

Knowledge acquisition and strategic action in "Mastermind" problems

JOHN B. BEST

Eastern Illinois University, Charleston, Illinois

In two experiments, the strategies used by subjects playing the logical-deduction game, Mastermind, were examined. In the first experiment, subjects showed improvement resulting from the continued use of a particular strategic action, and the data suggested that the subjects learned the strategy from their transactions with the task. In the second experiment, the question of changes in underlying strategic knowledge of Mastermind was examined. The accuracy and complexity of the subjects' deductions and their use of the previously identified strategy were used to generate a model of the cognitive operations involved in Mastermind. Although there were improvements in the accuracy and complexity of the subjects' deductions resulting from continued play, these improvements were unrelated to the use of the strategy. Moreover, the likelihood of making accurate and complex deductions was well accounted for by a Markovian model, suggesting that the deployment of the strategy was not driven by any change in the subject's underlying knowledge structures. Rather, the subjects seemed to use the strategy to create Mastermind situations whose interpretation was fairly easy. The implications for previous work on the issue of Mastermind strategies and the development of logical-deduction strategies are discussed.

One of the important legacies of the work of Bruner, Goodnow, and Austin (1956) was the notion that humans are strategic in their acquisition of knowledge. As is commonly known, Bruner et al. identified several strategies used by subjects, and their findings suggested that subjects who perform well on a task do so precisely because they use a particular approach. According to the view implicit in Bruner et al.'s research, and later solidified by subsequent findings (e.g., Anzai & Simon, 1979; Langley, 1985), individuals who wish to solve a problem first access their knowledge of the problem and then use this knowledge to assemble a strategy, meaning here a plan that specifies actions to be taken.

The influence of the action taken on the knowledge used to generate it has not been well understood. One common view is that information derived from the action taken, that is, from the outcome of the strategy, "loops back" onto the knowledge generating the action, resulting in a change in the amount or the organization of the information the person is able to access. According to this view (e.g., as seen in the research of Newell & Simon, 1972), "good" strategies are actions that are most informative in the sense of being most likely to loop back onto the underlying cognitive structure and produce a change therein. The resulting structural change may in turn drive a change toward even more superior actions, suggesting that structural change and strategic action may have a "bootstrapping" relationship with each other.

With this conceptual framework as their base, Laughlin, Lange, and Adamopoulos (1982) studied strategy use in the logical-deduction game, Mastermind. In the standard version of this game, the solver's task is the deduction of a left-to-right ordering of four color names, called the code. To deduce the code, the solver generates a hypothesis of four or fewer color names. Feedback is then supplied that informs the solver about the quality of the hypothesis. In general, feedback indicates the congruence between that particular ordering sequence and the code. Feedback can be black or white. Each unit of black feedback indicates that one of the color names in the just-advanced hypothesis matches a code member in color and location. Each unit of white feedback indicates that one of the color names in the hypothesis matches a code member in color, but not in location. In the standard game, each code consists of four colors drawn with replacement from a pool of six colors, thus affording 1,296 different codes. Each hypothesis that the subject makes, and its associated feedback, is displayed using a set of plastic tokens. A typical hypothesis and its associated feedback are shown below:

Code:	RD	GN	BL	WH
Hypothesis:	RD	YW	WH	BK
Feedback:		1 BK,	1 WH	

Thus, Mastermind seems to provide a useful situation for studying problem solving: it retains many of the formal mathematical properties of the Bruner et al. (1956) task, yet it has continued to be a popular commercial item for more than a decade.

The game tree for the standard game being intractably large, Laughlin et al. (1982) studied the performance of

Reprints may be obtained from John B. Best, Department of Psychology, Eastern Illinois University, Charleston, IL 61920.

—Accepted by previous editor, Alice F. Healy

subjects who played a reduced version of Mastermind. Their version had four dimensions or locations (i.e., first location, second location, etc.) that could each take on, with replacement, one of three numerical values (0, 1, or 2), thus producing a "universe" of 81 codes. This reduction resulted in four formally equivalent opening hypotheses that could be used by the solver: (1) all four dimensions having the same value (three instances: 0000, 1111, 2222); (2) three dimensions having the same value with the fourth dimension having a different value (24 instances: 0001, 0002, ... 2221); (3) two dimensions having the same value with two dimensions having a different value (18 instances: 0011, 0022, ... 2211); and (4) two dimensions having the same value with each of the remaining two dimensions having a different value (36 instances: 0012, 0021, ... 2210). These four classes of hypotheses were then designated by their lowest base-three instance, prefixed by a letter (F = first hypothesis, etc.): F0000, F0001, F0011, and F0012.

As a basis with which to compare human performance, Laughlin et al. (1982) evaluated the theoretically optimal efficiency of each of these four classes of opening hypotheses by using a Monte Carlo technique. Assuming an information processor who simply picked members of each class of hypothesis randomly, but who was capable of making all currently available deductions, simulations showed fairly small but intriguing differences. In the two extreme cases, when a hypothesis of the F0000 class was used as the initial hypothesis, a mean of 4.68 hypotheses were required for a solver to have the necessary and sufficient information to deduce the code. But when a hypothesis of the F0012 class was used as the initial hypothesis, a mean of 3.95 hypotheses were required to deduce the code, a reduction of 16% from the F0000 situation.

Laughlin et al. (1982) next identified plausible human strategies, and they compared the efficiency of these strategies to the theoretical optima generated by the Monte Carlo technique. For example, in Mastermind, a focusing strategy consists of beginning with one of the three F0000 hypotheses. By doing so, the solver will find out how many of the dimensions in the code have the same value as the hypothesis. Thus, if the feedback from the above hypothesis is "2 Black" units, the solver knows that two dimensions in the code will have 0 as their value. The use of the focusing strategy permits the solver to deduce the values of the code in two hypotheses, leaving only the task of assigning those values to their correct locations. This task can also be approached algorithmically, requiring an additional 2.5 hypotheses. Thus, the optimal use of the focusing strategy in the reduced version of Mastermind requires 4.5 hypotheses to deduce the code. This is less than the 4.68 hypotheses required by a random selection of hypotheses initiated with an F0000 hypothesis, but more than the 3.95 hypotheses required when the selections are initiated by an F0012 hypothesis.

