F. 112

8.

UMEÅ PSYCHOLOGICAL REPORTS

No. 122 1977

Department of Psychology
University of Umeå

SUBJECTS' SELECTION OF FEEDBACK INFORMATION IN AN
INDUCTIVE INFERENCE TASK

Lars-Åke Lindberg

Berndt Brehmer

SUBJECTS' SELECTION OF FEEDBACK INFORMATION IN AN INDUCTIVE INFERENCE TASK

Lindberg, L-Å., and Brehmer, B. Subjects' selection of feedback information in an inductive inference task. Umeå Psychological Reports No. 122, 1977. - Subjects' selection of outcome feedback instances in a probabilistic inference task with one linear and one nonlinear cue was investigated in a 2 (levels of task information: maximum vs. minimum) x 4 (blocks of trials) factorial experiment. Maximal information included the criterion distribution for each cue combination as well as aids for recording feedback instances, while minimal information consisted of standard instructions and no memory aids. Under both information conditions, subjects (1) examined one cue at a time, (2) gave priority to the more difficult quadratic cue, and (3) systematically changed their distribution and order of selections as a function of blocks. That is, subjects used the feedback information in a way that is not possible in the standard experimenter controlled reception paradigm. Furthermore, subjects made consecutive selections of the same cue combination more frequently in the maximum information condition than in the minimum information condition. These results are consistent with the assumptions that subjects under both conditions try to find the functional parameters of the task by means of a hypothesis testing activity, and that they learn to eliminate incorrect hypotheses with practice.

This is the second report on an experiment designed to investigate the reasons why subjects do not reach the optimal level of performance in multiple cue probability learning. Subjects performed the task in a series of alternating learning blocks with outcome feedback and test blocks without feedback. In the first report (Lindberg & Brehmer, 1977) results from the test stage are presented. In the present report, subjects' selection of feedback information in the learning stage is examined.

In the test stage, the effects of extensive task information and active feedback control were investigated in a 2 (maximum vs. minimum task information) by 2 (active vs. passive feedback control) by 4 (blocks of 25 learning trials). The learning task had one linear cue and one quadratic cue. Under the maximum information condition, subjects were informed about the distribution of criterion values, including the mean, for each cue value as well as for each combination of cue values. They were also provided with aids for recording feedback instances. Under the minimum information condition, subjects were given standard instructions. Under the active feedback condition, subjects selected cue combinations and feedback instances in whatever order they wanted, while under the passive feedback condition, this information was provided by the experimenter in a preestablished order.

Results from the test stage indicate that subjects were able to make use of the additional task information as well as of the possibility for active feedback control to improve their prediction strategies with respect to selection of inference rules and level of task control. In particular, subjects under the maximum information condition learned the optimal rules for both cues, while subjects under the minimum information condition learned to control the ~~linear~~ cue but not the quadratic cue. These results are consistent with Brehmer's (1974) hypothesis testing model.

In the learning stage of the active feedback condition, subjects' selections of cue combinations and feedback instances were recorded. In the present report, results from the analyses of subjects' selection strategies will be presented.

Problems

The first question to be investigated is whether the positive effects of active feedback control can be explained in terms of possible differences in how the subjects structure the feedback information in this condition compared to the condition where they receive feedback in an experimenter controlled sequence. Therefore, we will compare the subjects' selection strategies to the sequence of feedback information under the passive feedback condition.

According to the hypothesis testing model (Brehmer, 1974) subjects learn an inference task in two stages. In the first stage they generate and test hypotheses about the form of the cue-criterion function, and in the second, they learn to apply the inference rule consistently. According to this model, then, the subjects' performance in the test stage is a joint function of hypothesis selection and rule application. In the learning stage. In the present experiment, however, the subjects made no predictions. They could therefore focus on the testing of hypothesis without a concomittant requirement of task control. Therefore, the analyses of selection strategies should make it possible to distinguish subjects' hypotheses from their application of rules, and thus, give enable us to test the validity of the hypothesis testing model in a new way.

