NASA TECHNICAL NOTE

NASA TN D-5684

c. 1

NASA TN D-5684

PROBABILITY LEARNING: FIRST-ORDER MARKOV STRUCTURES OF QUARTERNARY EVENTS

by Edward M. Huff

Ames Research Center
Moffett Field, Calif.

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION • WASHINGTON, D. C. • FEBRUARY 1970

NASA TN D-5664

PROBABILITY LEARNING: FIRST-ORDER MARKOV STRUCTURES OF QUARTERNARY EVENTS

By Edward M. Huff

Ames Research Center
Moffett Field, Calif.

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

For sale by the Clearinghouse for Federal Scientific and Technical Information
Springfield, Virginia 22151 – CFSTI price $3.00

PROBABILITY LEARNING: FIRST-ORDER MARKOV STRUCTURES OF QUARTERNARY EVENTS

By Edward M. Huff

Ames Research Center

SUMMARY

In many control tasks, man is required to make decisions based on a subjective evaluation of the likelihood of future events. It is necessary to examine, therefore, the degree to which predictive behavior corresponds with the objective probabilities of events, and the degree to which man can learn to use information contained in event sequences.

In the present study, 15 groups of subjects were exposed to different sequentially dependent stimulus sequences in a four-alternative probability learning paradigm. Groups were found to differ in both learning rate and terminal response levels as a joint function of the stochastic structure and the degree of dependence in the stimulus sequence. Events with common structural properties in different sequences were learned in a similar fashion.

INTRODUCTION

The classic probability learning paradigm in which subjects predict successive events may, in principle, be used with any finite number of stimuli and any stimulus generator. Most research, however, has been restricted to binary situations in which events are sequentially independent (ref. 1). The few studies that have been reported with Markovian binary structures (refs. 2-5) have found that humans learn first-order event dependencies by responding to short stimulus-response sequences that precede their predictions, although all such sequences are not learned equally well. The manner in which subjects learn to predict four-alternative event sequences generated by first-order Markov processes (ref. 6) is reported here.

In one of the few studies using more than two dependent events, Bennett, Fitts, and Noble (ref. 7) found that a five-alternative Markov process that generated diagrams "concordant" with previously determined response preferences was learned more quickly than a similar stimulus process that generated "discordant" diagrams. These generators were of a special type in which only two nonzero (and in this case equal) transition probabilities were allowed in each row of the Markov matrix. The nonzero entries defined two "appropriate" choices following each stimulus. Although the authors concluded that response preferences are important in probability learning, their use of concordant and discordant diagrams was confounded with structural differences between the generators (ref. 8). Thus, the fact that the concordant generator created different kinds of event sub-sequences than the discordant generator (e.g., repetition and alternation runs) could have accounted for their results. It is difficult to evaluate this hypothesis since the response vector for the two appropriate alternatives was not reported; however, the discordant generator did not suppress the proportion of appropriate choices below chance level during the early stages of learning. As the authors themselves note, such suppression would be expected if response preference were the controlling variable. Their suggestion that between-group differences resulted in the discordant sequences being effectively "neutral" is clearly speculative.

DESIGN AND PROCEDURE

The quarternary Markov matrices used in the present study were restricted to those that are doubly stochastic and contain one high transition probability $\alpha(>0.25)$ in each row, the remainder of the probabilities being equal to $\beta\{=(1-\alpha)/3\}$. In such matrices, each event has a high probability (α) of being followed by some particular event (perhaps itself) and, provided that $\alpha < 1$, each event will asymptotically appear with equal relative frequency. Predicting the event specified by the α probability will be called a *profitable*[1] prediction.

