in the latter study, subjects were reminded throughout the session "to ignore the runs and to avoid the gambler's fallacy." To be fair, subjects are not quite as dense as these data would suggest. Nies (1962) reported that, although only 3% of his subjects optimized, 60% were able to verbalize the optimal strategy in a post session interview.

There are several reasons why many subjects fail to learn the optimal strategy and why most of those who do learn it fail to use it. First, as Siegel (1959) has suggested, there may be a certain utility in varying one's response, in attempting to outguess the experimenter. Second, subjects believe that patterns are present in the event sequence (for example, Nies, 1962) and exaggerate the likelihood of short runs (Tune, 1964). Third, these beliefs are essentially correct in many of the experiments in which instructions have been manipulated. Nies, who randomized in 50-trial blocks, reported that more short runs were present in his event sequence than would be expected for unconstrained random event sequences. This is undoubtedly even truer in studies by Goodnow (1955), McCracken *et al.* (1962) and Braveman and Fischer (1968), in which event proportions were fixed for blocks of 10, 20, and 30 trials, respectively. As Jones and Myers (1966) have demonstrated, when the sequence is randomized in such short blocks, subjects can outguess the experimenter, achieving considerably more correct responses than would be expected on a chance basis. With such constrained sequences, the "gambler's fallacy" (the other event is due) is not a fallacy, and the strategy of uniformly predicting the higher probability event is not necessarily optimal. Jones (1971) has pointed out that sequence structure is a form of instruction; if so, subjects in experiments such as those just cited receive conflicting messages.

c. Monetary payoffs. In the studies just considered there was no tangible incentive for subjects to optimize. Those who still have faith in the decision-making capabilities of human subjects might expect the introduction of monetary gains and costs to markedly increase probability of predicting the more frequent event. This expectation is confirmed. Comparing subjects who won one cent for correct predictions and lost one cent for incorrect predictions with subjects who had no monetary incentive, Myers *et al.* (1963) found $P(A_1)$ to be significantly higher for the one-cent group. The differences in terminal (Trials 301–400) response probabilities were .03, .12, and .06 at π values of .6, .7, and .8, respectively.

An additional expectation—that $P(A_1)$ would be a monotonic increasing function of incentive magnitude—at best receives weak confirmation. In the Myers *et al.* (1963) experiment cited above, the average value of $P(A_1)$ was .03 higher in ten-cent than in one-cent groups. More generally, a review of

eight experiments in which two nonzero payoff levels were compared reveals differences in terminal values of $P(A_1)$ ranging from essentially zero (Jones & Myers, 1966) to about .05 (Suppes & Atkinson, 1960; Castellan, 1960). One point is evident: while instructional and motivational manipulations can yield increased probability of predicting the more frequent event, subjects consistently fall short of the optimal strategy of always predicting that event. What is less clear is the extent to which this reflects a failure to learn the optimal strategy as opposed to a failure to use that strategy.

Rather substantial effects of payoff magnitude can be produced by employing a within-subject paradigm in which each subject makes predictions under two payoff levels. This has been done by randomly sequencing equal number of trials at each payoff level (Schnorr, Lipkin, & Myers, 1966; Schnorr & Myers, 1967) or by changing the payoff level partway through the sequence of trials (Castellan, 1969; Halpern, Schwartz, & Chapman, 1968; Swensson, 1965). Under either approach, choice data reflect a "negative contrast" effect. On high-payoff trials, $P(A_1)$ is at about the same level as in payoff groups in the studies cited above; however, on low-payoff trials, $P(A_1)$ is considerably depressed, typically below the probability matching level. Schnorr and Myers (1967) demonstrated that, on high-payoff trials $P_\infty(A_1)$ is independent of payoff magnitude whereas, on low-payoff trials, $P_\infty(A_1)$ decreases as the difference between high and low payoff decreases. Schnorr *et al.* (1966) and Swensson (1965) have also found that estimates of event probability, obtained at the end of the experimental session, are considerably under the true value of π for low-payoff trials. This may indicate that subjects based their estimates on their response sequences. Alternatively, incentive may influence the underlying probability learning process and negative contrast may reflect basic differences in learning, rather than in strategies, on high- and low-payoff trials.

Whatever the explanation, the negative contrast effect is a phenomenon of some generality, not a peculiar consequence of the probability learning task or the subject population used. The effects observed in experiments with humans rather neatly parallel those obtained with rats in runways and T mazes (Black, 1968), both response times and choice proportions being depressed on low-incentive trials as a function of the difference in magnitude of the two amounts of reward.

Several models have been proposed to account for the effects of payoffs upon human probability learning (Luce & Suppes, 1965). Three of these will be considered, chosen for discussion because they are capable of generating predictions for both the learning curve and the sequence of responses, have provided good fits to several data sets, and represent somewhat different assumptions about the choice process.

Both Siegel (1959, 1961) and Estes (1962) have proposed models which incorporate two independent processes, probability learning and decision. Both assume that the underlying learning process asymptotes at π and is described by some variant of SLT. With respect to the decision process, Siegel has assumed that subjects choose a strategy that maximizes the sum of the expected utility of payoffs and the expected utility of varying one's response. This quantity is maximized when

$$P_\infty(A_1) = (k_1 + k_2)\pi + (.5 - k_2), \qquad (1)$$

where the k_i are functions of utilities of payoffs and response variation. Siegel's key contribution, the notion that response variation has value for a subject, is consistent with both intuition and the finding that subjects who have not optimized have frequently indicated knowledge of the optimal strategy. Nevertheless, the results of direct attempts to test the validity of the concept of utility of response variation have been mixed (Messick, 1965; Halpern & Dengler, 1969; Halpern, Dengler, & Ulehla, 1968). Fits to asymptotic choice proportions support the model; for example, setting k_1 equal to k_2 in Equation (1), the Myers *et al.* (1963) data have been fit and an average absolute deviation of observed from predicted values of .022 obtained. Further direct tests of the basic assumptions would be useful, as would fits to choice proportions obtained with nonsymmetric payoff matrices. Of particular interest would be a test of a stochastic version; if SLT describes the learning of subjective probability (Siegel, 1961), sequential statistics and learning curves can, and should, be fit.

Random walk models (Bower, 1959; Estes, 1960, 1962) provide a more molecular analysis of choice behavior. In one such model, it is assumed that on any trial, n, the subject has probability $\phi_{i,n}$ of orienting toward A_i; once the subject is oriented toward a response, the probability of making it is $p_{i,n}$.

Consideration of the following diagram may make the model clearer. Let S denote the starting position on a trial, O_i indicate orienting toward response i, and A_i represents execution of response i. Then the process on a single trial may be represented by

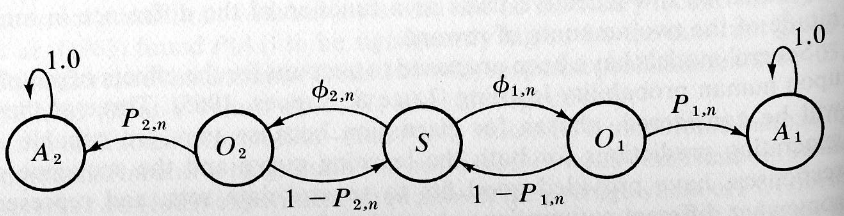

If $\phi_{i,n}$ and $p_{i,n}$ are independent functions of n described by SLT, then both have asymptotes at π_i and $P_\infty(A_1) = \pi^2/(1 - 2\pi(1 -))$. For π of .6, .7, and .8, predicted $P_\infty(A_1)$ are .692, .845, and .941, respectively; these values closely approximate those obtained by several investigators who used monetary payoffs. In the case in which no payoffs are at stake, $p_1 = p_2 = 1$ will yield $P_\infty(A_1) = \pi$.

The model has a certain appeal; it is conceptually simple and permits derivations of learning curves and sequential statistics (Cole, 1965, has fit such statistics), as well as response times (the expected number of orienting responses prior to reaching an A_i state is easily derived), and is thus subject to test in a variety of ways. However, the predicted lack of influence of payoff magnitude upon $P_\infty(A_i)$ is too strong a result; choice proportions are known to be a function of payoff magnitudes (markedly so when gains, or losses, depend upon the particular response), and the model should be modified to be responsive to this fact.

