INSTITUTE FOR THE

STUDY OF

INTELLECTUAL

BEHAVIOR

A Process Model for Hobbits-Orcs and Other River Crossing Problems

Robin Jeffries, Michael E. Atwood, Peter G. Polson and

Lydia R. Hooke

University of Colorado

Program* on Cognitive Factors in Human Learning and Memory Report. No. 54

Institute for the Study of Intellectual Behavior

University of Colorado

Boulder, Colorado 80309

August, 1976

ABSTRACT

We extend the model originally developed by Atwood and Polson (1976) for the water jug task to four versions of the Missionaries-Cannibals problem. Experiments showed that variations in the cover story produced no differences in legal moves to solution, but large differences in illegal moves. A three stage model incorporating means-ends heuristics, assumptions about the utilization of memory, and an illegal move detection process was able to account for both the legal and illegal move data from all four versions of the task.

---

*Supported in part by NSF Grant BNS 72 02084

A Process Model for Hobbits-Orcs and Other River Crossing Problems

In this paper, we develop and evaluate a process model for the Hobbits-Orcs (née Missionaries-Cannibals) river crossing problem. Two sets of issues motivate this research. Elsewhere, two of us have presented a General Problem Solver (GPS) - like model for water jug problems (Atwood and Polson, 1976). The primary purpose of this paper is to demonstrate the generality of the process assumptions underlying the water jug model by showing that another model derived from these same assumptions can account for the data from the Hobbits-Orcs problem and other similar river crossing tasks. Secondly, we will attempt to resolve some of the issues that have been raised by other studies of Hobbits-Orcs and related problems (Thomas, 1974; Greeno, 1974; Reed , Ernst, and Banerji, 1974; and Simon and Reed, 1976). We argue that a successful demonstration of the generality of the Atwood-Polson assumptions gives strong evidence for their validity.

We begin by describing in detail the Hobbits-Orcs problem and other related river crossing tasks. We then present a process model for this class of problems that is derived from the assumptions underlying the Atwood and Polson (1976) model for the water jug task. We follow with analysis of river crossing problems and discuss results and conclusions of other investigators in the context of our model. Then, we present two experiments that compare performance on different isomorphs of the Hobbits-Orcs problem. Finally, we discuss the results of comparisons between quantitative predictions obtained by simulation from the model and observed performance. Our overall objective is to demonstrate the validity of the assumptions made by Atwood and Polson by showing that it is possible to derive models from them that enable us to

account for the performance of naive subjects solving two quite different sequential problems: water jug tasks and river crossing problems.

The Task

The Hobbits-Orcs problem is one of a large number of river crossing tasks in which a collection of objects or group of travellers must be transported across a river from the left to right bank. The problematic aspects of these tasks are the limited capacity of the boat and restrictions that make illegal certain combinations of travellers on either side of the river. The basic version of the task used in this study involved moving six travellers, three Hobbits and three Orcs (Tolkien, 1937), across a river using a boat that only holds two of them. In addition, if the Orcs ever outnumber the Hobbits on either side of the river, the Orcs will kill the Hobbits. Moves that lead to such configurations are illegal. The graph of the problem space is shown in Figure 1.

Insert Figure 1 about here

A state of the problem is defined by the numbers of Hobbits and Orcs on either bank and the position of the boat. The boxes in Figure 1 are the legal states of the problem (Orcs do not outnumber Hobbits on either bank). The ovals are states in which Orcs outnumber Hobbits on one or the other bank; illegal moves are moves leading to these states. For each state, the number of Hobbits (H) and Orcs (O) and the position of the boat (*) is shown. We will refer to states by the numbers shown in Figure 1 and by triples: the number of Hobbits on the right bank, the number of Orcs on the right bank, and the position of the boat. We let the boat position equal L if the boat