The hypothesis testing model does, however, not make any predictions about how subjects learn other aspects of the task than the cue criterion functions. Consequently, only part of the subjects selection strategies can be explained by this model. Therefore, subjects' selection strategies were also analyzed with respect to assumptions made by the normative model implicit in Social Judgment Theory (Hammond, Stewart, Brehmer, & Steinman, 1975). Although the normative model does not give unambiguous criteria for an optimal selection strategy for the present tasks, it nevertheless shows what kind of hypotheses the subjects might be expected to generate and test by means of their selection strategies. According to this model, optimal performance is achieved when the subject find and make use of (1) the cue-criterion functions, (2) the cue weights, (3) the cue

integration rule, and if he (4) minimizes $(R-E)^2$, where R is the prediction made by the subject and E the criterion value. Under the maximum information condition, subjects have the alternative possibility to simply use the given criterion means as their predictions, without considering the structural aspects of the task.

In summary, subjects' selection strategies were analyzed to answer the following questions: (1) In what respects, if any, do subjects' selection strategies differ systematically from the structure of the feedback information given under the passive feedback condition? (2) Are these differences consistent with predictions from the hypothesis testing model? (3) Are these differences consistent with the assumptions made by the normative model? (4) Are subjects' selection strategies consistent with their performance in the test stage?

Analysis of selection strategies

Basic assumptions

To make it possible to draw any conclusions about the subjects' hypotheses from their selection strategies, the following basic assumptions are needed: (1) The subject's selection strategy maximizes the likelihood of relevant information with respect to the hypothesis under consideration, (2) the strategy involves direct tests of the hypothesis, and (3) it minimizes the cognitive strain on part of the subject. These assumptions are consistent with the results from studies on conceptformation (Bruner, Goodnow, & Austin, 1956).

Active vs. passive control

Under the passive feedback condition, every cue combination was given once in each learning block, and the order of presentation was random. Thus, all the 25 cue combinations were equally likely to occur, as were each row, column and principal diagonal of the cue matrix (see Figure 1).

		Cue 2				
		1	2	3	4	5
Cue 1	1	23 (1:1)	14	11	14	23 (1:5)
	2	26	17	14	17	26
	3	30	21	18 (3:3)	21	30
	4	34	25	22	25	34
	5	37 (5:1)	28	25	28	37 (5:5)

Figure 1. The cue-criterion matrix with criterion means.

Furthermore, the probability that two consecutive selections are drawn from the same row, or from the same column, is considerably less than the probability that they are drawn from different rows and columns. The probability of two selections from the same cell of the matrix is, of course, zero.

Under the active feedback condition, subjects may develop selection strategies which deviate from the above rules with respect to (1) order of selections within each block, (2) distribution of selections over the cue matrix in each block, and (3) changes over blocks.

Selection sequences. Two consecutive selections are classified into one of the following mutually exclusive and exhaustive categories:

Category 1: Two consecutive selections are made from the same column. Category 1 selections are defined to indicate that the subject is concerned with the linear cue 1, i.e., that he tests cue 1-hypotheses.

Category 2: Two consecutive selections are made from the same row. Category 2 selections are defined to indicate that the subject tests hypotheses about the quadratic cue 2.

Category D (diagonal): Two consecutive selections are made from different columns and different rows. According to the basic assumption that subjects will minimize the cognitive strain, it follows that they will not test several hypotheses simultaneously. Thus, category D selections are defined to indicate that subjects test hypotheses which do not involve separate cues.

Category C (cell): Two consecutive selections are made from the same cell indicating that the subject either examines the nature of the within cells variance, and/or that he tries to find the criterion mean for this cue combination.

For each subject and block in the learning stage, the 24 consecutive selections were classified into these four sequence categories. Each category was analyzed by means of a 2 (experimental condition: maximum vs. minimum information) x 4 (blocks of trials) analysis of variance (ANOVA) with repeated measures on the second factor.

Distribution of selections. A priori, the cue matrix could be categorized into subparts in a great number of ways. For the present purpose, two categorizations were made: (1) Subjects distribution of selections over the rows, columns and principal diagonals of the cue matrix, and (2) subjects distribution of selections over the cells of the cue matrix.

For each subject and block in the learning stage, the 25 selections were analyzed with respect to number of selections for each row, column, and principal diagonal. Data were analyzed by means of three ANOVAs. All ANOVAs had Experimental conditions (with 2 levels: maximum vs. minimum information) as the first factor, and Blocks (with 4 levels) as the second. The third factor was Rows (with 5 levels) in the first analysis, Columns (also with 5 levels) in the second, and Diagonals (with 2 levels) in the third analysis.