In contrast to the two models just considered, both of which postulate independent probability learning and decision processes, Myers and Atkinson's (1964) weak–strong model assumes that the response depends directly on the conditioning state of a sampled stimulus element; the element is either strongly or weakly conditioned to some A_1. Following an incorrect response, the sampled element changes conditioning state with probability δ; following a correct response, the transition to a new state occurs with probability μ. The model has yielded good fits to asymptotic choice proportions and sequential statistics in experiments employing a wide range of payoff matrices (Myers & Atkinson, 1964; Myers, Suydam, & Heuckeroth, 1966). Like the models considered previously, it is applicable to more than two-choices and, like the random walk models, it is also capable of describing choice latencies (Myers, Gambino, & Jones, 1967).

d. Discussion. Our review of asymptotic choice proportions raises two questions. First, is there a fundamental probability learning process that develops with practice, asymptotes at π, and is directly manifested in prediction behavior in early trials, given proper instructions and no tangible payoffs? If so, overshooting is merely the result of conditions in which the subject is led to incorporate such probabilities into complex response strategies and decision rules. The Estes and Siegel models for choice under payoff are two representations of this view. On the other hand, the Myers–Atkinson model assumes a single learning process the rate and asymptote of which is determined by payoff parameters. The duoprocess position would be enhanced if estimates of event probability were demonstrated to develop in accord with SLT under conditions in which asymptotic choice proportions deviated from π; for example, under extended practice, instructions to optimize, or monetary payoffs.

A second unresolved issue is the extent of the decision making capabilities of subjects. Is the optimal strategy as difficult to learn as the widespread failure to optimize would suggest? Or is it that subjects perform more intelligently than it appears, maximizing an expected utility based on something more than monetary payoffs, and responding to reliable sequential cues that we have unthinkingly built into our sequences? With respect to this last point, a clearer picture of decision making capabilities should evolve with the use of unconstrained random sequences rather than sequences randomized in very small blocks.

One final point is in order. Evaluation of models of choice has too often rested solely, or at least primarily, on fits to marginal response probabilities from terminal trial blocks. It was noted earlier that this approach has a problem; the predicted asymptote may be exceeded not because the model is wrong, but because extended practice represents some second phase of performance not reflected in the model. On the other hand, good fits to such terminal choice proportions are hardly grounds for euphoria. The true learning process may not yet have asymptoted, and the fits may be fortuitous. The same model that fits "asymptotic" probabilities at 300–400 trials may also fit, with a change in parameter values, a higher "asymptote" obtained with still more practice. There is no clear way to define the true asymptote; consequently, good fits to terminal trial block proportions are by themselves, at best, weak support for a model. Sequential statistics would appear to provide more information about the status of stochastic models. Such data are considered next.

2. *Sequential Dependencies*

Fine-grain analyses of the data that focus on the dependency of responses upon preceding patterns of responses and events are considerably more informative than marginal probabilities. Such sequential dependencies provide a measure of trial-to-trial changes in performance and are, therefore, more direct tests of basic assumptions about reinforcing effects of trial outcomes. Furthermore, where earlier there was one marginal probability to predict for a single trial block, there now is a large array of probabilities conditional on the preceding pattern of responses, or events, or both.[1]

a. Runs of events. The run curve, response probability conditional upon the length of the preceding run of events, most directly reflects the

[1] Not all of these quantities will be independently distributed, and we will often require more parameters to fit them than to fit marginal statistics. Nevertheless, the main point is important; the use of conditional statistics typically increases the number of free predictions.

validity of the assumption that event occurrence is the effective reinforcement. According to SLT, with each successive repetition of the preceding event, probability of predicting that event should increase by some fraction of the distance to the upper bound of unity. Thus, the run curve should be a monotonically increasing, negatively accelerated, function of run length with asymptote of 1.0. It has long been apparent that this prediction of positive recency is not usually supported by data; negative recency, descent of the run curve after several event repetitions, has been observed in numerous studies (for example, Anderson & Whalen, 1960; Jarvik, 1951). Indeed, positive recency is ordinarily obtained only when the average run length is long—a condition met by intentionally constructing sequences with this property (Derks, 1963; Jones & Myers, 1966), or by using π of .7 or higher—and subjects have had at least 100 trials in the task (Derks, 1962; Edwards, 1961) or have had several hundred trials of event observation prior to prediction trials (Reber & Millward, 1968). Even in these cases, the asymptote of the run curve is consistently 10–15% below unity.

Findings of negative recency are not quite as disastrous for SLT as they might at first appear. Recognizing that the "gambler's fallacy" is strongly entrenched in the belief systems of most college students, one might argue that it plays a role in the probability learning task; subjects enter the experiment with a firm conviction that events come in short runs, a conviction that is largely reinforced at low π levels and generally takes several hundred trials to negate. This line of reasoning led Friedman and five coinvestigators (1964) to run subjects for three sessions of 384 trials each. The first two sessions presumably provided an opportunity to wash out the gambler's fallacy; 48-trial blocks with π of .5 were alternated with randomly sequenced blocks with π values ranging from .1 to .9 in increments of .1. In the critical third session, subjects experienced 288 trials with π of .8, preceded and followed by a single block with π of .5. Fits to a wide array of statistics were generally good and, most importantly, positive recency was obtained for both the .8 and terminal .5 series in the third session.

These results, together with other data in which negative recency is replaced by positive recency during the course of the session (Derks, 1963; Edwards, 1961), are consistent with the proposition that extended practice eliminates preexperimental biases, permitting the basic conditioning process envisaged within SLT to be revealed. Anderson (1964) has put forth an alternative interpretation of the results of Friedman and his co-workers. He noted that the probability of an event repetition was greater than .5 over the first two sessions. The schedule, rather than extinguishing preexperimental biases, may have built in a bias to expect runs to continue. According to this interpretation, the good fits of the third session are not the product of an ongoing conditioning process, but rather reflect the subject's memory of the event sequence of previous days. That subjects can learn event repetition probabilities, and that such learning is readily transferred in the presence of new schedules, has been amply demonstrated in several experiments (Anderson, 1960; Witte, 1964).

This argument led Friedman, Carterette, and Anderson (1968) to run subjects for 25 350-trial sessions under a 50 : 50 schedule. Analyses carried out for individual subjects, week by week, revealed a high degree of variability among subjects and over time, as well as a tendency to alternate responses that was at odds with the limited-memory assumption of the linear and pattern models presented earlier. Few subjects exhibited consistent positive recency and run curves typically asymptoted below the predicted value of unity. These failures of SLT were evident within the first week of the experiment.

It is not clear that these data provide broad grounds for rejecting SLT. Rather, they might be viewed as further restricting its range of application. Both the experimenters and Estes (1972) have noted that the use of 50 : 50 schedule may yield unrepresentative results. Prolonged practice with a low degree of success may motivate the subject to seek alternative strategies that will be more productive; this would account for the marked variability in the Friedman *et al.* (1968) study. Furthermore, long runs are infrequent with such a schedule, and thus, the event sequence may have reinforced those biases with which subjects entered the experiment.

b. Runs of successes. In the studies reviewed thus far noncontingent event schedules have been employed. In such schedules, the event, E_1 or E_2, is independent of trial number or of any aspect of the preceding sequence of responses or events. In contrast, Yellott (1969) employed a noncontingent success schedule in which the trial outcome, success or failure of the prediction, was independent of previous outcomes. The probability of a successful outcome—that is, that the event on a trial matched the response—was referred to as δ and was constant over responses. The advantage of such a "noncontingent success" schedule is that it provides a unique opportunity to differentiate the pattern and linear models. In the pattern model, learning occurs only following errors; following a correct response, the sampled element cannot change state and so the proportion of elements conditioned to each response (and, accordingly, response probability) does not change. As a consequence, the pattern model predicts that alternation responses in the noncontingent success task will be independent of the length of the preceding run of successes. On the other hand, in the linear model, only the event and not the outcome, dictates changes in response probability. The linear model predicts that the probability of alternation responses in the noncontingent success task will decrease as the run of preceding successes increases. Using δ of .8 and 1, Yellott clearly demonstrated the correctness of the pattern model prediction. Furthermore, both variances and a variety of sequential statistics were well fit by the pattern model for the .8 block, and the estimate of the learning rate was .173, a value very close to that obtained in experiments using noncontingent event sequences.

These results clearly differentiate between alternative models and also provide strong support for the general proposition that a simple conditioning process underlies probability learning. There is, however, one departure from this pattern of support for SLT. Under the $\delta = 1$ schedule, some subjects exhibited very structured patterns of responding, consistent with the view that they had actively processed the event sequence, encoded it into structured chunks, and then had engaged in hypothesis testing. That sequential information is generally encoded in memory is also suggested by results of the few studies in which direct probes of memory have been employed (Millward & Reber, 1968; Vitz & Hazan, 1969). Perhaps one can best reconcile such evidence of memory processes with the generally excellent account of Yellott's data provided by the pattern model by concluding that his subjects utilized the information stored about the sequence only under the rather artificial $\delta = 1$ condition.