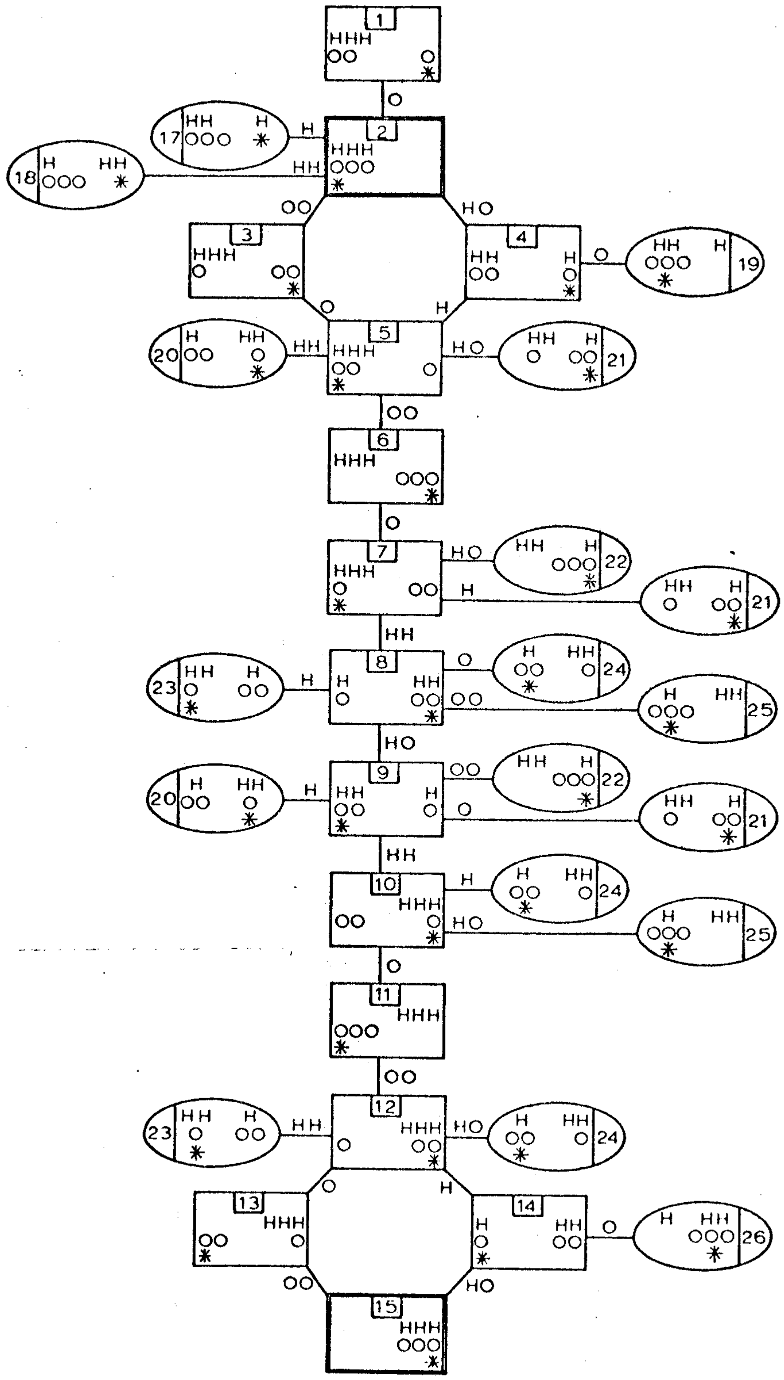

Figure 1. A graph of the problem space for the Hobbits-Orcs problem. The legal states are shown in the boxes, the illegal states in the ovals. The position of the boat is given by the *.

is on the left bank and R otherwise. Thus, (0H,00,L) is the start state and (3H,30,R) is the goal state.

The subject begins the problem in state 2:(0H,00,L) where he has a choice of three legal moves to states 3:(1H,10,R), 4:(0H,20,R) or 1:(0H,10,R). From either state 3 or 4, the subject can legally return to the start state or make a correct forward move to state 5:(0H,10,L). State 5 is a point on the graph where the subject's choices include two legal moves that don't lead toward the goal. The graph of legal moves is essentially linear from state 6:(0H,30,R) on. Only two legal moves lead from a majority of the states: the correct forward move and a return to the just previously visited state. Observe that the move from state 8:(2H,20,R) to state 9:(1H,10,L) is the only place in the problem where it is necessary to move two travellers away from the goal (right) bank back to the left bank. At state 12:(3H,10,R) there is again the choice of two forward moves.

We propose to evaluate the generality of the assumptions underlying the Atwood and Polson model by showing that we can derive a model for a problem that has a very different structure from the water jug task. A second approach to the issue of generality is to show that the model can account for the behavior of subjects who solve isomorphs or homomorphs of the original problem. An isomorph is a problem with a different cover story but exactly the same set of legal and illegal moves and transitions among them as the original problem. For example, Hobbits-Orcs is an isomorph of Missionaries-Cannibals. Hayes and Simon (1974) argue that the details of the cover story may change the augmented problem space (Newell and Simon, 1972) constructed by a subject to represent a problem and thus alter problem solving behavior. Greeno (1974)

used an isomorph of Hobbits-Orcs involving elves and men to investigate a hypothesis about difficulty at a particular point in the problem. Reed, Ernst, and Banerji (1974) employed a homomorph (a problem with the same pattern of legal moves, but with different illegal moves) in a study of transfer performance. We will show that our model can account for the similarities of subjects' problem solving behavior on four different isomorphs of Hobbits-Orcs.