In addition to this focusing strategy, Laughlin et al. (1982) also identified a powerful tactical strategy (Johnson, 1971, 1978). Essentially, the tactician conceptually

divides all hypotheses into two categories: those that will be eliminated as a function of anticipated feedback, and those that will continue to be viable. The difficulty of using this strategy stems in part from cognitive strain considerations, but the larger burden would seem to be the depth of knowledge required to implement the strategy appropriately. As noted by Laughlin et al., the solver must advance a series of maximally informative hypotheses over anticipated distributions of feedback. This means that the solver must have some knowledge of the patterns of feedback that are both possible and plausible given a number of as-yet-undeployed hypotheses, and must be able to deduce from this information the hypotheses that remain plausible and those that have been eliminated. Such knowledge would appear to be based upon a deep understanding of the game's structure and upon a thoroughgoing comprehension of the principles of logical deduction. However, if done correctly, the tactical strategy is powerful. Collapsing across the four classes of initial hypotheses, Laughlin et al. found an expected 3.58 hypotheses to solution when the tactical strategy was initiated by an F0000 hypothesis, 3.15 for an initial F0001, 3.08 for an initial F0011, and 3.02 for an initial F0012. As these figures suggest, the tactical strategy is most powerful when it is initiated by an F0012 hypothesis, and on that basis Laughlin et al. concluded that the tactical strategy was most likely to be initiated by an F0012 hypothesis. Subjects using the less demanding and less powerful focusing strategy, on the other hand, should be most likely to initiate it with an F0000 hypothesis.

In their first experiment, Laughlin et al. (1982) found that these two strategies apparently predominated, with 36% of all problems begun with an F0000 hypothesis (i.e., the focusing strategy) and 31% begun with an F0012 hypothesis (i.e., the tactical strategy). Subjects required a mean of 5.27 hypotheses to solve problems initiated with the F0000 hypothesis type, 6.47 hypotheses for F0001, 6.1 hypotheses for F0011, and 5.98 hypotheses for F0012. Thus, as expected, subjects using the focusing and tactical strategies performed better than subjects using some other approach. In addition, subjects using the focusing strategy outperformed those using the more powerful tactical strategy, in part because of the demands imposed by the latter. In Laughlin et al.'s second experiment, subjects were induced to use a particular strategy by the problem instructions, which limited the subject's first hypothetical choice to a member of one of the four classes of hypotheses. Subjects who were induced to use the focusing strategy required 5.87 hypotheses to deduce the code; those who were induced to use the tactical strategy required 6.23 hypotheses. Both groups outperformed subjects who were instructed to begin with a hypothesis of another class (6.47 hypotheses for F0001, 7.29 hypotheses for F0011).

In summary, the Laughlin et al. (1982) experiments seem to make the following principal points:

1. An ideal information processor using an algorithmic approach in a reduced version of Mastermind performs

better or worse depending upon the class of the opening hypothesis used.

2. Some Mastermind strategies seem to correspond to the strategies observed in the selection tasks used by Bruner et al. (1956), and a Mastermind focusing strategy can be categorized on the basis of the initial class of hypothesis deployed. A Mastermind scanning strategy offers the possibility of superior performance, relative to a focusing strategy, particularly when it is initiated by a hypothesis of a particular class.

3. Subjects who use either the focusing or the scanning strategy should outperform solvers using other strategies.

4. The Mastermind scanning strategy requires a deep comprehension of the situation and of logical deduction, and for these reasons, it is a difficult strategy to use consistently.

5. Solvers who are induced to begin a game with either strategy outperform subjects who are induced to begin a game with a hypothesis of another category.

However, Laughlin et al. (1982) did not take into account what seems to be an important distinction between strategic behavior and underlying structural knowledge. Consequently, solvers who use what seems to be a fairly sophisticated strategy may not necessarily have a deep knowledge of the task's underlying structure. Thus, strategic actions may be driven by mechanisms other than underlying structural knowledge. In addition, Laughlin et al. have suggested that a single exposure to a particular starting sequence is sufficient experience to induce a complex strategy. But it seems paradoxical to maintain that a complex strategy can be induced by a single exposure, while claiming, as Laughlin et al. did, that successful use of the tactical strategy requires a deep knowledge of the logical structure of the game. Some researchers (e.g., Chi, Glaser, & Rees, 1982) have contended that underlying structural knowledge is built up only as a result of lengthy personal experience with the task, perhaps as the result of trying several strategies. If such contentions are accurate, then structural knowledge, if it emerges at all, should be seen only after fairly lengthy experience with the task.

Related to the notion of lengthy experience as a necessary underpinning of structural knowledge is the concept of *dynamism*, or flexibility in the use of strategic actions. Although Laughlin et al. (1982) did not address this issue explicitly, their findings suggest a certain fixity to strategic actions: once the subject adopts a strategy, his or her rate of use of that strategy seems to be more or less constant. However, strategic actions may instead follow a developmental course. If such is the case, then it would seem reasonable to expect some development of solvers' strategic actions. It seems unlikely that solvers can immediately comprehend the effectiveness of the most efficient strategy on any but fairly trivial tasks. Moreover, even when strategic action emerges quickly on a complex task, the rate of use of this strategic action should change as a function of the solver's experience.

These issues were explored in the following experiments. It was contended that:

1. The strategic actions of individuals are only weakly specified by their initial choice of hypothesis. Rather, subjects' strategies are better seen in their choice of a subsequent hypothesis after feedback has been given. In the present study, it will be argued that the deployment of different classes of hypothesis in Mastermind is almost completely driven by the feedback to which such hypotheses are in response, rather than by the subject's adherence to a focusing or tactical approach. Moreover, the principal effect of using such a strategy is to alter the likelihood of forthcoming feedback events to reduce the likelihood of uninformative feedback patterns. Consequently, subjects who perform well overall should do so because they get less of the prevalent feedback in the first place, and/or because they learn a strategy that reduces the likelihood of such feedback.

2. The strategic actions of individuals are not necessarily fixed; rather, subjects learn to deploy particular classes of hypotheses in response to feedback. This means that the deployment of a particular class of hypothesis is related in a principled way to the subject's previous deployment of that class of hypothesis. Thus, we can expect solvers to become increasingly likely to deploy particular classes of hypotheses in response to particular feedback patterns.