For each subject and block in the learning stage, the distribution of selections over the 25 cells of the cue matrix was also recorded. Specifically, the most frequently selected cells were registered as well as the number of cells from which no selection was made.

Systematic changes over blocks. In addition to the above analyses, the number of selections in the most frequently used category (as defined above) was recorded for each subject and block in the learning stage, and subjected to a 2 (experimental conditions: maximum vs. minimum information) x 4 (blocks of trials) ANOVA with repeated measures on the second factor.

Individual strategies. The above analyses give information about the behavior of the "mean subject". In addition, the number of subjects following each strategy was calculated for each experimental condition and block in the learning stage. The results from these analyses will be reported when they are not consistent with the results obtained for the "mean subject".

Hypotheses

Minimum information condition

1. According to the hypothesis testing model, subjects generate hypotheses about the separate cue-criterion functions. Prediction 1: Sequence categories 1-2 will be used more frequently than category D. Under the passive feedback condition, sequence category D is more probable than categories 1-2.

2. According to the hypothesis testing model non-optimal hypotheses are rejected with practice. Prediction 2: The frequency of category D selections will decrease as a function of blocks. Under the passive feedback condition, the probability of category D selections does not change over blocks.

3. According to the hypothesis testing model, subjects will test the hypothesis that the cue-criterion function is linear before they test the hypothesis that it is quadratic. Furthermore, linear rules are easier to detect than are quadratic rules (Brehmer, Kuylenstierna, & Liljergren, 1974). The first hypothesis cannot be tested directly since subjects may test both kind of hypotheses through continuous scanning of cells within rows or within columns. However, if subjects develop different strategies for the two cues as a function of blocks, the subjects should detect that cue 2 is non-linear and they should therefore start testing non-linear hypotheses with respect to this cue. Since non-linear relations are more difficult to detect, the relative number of tests of hypotheses about cue 2 will increase over blocks. Prediction 3: The relative frequency of category 2 selections will increase as a function of blocks, compared to the frequency of category 1 selections. Under the passive feedback condition, the two categories are equally probable in all blocks. Prediction 4: The distribution of selections over rows will differ from the distribution of selections over columns, when subjects detect that cue 2 is non-linear. These distributions do not differ under the passive feedback condition.

4. According to the hypothesis testing model, subjects first learn the form of the cue-criterion functions. After that they learn the parameters of the functions. Since a lower number of cells is necessary for the latter stage, subjects will restrict their distribution of selections over the cue matrix with practice. Prediction 5: The frequency of cells from which no selections are made will increase as a function of blocks. For the passive feedback condition, this number is zero in all blocks. Alternatively, an increased number of such cells may be due to an increase in the number of cells from which more than one selec-

tion is made. Therefore, the number of cells from which no selections are made should be compared with the number of cells from which more than one selection is made. Finally, an initial search over the entire cue matrix, may also indicate that the subjects try to determine whether the cues are independent, i.e., that they try to investigate the cue integration rule. If so, the number of cells from which no selections are made should increase as a function of blocks, provided that the subjects detect the correct additive combination rule.

5. According to the normative model, subjects also have to learn the cue weights. This could be achieved by means of the same strategy as that for finding the function forms (prediction 1). Furthermore, subjects have to detect and account for the within cells variance. The model does not, however, predict how this could be achieved, and the extent to which the subjects should focus on the within cells variance compared to the other aspects of the task. In accordance with a hypothesis put forward by Brehmer and Kuylenstierna (1974), we assume that subjects consider the task as deterministic. Prediction 6: Subjects do not make category C selections. Under the passive feedback condition, there were no category C sequences.

Maximum information condition

1. Although subjects under the maximum information condition are informed about the $\bar{E}$ - values, they do not know that these values are considered to be the optimal predictions. Therefore, subjects are free to test different hypotheses about possible criteria for optimal performance for the task. This may be one of the reasons why these subjects change their performance in the test blocks, as well as why they should be expected to change their selection strategies with practice.