The Hobbits-Orcs problem, or variations of it, has been used extensively in investigations of problem solving behavior. Thomas (1974) examined the behavior of naive subjects on this problem. Greeno (1974) attempted to describe the cognitive changes that are induced by repeated solution of the problem. Reed, et al. (1974) examined possible transfer effects of solving a homomorph on subsequent solution of this problem. All three papers concluded that subjects can be described as progressing through several "cognitive states" in the process of solving the problem. Each of these cognitive states was assumed to encompass a multiple move sequence and is characterized by application of a particular strategy. Furthermore, forward planning occurs in each of the cognitive states. Our model, to be presented in the following section, makes no such assumption. We argue that memory resource limitations (Norman and Bobrow, 1975) confine subjects to the consideration of one-step sequences.

The Model

The model for Hobbits-Orcs that we develop in this section is derived from a theory for MOVE problems originally proposed by Atwood and Polson (1976). The theory assumes that move selection involves the interaction between General Problem Solver (GPS) - like means-ends heuristics and information stored in memory about previously occupied problem states. Atwood and Polson, in their presentation of the water jug model, give a detailed rationale for the assumptions underlying the theory. We will focus on the presentation of a model for Hobbits-Orcs and related problems. The reader is referred to Atwood and Polson (1976) for a more detailed discussion of the rationale for the assumptions underlying the model.

Following Atwood and Polson (1976), we propose that a naive subject's attempts to solve MOVE problems can be characterized as a series of episodes, each involving the selection of a successor to the current problem state. Move selection involves the interaction of two sets of processes: 1) the means-ends processes, in which a subject attempts to find a successor to the current state that he evaluates as being closer to the goal state, and 2) the memory processes, which cause the subject to tend to reject moves that lead to previously occupied states. The model describes how information from the results of the means-ends evaluation and data from memory are integrated to select a next move. Before describing the complete model we will describe the means-ends processes and memory assumptions.

Means-Ends Processes. The components of the means-ends processes are: (1) the evaluation function, (2) the move acceptability criteria, and (3) the choice of an optimal or best move.

The evaluation function is a numerical representation of the subject's state evaluation strategies. We propose that the evaluation function is the major task specific component of any given model derived from the theory. The major difference between models for different tasks (water jugs vs. Hobbits-Orcs) or for different isomorphs of a problem is in the structure of the specific evaluation function. We assume that the evaluation function represents a subject's encoding of the task instructions and his decision concerning overall strategies. The evaluation function incorporated into the simulation model is numerically valued with arguments that include the number of travellers on each bank and configurations of travellers, e.g., the number of pairs of Hobbits and Orcs on the goal bank. The evaluation functions incorporated into models for Hobbits-Orcs and its isomorphs will be presented in the discussion of the simulation results.

The means-ends heuristics classify any given move as acceptable or unacceptable. Let $\underline{e}_i$ be the value of the evaluation function for the current state; let $\underline{e}_j$ be the value of the state that would result if the move under consideration were taken. We have defined the evaluation functions such that larger values are associated with more desirable states, i.e., more travellers on the right (goal) bank. Thus, with rare exceptions, $\underline{e}_i < \underline{e}_j$ for any move from left to right, a move taking one or more travellers to the goal bank. Conversely, moves to bring the boat back, right to left, all result in states with worse evaluations, $\underline{e}_i > \underline{e}_j$, since at least one traveller must be removed from the goal bank. Move evaluation involves the comparison of the values of $\underline{e}_i$ and $\underline{e}_j$. Let δ_d, d = L or R, be values that define acceptability criteria. When moving from left to right, a move is acceptable if $\underline{e}_j \geq \underline{e}_i + \delta_L$. A move from right to left is acceptable when $\underline{e}_j \geq \underline{e}_i - \delta_R$. If a move does not satisfy the appropriate criterion, we say it is unacceptable. A subject, when moving from left to right, is attempting to find a move that leads to a state with a better evaluation. When moving from right to left, δ_R defines a indifference criterion.

In the final stage of the move selection process, there are circumstances in which a subject will attempt to choose an optimal move. When moving in either direction, the optimal move is the move that leads to the state with the largest value of the evaluation function.

Memory Processes. The memory processes incorporated into our Hobbits-Orcs model are identical to those proposed by Atwood and Polson (1976) for the water jug task. We assume a multi-store model of memory (Bower, 1975). Information generated during the move selection process is stored in short-term memory (STM). Information about states entered during previous episodes is stored in long-term memory (LTM).

During the first stage of the move selection process, a subject computes and stores the following information in STM for each of the successors of the current state: (1) the move, (2) the resulting state, (3) the value of the evaluation function for the resulting state, and (4) information retrieved from LTM about previous entries into the state. Information about at most r successors can be stored in STM. It is assumed that if the current state has more than r successors, information will be lost from STM.