3. The consensual view holds that such strategy use is supported or enabled by an underlying cognitive structure which, in turn, is modified and improved by the continued use of the strategy. In the case of Mastermind, it seems that this cognitive structure must include some knowledge of logical deduction. Consequently, the consensual view would argue that subjects who improve on a logical-deduction task probably learn something about logical deduction as an outcome of their successful use of a particular strategy. In this paper, it will be contended that the repeated deployment of particular classes of hypotheses in response to particular feedback patterns (i.e., the use of particular strategies) does not necessarily produce an improvement in cognitive structure. That is, solvers who use a strategy do not necessarily know or learn much about logical deduction per se. This means that the continued use of a strategy must be supported by some other kind of event, and in this paper it will be argued that most solvers who play the game several times over a several-week period maintain particular strategies because such strategies are likely to produce an increase in the total amount of feedback.

EXPERIMENT 1

Method

Subjects. Fifty-five undergraduates at Eastern Illinois University who were unfamiliar with the experimental task received extra credit in their upper-division psychology classes in exchange for their participation.

Experimental Task. The subjects played a version of the standard Mastermind game, using the board and its plastic tokens. Thus, in each game, the subjects were able to see all previously made hypotheses and the feedback associated with them. To establish an intermediate level of difficulty on the task, the subjects were told that each code consisted of four different colors, randomly drawn from a pool of six colors. There were thus 360 different possible codes in this version of the task.

Procedure. The experimenter explained the rules of the task and answered any questions the subject may have had. When the subject seemed familiar with the task, the problem-solving session began. The subjects were informed that they would play three consecutive games, or until a 1-h time limit had elapsed. In each game, the subjects were permitted to make 10 consecutive hypotheses. If the subject had not deduced the code after making 10 hypotheses, play on that game was discontinued. The subjects were not overtly compelled to play rapidly.

Results

A performance criterion was established and applied to the subjects' responses. Subjects were eliminated from the analysis if they failed to complete three games, or if they made 10 hypotheses in either their second or third game without deducing the code. Failure to meet the performance criterion resulted in the elimination of 11 subjects from further analysis.

The data were scored by counting the number of hypotheses each subject made in each game, with fewer hypotheses indicating greater skill in deduction. The mean number of hypotheses required to deduce the code on the initial attempt was 6.8. In Games 2 and 3, 5.8 and 5.2 hypotheses, respectively, were required. Planned comparisons indicated that the mean number of hypotheses required on Game 2 was significantly lower than that required on Game 1 [$t(43) = 2.77, p < .01$]. The mean number of hypotheses produced in Game 3 was not significantly lower than the mean number required in Game 2 [$t(43) = 1.76, p = .08$].

One of the contentions of the present experiment was that the deployment of different classes of hypotheses is almost completely driven by the feedback that precedes the hypothesis. Demonstrating this point involves the development of a procedure capable of categorizing Mastermind hypotheses and the application of this procedure to the hypotheses made by the subjects. To generate a hypothesis, solvers typically review their immediately previous hypothesis and the feedback it received. Usually, subjects modify the previous hypothesis in fairly typical ways. For example, a solver may change one of the colors in a given hypothesis, selecting a color from the pool that had not been used on the previous hypothesis. If all other colors from the previous hypothesis were played again in the same locations, then such a hypothesis was coded a "one color" or "1C" hypothesis. In addition to, or instead of, changing colors, a solver may elect to change only the locations of previously used colors on a given hypothesis. If, for example, a subject produced a hypothesis by shifting two of the colors used in the just-previous hypothesis, then such a hypothesis was coded a "two location" or "2L" hypothesis. This procedure does not per-

mit the categorization of the subjects' opening hypothesis of each game, but such hypotheses were fairly uniform; 88% of the opening hypotheses consisted of four different colors. All hypotheses subsequent to the opening one were codable using this procedure.

The feedback-driven nature of the subjects' hypothesis-making can be seen by looking at the hypotheses made by the subjects in relation to the feedback they had just been given. The 44 subjects passing the performance criterion made a total of 651 codable hypotheses. Table 1 shows the relationship between color-change hypotheses (collapsed across all location changes) and the total number of feedback pins awarded. As Table 1 shows, the number of color changes was inversely related to the number of feedback pins awarded on the previous hypothesis. When no feedback or one feedback pin was awarded, the subjects introduced an average of 2.4 new colors into their next hypothesis. When two feedback pins were given, the subjects introduced 2.13 new colors; for three feedback pins, only one new color was played. Finally, when four feedback pins were awarded, essentially no color changes were made. What this suggests is that, in deciding whether to include a new color in a forthcoming hypothesis, the subjects simply reviewed the total number of feedback pins they had just been awarded. The subjects tended to produce hypotheses with many color changes if few pins had been awarded, and few color changes if many pins had been awarded.

Table 2 shows the 651 codable hypotheses retabulated to show the relationship between location changes (collapsed across all color changes) and the number of black feedback pins awarded. As Table 2 shows, the number of location changes was inversely related to the number of black feedback pins awarded to the previous hypothesis. When the subjects received no black feedback pins, they made an average of 2.9 location changes. However, when three black feedback pins were awarded, the subjects made

Table 1
Frequency of Color-Change Hypotheses as a Function of the Total Number of Feedback Pins Awarded

Type of Hypothesis	Total Feedback Pins			
	0-1	2	3	4
No Color Change	0	11	26	187
1 Color Change	1	55	223	1
2 Color Changes	2	117	26	0
3-4 Color Changes	2	0	0	0

Table 2
Frequency of Location-Change Hypotheses as a Function of Black Feedback Pins Awarded

Type of Hypothesis	Black Feedback Pins			
	0	1	2	3
No Location Change	9	16	38	36
1 Location Change	8	53	35	2
2 Location Changes	64	88	48	2
3 Location Changes	77	69	10	0
4 Location Changes	83	9	4	0

only 0.15 location changes. The findings shown in Table 2 suggest that, in trying to decide whether to change the location of a color in a forthcoming hypothesis, the subjects reviewed the number of black feedback pins they had just been awarded. Tables 1 and 2 suggest that the subjects were not using either a focusing or a tactical strategy. Rather, the elements of a forthcoming hypothesis seem to be well accounted for by an analysis of the most recently given feedback.

An additional contention of the current experiment was that the principal effect of using strategic actions in Mastermind is to alter the distribution of forthcoming feedback events to reduce the occurrence of prevalent (and perhaps uninformative) feedback patterns and to increase the occurrence of less likely, but perhaps more informative, feedback. Consequently, subjects who perform well overall should do so because they get less of the popular feedback in the first place, and/or because they acquire the strategic action more quickly than do other subjects.