2. Another reason may be that subjects try to learn the structural aspects of the task, not explicitly given in the task information provided. Subjects' knowledge of the task is "blind" in the specific sense that they cannot solve the prediction problem if the task information is removed. According to this argument, subjects under the maximum information condition may adopt a problem solving approach

to the task in the same way as they do in the minimum information condition. Such a strategy would also be helpful for the subjects with respect to the problem to determine the optimal performance criteria, discussed above.

3. However, in contrast to the minimum information condition, subjects in the maximum information condition do not have to examine the within-cells variance or to detect the $\bar{E}$ values, and they do not have to scan the entire matrix in order to find out whether cues are independent. It is also assumed that the form of the cue-criterion functions, as well as the parameters, are easily detected by means of visual inspection of the cue-criterion matrix. Prediction 7: The proportion of cells, from which no selections are made, is initially greater under the maximum information condition than under the minimum information condition. Prediction 8: In contrast to the minimum information condition, the proportion of cells from which no selections are made will not increase as a function of blocks in the maximum information condition.

4. It is assumed that subjects under the maximum information condition try to integrate separate kinds of information about the task into a selection strategy which enables them to test statistical hypotheses, that is, hypotheses which combine the structural aspects with the probabilistic aspects of the task. If so, subjects in the maximum information condition should distribute their selections over rows and columns in the same way as subjects in the minimum information condition. At the same time however, they should have a greater proportion of category C selections. Prediction 9: With respect to rows and columns of the cue matrix, subjects in the maximum information condition distribute their selections in the same way as subjects in the minimum information condition. Prediction 10: The proportion of category C selections will be greater in the maximum information condition than in the minimum information condition.

Method

Subjects

Thirty-two undergraduates from the University of Umeå participated as paid volunteers. Subjects were randomly assigned to the four experimental groups.

Design

Subjects went through the experiment individually in a series of alternating learning blocks with outcome feedback and test blocks without feedback. The learning stage as well as the test stage consisted of four blocks of trials. Each learning block included twenty-five trials while each test block included fifty trials, comprising two replicates of each one of twenty-five different cue combinations.

In the learning stage, two factors were combined factorially, each factor had two levels. For the first factor, task information, the levels were maximum vs. minimum task information and for the second factor, the feedback control factor, the levels were active vs. passive feedback control. The learning stage, then, follows a 2 (maximum vs. minimum task information) x 2 (active vs. passive feedback control) x 4 (blocks) factorial design.

Learning task

The learning task was a two-cue task with metric cue and criterion variables. Cue values, from 1 to 5, were presented on cards as bars of different height, each card containing one value for each cue. On the back side of each card the criterion value was given as a number which could take on values from 5 to 43. Cue 1 was linearly related to the criterion while cue 2 was quadratically related to the criterion. The criterion variable was an additive function of the two cue variables. The cues were uncorrelated, and the cues had equal validities (weights). The statistical characteristics of the task are given in Figure 1 and Table 1.

Table 1. Statistical characteristics of the task.

Criterion variable (E), integers, 5-43
Linear cue (C_1), bar varying in height 1-5
Quadratic cue (C_2), bar varying in height, 1-5
$E = .60 (C_1 - \bar{C}_1) + .60 (C_2 - \bar{C}_2)^2$
$r_e = .85 \pm .05$, the multiple correlation for the cue-criterion relation
$r_{C_1 C_2} = .00 \pm .05$
$r_{C_1 E} = r_{C_2 E} = .60 \pm .05$

Experimental conditions

Group 1: Minimum task information/active feedback control. Subjects were instructed that their task was to learn to predict numbers on the back side of the cards from the values of the two bars on the front side of the cards. They were informed that there was a systematic relation between each bar and the number, that they could not expect to hit the correct number on every trial but that they could improve by training, and that the experiment consisted of alternating training blocks and test blocks.

Cards used in the learning blocks were placed on a table in front of the subject in a 5 x 5 matrix with eight cards in each cell. The subject was asked to select cards, one at a time, from any cell and in whatever order he wanted, and to check the number on the back side. He was informed that each learning block consisted of 25 selections, after which would come a test block of 50 trials, and that he was free to distribute his selections as he wished over the total of 25 x 4 training trials, but that he would not be given any new cards if one of more cells became empty. For each learning trial the subject (1) selected a card, (2) checked the criterion number, (3) put the card aside (no replacement), and (4) drew a new card, and so on. In the test blocks, the subject was given a deck of cards by the experimenter,

consisting of 50 cards, with each cue combination occurring twice. These cards did not have criterion numbers and were given in a different order for each subject. For each test trial, the subject (1) looked at the cue values, (2) wrote down his prediction of the criterion value on an answer sheet, and (3) looked at the next card in the deck, and so on.