What will be referred to as the *structure* of the process is the general arrangement of the α probabilities within the 4×4 matrix. Since response preferences for individual stimulus diagrams were not considered, exactly five structures were of interest. These correspond with the unordered partitions of four indistinguishable events:

$$S_1 = (E_i)(E_j)(E_k)(E_\ell)$$

$$S_2 = (E_i)(E_j)(E_k, E_\ell)$$

$$S_3 = (E_i)(E_j, E_k, E_\ell)$$

$$S_4 = (E_i, E_j)(E_k, E_\ell)$$

$$S_5 = (E_i, E_j, E_k, E_\ell)$$

where $i \neq j \neq k \neq \ell$, and the subscripted S's and E's refer to the structures and events, respectively. For each structure, each subset of events tends to generate a most probable sub-sequence in which the elements cycle among themselves. For example

$$S_3 = \begin{bmatrix} \alpha & \beta & \beta & \beta \\ \beta & \beta & \alpha & \beta \\ \beta & \beta & \beta & \alpha \\ \beta & \alpha & \beta & \beta \end{bmatrix}$$

[1]Brunswik's terminology (ref. 9) is expanded here.

corresponds with *any* matrix in which one of the events, E_i, has a high probability of being followed by itself, and each of the remaining three events, E_j, E_k, and E_ℓ, has a high probability of being followed by a different event within the subset.

Since all the events are members of some structural subset, each event may be identified with a class C_n ($n = 1, 2, 3, 4$), where n is the number of elements in the subset and, hence, the number of elements in a most probable sub-sequence cycle. Structure S_3 may alternatively be regarded, then, as containing two classes of events, C_1 and C_3, which tend to create the most probable sub-sequences ... $E_iE_iE_i$ ... and ... $E_jE_kE_\ell E_jE_k$ ..., respectively. Inspection will reveal that each S_1, S_2, and S_3 contains C_1 events, and each S_2 and S_4 contains C_2 events. Structures S_3 and S_5 alone contain events in classes C_3 and C_4, respectively.

The advantage of using these particular restrictions (ref. 10) is that for fixed α, the redundancy (ref. 11) of all five structures is equal. Further, independent of α, the marginal (zero-order) uncertainty of all structures is maximal (i.e., two bits). Hence, not only must the subject learn first-order dependencies for performance to improve, but differences in behavior as a function of the structure variable are isolated from the degree of redundancy in the stimulus sequence.

Three α values, 0.4, 0.55, and 0.7, were used in the present study because they are sufficiently different from 0.25 and 1.0 to allow departures from either change or maximizing[2] behavior to be assessed; they also allow a reasonable range of stimulus dependency to be examined. At each level of α a single sequence of 500 events was generated according to the rules of each stochastic structure. Sequences were required to have marginal event frequencies and transition frequencies governed by the α probabilities, within 4 and 5 percent of their expected values, respectively. Transition frequencies governed by the β probabilities were allowed to vary from 10 to 20 percent as an inverse function of α. In all cases, the empirical matrices were judged to be highly representative of the corresponding theoretical process.[3]

Each of the 15 event sequences (five structures for each of the three α values) was recorded on paper tape and used to control the temporal sequence of four visual stimuli (+, -, o, x) to a different group of subjects. Ten college students between the ages of 17 and 28 were randomly assigned to each group and were paid at a fixed hourly rate. In order to distribute the effects of possible symbol diagram preferences, the four events recorded on the control tape were randomly identified with a different symbol ordering for each subject.

[2]Such behavior would be obtained if subjects optimized the number of profitable predictions by invariably choosing alternatives specified by the α probabilities.

[3]As β diminishes the expected marginal frequency for each unlikely alternative approaches zero. The proportional change due to even a single frequency inversion, therefore, becomes quite severe and requires special treatment. The Anderson-Goodman contingency test (ref. 12), nevertheless, did not reveal significant departures from the theoretical matrices at the 0.01 level of confidence.

Subjects were tested two at a time in separate booths and instructed to predict which symbol would appear next on a small rear projection console screen. They made their predictions by pressing one of four buttons, each identified with one of the symbols, during a 1.5 second interstimulus interval identified by the onset of a yellow panel light. Each symbol appeared for 4 seconds. The selected response button lit up following a prediction and remained on until the next symbol appeared.[4]

RESULTS

A separate learning curve was first computed for each event, E, in each structure by determining the average percentage of profitable predictions made by the 10 subjects during each block of 12 successive E occurrences.[5] Since all events occurred at least 120 times in a given sequence, this procedure guaranteed a minimum of 10 blocks of equal statistical precision for each curve.