Some types of feedback were indeed given more frequently than other types. Of the 13 combinations of feedback possible, the five most frequently occurring types in Game 1 accounted for 65% of all feedback given. These five "popular" types of feedback also accounted for a substantial proportion of the feedback awarded in Games 2 and 3 (59% and 49%, respectively).

Table 3 shows the frequency of each hypothesis type made in response to the five most popular feedback types, aggregated across Games 1-3. As Table 3 suggests, the subjects usually produced a particular hypothesis type in response to each feedback type. Of the 13 feedback types, 10 were given at least 30 times to different subjects, and each of these 10 feedback types was followed by a different predominant response. This predominant response, called the *modal hypothesis*, accounted for between 48% and 96% of all responses to each of the feedback types. The mean response rate of all modal hypotheses was 72%.

The relationship between modal hypothesis-making in response to popular feedback and performance is clear. The 11 subjects whose overall performance was poor (i.e.,

at the 25th percentile or lower, M hypotheses/game = 7.6) typically received some form of popular feedback (55% of all hypotheses) and responded to popular feedback with the modal hypothesis 51% of the time. For the 13 subjects whose performance was significantly better than this group [i.e., subjects at the 70th percentile or higher, M hypotheses/game = 4.4, $t(22) = 3.07$, $p < .001$], much less popular feedback was received (38% of all hypotheses), and these subjects responded to popular feedback with the modal hypothesis 86% of the time, which is significantly more frequent than the 38% of the low-performance group ($z = 3.88$, $p < .001$). This finding suggests that subjects who respond with the modal hypothesis to popular feedback should perform better than subjects who do not respond modally to such feedback. This expectation was confirmed: the 10 subjects who had the highest modal response rate, regardless of the amount of popular feedback they were given (M modal response rate = 100%), performed significantly better (M hypotheses/game = 4.53) than the 10 subjects whose rate of modal response was closest to the median modal response rate [M modal response rate = 67%, M hypotheses/game = 6.46, $t(18) = 4.33$, $p < .001$]. However, the performance of the 10 subjects whose rate of modal response was closest to the median was not significantly different than that of the 10 subjects who had the lowest modal response rate [M modal response rate = 32%, M hypotheses/game = 6.67, $t(18) = .42$, $p > .05$].

As contended, popular feedback seems to be rather uninformative. Consequently, subjects who, for whatever reason, receive substantial amounts of popular feedback seem to perform worse overall than subjects who receive moderate amounts of it. Furthermore, the results suggest that the practical effect of deploying the modal hypothesis in response to popular feedback is to decrease the likelihood of getting another round of popular feedback in response to the hypothesis.

In the current experiment, it was also contended that the strategic actions of individuals are not necessarily fixed. Rather, subjects learn to deploy particular classes

Table 3
Number of Hypothesis Types Made in Response to Popular Feedback Types

Hypothesis Type	Feedback Type				
	2 White	1 Black, 1 White	3 White	1 Black, 2 White	2 Black, 1 White
1 Color	2	2	1	5	6
1 Color, 1 Location	2	4	3	7	27
1 Color, 2 Locations	5	16	4	65	4
1 Color, 3 Locations	14	0	58	3	2
2 Colors	3	6	0	1	5
2 Colors, 1 Location	3	29	0	13	1
2 Colors, 2 Locations	49	1	5	1	0
3 Colors	0	0	0	0	0
3 Colors, 1 Location	0	0	0	0	0
4 Colors	0	0	0	0	0
1 Location	0	0	0	0	0
2 Locations	1	2	0	1	3
3 Locations	0	1	2	6	3
4 Locations	6	0	8	0	1

of hypotheses in response to feedback, implying that the deployment of a particular class of hypothesis should be related in a principled way to the subject's previous deployment of that class of hypothesis. As an outcome of this learning, we can expect solvers to become increasingly likely to deploy particular classes of hypotheses in response to particular feedback patterns as they gain experience with the task.

The outcome of the following analysis supports this contention. Each occurrence of a particular type of feedback was considered an "encounter" with that feedback type. The proportion of modal hypotheses made by individuals at each encounter was computed for subjects who had, and for those who had not, produced a modal hypothesis on the previous encounter. This computation was carried out separately for each of the five popular feedback types, and the computation continued for each of the five types of popular feedback as long as a majority of the subjects had such an encounter. After the proportions were computed for each encounter, for each feedback type, the resulting proportions were summed across encounters to produce five transition matrices showing the overall proportion of nonmodal-to-modal transitions for each of the popular feedback types. If any changes in the subjects' hypothesis-selection procedure occurred as a function of their experience with the task over Games 1-3, then such changes should be reflected in changes in these transition matrices—that is, in changes of the frequency with which hypotheses were selected.

Finally, subjects' choices may be compared with a theoretically plausible model. If the act of producing a modal hypothesis is completely unrelated to earlier productions, then such a selection process can be described as *Markovian*. A Markovian process is said to be path-independent, in this case meaning encounter-independent. Thus, the transition matrix of a Markovian process describes the process almost completely: if the Markovian property is observed, then the probability of a solver's entry into a particular state (i.e., choice of a particular hypothesis) is independent of all but the immediately previous encounter with a particular feedback type. Rejection of the Markovian hypothesis implies that the transition matrices for each encounter, and for each form of popular feedback, are significantly different from one another.

The transition matrices were then used to generate expected frequencies of modal-hypothesis selection by encounter, for each type of popular feedback. The expected frequencies can be compared with those observed using a Wald-Wolfowitz runs test (Siegel, 1956; Swed & Eisenhart, 1943). When this is done, the Markovian hypothesis is rejected (critical value = 24, obtained value = 16, $p < .05$). The rejection of the Markovian hypothesis means that subjects' responses to popular feedback are not path- or encounter-independent. Rather, the likelihood of a subject producing a modal hypothesis on any given encounter is a function of the number of times the subject has encountered that feedback. Specifically, the sub-

jects responded to popular feedback with the modal hypothesis significantly more frequently on each encounter with popular feedback than would be expected if the Markovian hypothesis were true.