Group 2: Minimum task information/passive feedback control. This condition differed from condition 1 above in that the cards in the training blocks were presented by the experimenter, one at a time, in a predetermined order which was varied among subjects. In each training block, each subject was given one card from each of the 25 cells.

Group 3: Maximum task information/active feedback control. This condition differed from condition 1 in that subjects were provided with the following task information and aids for data recording: (1) Throughout the experiment, subjects could inspect an information sheet which gave the information in Table 2, and where also all criterion values in each cell were given; (2) For the learning blocks, subjects received a recording sheet in the form of the cue-criterion matrix, on which they recorded each feedback (criterion) value.

Group 4: Maximum task information/passive feedback control. This group was treated as group 2 with respect to procedure, and as group 3 with respect to task information and recording aids.

Under all conditions, subjects finished the experiment within two hours.

Results

Selection sequences

The number of selections for each of the four sequence categories are given in Figures 2-5 as functions of experimental condition and blocks of trials.

Figure 2. Mean number of Category 1 selections as a function of experimental conditions and blocks.

Figure 3. Mean number of category 2 selections as a function of experimental conditions and blocks.

Figure 4. Mean number of Category D (diagonal) selections as a function of experimental conditions and blocks.

Figure 5. Mean number of Category C (cell) selections as a function of experimental conditions and blocks.

In the first block, there are only minor differences among groups as well as among categories. As compared to the passive feedback condition, subjects in the active feedback condition make a considerable number of category C choices, but the number of category D selections is lower.

In the last block, categories 1-2 dominate in the minimum information condition. This is consistent with prediction 1. Category C dominates in the maximum information group. For both groups, the number of category D selections decreases as a function of blocks ($F_{3/42} = 5.26$, $p < .01$), supporting prediction 2. The ANOVA also yielded an ~~condition~~ $\times$ blocks interaction for category C selections ($F_{3/42} = 5.46$, $p < .01$), a result which partly supports prediction 10, that the number of category C selections will be greater in the maximum information condition than in the minimum information condition. Under both conditions, the number of category 2 selections is higher than the number of category 1 selections, a difference which does not change over blocks. This result indicates that subjects detect the quadratic cue 2 rule already in the first block, which would explain why prediction 3, that subjects will increase their relative focussing on cue 2, is not supported. Finally, prediction 6 is contradicted in that subjects, in fact, made category C selections.

The number of selections in the most frequently used category increases as a function of blocks ($F_{3/42} = 3.58$, $p < .05$), indicating that subjects become more consistent with practice, i.e., that they restrict their range of hypotheses, regardless of which hypotheses are dominating for each subject. The most frequently used category for each subject and block is given in Table 2, and in Table 3 the number of subjects per category is summarized for each condition and block.

Although these results are consistent with the above results for the "mean subject", an examination of individual selection patterns also shows that there are great individual differences with respect to selection strategies. It is of some interest that subjects 2 and 4 in the minimum information group, and subject 4 in the maximum information group, all of which have chosen category C as the most frequently used category in the first block, in fact search through the principal diagonal 1 (1:1-5:5) as is shown in Figure 6.

Table 2. Selection sequences: The most frequently used category for each subject, condition, and block.

Condition	Subject	Block:	1	2	3	4
Minimum	1		1	2	2	2
	2		C(D)	C(D)	C(1)	C(2)
	3		D	D	1	2
	4		C(D)	2	2	2
	5		C(2)	C(D)	C(D)	1
	6		D	D	1	1
	7		2/C	C(1)	2	2
	8		2	2	D	1
Maximum	1		D	1/C	C(1/D)	C(1/D)
	2		D	C(2)	C(D)	C(D)
	3		2	C(D)	C(2)	C(2)
	4		C(D)	C(1/D)	C(2)	C(1)
	5		2	2	2	2
	6		D	C(2)	C(D)	C(D)
	7		D	D	C(1/D)	C(1/D)
	8		C(D)	D	D/C	C(D)

Table 3. Selection sequences: Number of subjects choosing different categories as the most frequently used category for each condition and block.