Figure 1 presents learning curves for stimulus classes C_n (n = 1, 2, 3, 4) averaged by block over structures and subjects. Figure 2 separates the learning curves by structure for each of the two event classes which are combined over structures in figure 1, that is, C_1 and C_2. The statistical precision characteristically varies from curve to curve as a result of the unequal class membership distribution between the structures. In all cases where two or more events within a structure belonged to the same class, however, no meaningful differences between the individual learning curves were noted.

The present data support the hypothesis presented earlier concerning the importance of the structural variable in probability learning. They show that quarternary structures, and in all likelihood *n*-ary structures in general, involve multiple learning rates and asymptotic levels (see fig. 1) which are a joint function of the structural properties and the redundancy of the controlling stimulus generator. Furthermore, although some degree of within structure interaction is reflected in the learning rates (see fig. 2), for a given α the terminal response level for an event is mainly determined by the sub-sequence class of which it is a member, rather than the particular stimulus structure in which it is imbedded. Indeed, these asymptotic levels are,

---

[4]The presentation interval was made quite long relative to the response interval in order to induce the subject to prepare his response while examining the prior stimulus. The subjects, therefore, may have found it easier than subjects in previous studies (ref. 13) to learn first-order associations.

[5]The common analytic technique with binary sequences is to examine blocks of arbitrary event occurrences. This procedure is most convenient when the diagram frequencies have also been forced into theoretical agreement for each block, and it is identical with the present procedure when marginal frequencies are equal. Here (in order to avoid unnecessary higher-order dependencies), diagram frequencies were not restricted within blocks.

as a rule, different from α and support Hake and Hyman's contention (ref. 2, pp. 72-73) that probability matching, when observed, is the fortuitous result of a complex learning process. No tendency for matching was found in the present data. There is some indication, however, that with sufficient stimulus dependency. The terminal proportions of profitable predictions would converge to a common level for all event classes (see fig. 1(c)).

DISCUSSION

Certain interrelated factors bear on the present findings. These are mentioned briefly in order to point up similarities with other classical learning situations. First, as the number of stimuli in an event class diminishes, the occurrences of each component stimulus become clustered within the sequence. The results, therefore, may reflect merely a relative advantage of massed versus distributed practice in a complex task (ref. 14). Second, as the number of stimuli involved in a structurally preferred sub-sequence is reduced, the potential of each component stimulus to develop conflicting remote associations is diminished, provided that associative strength is inversely related to distance within the sequence (ref. 15). Both factors are complicated, however, by the fact that most probable sub-sequences involving the greatest number of events also have the greatest number of starting (entry) points, as well as the greatest potential for atypical reentry into their own event set. Hence, the present results would be expected if one posits, as Hake and Hyman (ref. 2) suggested, that specific stimulus n-tuples, rather than the probabilistic structure of the stimulus generator, are learned. In this case a different n-tuple would be associated with each starting point in the sub-sequence. It is also true, however, that behavioral adjustment to unexpected contingencies is often difficult because of mental or motor *set* (refs. 16, 17, 18). Even if subjects did perceive the probabilistic properties of the generator, then, one would expect a differential performance decrement between event classes as a function of the number of atypical sub-sequence reentries.

These arguments may be somewhat difficult to untangle, but it is clear that many of the same variables are operating here as in the more deterministic serial and paired-associates learning paradigms. It may be noted, moreover, that in the extreme ($\alpha = 1.0$), each structural subset represents a serial learning task. The present findings would appear to be, therefore, the probabilistic analog of the well documented inverse relationship between serial list length and learning speed (ref. 14).