The overall implications of sequential dependency data for SLT are mixed. The form of event–run curves in the Friedman *et al.* (1964) study and of success-run curves in Yellott's study, and quantitative fits to these statistics support SLT. On the other hand, repeated findings of negative recency, the influence of event alternations in the Friedman *et al.* (1968) study, results of memory probes, and Yellott's $\delta = 1$ data also make it clear that information about the pattern of prior events is stored in memory and, except under carefully restricted conditions, influences predictive behavior.

C. Discussion

The broad impact of two decades of theoretical work on predictive behavior is evident. The conceptual and mathematical apparatus developed during the course of this work comprises a major general contribution to psychological theorizing. Regardless of evaluation of the success of SLT in explaining probability learning, the theoretical enterprise demonstrated the viability of an approach in which predictions of the fine details of the course of learning can be derived from a relatively small and basic set of assumptions.

It is more difficult to assess the outcome of the specific issue that generated the bulk of the prediction research. Is binary prediction the result of "the pervasive operation of a rather simple form of conditioning" (Estes, 1964, p. 121)? Certainly, the conditioning framework accounts for much of the data under a wide range of reinforcement schedules, and in several extensions of the simple noncontingent-event task. Nevertheless, fundamental predictions do frequently fail. Most critical for a simple stimulus–response theory is evidence that subjects remember (Millward & Reber, 1968) and respond to (Friedman *et al.* 1968) local event patterns, learn something about the probabilities of such patterns and transfer this knowledge to new sequences (Anderson, 1960; Witte, 1964), and formulate and test hypotheses based on event patterns (Yellott, 1969). In view of the successes of the theory, particularly in its application to the Friedman *et al.* (1964) data and to the early phases of the Yellott (1969) study, it may be that the failures cited serve not to invalidate the theory, but rather to circumscribe its range of application. In short, there is evidence of an underlying probability learning process which may best be characterized as associative learning, but whose operation is easily obscured by the propensity of subjects to seek sequential cues and to incorporate such information into their decision making process.

In retrospect, the need to distangle processing of event probabilities and sequential information is evident. Probability learning may be more directly investigated by collecting subject's estimates of event probabilities. There is already some evidence (Shanteau, 1970) that such estimates are described by Anderson's (1968) linear integration model, which simply assumes that the response is a weighted average of the preceding set of events. Since both Estes' linear model and a Bayesian revision model (Beach *et al.*, 1970) are special cases of the linear integration model (Anderson, 1968; Messick, 1970), they are also tenable. Further investigation of probability estimates may provide a better understanding of the process by which we develop subjective probabilities. Furthermore, such research may well converge with recent work in verbal learning addressed to the general issue of how information about stimulus frequency and recency are represented in memory (for example, Hintzman, 1969; Howell, 1973).

Sequential information processing may be studied profitably in the binary prediction situation with the use of considerably more constrained sequences than have been employed in the research described thus far. By systematically manipulating salient characteristics of the event sequence, such as length of runs or alternations, one is more likely to discover how subjects learn to respond to such sequential stimuli than if one employs random sequences and post hoc analyses of responses to patterns which chanced to occur. Such experimentation has been carried out concurrently with the development of models in which constructs such as encoding, memory, and hypothesis testing play a central role. The remainder of this chapter will deal with such developments.

III. SEQUENCE LEARNING: CONDITIONAL EVENT SCHEDULES

In contrast to earlier research with noncontingent event schedules, no single theoretical position has dominated the research considered next. Consequently, several alternative models of prediction will be described first. These differ in orientation and detail but have the common goal of accounting for such sequential effects as negative recency and the common view that subjects remember event patterns and base their predictions upon that memory. The account of these models will be followed by a summary of relevant experimental findings. Then, armed with the facts, or some approximation thereof, I will attempt to evaluate the models under consideration.

A. Models

1. *Fixed Memory-Span Models*

In the simplest such model (Burke & Estes, 1957), the trace of the immediately preceding event is represented as a unique set of stimulus elements. Thus, on each trial, the subject samples from one of two distinct sets of elements. Application of the conditioning axioms of SLT reveals that the asymptotic probability of a repetition response, $P(A_{i,n}|E_{i,n-1})$, should slightly exceed the true probability that an event will be repeated on the next trial.

Restle (1961, pp. 109–111) has suggested an extension of the event-trace model capable of accounting for the ability of subjects to learn to respond to patterns spanning several trials. He assumes that the last k events are in memory; the response is based on a sample of elements drawn from a set corresponding to that pattern of k events. One strong prediction of the model is that subjects learn to respond without error to any pattern of length k or less that perfectly forecasts the next event. For example, if runs longer than length three never occur, the subject should learn not to make a repetition response following three events of the same type.

2. *Determining the Structure of Sequences*

Feldman and Hanna (1966) have assumed that subjects keep track of the changing conditional probabilities of E_1 following all possible event patterns of length k (five in their paper) or less. The subject first learns to discriminate the longest pair of subsequences S and S' such that the two differ only in the first, or earliest, position (for example, $E_1E_2E_1E_1E_2$ and $E_2E_2E_1E_1E_2$) and $P(E_1|S)$ differs significantly from $P(E_1|S')$. Discrimination is equivalent to probability matching at the level of subsequences; thus, $P(A_1|S)$ and $P(A_1|S')$ would match $P(E_1|S)$ and $P(E_1|S')$, respectively. Shorter subsequences would be discriminated later in learning. This structural-analysis model appears to credit subjects with more memory and computational power than they possess, but it may prove a useful baseline against which to evaluate performance.

3. *Encoding Event Runs*

The models just described place a considerable strain on subject's memories and one's credulity. The strain on memory at least is reduced if one takes a very different tack and assumes that subjects encode only certain types of patterns. In view of the considerable evidence attesting to the salience of event runs (Myers, 1970), a reasonable assumption is that the subject remembers only the type of event he has just seen and the number of times in a row that it has just been presented. The two models presented next assume just such a short-term memory as well as some long-term information about the distribution of run lengths.

a. Restle's (1961) schema model. Restle assumed that the subject attempts to match the run in progress against some schema stored in memory. For example, if the events for the last four trials were $E_2E_1E_1E_1$, the probability of an A_1 response is the probability that the schema $E_2E_1E_1E_1E_1$, rather than $E_2E_1E_1E_1E_2$, is found in long-term memory. The probability of finding a particular schema was assumed to be a function of both its relative frequency of occurrence in the past and its length; Restle assumed that long runs are more salient than short runs in memory. Depending upon the distribution of runs in the event sequence, the schema model can predict either positive or negative recency. The overshooting observed in noncontingent event experiments is also predicted, particularly the pronounced overshooting observed with more than two choices. The schema model makes two qualitative predictions that were also noted for the *k*-span model; first, the probability of a repetition response will exceed the actual probability of an event repetition because of the weight given long runs, and second, subjects will learn always to predict the continuation of a run that has always continued, and the breaking off of a run that has always broken off in the past.

b. Gambino and Myers' (1967) generalization model. These investigators assumed that the subject has a set of expectancies, one for each possible run length that he might encounter in the experiment. If the subject has just seen *m* consecutive events of a particular type, his expectancy that runs of length *m* will continue is increased or decreased by a fraction, depending upon whether the current run does continue. Furthermore, the

continuation (breaking off) of a run of length *m* results in a generalized increase (decrease) in expectancies that runs of other lengths will continue; as one might expect, generalization is assumed to be greatest for run lengths close to *m* in value. The generalization assumption is the critical difference between this and other models capable of predicting negative recency. It provides a mechanism for generating errors at points in the sequence at which the next event is perfectly predictable. For example, suppose the event sequence contains no runs longer than five; the generalization model, in contrast to the models previously described, predicts a greater-than-zero error probability because of generalized expectancies resulting from the continuation of runs of lengths one through four.

B. Results

1. *Responding to Event Contingencies*

A direct test of the Burke–Estes trace model is provided by manipulating π_{11}, the probability that an event occurrence is repeated on the next trial. As the model predicts, the asymptotic probability of a repetition response is close to, but slightly above π_{11} for values of .3 and greater (Anderson, 1960; Engler, 1958; Witte, 1964); this result can also be predicted by the Restle and Gambino–Myers models. A result contrary to the prediction of overshooting was obtained by Anderson who found that the probability of predicting an event repetition was below the event repetition probability for π_{11} of .1 and .2. Further evidence against the Burke–Estes and Restle models is provided by the run curves in Witte's experiments; the Burke-Estes model generally predicted too much positive recency, and the Restle model predicted too much negative recency. These data also are inconsistent with the Feldman–Hanna model; since the occurrence of an event depends only upon the preceding event, the structural analysis model would incorrectly predict flat run curves.

