Reprinted from Science, April 30, 1965, Vol. 148, No. 3670, pages 661-662

Observable Changes of Hypotheses under Positive Reinforcement

Abstract. *In mathematical models of concept learning it has consistently been assumed that positive reinforcement cannot lead to a change of the hypothesis determining the overt response. When hypotheses are experimentally identified and recorded along with positive and negative reinforcements of stimulus-response pairs, it can be shown that hypotheses may change after a positive reinforcement. Positive reinforcement has an information content for subjects that has not yet been adequately recognized in concept formation studies.*

Various stochastic theories of learning account for the classification or conceptualization of behavior in terms of hypothesis testing. The learning process is described as follows: the subject randomly samples from a population of hypotheses, and when a stimulus is presented he makes the response determined by the hypothesis sampled. In all hypothesis models proposed so far, it has been assumed that the same hypothesis is kept until information is received which infirms it (*1*). As a consequence, the probability is zero that a subject will shift from an incorrect hypothesis to the correct hypothesis on correct trials. The critical assumption of "no-change-if-no-error" has not been checked directly. Indeed, in most of the experiments designed to test these models of learning, it has not been possible to identify the hypotheses used on successive trials. Levine (*2*) attempted such identification by inferring hypotheses from sequences of responses given on blank trials (that is, without outcomes). However, since his procedure did not rule out the possibility that the hypotheses could change in the course of the test responses themselves, identification of the hypotheses could be impaired. This objection may be important, as suggested by the following results.

Three experiments were conducted with three different groups of college students. In experiment I, each subject was given a set of eight white cards; on each card was printed a string of three letters, each letter being either D or K, with all possible combinations represented once in the group of cards. The subject was requested to classify the cards correctly by placing them in two columns, headed by a pink and a blue label, respectively. After the subject had classified the cards once without any clue concerning the "correct" criterion, the experimenter replaced one of the white cards by a colored card, either pink or blue, having exactly the same pattern of letters as that on the white card removed. The subject was instructed to place this card in the group to which it belonged (identified by the color) and, furthermore, to make as many changes in the placement of the other cards, or none, as appeared necessary in order to achieve the correct classification of all cards. When the subject indicated that he was satisfied with his new classification, which might be the same as the previous one, he was given a second colored card. The presentation of each additional colored card started a new trial. The procedure was repeated until all white cards were replaced. As the subject kept all cards before him, the "correct" classification, defining the conceptual problem to be solved, was completely shown after eight trials. In this experiment the problem was the same for all subjects. It is described by the information cards which were given in this fixed order: DKD(blue), KKK(pink), DDD(pink), KDK(blue), DKK(pink), KDD(pink), DDK(blue), KKD(blue).

In experiment II, exactly the same cards and essentially the same procedure were used as in experiment I. But the subjects were given successively six problems to solve, with the same cards, ranging from the simplest (one-dimensional) problem to complex ones (disjunctive three-dimensional). The order of the problems and the sequence of information were constant.

In experiment III, the patterns to be classified included strings of 1, 2, 3, 4, and 5 letters, all 62 possible combinations of D's and K's being represented once. The strings were typed together on single sheets of paper, in fixed order for the individual, and in randomized order for the group of subjects, one sheet being used per trial. The patterns were to be classified by either circling or crossing the strings of letters. After each classification the subject selected one string of letters for which the experimenter indicated the "correct" response (circling or crossing). Each new sheet contained the past as well as the new information given by the experimenter. The single problem, given to all subjects, resembled that of experiment I: circling was the "correct" response requested for all strings of letters ending in DD or KK. Except for the differences in the experimental procedure, the instructions were similar to those given in experiments I and II.

Thus, the three experiments had the following common features: (i) on each trial the subject made a binary classification of the entire set of stimulus patterns (strings of letters); (ii) on each trial the classification of a single stimulus card was reinforced; (iii) the subject had before him a complete record of past information and was prevented from making any classification response inconsistent with this information.

Figure 1*A* shows, for successive trials of the three experiments, the proportions of subjects who did change their classifications although not forced to do so by the information received—that is, the information card given them agreed with their own classification. The values plotted for experiment II are averages computed over the six problems.

With exception of one peak value (trial 4, experiment I), the proportions of changes on a correct trial fluctuated around 0.20 for experiments I and II, in which the same stimuli and the same procedure were used. The proportions were much higher, ranging between 0.67 and 0.48, in experiment III, which differed from the other two mainly in the complexity of the stimuli and the subject's freedom to select the sequence of information. Inspection of individual data in experiments I and II reveals that changes of hypothesis on a correct trial is not a characteristic either of a few subjects only, or of a particular trial or problem. Actual transitions from specific hypotheses to other hypotheses were analyzed in detail for experiment II. They were quite similar on correct and on incorrect trials. In particular, the proportion of transitions from simple hypotheses to classifications corresponding to no obvious hypotheses remained equally low in both cases. In experiment III the complexity of the task introduced the possibility of clerical errors or inconsistencies that could artificially inflate the proportions of change. But, in fact, changes involving 10 or more individual strings of letters represent more than 90 percent of the changes on correct trials, up to trial 5.