Ames Research Center
National Aeronautics and Space Administration
Moffett Field, Calif., 94035, April 8, 1969
127-51-09-01-00-21

REFERENCES

1. Estes, W. K.: Probability Learning. Categories of Human Learning, A. W. Melton, ed., Academic Press (New York), 1964, p. 89.

2. Hake, H. W.; and Hyman, R.: Perception of the Statistical Structure of a Random Series of Binary Symbols. J. Exp. Psychol., vol. 45, 1953, pp. 64-74.

3. Anderson, N. H.: Effect of First-Order Conditional Probability in a Two-Choice Learning Situation. J. Exp. Psychol., vol. 59, 1960, pp. 73-93.

4. Chapman, Jean P.: The Spacing of Sequentially Dependent Trials in Probability Learning. J. Exp. Psychol., vol. 62, 1961, pp. 545-551.

5. Anderson, N. H.: An Evaluation of Stimulus Sampling Theory: Comments on Professor Estes' Paper: Categories of Human Learning, A. W. Melton, ed., Academic Press (New York), 1964, p. 129.

6. Feller, W.: An Introduction to Probability Theory and Its Applications. John Wiley & Sons, Inc., 1957, p. 338.

7. Bennett, W. F.; Fitts, P. M.; and Noble, M.: The Learning of Sequential Dependencies. J. Exp. Psychol., vol. 48, 1954, pp. 303-312.

8. Garner, W. R.: Uncertainty and Structure as Psychological Concepts. John Wiley & Sons, Inc., 1962, p. 287.

9. Brunswik, E.: Probability as a Determiner of Rat Behavior. J. Exp. Psychol., vol. 25, 1939, pp. 175-197.

10. Evans, S. H.: VARGUS 7: Computed Patterns From Markov Processes. Behav. Sci., vol. 12, 1967, pp. 323-328.

11. Shannon, C. E.; and Weaver, W.: The Mathematical Theory of Communication. Univ. of Illinois Press, Urbana, 1949.

12. Anderson, R. W.; and Goodwin, L. A.: Statistical Inference About Markov Chains. Ann. Math. Stat., vol. 28, 1957, pp. 89-110.

13. Tune, G. S.: Response Preferences: A Review of Some Relevant Literature. Psychol. Bull., vol. 61, 1964, pp. 286-302.

14. Hovland, C. I.: Human Learning and Retention. Handbook of Experimental Psychology, S. S. Stevens, ed., John Wiley & Sons, 1951, Ch. 17.

15. Woodworth, R. S.; and Schlosberg, H.: Experimental Psychology. Holt, Rinehart, and Winston, N. Y., 1954, p. 695.

16. Dashiell, J. F.: A Neglected Fourth Dimension in Psychological Research. Psychol. Rev., vol. 47, 1940, pp. 289-305.

17. Gibson, J. J.: A Critical Review of the Concept of Set in Contemporary Environmental Psychology. Psychol. Bull., vol. 38, 1941, pp. 781-817.

18. Bugelski, B. R.; and Huff, E. M.: A Note on Increasing the Efficiency of Luchins' Mental Sets. Am. J. Psychol., vol. 75, 1962, pp. 665-667.