The runs model stimulated several studies in which the event sequences were constructed of a limited number of run lengths. In a typical sequence, equal numbers of runs of lengths 2 and 5 might be randomly ordered; following any run length except two, the next event is perfectly predictable. At the uncertainty point, that is, after a run of length 2 in the example, the probability of a repetition response exceeds the probability of an event repetition (Gambino & Myers, 1966; Restle, 1966), a result predicted by both run models. In addition, the probability of a repetition response increases with the number of long runs (for example, length 5) and decreases with the number of short runs (for example, length 2) immediately preceding the run in progress (Butler, Myers, & Myers, 1969), a result demanded by the reinforcement axioms of the Gambino–Myers model.

However, when contingencies among run lengths are not random, the results are more difficult for any runs model. For example, if the probability is high that long and short runs will alternate, repetition response probability will be higher when the run preceding the current run was short than when it was long (Butler *et al.*, 1969).

The event contingencies considered thus far only scratch the surface of the sequential processing capacity of human subjects. In an attempt to evaluate their structural analysis model, Feldman and Hanna generated a sequence with some contingencies complex enough to defy brief description here. The subjects learned to respond differentially to different patterns; in fact, the probability of predicting an event following each of 62 different patterns of length 5 or less is related to the true probability of the event following that pattern by a linear regression line having slope of .998 and intercept of .013. While Feldman and Hanna have taken this result as evidence of contingency matching, a prediction of their model, it should be noted that there is considerable scatter about this best-fitting function. Nevertheless, the main point is well taken; while runs are extremely salient sequential cues, other patterns can also serve as a basis for prediction.

2. *Errors of Prediction*

All of the models described in the preceding section, except the Gambino–Myers generalization model, predict that the subject will learn to eliminate errors at those points at which the sequence is perfectly predictable. Thus, in a sequence resulting from the random ordering of runs of lengths 2 and 5, subjects should learn to always predict the preceding event if the ongoing run is of length 1, 3, or 4, and to predict the alternate event if the current run is of length 5. Inappropriate failures to make a repetition response will be referred to as anticipatory errors while inappropriate repetition responses will be referred to as perseverative errors. While such errors are clearly less frequent than would be expected if subjects were guessing or merely matching the overall probability of an event repetition, they continue to occur after as many as 700 trials with no indication that further practice would result in improvement; indeed, error rates appear to asymptote within the first 100 trials (Gambino & Myers, 1966; Restle, 1966; Rose & Vitz, 1966). Typically, anticipatory errors increase as mean run length decreases and perseverative errors are an increasing function of the number of run lengths present and of the difference in lengths when only two are present. These error data, as well as repetition responses at the uncertainty point, are well fit by the generalization model (Gambino & Myers, 1967).

Contrary to the predictions of the generalization model, or of any run model, rules about patterns other than runs can also be learned although, again, not to an errorless criterion. Rose and Vitz (1966) found better than chance acquisition of such rules as "if the current pattern is 1121, predict a 2." Wolin, Weichel, Terebinski, and Hansford (1965) have also reported some learning of responses to patterns other than runs.

3. *The Acquisition of Sequential Information*

Of the models considered, only the Feldman–Hanna structural analysis model predicts differential rates of learning different contingencies; the schema and generalization models envisage the encoding only of runs and no one has hypothesized any relationship between learning rates and event patterns for the *k*-span model. Hanna and Feldman's assumption that long subsequences are discriminated first appears to be incorrect; their own data analyses revealed that subjects first responded differentially to patterns which span fewer events. For patterns of the same length, the order in which subjects learned to respond with different probabilities appeared to depend upon how much the patterns differed with respect to probability of the next event.

In contrast to the precise pattern analysis assumed by Feldman and Hanna, Wolin *et al.* (1965) characterize their subjects as first learning general, inexact, aspects of the sequence and later more specific rules. For example, a subject might learn first that the sequence was composed of runs and single alternations and later learn that runs of E_1's had to be an odd-numbered length or even some specific length. Still another view of the learning process is provided by Butler *et al.* (1969). In their data, it appeared that learning of recurrent sequential units (runs of different lengths) and learning of contingencies between such units proceeded concurrently; however, learning the units, as indicated by anticipatory and perseverative errors, appeared to stabilize well before contingency learning was complete.

4. *Transfer*

Once probabilistic contingencies are learned, their effects persist for many trials after the sequence has been changed. At the end of 200 trials of π_{11} of .5, Anderson (1960) observed that the probability of a repetition response was approximately .7 for subjects originally trained with π_{11}'s ranging from .6 to .9 and approximately .5 for subjects originally trained with π_{11}'s of from .1 to .4; there was no indication that the two groups of curves were converging. Still more impressive evidence of transfer was obtained by Witte (1964), who found that run curves still differed as a function of original training after four once-a-week sessions at π_{11} of .5. Such persistent effects of original training are difficult for any of the models under consideration.

Using a four-choice prediction task, Jones and Erickson (1972) have demonstrated transfer of still other types of contingencies. Groups trained to attend to the length and event class of two preceding runs performed better in a transfer task incorporating multirun contingencies than did groups trained only to attend to current run length or given no training. None of the four models presented thus far can account for these results because what was transferred was not knowledge of the consequences of specific patterns—be they runs or various other configurations over several event positions—but rather an abstract rule, perhaps "pay attention to the last two run lengths and the positions in which they appeared."

Following Yellott's (1969) lead, Colker and Myers (1971) employed a transfer phase in which all predictions were correct. Prior to this all-correct phase, subjects experienced one of four types of sequences composed of two run lengths; these varied with respect to length of the two event runs and probability that a run of a particular length would be followed by a run of the same length. Protocols for the all-correct phase were divided into two categories—simple periodic solutions which required that the subject remember only the run in progress and the length of the immediately preceding run (for example, 2/5/2/5/2 . . . , the digits represent run lengths) and complex solutions which encompassed all response patterns that required the subject to remember more than the immediately preceding run length (for example, 2/2/5/5/2//5/5/2/2/5//2/2/ . . .). Within each experimental group, there were significantly more errors in the training phase for subjects who subsequently exhibited complex solutions. Furthermore, those groups with higher mean error rates in training had more complex solutions in the transfer phase. None of the models we have considered encompass such results. They seem to require us to conceptualize a limited-capacity information processing organism; when, either because of preexperimental biases or experimental influences, strategies are pursued which place a heavy load on short-term memory, more predictive errors are made. Butler's (1969) finding that a display of prior events beyond the current run results in more errors than a display of only the current run is also consistent with the limited capacity hypothesis; give a subject information and he will try to use it and, if the information is irrelevant, as prior runs were in Butler's study, the subject's processing will be impaired.

5. *Memory for Binary Patterns*

It is clear that subjects do not learn perfectly in the face of patterns that are perfect predictors of the next event, at least in the sequences under consideration, all of which contain some points at which there is stimulus uncertainty. The Gambino–Myers model provides one mechanism for errors; an alternative explanation is simply that short-term memory is not perfect under the processing load placed upon the subject with such sequences. This possibility suggests that one look at memory for event patterns.

Glanzer and Clark (1963) provided subjects with brief exposures to sets of eight binary figures. Accuracy of subsequent reproduction was best for runs of eight like figures, deteriorated as the number of runs increased, but was considerably improved again for single alternations. Millward and Reber (1968) obtained similar results in the context of a binary prediction task, probing on each trial for a specific event k trials back. As would be expected, recall accuracy decreased as a function of depth of probe. Of somewhat greater interest is the finding that, with depth held constant, recall accuracy decreased as a function of the number of runs intervening between the target item and the probe, with one exception: when a single alternation intervened, there was a decided upturn in the recall function. The results of these two experiments suggest that the amount remembered depended upon the patterning of events, a result at odds with both the fixed-span and Feldman–Hanna models. Furthermore, as Restle and Gambino and Myers have assumed, runs were basic units of information storage. On the other hand, other patterns were recalled above a chance level.

Vitz and Hazan (1969) also probed memory in the context of the typical binary prediction task. They, however, probed only three times during the session, asking subjects to recall as much as they could of the preceding event series. Consistent with the implications of several transfer studies cited earlier (Anderson, 1960; Witte, 1964), rather accurate long-term memory was exhibited.

The relationship between prediction and memory is clearer in another study (Myers, 1970). Probes of memory for recent events yielded more errors for those experimental conditions in which errors of prediction were greatest. Thus, contrary to all of the models considered, short-term memory is fallible and, in fact, correlated with predictive errors. Of additional interest is the distribution of remembered run lengths. In more than 97% of 1,440 probes of memory, the preceding run length was either remembered correctly or was remembered as shorter than it actually was. As might be expected, the distribution function was monotonic; the correct run was reported most often and the probability of a particular run length being reported decreased as a function of distance from the correct value.