Discussion

It seems that subjects improve quickly and substantially on this difficult task. In addition, as contended, the subjects' strategies were only weakly specified by their initial choice of hypotheses. Rather, the subjects' strategic actions were apparently almost completely driven by the feedback they were awarded. Moreover, the data suggest that the principal effect of using certain strategies in Mastermind is to alter the distribution of forthcoming feedback events in such a way as to reduce the occurrence of prevalent (and perhaps uninformative) feedback patterns and increase the occurrence of less likely, but more informative, ones. Thus, subjects who performed well overall apparently did so because they got less of the popular feedback in the first place, and because they acquired the appropriate strategic action more quickly than did other subjects. The subjects also learned to deploy particular classes of hypotheses in response to feedback. Evidence for this conclusion was provided by evaluating a path-independent Markovian hypothesis, according to which the deployment of particular classes of hypotheses was not related in any principled way to the subject's previous deployment of such hypotheses. Contrary to the Markovian model, the subjects became increasingly likely to deploy the modal hypothesis as a function of encounters with popular feedback.

The rejection of the Markovian hypothesis is particularly interesting because Markovian models have been used as theoretical statements of learning in concept-formation tasks (Coombs, Dawes, & Tversky, 1970; Wickens, 1982). Markovian models typically assume that concept learners in a selection paradigm make more-or-less unprincipled choices of the next exemplar they wish to examine until certain conditions are fulfilled that produce insight. In such a context, insight should be understood as complete, or almost complete, knowledge of the rule that links the exemplars. The achievement of insight in concept-learning tasks is indicated in terms of the subject's choice of an exemplar to examine, and by the subject's correct prediction concerning the inclusion or exclusion of that exemplar from the concept. Thus, when applied to concept-learning tasks, Markovian models argue that subjects initially make unprincipled selections because they learn little or nothing about the rule linking the exemplars until they achieve insight. As seen in this experiment, the rejection of a Markovian model suggests that the subjects realized how to deal with different feedback patterns on their initial encounter with such feedback.

EXPERIMENT 2

The findings of the first experiment suggest that the hypotheses produced by subjects in Mastermind are es-

entially responses to the most recently awarded feedback, rather than instantiations of a multiple-hypothesis focusing or scanning strategy. In addition, the findings raise the possibility that such strategic actions may be produced by cognitive processes that are at least somewhat independent of the operations underlying logical deduction per se. That is, the findings of the previous experiment raise the possibility that strategic actions and underlying structural knowledge of Mastermind, and of logical deduction more generally, may not necessarily have a mutually enhancing relationship. Such a view can be contrasted with what could be called the *consensual* position.

The consensual view (e.g., Anzai, 1987; Ohlsson, 1987) holds that such strategy use is supported or enabled by an underlying cognitive structure which, in turn, is modified and improved by the continued use of the strategy. In the case of Mastermind, one might argue that, at a minimum, this cognitive structure must include some knowledge of logical deduction. Consequently, proponents of the consensual view might argue that subjects who improve on a logical-deduction task must learn something about logical deduction as an outcome of their successful use of a particular strategy.

The purpose of Experiment 2 was to demonstrate that the consensual view is not particularly correct with regard to the use of strategies in Mastermind. Specifically, this experiment shows that the repeated deployment of particular classes of hypotheses in response to particular feedback patterns (i.e., the use of particular strategies) does not necessarily produce an improvement in the cognitive structure that, according to the consensual view, is responsible for generating a strategic action and for interpreting its outcome. In other words, solvers who become adept in the use of a particular strategy do not necessarily know or learn much about applying logical-deductive principles to Mastermind. Moreover, this experiment presents evidence that most solvers who play the game several times over a several-week period maintain their use of particular strategies because such strategies appear likely to produce an increase in the total amount of feedback given.

Method

Subjects. Thirteen undergraduates at Eastern Illinois University received extra credit in an upper-division psychology class in exchange for their participation. No subject participated in both experiments reported here.

Procedure. The procedure was similar to that used in the first experiment. The subjects were informed that they would play two consecutive games, or until a 30-min time limit had elapsed. The subjects scheduled themselves for one such problem-solving session per week, for a period of 5 weeks.

My contention is that subjects who become adept in the use of a particular strategy do not necessarily know or learn much about applying logical-deductive principles to Mastermind. The subjects' abilities to make such deductions were assessed in the following way. Immediately prior to making every hypothesis except the first one of each game, the subjects were asked to indicate the extent of their knowledge at that point by filling out a form. The subjects

used a shorthand notation to indicate the deductions they had made about each color in the pool.

There were three broad categories of deductions. The subjects marked "consideration" deductions if they believed that a particular color might be a code member. There were three levels of consideration deductions, corresponding to the extent to which the subject had narrowed the possible locations for the color, if the color was indeed a code member. Level 1 considerations (the narrowest) were marked by subjects who had determined that if a color was in the code, then it must be located at one particular location. Level 2 considerations were marked by subjects who had determined that a putative code member must be located at two or three particular locations. Level 3 considerations were marked by subjects who had not narrowed the putative code member's possible locations.

"Inclusions/exclusions" was the second category of deductions, and there were two types. Inclusions were marked by a subject who had determined that a particular color was definitely a code member, although no determination could be made about its location. Exclusions were marked by a subject who had determined that a color was definitely *not* a code member.

"Assignments" was the third category of deductions, and there were two types of these. Assignment to more than one location was marked by subjects who had deduced that a particular color must be a code member, and must be correctly located in one of two or three indicated locations. Assignment to one location was marked by subjects who had narrowed the location of a positive code member to a single position.

The subjects were not required to make deductions about each color after each hypothesis, but they were asked to record the extent of their current deductions as completely as possible. A sheet summarizing the shorthand notation was visible during the entire problem-solving session. After the first game or two, the subjects asked no questions about the use of the notation. After the subjects were conversant with the notation, filling out the form required no more than 5 sec in the majority of cases.

Results

Performance of subjects. The subjects' performance improved over the course of the 5-week test interval. The mean number of hypotheses required to deduce the code in the initial week of play was 7.38; by Week 4 of play, 5.71 hypotheses were required. The difference between these means is significant [$t(12) = 3.63, p < .01$]. Performance worsened in Week 5 of play, which coincided with the final week of classes in the Spring semester ($M = 6.33$ hypotheses/game).