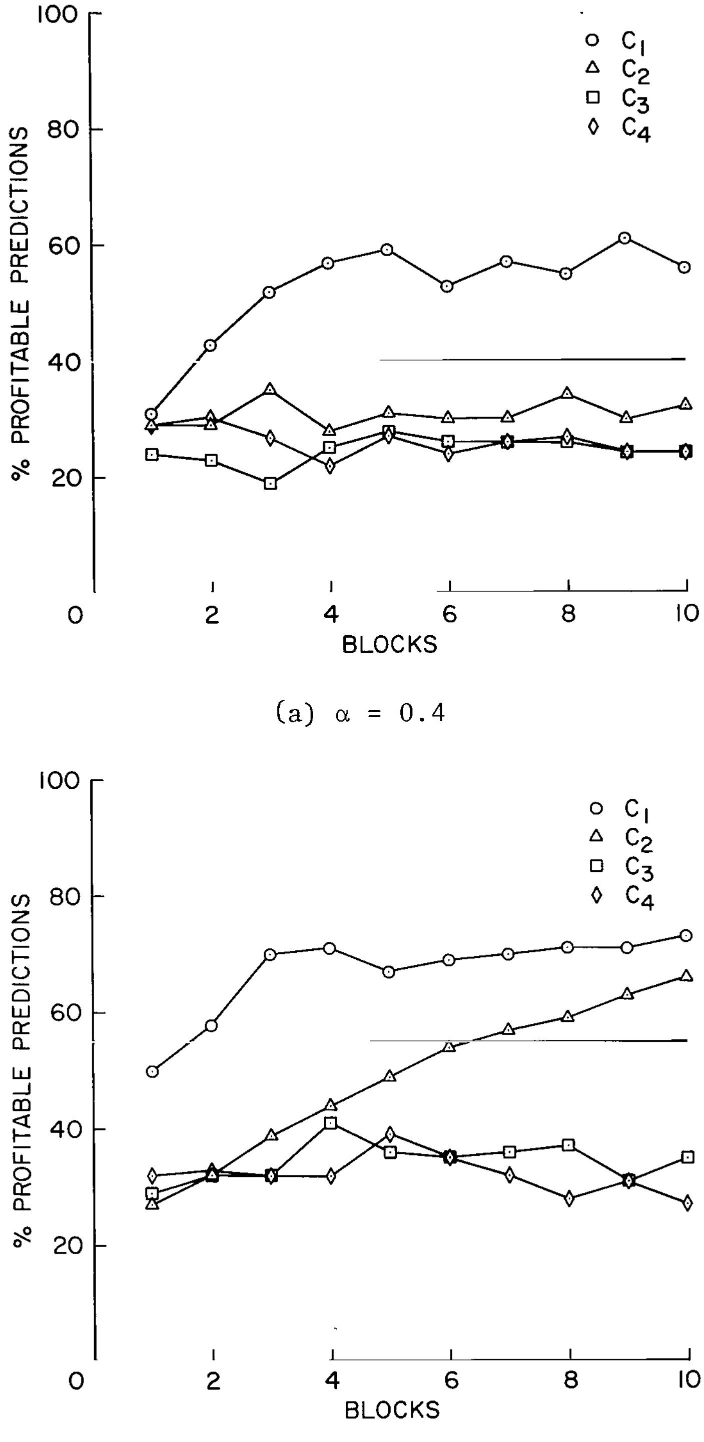

(a) α = 0.4

(b) α = 0.55

Figure 1.- Learning curves for event classes C_n (n = 1, 2, 3, 4) averaged over structures and subjects for each of three α levels.

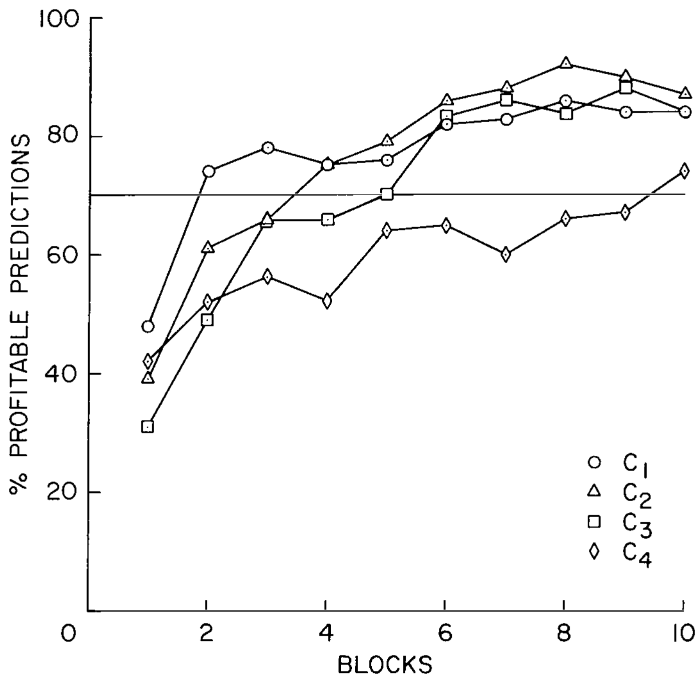

(c) α = 0.7

Figure 1.- Concluded.