C. Discussion

It is clear that, as Restle, and Gambino and Myers, have hypothesized, event runs are units of encoding and information about the events following each run length is available in long-term memory. It is equally clear that, contrary to the runs models, other patterns of events are encoded. It appears that an appropriate model will incorporate the assumption-that what is remembered and responded to is not a fixed-span of events but rather a variable event span whose length depends upon the particular pattern of events. Furthermore, such a model will be all too human; unlike the current model, it will have a fallible short-term memory.

Two possible mechanisms for memory failure are suggested by the Myers (1970) data, in particular, by the finding that run lengths were almost always remembered as shorter than they actually were if they were remembered incorrectly. First, errors may occur at event input; with some probability, a counter may fail to register the incoming event. Second, errors may occur at retrieval. For example, suppose the subject has correctly registered a run of length five. Then, he has recently registered a run of length four, before that, a run of length three, and so on. Assuming that traces of inputs vary in strength as a function of proximity to the probe, and that the probability of retrieving a trace when probed is proportional to strength, something very much like the observed distribution of reported run lengths would be obtained. Distinguishing between these two positions will be of interest. In addition, memory probe experiments employing other sequential constraints are required to determine whether the apparent information loss is a general phenomenon, or is peculiar to the run-structured sequences employed in the experiment described.

Transfer studies demonstrate long-term perseverative effects of training with event contingencies, effects that seem to be beyond the scope of the models considered. It appears that, having once formed hypotheses about the sequences, subjects give greater weight to those event patterns that support their hypotheses, only slowly changing their response pattern as evidence accumulates that the sequential structure has changed. In line with this view, probes of memory during transfer might reveal that subjects trained differently would differ in what they recall of the common transfer sequence; subjects trained with high event repetition probabilities might be more accurate in recalling prior run lengths, whereas subjects trained with low event repetition probabilities might be more accurate in recalling event alternations.

The Colker and Myers (1971) results suggest a limited-capacity information processing system. If this is so, failures in short-term memory, and perhaps the amount of experience required for appropriate revision of hypotheses during transfer, will be a function of the amount of material stored in short-term memory and the complexity of rules for response selection. This suggests that part of the subject's problem in eliminating anticipatory and perseverative errors may lie in striving for perfection. As numerous investigators have noted, even in the face of more clearly random sequences subjects believe that there is a solution, a strategy that will

completely eliminate errors. Presumably, the subject develops and tests progressively more complex hypotheses, and, therefore, stores more information, in an attempt to find the solution. The result is a loss of information from an overburdened processing system with the consequence that errors which could be eliminated are not. One further implication of this view of the subject as a limited-capacity information processing system is that a sufficient model cannot merely incorporate a mechanism for memory failures (whether at input or retrieval) but must also specify how such failures depend upon the information processing load. This implies, in turn, a need for a more precise definition of such terms as "processing load" and "sequence complexity."

A clearer picture of sequential processing has emerged from studies in which the number and complexity of hypotheses has been reduced through the use of repeated event patterns. In particular, such studies have provided further data relevant to two related questions: What determines the subject's response at any position within the pattern? What determines the relative level of difficulty of responses at different positions within a pattern? The next section provides a discussion of models and data within the context of such deterministic sequences.

IV. SEQUENCE LEARNING: DETERMINISTIC SCHEDULES

In this section, how subjects learn to predict sequences that consist of repetitions of a single pattern is considered. The pattern could be a simple binary one, for example, *aaabb*. At the other extreme, it may involve more than two events and be generated by a relatively complex set of rules. Regardless of the complexity of the sequence, interest here will focus on the serial position function. Of primary concern will be the ability of different models to predict variation in error rate over positions in the pattern as a function of pattern structure.

A. Models

1. *An Association Model*

Vitz and Todd (1967) have proposed a model to account for prediction of repeated simple binary patterns of the form m a's followed by n b's, for example, *aaabb*. They assume that the pattern can be viewed as a set of stimuli where each stimulus is the run preceding a position. Thus, if the pattern *aaabb* is repeatedly presented, the stimuli and responses are: $a \rightarrow a$, $aa \rightarrow a$, $aaa \rightarrow b$, $b \rightarrow b$, $bb \rightarrow a$. The stimulus–response connections are assumed to be learned all-or-none with probability c, as in Bower's (1961) one-element model of paired-associate learning. If the association is not learned, the subject guesses the correct response with probability g. Such a model predicts stationarity, constant probability of a correct response over trial blocks prior to the last error to a stimulus, as well as independence of responses. Furthermore, assuming that c is constant over positions, error rates should not vary significantly with position in the sequence or with length of runs of a's and b's.

The scope of the model is quite limited; for example, it fails to account for patterns such as *aabab* because the stimulus a is followed by an a at Position 2, and by a b at Position 5. Thus, the preceding run does not contain sufficient information to allow the subject to learn the entire pattern. Nevertheless, if the model successfully accounts for learning of the simple patterns for which it was designed, it might be elaborated to deal with the more complex patterns much in the way that Bower's one-element model for paired-associate learning was extended to account for additional stages such as stimulus differentiation and response integration.

2. *A Two-Stage Model*

Restle (1967) has provided one possible elaboration of the model sketched above. He distinguishes between "mandatory" and "optional" positions; in the pattern *aabab*, Positions 2 and 5 are optional in the sense that the subject's response is not completely determined by the preceding run length. Restle's account of the learning of mandatory positions is essentially the same as Vitz and Todd's although he prefers to speak of learning mandatory rules rather than forming associations. Responses at optional positions are assumed to require a second all-or-none stage to be learned; thus, error rates should be higher at optional than at mandatory positions.

3. *Hierarchical Rule Learning Model*

The two models described above deal with prediction of binary sequences. A richer conceptualization incorporating more, and more complex, hypotheses, becomes possible when one considers multichoice tasks. Restle (1970), and Restle and Brown (1970), have analyzed the learning of such tasks. They have assumed that subunits are learned first and then higher-order rules, which integrate such subunits. The model has great scope, predicting the relative difficulty of sequences, the relative difficulty of positions within sequences, the effects of manipulations within a few positions in sequences, and transfer effects. Three operators are defined (others are possible); $T(X)$ is a transposition operator and implies incrementing

each element in X by one; if X is the single element 4, $T(X)$ implies 45, and if X is the subunit 123, $T(X)$ is 123234. The second operator, $R(X)$ is a repetition operator; thus $R(5)$ implies 55 and $R(23)$ implies 2323. The third operator is a mirror-image operator which implies subtracting each element in X from 7 (or in general, from one more than the total number of events); thus, $M(2)$ implies 25 and $M(14)$ implies 1463. Now consider the sequence 12126565. This may be generated by letting $X = 1$, applying T, then R, then M; one may represent this sequence of operations as $M(R(T(1)))$. Restle and Brown assume that error rates will be a function of the level of the rule applied to generate a response for a position in the sequence; level corresponds to distance to the left of X in this notation. Thus, Position 5, which is generated by the leftmost, or M, operation, should prove most difficult. Positions 3 and 7, which require the R operation to generate a correct response, should be next in error rate. Events at Positions 2, 4, 6, and 8 are generated by the transposition operation which is the lowest level (rightmost) operation, and should, therefore, be easiest to learn. Two additional predictions are immediately evident. First, if one divides the pattern into subunits and rearranges these so that the new sequence is no longer generated by a single rule hierarchy, the sequence should be considerably more difficult to learn. Second, transfer to a new sequence employing the same rule hierarchy but a different element X should be positive.

B. Results

1. *Binary Event Sequences*

a. Mandatory positions. Both Vitz and Todd (1967) using patterns with only mandatory positions, and Restle (1967) using more complex patterns, have found that an all-or-none learning model provides a reasonable fit to data at such positions. Furthermore, Vitz and Todd tested two critical predictions of the all-or-none model, stationarity of error probabilities and independence of responses, and obtained no significant difference. Nevertheless, there is at least one problem with the model. Both Vitz and Todd, and Derks and House (1965), have noted higher error rates at the first and last positions in a run than at other positions. One not very plausible interpretation of such a result is that c, the conditioning probability, varies as a function of position in the run. It is more likely that subjects lose track of the preceding run length; several of Vitz and Todd's subjects reported miscounting, and data considered in Section III clearly demonstrate that this happens in nondeterministic sequences.