As suggested by the curve obtained for experiment III, the probability of change on a correct trial gradually decreased over trials. If this decline were a simple function of the accrued information or of the shrinkage of the set of hypotheses theoretically available, it should have been more pronounced in experiments I and II which, unlike experiment III, were continued until all information was given. On the contrary, if trials 2 and 3 are compared with trials 6 and 7 the decline, if any, is hardly noticeable in experiments I and II.

There were two peculiarities in trial 4 of experiment I, for which the proportion of changes appeared surprisingly high: it was the first trial on which the preceding information eliminated all simple (one-dimensional) types of classification; it was also the trial on

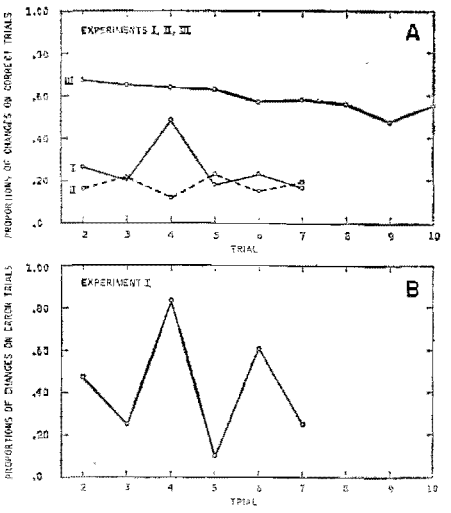

Fig. 1. (*A*) Plot of proportions of changes of hypotheses on correct trials, for experiments I (157 subjects), II (30 subjects), and III (155 subjects). (*B*) Plot of proportions of subjects forced to change their hypotheses by the information received, for experiment I.

which the proportion of subjects forced by the information to change their classifications was the highest. In fact, the curves of Fig. 1*B* and Fig. 1*A* for experiment I suggest that a given piece of information had a particular effect whether it agreed or disagreed with the subject's hypothesis. The peaks and valleys match exactly for the two curves of experiment I.

A tentative interpretation is that, when the set of hypotheses is large (*3*), the subject "samples" or attends to several hypotheses simultaneously, only one of which will be overt in terms of classification responses. (i) The information received may agree with the overt hypothesis and disagree with the covert hypotheses. In this case the subject may be led to alter his subset of hypotheses; this may, in turn, disturb the dominant status of the overt hypothesis. To account for a relation between the probability of change on a correct trial and the overall probability of error on that trial, it must be further assumed that similar subsets of hypotheses were sampled by subjects receiving the same information. This assumption seems supported by the similarity of the transitions between hypotheses, observed on correct and on error trials separately. (ii) If the information received on a given trial agrees with both an overt and a covert hypothesis, it may reinforce either one with a different probability, the reinforcing effect of specific information depending on the hypotheses sampled. This is, indeed, suggested by our data. For example, in experiment I, after the information card DKD was given (*4*), the conditional probability of a classification based on the middle letter of the pattern was .15, given a previous classification based on the first letter, but it was .00, given a previous classification based on the number of D's (or K's) in the pattern. (iii) It is also conceivable that a subject might sample spontaneously, at any time, or under stimulations other than those planned by the experimenter. A more detailed exploration of these ideas, including a test of Bayesian approaches to information processing, is now being made.

The divergence of our results, for which the assumption "no-change-if-no-error" is clearly wrong, from other studies in which it seems to work, even if not directly checked, may have been enhanced by our procedure of keeping the subject constantly informed of all past reinforcements (*5*). Our use of a larger universe of hypotheses and a more complex situation yields another source of divergence. At any rate, the present results show that this assumption is unsound for a significant class of concept learning experiments. It would seem that the assumption that the probability of change on a correct trial is zero should be made a special case rather than a postulate in a comprehensive theory of concept learning.

PATRICK SUPPES
MADELEINE SCHLAG-REY
Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford, California

References and Notes

1. G. H. Bower and T. R. Trabasso, in *Studies in Mathematical Psychology*, R. C. Atkinson, Ed. (Stanford Univ. Press, Stanford, 1964), pp. 32–94; M. Levine, H. Leitenberg, M. Richter, *Psychol. Rev.* **71**, 94 (1964); F. Restle, *ibid.* **69**, 329 (1962); ———, in *Studies in Mathematical Psychology*, R. C. Atkinson, Ed. (Stanford Univ. Press, Stanford, 1964), pp. 116–172.
2. M. Levine, "Hypothesis Behavior by Humans during Discrimination Learning," report presented at the Psychonomic Society Meeting (Oct. 1964).
3. The size of the intial set was $2^8 - 2$ in experiments I and II, and $2^{32} - 2$ in experiment III.
4. Note that this information itself did not infirm any criterion of classification because the subject could merely reverse his response assignment if necessary.
5. See J. S. Bruner, J. J. Goodnow, G. A. Austin, *A Study of Thinking* (Wiley, New York, 1956), for a discussion of the load put on memory, respectively, by simultaneous scanning and successive scanning of hypotheses.
6. Supported by the Carnegie Corporation of New York and U.S. Office of Education contract OE 3-10-009. These experiments were performed in the Human Communication Laboratory, Department of Psychology, University of California, Los Angeles. We are grateful to Dr. E. C. Carterette for making the facilities of the Laboratory available to us.

12 February 1965