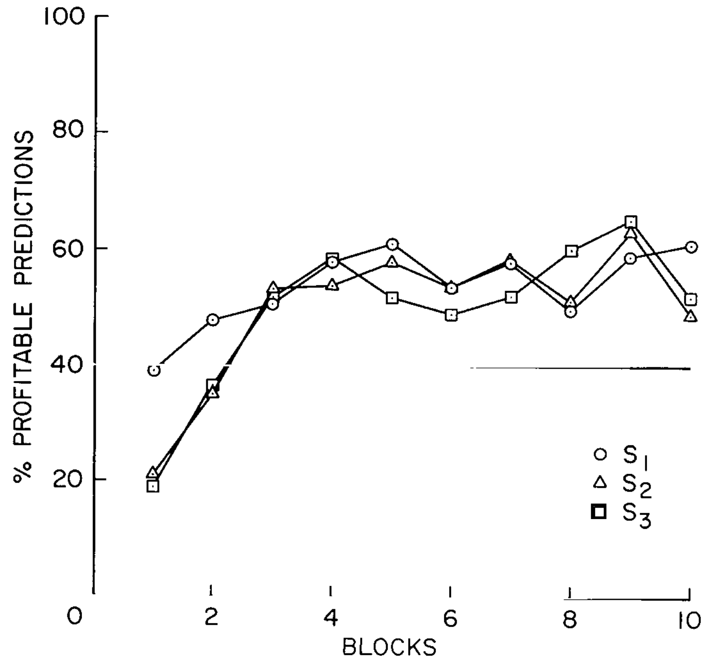

(a) C_1 events, α = 0.4

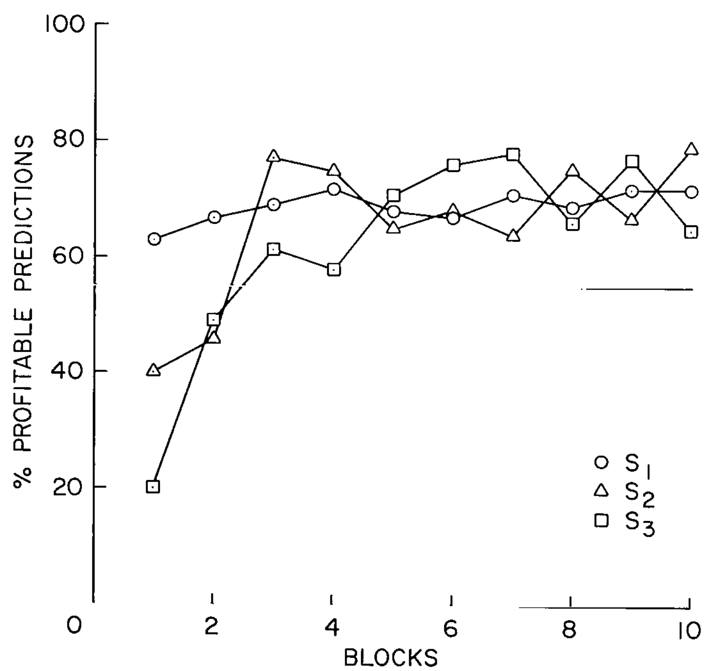

(b) C_1 events, α = 0.55

Figure 2.- Learning curves for event classes C_1 and C_2 averaged separately by structures for each α level.