Another possible source of difficulty, even in simple patterns, is suggested by a result obtained by Garner and Gottwald (1967). Using a simple *aaabb* pattern, they found that prediction errors were most frequent following a run of two *a*'s. The pattern provided greater difficulty when the second *a* was presented on trial 1, that is, *aabbaaa* . . . On the other hand, *aaabbaaa* . . . was relatively easy for subjects. Apparently, subjects develop expectancies which conform to some simple hypothesis about the structure of the pattern; in this case, subjects appear to expect a double alternation. Errors pile up at positions that deviate from the expected structure. Furthermore, the extent to which early trials confirm the subject's expectancy influences difficulty of learning the pattern.

b. Optional positions. As Restle (1967) hypothesized, optional positions, those at which the response is not completely determined by the immediately preceding run length, do have higher error rates than mandatory positions. This is evident in both his own data and those obtained by Derks and House (1965). Furthermore, the frequency distribution for errors is well fit by Restle's two-stage model. However, in both studies there is variation in error rates among optional positions and, in some of the Derks and House sequences, some mandatory positions have high error rates. These are generally at the end of runs longer than length three, suggesting miscounting. In short, there is support for the hypothesis of all-or-none learning of stages but Restle's model, like Vitz and Todd's, fails to provide a complete account of variations in error rate over positions in the pattern.

A closer look at how the second, optional position, stage is learned is instructive. Consider one sequence employed by Restle:

Event:	*a*	*a*	b	*b*	a	*b*	*b*	a	*b*
Position:	1	2	3	4	5	6	7	8	9

Italicized letters denote optional positions. Note that Positions 1 and 7 demand different responses, but the preceding five events, or three runs, are identical. No other pairs of positions require this much information in order to be discriminated. Nevertheless, only Position 1 was clearly more difficult than other optional points; there were actually slightly more errors at Positions 2 and 9 than at 7. These results argue against an associative theory in which combinations of preceding events or runs are a discriminative cue for prediction. If such a theory were correct, one would expect both Positions 1 and 7 to be considerably more difficult than other positions on two grounds. First, it is reasonable to assume that it takes more trials to integrate a longer discriminative cue; if this were not true, there is no reason for the consistent differences in error rate between mandatory and optional positions. Second, longer discriminative cues place a greater burden on memory and should result in more miscounting and thus more errors. That this relationship between memory load and processing exists was amply documented in Section III. Restle has proposed an alternative account of Stage 2 learning. He assumes that following the learning of mandatory positions the subject develops some general rule that integrates the individual runs. Errors should develop at positions which depart from the general rule since such positions have to be learned as exceptions. In the example under consideration, the sequence can be described as a repetition of the simple pattern *bba* except for Position 1, which has a notably higher error rate than other positions. This analysis is consistent with Garner and Gottwald's account of simple pattern learning and confirms their emphasis on the role of expectancies.

2. *N-ary Sequences* ($N > 2$)

Restle and Brown (1970) have reported a series of ten experiments in which a six-choice task was employed. Such experiments are an advance over binary choice experiments because they permit greater variation in the rules used to construct sequences and because errors are more informative; one notes not only where errors occur, but also which erroneous response occurred.

Learning of such sequences is dictated by abstract properties of the sequence, such as scales (for example, 2345) and trills (for example, 2323), rather than by specific events or numerical intervals between events. Restle and Brown first demonstrated this by comparing learning of an initial sequence with learning of three other sequences derived from it by either transposing the initial sequence by one event (for example, 123 becomes 234), inverting it (123 becomes 654, or by transposing and inverting it (123 becomes 543). Sequence had little effect on error rates. More important to the point at issue, sequences did not differ with respect to the profiles of error rates over positions.

Errors in this study were typically most frequent at the beginning of the scale and trill subunits and were the result of overextension of the previous subunit. Thus, in a sequence beginning 234654, the fourth position would have a high error rate and the error would frequently be the prediction of a 5 instead of 6. Fritzen and Johnson (1969) have shown that at least one cause of such errors is the subject's failure to recognize the end of the subunit; error rates were lower when all subunits in a sequence were of the same length or when subunits could not physically be continued (for example, 321).

Additional evidence for the importance of subunits such as runs and trills was obtained by using an ambiguous sequence. For example, 6543432345 may be chunked either as 6543/432/345 (scales) or as 65/434/323/45 (trills). Different error profiles can be produced by pretraining on sequences clearly composed of scales or of trills, or by pretraining on an ambiguous sequence with pauses introduced which chunk the sequence into scales or into trills. It is quite evident that the errors produced in transfer are a result of the induced expectancy. For example, in the ambiguous sequence presented above, subjects trained to expect scales should make more errors at Position 5 than subjects trained to expect trills, and the error for the scales group should usually be a 2; this is what was observed.

Given that subjects learn rules which guide them through subunits of the overall pattern, how are the subunits integrated? Restle's (1970) hierarchical rule learning model provides an answer which receives considerable support from the data of the last five Restle and Brown experiments. Several results are of particular relevance. First, the error rate at a position in a sequence appears to depend not upon the particular rule required to generate the response at that position, but rather upon where that rule lies in the rule hierarchy. For example, given the sequence 12122323121223236565545465655454, which is so much more elegantly represented by $M(R(T(R(T(1)))))$, we find most difficulty at the seventeenth position, where the M operation must be applied, somewhat less at Positions 9 and 25, where the higher-order (leftmost) R operation must be applied, and so on. The sequence generated by $T(M(R(T(R(1)))))$ has much the same error profile indicating that position in the rule hierarchy, rather than the specific rule applied, is critical.

Second, there is some evidence that when a sequence can not be described by a hierarchy of rules, performance at transition points between subunits is poor. Restle and Brown found that a sequence obtained by randomly ordering subunits of a sequence generated by a rule hierarchy had more errors at the first position within subunits; typically, these were due to overextension of the previous subunit. This result not only confirms the importance of the rule hierarchy in integrating subunit learning but also indicates that lower-order regularities can be learned in the absence of higher-order rules.

Third, there is evidence that rule hierarchies are learned from the bottom up, or from right to left in terms of the operator notation we have used. In essence, subjects appear to first determine the lowest element X, then to learn the operation on it, for example $T(X)$, gradually building up the hierarchy. At each level, the repetition operator R is tried out first, resulting in errors. This is natural since the subject knows that the whole pattern is repeated and does not know the pattern length. This interpretation suggests that subjects should learn a complex pattern most readily if small subunits are first learned and then the pattern is extended by introducing rules gradually in a lower- to higher-order direction. The conjecture was borne out; subjects had less difficulty in learning the sequence $T(M(T(R(T(1)))))$ when preceded by a series of training blocks in the order $T(1)$, $R(T(1))$, $T(R(T(1)))$, $M(T(R(T(1))))$ than when preceded by training blocks in the order $T(1)$, $T(M(1))$, $T(M(T(1)))$, $T(M(T(R(1))))$.

It is possible that the results considered are specific to the prediction task, that they depend upon verbal rehearsal and a conscious and deliberate encoding and storing of information. Two experiments, employing rather different methodologies, have addressed this issue. Garner and Gottwald (1968) required subjects to observe repeating patterns of eight binary events, presented either visually or auditorily, and at one of five rates of presentation. Subjects stopped the presentation when they felt able to provide a description of the pattern. Stopping point and accuracy of description were differently affected by the independent variables at high and low speeds, a result which led the authors to conclude that processing at low speeds ("learning") is an active intellectualized process, while at high speeds ("perception") it is a passive experience of an integrated sequence. Restle and Burnside (1972), who required subjects to track six-event sequences by pushing the appropriate button when a light came on, reached a very different conclusion. Redoing several of the Restle and Brown prediction experiments with the tracking task, they found similar effects on error profiles. In contrast to Garner and Gottwald, they concluded that learning and perception of serial patterns are closely related, that both involve a rather rapid organization of serial information, an organization controlled by rule hierarchies.

C. Discussion

The data present difficulties for associative models. The generally intricate profiles of errors argue against simple conditioning of responses to serial positions. Simple conditioning on the basis of last event, the models with which we began this chapter, also clearly does not stand up to the results. Conditioning of responses to the immediately preceding run, essentially the Vitz–Todd model, does not completely describe even data from simple binary patterns, although the addition of some forgetting mechanism might suffice to resolve the problem. Nevertheless, it is evident that a complex conditioned stimulus must be assumed if one is to cope with optional positions in binary sequences, or with the results obtained in multichoice prediction tasks.

Conceivably, the conditioned stimulus may be some preceding set of events or runs. However, as was noted in considering the Restle (1967) data, and as Restle and Brown (1970) have also found in their six-choice experiments, there is no clear relationship between errors at a position and the memory load required for accurate response discrimination. Furthermore, Restle and Brown have also noted that changes in the sequence may have effects at some position quite distant from the locus of change. Developing an adequate account of the data from an associative framework may prove quite challenging.