Type, accuracy, and complexity of deductions. The nature of the deductions made by the subjects also changed over the 5-week period. In Week 1, 72% of all deductions were considerations. By Week 4, this proportion had declined significantly, to 56% ($z = 5.33, p < .01$), and by Week 5, to 45% of all deductions. In contrast, in Week 1, 9% of all deductions were assignments. This proportion increased significantly by Week 4, to 22% ($z = 5.65, p < .01$), and by Week 5, to 37%. The percentage of inclusions/exclusions did not change significantly over the time period.

Accurate deductions were those in which a subject correctly considered, included, eliminated, or assigned a color. The proportion of accurate deductions, relative to all deductions made, increased over the course of the 5-week interval. In Week 1, 44% of each subject's deduc-

tions were accurate. By Week 4, this percentage had increased significantly, to 63% [$F(3,36) = 7.52, p < .001$], and in Week 5, 60% of each subject's deductions were accurate.

It would seem that some deductions are harder to make than others. For example, a Level 3 consideration was always available to the subject for each color, but the assignment of a color to a specific position requires determining that a color is definitely included in the code, as well as deducing the color's location. Considering the elimination of a color, and the assignment of a color to a single location as "complex" deductions, a higher percentage of the deductions made were complex following a modal hypothesis (25%) than following a nonmodal hypothesis (15%) [$F(1,12) = 15.9, p < .005$].

Modeling the acquisition of knowledge. The subjects were more likely to make complex deductions after making a modal hypothesis than they were after making a nonmodal hypothesis. However, the fact that accurate deductions were equally likely after modal and nonmodal hypotheses suggests that the cognitive processes used to deduce information from the problem array were not necessarily linked to the cognitive processes used to determine the nature of the next hypothesis. This distinction is shown in Figure 1, which is a representation of some of the putative operations in which a problem-solver faced with this task might engage. Note that some of the operations result in an action such as production of a modal hypothesis. Other operations, such as those involved in making accurate and complex deductions, are internal.

Superimposed over these operations are the findings that result when subjects engage in the operations in certain "sequences." These sequences can be described in terms of (1) the nature of the feedback (popular or nonpopular) that initiates the sequence, (2) the timing of the deductions (immediately or following the production of a hypothesis), and (3) the nature of the hypothesis (modal or nonmodal) generated following the awarding of feedback.

For example, Sequence 1 describes the situation in which the solver, having received popular feedback, attempts to make accurate and complex deductions immediately from the feedback. Solvers who are successful at this can be described as having entered a particular mental state—that of having made such a deduction. By contrast, Sequence 3 describes an operation initiated by popular feedback, followed by making a modal hypothesis, followed by an attempt at making accurate and complex deductions. Solvers who succeed at this can be described as having entered a mental state, too, but perhaps a somewhat different one than that entered by solvers operating in Sequence 1, because making a modal hypothesis has preceded the deductive processes in the latter case.

As Figure 1 suggests, when the subjects received nonpopular feedback, they made significantly more deductions ($M = 4.11$) than when they received popular feedback ($M = 3.76$) [$t(12) = 3.5, p < .01$]. Moreover, as shown in Figure 1, the likelihood of a subject's making a deduction that was both accurate and complex was four to five times greater when the subject was operating in a sequence initiated or terminated by nonpopular feed-

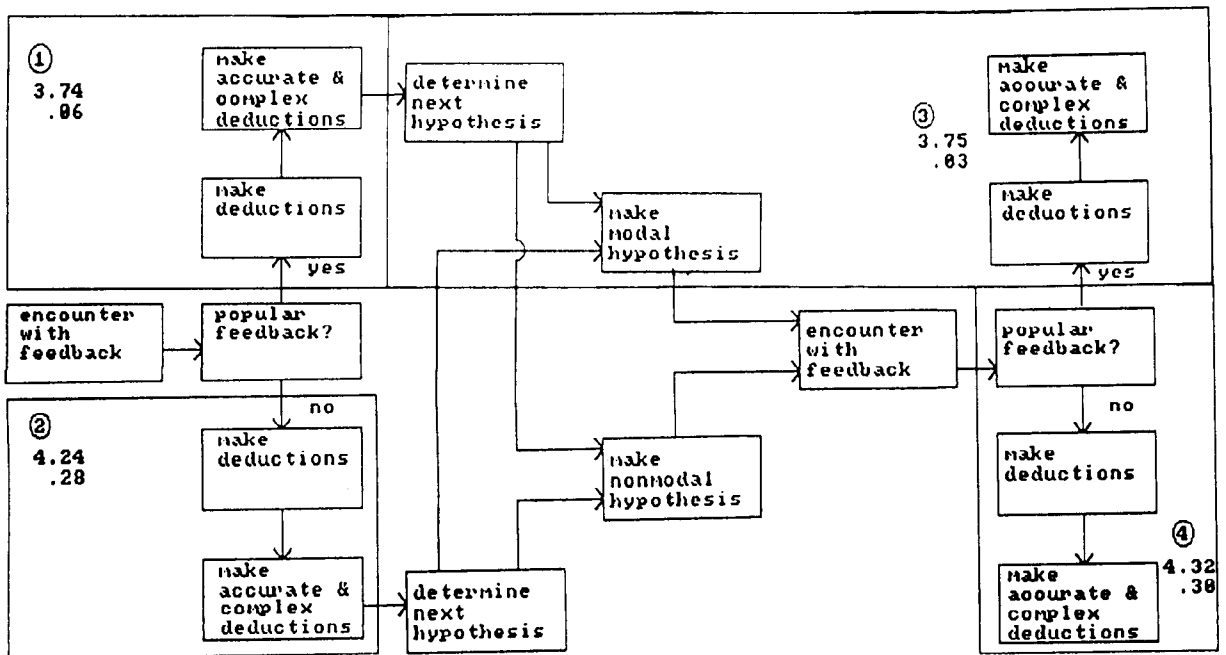


Figure 1. Putative operations engaged in by Mastermind solvers. Below each of the four circled sequence numbers are two other numbers, the first showing the mean number of deductions made when that route of information processing was taken, the second number, immediately below the first, showing the proportion of all deductions made that were both accurate and complex.

back than when he/she was operating in a sequence initiated or terminated by popular feedback.

One of my contentions is that the repeated deployment of particular classes of hypotheses in response to particular feedback patterns (i.e., the use of particular strategic actions) will not necessarily produce an improvement in the cognitive structure that, according to the consensual view, is responsible for generating the strategic action in the first place. Thus, solvers who become adept in the use of a particular strategy do not necessarily have deep knowledge about the logical structure of Mastermind, nor do they learn much about such structures despite their apparent improvement in performance.