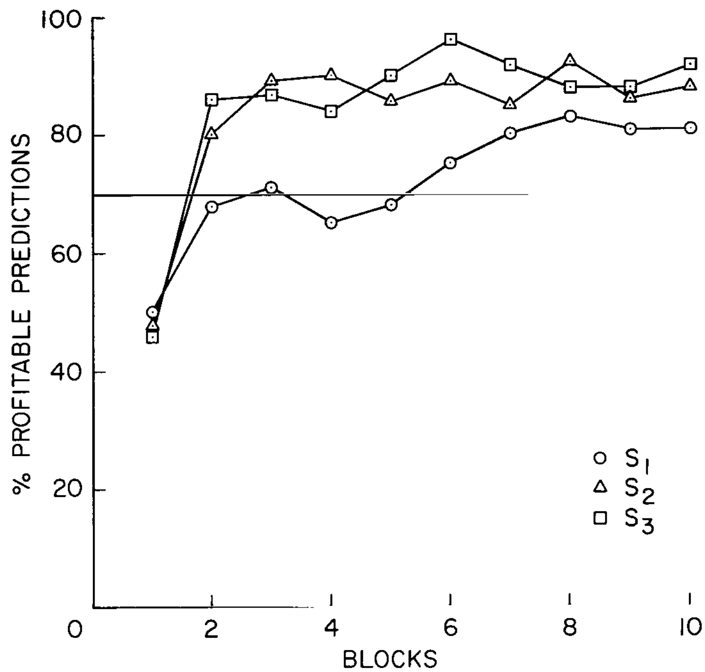

(c) C_1 events, $\alpha = 0.7$

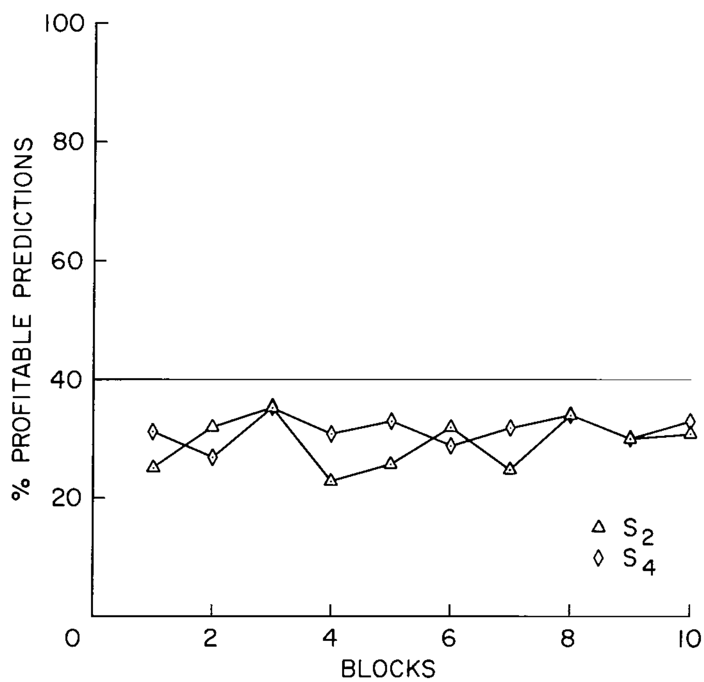

(d) C_2 events, $\alpha = 0.4$

Figure 2.- Continued.