Restle and Brown's (1970) rule hierarchy model provides an alternative frame of reference. It is assumed that subjects discriminate some basic element, learn to operate on this, and proceed to concatenate operators until the sequence is learned. Operators that are applied to larger subsequences are learned later. Errors occur at the beginning of subunits and these are erroneous extensions of the currently applied operator. Errors also apply at points where there is ambiguity about the nature of the subunit and the subject makes the wrong decision on the basis of prior expectations. Finally, errors occur because of premature attempts to apply the repetition operator, presumably owing to a lack of definition of the pattern length.

Several issues merit consideration. First, why do errors accumulate at initial positions in subunits? Fritzen and Johnson (1969) have noted several possible sources of such errors. The subject may not know when to terminate the preceding subunit. Their data indicate that this is a problem for the subject. It is not clear whether this is because the subject has miscounted the length of the preceding subunit or has not learned what its length should be. Furthermore, some errors may occur when subjects know that a new operator is to be applied but are not sure which operator or know which operator but are not sure of which number to begin with. The relative contributions of these factors should be assessed.

Second, what is the nature of the learning process? How is each new rule acquired? Is the process all-or-none, as Restle (1967) at one time hypothesized? Are all rules equally easy to learn? Or does the ease of acquiring a rule depend upon the rule and perhaps upon the nature of the rules that have already been learned? There is little information on this, but it would be surprising if all rules were really equivalent in difficulty and if difficulty of a particular rule did not depend upon the overall configuration of rules.

Third, the Garner and Gottwald (1968) and Restle and Burnside (1972) results on perception versus learning of patterned information lead to discrepant conclusions. These studies differ in many ways—in the sequences, independent variables, and dependent variables. Determining the conditions under which the same sequences yield equivalent results with fast and slow event input rates seems fundamental to delineating the generality of any theory of sequential information processing and to defining the boundary conditions for its application. With respect to generality of such theories, Jones' (1974) discussion of serial patterns provides a useful consideration of a wide range of theories and their validity for tasks ranging from prediction to perception to recall of sequences.

Finally, given some representation of rules, how does the subject generate his response? Greeno and Simon (1974) have noted that the same pattern description can be a base for very different procedures for deriving the appropriate response. These procedures differ in short-term memory requirements—how many and which elements must be immediately available—and in the number of operations which must be applied in order to arrive at the appropriate response. Experimental investigation of the process of response generation appears to be a logical next step for investigators in this field.

V. CONCLUSIONS

The course of prediction research and theorizing has been a steady progression through three overlapping stages. The first was marked by the use of random sequences in experimentation and very simple associative models in theorizing. It served to set boundary conditions on the applicability of the models under consideration. That early work has generated two paths—more direct studies of how event probabilities are represented and of how patterns of events are processed.

The second stage was marked by the introduction of constrained sequences in which stimulus uncertainty was reduced but not eliminated. Associative models became more complex than previously, with the run, rather than the event, being a candidate for the discriminative stimulus. In addition, models were introduced that were more oriented toward the information processing approach developing in other fields of learning and in perception. The importance of a run of events as a salient stimulus became clear but at the same time it was also evident that subjects could encode, remember, and respond to considerably more complex patterns. In addition, the need to consider the capacity of the information processing system became apparent.

In the current, third, stage of investigation, subjects have been faced with recurrent patterns of events. The resulting data place the difficulties encountered by an associative theory in full perspective. Progress toward an information-processing model has been made. Prediction of future stages of theoretical development is a risky affair, but it would seem that such work will center about delineation of rules not yet considered, the relative difficulty of learning different rules, the nature of the rule-learning process, the generality across paradigms of various models, and a clearer definition of the information processing stages, with particular emphasis on the relationship between pattern representation and response generation.

REFERENCES

Anderson, N. H. Effect of first-order conditional probability in a two-choice learning situation. *Journal of Experimental Psychology*, 1960, **59**, 73–93.

Anderson, N. H. An evaluation of stimulus sampling theory. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press, 1964.
Anderson, N. H. A simple model for information integration. In R. P. Abelson *et al*. (Eds.), *Theories of cognitive consistency: A source book*. Chicago: Rand-McNally, 1968.
Anderson, N. H., & Whalen, R. E. Likelihood judgments and sequential effects in a two-choice probability learning situation. *Journal of Experimental Psychology*, 1960, **60**, 111–120.
Atkinson, R. C., & Estes, W. K. Stimulus sampling theory. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology*. Vol. II. New York: Wiley, 1963.
Beach, L. R., Rose, R. M., Sayeki, Y., Wise, J. A., & Carter, W. B. Probability learning: Response proportions and verbal estimates. *Journal of Experimental Psychology*, 1970, **86**, 165–170.
Beach, L. R., & Swensson, R. G. Instructions about randomness and run dependency in two-choice learning. *Journal of Experimental Psychology*, 1967, **75**, 279–282.
Black, R. W. Shifts in magnitude of reward and contrast effects in instrumental and selective learning. *Psychological Review*, 1968, **75**, 114–126.
Bower, G. H. Choice-point behavior. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory*. Stanford, California: Stanford University Press, 1959.
Bower, G. H. Application of a model to paired-associate learning. *Psychometrika*, 1961, **26**, 255–280.
Braveman, N. S., & Fischer, G. J. Instructionally induced strategy and sequential information in probability learning. *Journal of Experimental Psychology*, 1968, **76**, 674–676.
Burke, C. J., & Estes, W. K. A component model for stimulus variables in discrimination learning. *Psychometrika*, 1957, **22**, 133–146.
Butler, P. A. The role of information in choice behavior. Unpublished doctoral dissertation, University of Massachusetts, Amherst, 1969.
Butler, P. A., Myers, N. A., & Myers, J. L. Contingencies among event runs in binary prediction. *Journal of Experimental Psychology*, 1969, **79**, 424–429.
Castellan, N. J., Jr. Effect of change of payoff in probability learning. *Journal of Experimental Psychology*, 1969, **79**, 178–182.
Cole, M. Search behavior: A correlation procedure for three-choice probability learning. *Journal of Mathematical Psychology*, 1965, **2**, 145–170.
Colker, R., & Myers, J. L. Effects of sequential structure upon binary prediction under an all-correct procedure. *Journal of Experimental Psychology*, 1971, **89**, 416–418.
Cotton, J. W., & Rechtschaffen, A. Replication report: Two- and three-choice verbal conditioning phenomena. *Journal of Experimental Psychology*, 1958, **56**, 96–97.
Derks, P. L. The generality of the "conditioning axiom" in human binary prediction. *Journal of Experimental Psychology*, 1962, **63**, 538–545.
Derks, P. L. Effect of run length on the "gambler's fallacy." *Journal of Experimental Psychology*, 1963, **65**, 213–214.
Derks, P. L., & House, J. I. Effect of event run structure on prediction of recursive binary sequences. *Psychological Reports*, 1965, **17**, 447–456.
Detambel, M. H. A test of a model for a multiple-choice behavior. *Journal of Experimental Psychology*, 1955, **49**, 97–104.
Edwards, W. Probability learning in 1000 trials. *Journal of Experimental Psychology*, 1961, **62**, 385–394.
Engler, J. Marginal and conditional stimulus and response probabilities in verbal conditioning. *Journal of Experimental Psychology*, 1958, **55**, 303–317.
Estes, W. K. Toward a statistical theory of learning. *Psychological Review*, 1950, **57**, 94–107.
Estes, W. K. A random walk model for choice behavior. In K. J. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences*. Stanford, California: Stanford University Press, 1960.