The issue of learning from the feedback can be studied by considering three of the sequences of operations shown in Figure 1 on an encounter-by-encounter basis. Thus, for Sequence 2 for example, the number of accurate and complex deductions made by each subject from the initial awarding of nonpopular feedback was computed and expressed as a percentage of all deductions made on that encounter with nonpopular feedback. This computation was made for each subsequent encounter of that type, for each subject, as long as a majority of the 13 subjects actually had an encounter of that type. The sequences labeled 3 and 4 in Figure 1 were treated similarly. That is, the number of accurate and complex deductions made by each subject when the subject initially operated in Sequences 3 and 4 was computed and expressed as a percentage of all deductions made on the initial operation of that sequence. Each subsequent operation of Sequences 3 and 4 was treated the same way.

Given that each encounter of each of the types of sequences should be understood as a learning trial of that sequence, then the percentage of accurate and complex deductions, relative to all deductions made on that trial, should be a function of trials (encounters) and previous deductions. That is, if the subjects were truly getting better at making logical deductions, then the percentage of all deductions that were accurate and complex should increase as a function of the subjects' previous success in making such deductions. If this were true, then the percentage of subjects making greater than the overall mean percentage of accurate and complex deductions should be low on initial encounters and should increase as a function of the number of encounters. Thus, the more successful the subject had been at making such deductions, the more successful we would expect that subject to be on any subsequent encounter.

An alternative hypothesis is that the percentage of each encounter's deductions that are accurate and complex is independent of the subject's previous history of such deductions. If observed, this pattern, which would be consistent with a Markovian position, would suggest that the subjects were not actually learning how to make logical deductions from their previous experience.

Figure 2 shows the observed percentages (based on the number of subjects who had such an encounter) of subjects who had higher than the overall mean proportion

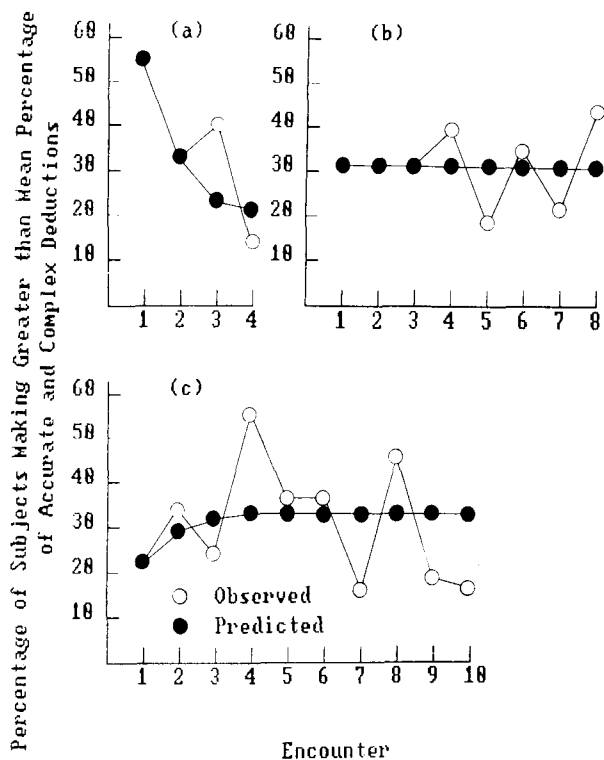


Figure 2. Percentage of subjects making greater than the mean percentage of accurate and complex deductions as a function of sequence type and encounter number. (a) Sequence type initiated by the award of popular feedback, followed by a modal hypothesis, followed by nonpopular feedback. (b) Sequence type initiated by popular feedback, followed by a modal hypothesis, followed by popular feedback. (c) Deductions made directly from nonpopular feedback. The observed percentages are shown relative to those that would be generated by a Markovian model of the subjects' underlying knowledge.

of accurate and complex deductions on that encounter for three sequences. The number of observed subjects was contrasted with the number of subjects who would be expected to have had higher than the overall mean proportion of accurate and complex deductions on that encounter, if the Markovian property were operative. For each of the three trend lines shown in Figure 2, none of the resulting chi-square analyses were significant [for Sequence 2, "immediate from nonpopular feedback," $\chi^2(9) = 5.49$; for Sequence 3, "popular feedback-modal hypothesis-popular feedback," $\chi^2(7) = 1.36$; for Sequence 4, "popular feedback-modal hypothesis-nonpopular feedback," $\chi^2(3) = 1.07$; all $ps > .05$]. The trend lines have the Markovian property: the subjects appear to have learned little or nothing about making accurate and complex deductions from the encounters in which they succeeded in making a greater than mean percentage of such deductions. Furthermore, the sequence of cognitive activity engaged in by the subjects seems to have had little influence on this lack of learning.

Properties of feedback and hypotheses. Finally, I have hypothesized that most solvers who play the game

over a several-week period maintain their use of particular strategies, because such strategies appear likely to produce an increase in the total amount of feedback given to the subjects. A count was made of the number of pins awarded each time a subject received any of the five types of popular feedback. On the 338 times such feedback was issued, the subjects received 887 feedback pins (2.62 pins/encounter with popular feedback). A similar count was made of the number of pins awarded on each issuance of nonpopular feedback. On the 218 times such feedback was issued, the subjects received 757 feedback pins (3.47 pins/encounter with nonpopular feedback). Thus, when a subject received nonpopular feedback, it typically meant that more information had been given than when popular feedback was awarded.

In addition, it appears that a modal hypothesis is more likely than a nonmodal hypothesis to produce an increase in the number of just-previously awarded feedback pins, an effect that can be seen in the following analysis. Assume that a subject who has produced a hypothesis consisting of four different colors is awarded the feedback type "1 Black, 2 White." The modal hypothesis for this feedback is "1 color, 2 locations." Under these conditions, there are 32 specific hypotheses that would be characterized as modal. Of these, nine hypotheses (0.28 of all modal hypotheses) would lead to the award of a total of four feedback pins. For the feedback in question, there are 1,264 possible nonmodal hypotheses (the total number of hypotheses, $6^4 - 32$). None of the hypotheses characterized as "2 colors," "3 colors," or "4 colors" will produce an increase in the total number of feedback pins, nor will any of the hypotheses characterized as "1 location," "2 locations," "3 locations," or "4 locations" do so. Of the eight possible "1 color" hypotheses, one will lead to an increase in the total number of feedback pins awarded. Similarly, of the eight possible "1 color, 1 location" and "1 color, 3 locations" hypotheses, one of each type will lead to an increase in the total number of feedback pins. Thus, three of the 1,264 possible nonmodal hypotheses (0.0023) will lead to an increase in the total number of feedback pins awarded. Similar properties were observed for each of the other popular feedback types.