Condition	Category	Block:	1	2	3	4
Maximum	1		1 } 3 ¹⁾	0 } 3	2 } 5	3 } 7
	2		2 }	3 }	3 }	4 }
	D		2	2	1	0
	C		4	3	2	1
Minimum	1		0 } 2	1 } 2 ¹⁾	0 } 1 ¹⁾	0 } 1
	2		2 }	1 }	1 }	1 }
	D		4	2	1	0
	C		2	5	7	7

1) One subject has used two categories equally often (see Table 1).

Subject 2/min	Subject 4/min	Subject 4/max
2 0 0 0 0	8 0 0 0 0	6 0 0 0 1
2 7 0 0 0	0 0 0 0 0	0 0 0 0 0
0 0 7 0 0	0 0 8 0 0	0 0 7 0 0
0 0 0 4 0	0 0 0 0 0	0 0 0 0 0
0 0 0 0 3	1 0 0 0 8	5 0 0 0 6

Figure 6. Distribution of selections for three subjects in Block 1.

On the other hand, there are subjects who distribute their selections over the entire cue matrix. For these subjects, the number of category D selections does not express examination of the principal diagonals.

Distribution of selections

Subjects' distributions of selections over the rows, columns, and principal diagonals of the cue matrix are illustrated in Figures 7-10, for the first block, and for the entire learning stage, respectively.

In the first block, rows 1 and 5, and columns 1, 3, and 5, are more frequently used than the other rows and columns, supporting prediction 4.

This pattern does not disappear when subjects' selections are summed over all four blocks. The same pattern appears in both groups, supporting prediction 9. The ANOVAs yielded a non-significant effect of columns ($F_{4/56} = 2.50$, $p < .10$), but no effects of rows or blocks.

However, as can be seen from these figures the principal diagonal 1 is used more often than any single row or columns for both groups. This result is not consistent with the hypothesis testing model or with the normative model. The ANOVA yielded an effect of diagonals

Figure 7. Distribution of selections over diagonals, columns, and rows for the minimal information group in Block 1.

Figure 8. Distribution of selections over diagonals, columns, and rows for the maximum information group in Block 1.

($F_{1/14} = 10.74$, $p < .01$) as well as an diagonals x blocks interaction ($F_{3/42} = 3.24$, $p < .05$). The latter effect is due to the fact that the number of diagonal selections decreases over blocks. As a consequence, the difference between the two diagonals is reduced.

Results with respect to the mean number of cells from which no selection is made, are given in Table 4, for each condition and block.

Table 4. Mean number of cells from which no selection is made for each condition and block.

Condition	Block:	1	2	3	4
Maximum		10.6	13.0	16.9	16.9
Minimum		11.0	10.5	7.4	6.7

Table 4 shows that the mean number of cells not selected for the minimum information group increases as a function of blocks, in contradiction to prediction 5, and that it decreases for the maximum information group, in contradiction to prediction 8. Also prediction 7, that the number of cells not selected will be greater initially under the maximum information group is contradicted. This result suggests that subjects in the minimum information group become more concerned with the cue integration rule with practice, and that they are unable to decide whether the cues are independent. In part, this result may also be an effect of the fact that the number of category C selections decreases for these subjects.

Discussion

The analyses of subjects' selection strategies show that subjects make use of active feedback control to structure the task information in a way that systematically differs from what is possible under normal, experimenter controlled, feedback presentation. This result supports the suggestion by Lindberg and Brehmer (1977), that negative effects

of outcome feedback demonstrated in previous experiments might be due to the order of feedback instances and the distribution of feedback over the cue matrix, and not only to the fact that outcome feedback contains error variance.