Estes, W. K. Theoretical treatments of differential reward in multiple-choice learning and two-person interactions. In J. H. Criswell, H. Solomon, & Suppes (Eds.), *Mathematical models in small group processes*. Stanford, California: Stanford University Press, 1962.
Estes, W. K. Probability learning. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press, 1964.
Estes, W. K. Research and theory on the learning of probabilities. *Journal of the American Statistical Association*, 1972, **67**, 81–102.
Estes, W. K., & Straughan, J. H. Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal of Experimental Psychology*, 1954, **47**, 225–234.
Feldman, J., & Hanna, J. F. The structure of responses to a sequence of binary events. *Journal of Mathematical Psychology*, 1966, **3**, 371–387.
Flood, M. M. Environmental non-stationarity in a sequential decision-making experiment. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes*. New York: Wiley, 1954.
Friedman, M. P., Burke, C. J., Cole, M., Keller, L., Millward, R. B., & Estes, W. K. Two choice behavior under extended training with shifting probabilities of reinforcement. In R. C. Atkinson (Ed.), *Studies in mathematical psychology*. Stanford, California: Stanford University Press, 1964.
Friedman, M. P., Carterette, E. C., & Anderson, N. H. Long-term probability learning with a random schedule of reinforcement. *Journal of Experimental Psychology*, 1968, **78**, 442–455.
Fritzen, J., & Johnson, N. F. Definiteness of pattern ending and uniformity of pattern size: Their effects upon learning number sequences. *Journal of Verbal Learning and Verbal Behavior*, 1969, **8**, 575–580.
Gambino, B., & Myers, J. L. Effect of mean and variability of event run length on two-choice learning. *Journal of Experimental Psychology*, 1966, **72**, 904–908.
Gambino, B., & Myers, J. L. Role of event runs in probability learning. *Psychological Review*, 1967, **74**, 410–419.
Gardner, R. A. Probability-learning with two and three choices. *American Journal of Psychology*, 1957, **70**, 714–185.
Gardner, R. A. Multiple-choice decision behavior. *American Journal of Psychology*, 1958, **71**, 710–717.
Garner, W. R., & Gottwald, R. L. Some perceptual factors in the learning of sequential patterns of binary events. *Journal of Verbal Learning and Verbal Behavior*, 1967, **6**, 582–589.
Garner, W. R., & Gottwald, R. L. The perception and learning of temporal patterns. *Quarterly Journal of Experimental Psychology*, 1968, **20**, 97–109.
Glanzer, M., & Clark, W. H. The verbal loop hypothesis: Binary numbers. *Journal of Verbal Learning and Verbal Behavior*, 1963, **2**, 301–309.
Goodnow, J. J. Determinants of choice distribution in two-choice situations. *American Journal of Psychology*, 1955, **68**, 106–116.
Greeno, J. G., & Simon, H. A. Processes for sequence production. *Psychological Review*, 1974, **81**, 187–198.
Halpern, J., & Dengler, M. Utility and variability: A description of preferences in the uncertain outcome situation. *Journal of Experimental Psychology*, 1969, **79**, 249–253.
Halpern, J., Dengler, M., & Ulehla, Z. J. The utility of event variability in two choice probability learning. *Psychonomic Science*, 1968, **10**, 143–144.
Halpern, J., Schwartz, J. A., & Chapman, R. Simultaneous and successive contrast effects in human-probability learning. *Journal of Experimental Psychology*, 1968, **77**, 581–586.
Hintzman, D. L. Apparent frequency as a function of frequency and the spacing of repetitions. *Journal of Experimental Psychology*, 1969, **80**, 139–145.
Howell, W. C. Representations of frequency in memory. *Psychological Bulletin*, 1973, **80**, 44–53.

Humphreys, L. G. Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology,* 1939, **25,** 294–301.
Jarvik, M. E. Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology,* 1951, **41,** 291–297.
Jones, M. R. From probability learning to sequential processing: A critical review. *Psychological Bulletin,* 1971, **76,** 153–185.
Jones, M. R. Cognitive representations of serial patterns. In B. R. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition.* Hillsdale, New Jersey: Lawrence Erlbaum Assoc., 1974.
Jones, M. R., & Erickson, J. L. A demonstration of complex rule learning in choice prediction. *American Journal of Psychology,* 1972, **85,** 249–259.
Jones, M. R., & Myers, J. L. A comparison of two methods of event randomization in probability learning. *Journal of Experimental Psychology,* 1966, **72,** 909–911.
Luce, R. D., & Suppes, P. Preference, utility, and subjective probability. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology.* Vol. III. New York: Wiley, 1965.
McCracken, J., Osterhout, C., & Voss, J. F. Effects of instructions in probability learning. *Journal of Experimental Psychology,* 1962, **64,** 267–271.
Messick, D. M. The utility of variability in probability learning. *Psychonomic Science,* 1965, **3,** 355–356.
Messick, D. M. Learning probabilities of events: A discussion. *Acta Psychologica,* 1970, **34,** 172–183.
Millward, R. B., & Reber, A. S. Event-recall in probability learning. *Journal of Verbal Learning and Verbal Behavior,* 1968, **7,** 980–989.
Myers, J. L. Sequential choice behavior. In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory.* Vol. 4. New York: Academic Press, 1970.
Myers, J. L., & Atkinson, R. C. Choice behavior and reward structure. *Journal of Mathematical Psychology,* 1964, **1,** 170–203.
Myers, J. L., Fort, J. G., Katz, L., & Suydam, M. M. Differential monetary gains and losses and event probability in a two-choice situation. *Journal of Experimental Psychology,* 1963, **66,** 521–522.
Myers, J. L., Gambino, B., & Jones, M. R. Response speeds in probability learning. *Journal of Mathematical Psychology,* 1967, **4,** 473–488.
Myers, J. L., Suydam, M. M., & Heuckeroth, O. Choice behavior and reward structure: Differential payoff. *Journal of Mathematical Psychology,* 1966, **3,** 458–469.
Neimark, E. D. Effects of type of nonreinforcement and number of alternative responses in two verbal conditioning situations. *Journal of Experimental Psychology,* 1956, **52,** 209–220.
Neimark, E. C., & Shuford, E. H. Comparison of predictions and estimates in a probability learning situation. *Journal of Experimental Psychology,* 1959, **57,** 294–298.
Nies, R. C. Effects of probable outcome information on two-choice learning. *Journal of Experimental Psychology,* 1962, **64,** 430–433.
Peterson, C. R., & Ulehla, Z. J. Sequential patterns and maximizing. *Journal of Experimental Psychology,* 1965, **69,** 1–4.
Reber, A. S., & Millward, R. B. Event observation in probability learning. *Journal of Experimental Psychology,* 1968, **77,** 317–327.
Restle, F. *Psychology of judgment and choice.* New York: Wiley, 1961.
Restle, F. Run structure and probability learning: Disproof of Restle's model. *Journal of Experimental Psychology,* 1966, **72,** 382–389.
Restle, F. Grammatical analysis of the prediction of binary events. *Journal of Verbal Learning and Verbal Behavior,* 1967, **6,** 17–25.

Restle, F. Theory of serial pattern learning: Structural trees. *Psychological Review,* 1970, **77,** 481–495.
Restle, F., & Brown, E. R. Organization of serial pattern learning. In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory.* Vol. 4. New York: Academic Press, 1970.
Restle, F., & Burnside, B. L. Tracking of serial patterns. *Journal of Experimental Psychology,* 1972, **95,** 299–307.
Rose, R. M., & Vitz, P. C. The role of runs in probability learning. *Journal of Experimental Psychology,* 1966, **72,** 751–760.
Schnorr, J. A., Lipkin, S. G., & Myers, J. L. Level of risk in probability learning: Within- and between-subjects designs. *Journal of Experimental Psychology,* 1966, **72,** 497–500.
Schnorr, J. A., & Myers, J. L. Negative contrast in human probability learning as a function of incentive magnitudes. *Journal of Experimental Psychology,* 1967, **75,** 492–499.
Shanteau, J. C. An additive model for sequential decision making. *Journal of Experimental Psychology,* 1970, **85,** 181–191.
Siegel, S. Theoretical models of choice and strategy behavior: Stable state behavior in the two-choice uncertain outcome situation. *Psychometrika,* 1959, **24,** 203–216.
Siegel, S. Decision making and learning under varying conditions of reinforcement. *Annals of the New York Academy of Science,* 1961, **89,** 766–782.
Suppes, P., & Atkinson, R. C. *Markov learning models for multiperson interaction.* Stanford, California: Stanford University Press, 1960.
Swensson, R. G. Incentive shifts in a three-choice situation. *Psychonomic Science,* 1965, **2,** 101–102.
Tune, G. S. Response preferences: A review of some relevant literature. *Psychological Bulletin,* 1964, **61,** 286–302.
Vitz, P. C., & Hazan, D. N. Memory during probability learning. *Journal of Experimental Psychology,* 1969, **80,** 52–58.
Vitz, P. C., & Todd, T. C. A model for simple repeating binary patterns. *Journal of Experimental Psychology,* 1967, **75,** 108–117.
Witte, R. S. Long-term effects of patterned reward schedules. *Journal of Experimental Psychology,* 1964, **68,** 588–594.
Wolin, B. R., Weichel, R., Terebinski, S. J., & Hansford, E. A. Performance on complexity patterned binary event sequences. *Psychological Monographs,* 1965, **79,** 1–18.
Yellott, J. I., Jr. Probability learning with noncontingent success. *Journal of Mathematical Psychology,* 1969, **6,** 541–575.