Discussion

As in the first experiment, the subjects improved substantially during the test period, and this improvement seemed to result from the subjects' choices of hypotheses in response to certain types of feedback. Moreover, the improvement in performance was accompanied by other changes. Predominating early in the 5-week interval were the fairly easy to make consideration deductions. However, by the conclusion of the experiment, a substantial proportion of all deductions were assignments or eliminations. Furthermore, the subjects' accuracy also improved across the experimental period. Although the subjects were no more likely to make accurate deductions following modal hypotheses than following nonmodal hypotheses, they were more likely to make complex deduc-

tions in response to modal hypotheses than in response to nonmodal hypotheses.

However, the feedback given seemed to play a strong role in the subjects' ability to make deductions. When the subjects received nonpopular feedback, they made significantly more deductions than when popular feedback was given. In addition, the likelihood of a subject's making an accurate and complex deduction was four to five times greater when he or she was operating in a sequence initiated or terminated by nonpopular feedback than in a sequence initiated or terminated by popular feedback. However, subjects who made accurate and complex deductions on a given encounter were no more likely to make such deductions on their next encounter than were other subjects, regardless of the nature of the hypothesis advanced, or the feedback given in the interim. This finding suggests that the solvers learned little or nothing about the underlying structure of this logical-deduction task despite their use of an effective strategy.

Rather, there are some findings that suggest that the use of the modal hypothesis strategy was maintained by the subjects for other reasons. Modal hypotheses were more likely than nonmodal hypotheses to produce an increase in the number of feedback pins given. This finding suggests that the modal hypothesis strategy was maintained simply because it led to an increase in feedback.

GENERAL DISCUSSION

These studies confirmed at least two well-known principles in the psychology of thinking: (1) strategic actions emerge very quickly on almost any task involving human problem solving, and (2) such actions continue to show development over the interval in which they are measured, suggesting that on all but the most trivial problems, humans continue to learn about the task. In addition, however, the findings of Experiment 2 suggest that the performance of such strategic actions does not necessarily have a principled relationship with a person's underlying knowledge of the task's structure.

The conventional view of improvement in the depth of knowledge as problem solving proceeds is that such improvement is all but inevitable given a subject's motivation and continued involvement with the task (Carbonell, 1986; Larkin, McDermott, Simon, & Simon, 1980). However, the present findings suggest that this course of events may not always be observed. In *Mastermind*, it seems clear that the solver perceives some sort of connection between the pattern of feedback given and the appropriate hypothesis to be ventured. Moreover, knowledge of this relationship is acquired rather quickly. This perceived connection, rather than any deep structural knowledge, seems to drive most nonexpert *Mastermind* play. It is true that the subjects in the present study were not expert players, so it is not completely clear what might take place were they to acquire expertise with time.

Nevertheless, the present findings indicate that these subjects could be proficient at playing *Mastermind* without gaining much in the way of underlying structural knowl-

edge. This, in turn, suggests that expertise may not always consist of deriving some structural improvement from the repeated deployment of a particular strategic action. Rather than getting better at interpreting the feedback, which in a more general sense corresponds to improvements in structural knowledge, the present experiments suggest that the solvers learned something different: How to convert the existing situation, however complicated and uninterpretable, into one of which they are able to make sense.

In summary, the imposition of a single strategic action may not always be sufficient to produce a dramatic reorganization of a solver's underlying strategic knowledge. In addition, the effective strategic action that subjects acquire fairly spontaneously in Mastermind does not necessarily promote any genuine comprehension of the variables involved in logical deduction in that game.

REFERENCES

- ANZAI, Y. (1987). Doing, understanding, and learning in problem solving. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development* (pp. 55-98). Cambridge, MA: MIT Press.
- ANZAI, Y., & SIMON, H. A. (1979). The theory of learning by doing. *Psychological Review*, **86**, 124-140.
- BRUNER, J. S., GOODNOW, J. J., & AUSTIN, G. A. (1956). *A study of thinking*. New York: Norton.
- CARBONELL, J. G. (1986). Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 137-161). Los Altos, CA: Morgan Kaufman.
- COOMBS, C. H., DAWES, R. M., & TVERSKY, A. (1970). *Mathematical psychology: An elementary introduction*. Englewood Cliffs, NJ: Prentice-Hall.
- CHI, M. T. H., GLASER, R., & REES, E. (1982). Expertise in problem solving. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 1, pp. 7-76). Hillsdale, NJ: Erlbaum.
- JOHNSON, E. S. (1971). Objective identification of strategy on a selection concept learning task. *Journal of Experimental Psychology Monographs*, **90**, 167-196.
- JOHNSON, E. S. (1978). Validation of concept learning strategies. *Journal of Experimental Psychology: General*, **107**, 237-266.
- LANGLEY, P. (1985). Learning to search: From weak methods to domain-specific heuristics. *Cognitive Science*, **9**, 217-260.
- LARKIN, J. H., McDERMOTT, J., SIMON, D. P., & SIMON, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, **208**, 1335-1342.
- LAUGHLIN, P. R., LANGE, R., & ADAMOPOULOS, J. (1982). Selection strategies for "Mastermind" problems. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **8**, 475-483.
- NEWELL, A., & SIMON, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- OHLSSON, S. (1987). Truth versus appropriateness: Relating declarative to procedural knowledge. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development* (pp. 287-328). Cambridge, MA: MIT Press.
- STEGEL, S. (1956). *Nonparametric statistics for the behavioral sciences*. New York: McGraw-Hill.
- SWED, F. S., & EISENHART, C. (1943). Tables for testing randomness of grouping in a sequence of alternatives. *Annals of Mathematical Statistics*, **14**, 66-87.
- WICKENS, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco: W. H. Freeman.

(Manuscript received December 10, 1987;
revision accepted for publication May 3, 1989.)