For the minimum information condition results support the assumptions made by the hypothesis testing model (Brehmer, 1974) that subjects focus on separate cues, except for the first block of learning, and that they give priority to quadratic rules as compared to linear rules, after both kind of rules have been detected. That is, the less efficient utilization of the quadratic cue in the test stage is, presumably, due to the fact that the quadratic rule is more difficult to apply, than the linear rule. The present results do not, however, support the notion that subjects first find the form of the cue-criterion functions and then detect the parameters of the functions. This assumption is contradicted by the fact that the number of cells from which no selections are made increases as a function of blocks. On the other hand, this result could also be interpreted to mean that subjects become increasingly concerned with the cue integration rule, i.e., that they start trying to find out whether the two cues are independent.

In the first block subjects in both task information conditions make use of selection strategies which are not consistent neither with the hypothesis testing model nor the normative model (Hammond et al., 1975). In short, subjects do not focus on single cues. Instead, they make consecutive selections from different rows and columns of the cue matrix, they examine the within cell variance, and/or they make their selections from the principal diagonal. These results indicate that, for this task, hypotheses about separate cue-criterion functions are not the strongest initial hypotheses, as is assumed by the hypothesis testing model.

In the maximum information condition, subjects distribute their selections over the cue matrix in the same way as the subjects in the minimum information condition, but they have an increasingly greater number of consecutive selections from the same cell. This result supports the hypothesis that these subjects try to test statistical hypotheses, and that they try to learn to control the task in such a way that they do not need to use the task information provided in advance.

In the test stage, subjects in the maximum information condition increase their task control with practice. In the learning stage, these subjects change their selection strategies as a function of blocks. Specifically, they increase their relative focus on the within cell variance. This may indicate that the subjects investigate different-performance criteria more directly, which they cannot do by inspection of the cue-criterion matrix. For example, a trial-by-trial examination of the task give the subjects empirical evidence that the mean prediction error (R-E) is not reduced if they try to match the within cell criterion variance in their predictions. It is possible, then, that the improvement in task control over blocks in the test stage is due to the fact that the subjects are able to eliminate certain nonoptimal hypotheses about performance criteria, by trial-by-trial examination of the within-cell variance in the learning blocks, and that they, accordingly, approach the least squares criterion assumed by the normative model.

However, since the subjects in the maximum information condition did not reach the maximal level of task control in the test stage, they probably did not find the optimal performance criterion. If the subjects are unable to find this criterion here, they probably do not do so in ordinary experiments either and we may therefore assume that at least part of their inconsistency is due to this factor in standard experiments on multiple-cue probability learning. That is, subjects may not follow the optimal prediction strategy even if they know everything about the structure of the task, if they are not informed that the normative model defines the $\bar{E}$ - values as the optimal predictions. It should be observed that this criterion could not be derived solely on the basis of an examination of the task structure. If subjects choose alternative performance criteria, therefore, it may be misleading to attribute all of their lack of consistency to lack of cognitive control (cf. Hammond, et al., 1975).

In conclusion, the active feedback condition, used in the present experiment, has proved to be a valuable device for examining the subjects' hypothesis testing activity. To increase the efficiency of this method, it should be combined with the "think aloud" technique, however. It may also be necessary to introduce certain restrictions on subjects' selection strategies, by means of instructions, or other methods to make it possible to eliminate some hypotheses about what the subjects actually try to achieve when they select task information.

This study was supported by a grant from the Swedish National Defence Research Institute (FOA). The authors are indebted to Jan Kuylenstierna and Lars-Erik Warg for assistance in the experimental and computational work.

References

- Brehmer, B. Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. Organizational Behavior and Human Performance, 1974, 11, 1-27.
- Brehmer, B., Kylenstierna, J., & Liljengren, J. Effects of function form and cue validity on subjects' hypotheses in probabilistic inference tasks. Organizational Behavior and Human Performance, 1974, 11, 338-354.
- Brehmer, B., & Kylenstierna, J. Evaluation of performance in the learning of probabilistic inference tasks. Umeå Psychological Reports No. 91, 1976.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. A study of thinking. New York: Wiley, 1956.
- Hammond, K. R., Stewart, T. R., Brehmer, B., & Steinman, D. D. Social judgment theory. In M. Kaplan and S. Schwartz (Eds.), Human Judgment and Decision Processes: Formal and Mathematical Approaches. New York: Academic Press, 1975.
- Lindberg, L-Å., & Brehmer, B. Effects of task information and active feedback control in inductive inference. Umeå Psychological Reports No. 123, 1977.