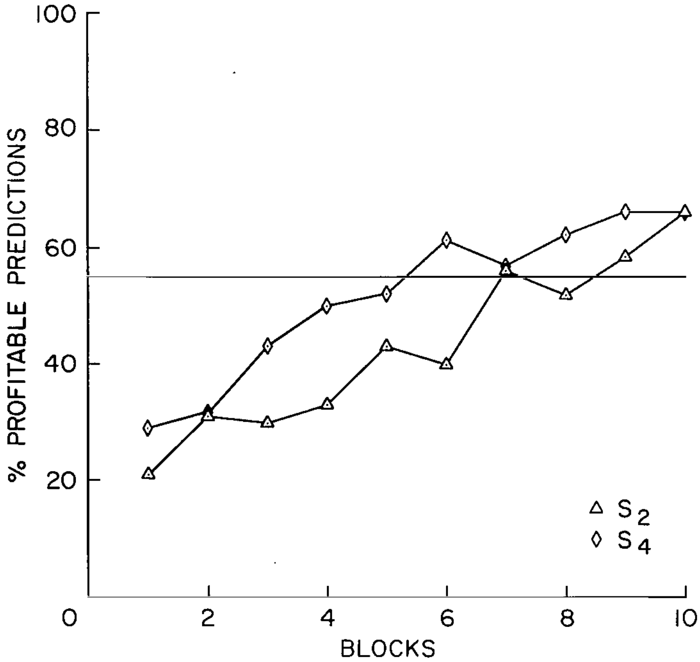

(e) C_2 events, α = 0.55

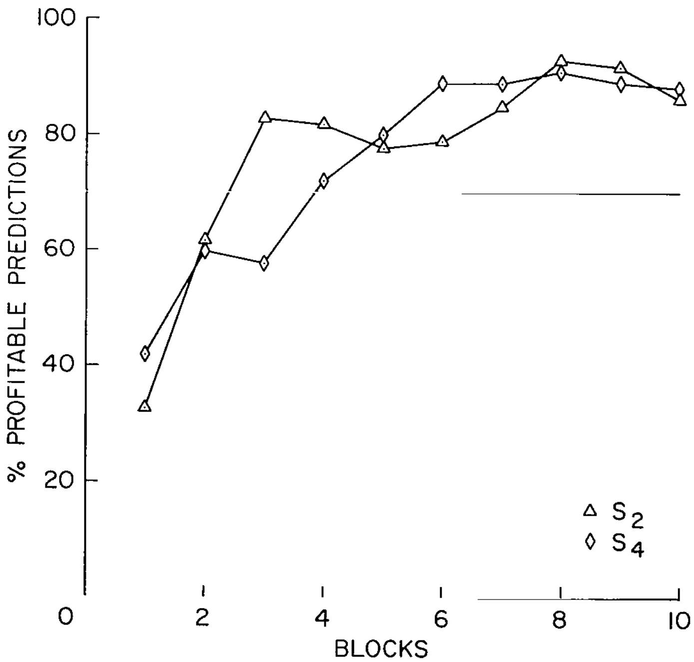

(f) C_2 events, α = 0.7

Figure 2.- Concluded.

"The aeronautical and space activities of the United States shall be conducted so as to contribute . . . to the expansion of human knowledge of phenomena in the atmosphere and space. The Administration shall provide for the widest practicable and appropriate dissemination of information concerning its activities and the results thereof."

—NATIONAL AERONAUTICS AND SPACE ACT OF 1958

NASA SCIENTIFIC AND TECHNICAL PUBLICATIONS

TECHNICAL REPORTS: Scientific and technical information considered important, complete, and a lasting contribution to existing knowledge.

TECHNICAL NOTES: Information less broad in scope but nevertheless of importance as a contribution to existing knowledge.

TECHNICAL MEMORANDUMS: Information receiving limited distribution because of preliminary data, security classification, or other reasons.

CONTRACTOR REPORTS: Scientific and technical information generated under a NASA contract or grant and considered an important contribution to existing knowledge.

TECHNICAL TRANSLATIONS: Information published in a foreign language considered to merit NASA distribution in English.

SPECIAL PUBLICATIONS: Information derived from or of value to NASA activities. Publications include conference proceedings, monographs, data compilations, handbooks, sourcebooks, and special bibliographies.

TECHNOLOGY UTILIZATION PUBLICATIONS: Information on technology used by NASA that may be of particular interest in commercial and other non-aerospace applications. Publications include Tech Briefs, Technology Utilization Reports and Notes, and Technology Surveys.

Details on the availability of these publications may be obtained from:

SCIENTIFIC AND TECHNICAL INFORMATION DIVISION

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

Washington, D.C. 20546