Representations of moves actually taken are stored in LTM. This information is retrieved using a recognition process. Upon entry into a new state, a representation of that state is stored with probability s in LTM. New states are those states that have never been entered during the course of attempting to solve the problem or whose representations were not successfully stored during previous entries into the state. A state whose representation has been stored will always be recognized. Finally, information about the start state (0H, 0O, L) is stored in LTM with probability 1.0, i.e., a primacy effect.

The Move Selection Process. The theory postulates a three stage process for move selection. In Stage I, each possible successor is evaluated using the means-ends processes. If no move is chosen in Stage I, a subject attempts to find a move that leads to a new state. The third stage is entered only if representations of all successors of the current state have been stored in LTM.

Stage I. A major difference between the Atwood and Polson (1976) model for the water jug task and our model for Hobbits-Orcs is that our model incorporates a fixed noticing order for the evaluation of successors of the current state. When the boat is on the left bank, moves are evaluated in the following order: 1) one Hobbit and one Orc (HO), 2) two Hobbits (HH), 3) two Orcs (OO), 4) one Hobbit (H), and 5) one Orc (O). We define the index, $n(x)$, of a move, x, is its position in the noticing order,

$$n(x) = i, \quad i = 1,2, \ldots, k \qquad (1)$$

where k is the number of possible successors. That is, for a state that has five successors and the boat on the left bank, $n(HO) = 1$, $n(HH) = 2$, etc. For states with fewer possible moves, $n(x)$ still increases linearly (e.g., for state 11:(3H, 0O,L), $n(OO) = 1$, $n(O) = 2$). When the boat is on the right bank, the noticing order and associated indices are inverted.

The noticing order embodies some of the subjects' general strategies for attacking this problem. Moves that result in more travellers moving to or staying on the right bank are preferred; as a sub-strategy moves that put more Hobbits on the right are preferred to those that increase the number of Orcs. The "pair" move (HO) is preferred to all other moves going left to right.

The Stage I process involves the following sequence of operations: 1) the next move in the noticing order is evaluated using the means-ends processes. Information about previous entries into the resulting state is retrieved from LTM. The move, the resulting state, its evaluation, and whether or not the state was recognized is stored in STM. 2) An unacceptable move is never taken, and the Stage I process continues with the next move on the noticing order. 3) If the move is acceptable, the probability that it will be taken depends on whether or not the resulting state is recognized. If it is not recognized, the move is taken with probability α^i, where i is the index of the move. If it is recognized, it is taken with probability β. 4) Moves returning to the immediately preceding state are not taken. 5) If a move leads to the goal state, it is taken with probability 1. 6) If the move is not taken, the next move in the noticing order is evaluated. The cycle continues until a move is taken or all possible moves have been evaluated and rejected.

Stage II. If a move is not selected during the Stage I process, a subject enters Stage II. In Stage II, successor states are considered in the order defined by the noticing order. A subject takes the first move that leads to a "new" state. If there are no new successors, the subject enters Stage III.

Stage III. In this final stage, a subject attempts to select an optimum move on the basis of information generated during the execution of Stage I. He attempts to choose a move that leads to the state with the largest value of the evaluation function. However, STM can only retain reliable information about r successors. If the current state has more than

r successors, it is assumed that an accurate record of the information about each move is no longer stored in STM. In this case, the subject randomly selects a possible move. If the current state has r or fewer successors, the subject will select the optimal move with probability α_{avg}, where $\alpha_{avg} = (\alpha^5 + \alpha^4 + \alpha^3 + \alpha^2 + \alpha)/5$. If the optimal move is not taken, the subject randomly selects a move.

Errors and Illegal Moves. The Atwood and Polson (1976) model for the water jug task has no mechanisms for dealing with errors. They were viewed as noise in the data since the only possible errors were physically impossible moves, pouring water from an empty jug or into a full one. For the Hobbits-Orcs task, however, there are two kinds of errors: nonsense moves (putting three travellers in the boat, not putting anyone in the boat, putting a traveller into the boat who is not available on that bank), and moves that let Orcs attack Hobbits. The first type are violations of the real-world constraints defined by the problem and are like attempts to pour water from an empty jug. The second type are moves that violate arbitrary constraints of the particular problem. In the following, we will refer to the first type as errors and the second type as illegal moves. We will add a mechanism to an Atwood-Polson type of model to attempt to deal with illegal moves.

The pattern of possible illegal moves is not uniform throughout the Hobbits-Orc problem. As can be seen in Figure 1, the number of possible illegal moves varies from zero to three. If illegal moves were generated by a random process, the number of illegal moves chosen at a particular state would be highly correlated with the number of illegal moves that are possible at that state. We will demonstrate, however, that subjects do not randomly select illegal moves; high frequency illegal moves are those that lead to states with high evaluations. Our illegal move process is based on two assumptions. First, the move selection processes associated with each of the three stages, discussed earlier, make no distinctions between illegal and legal moves. Second, after a given process selects a move, it is filtered through an illegal move detector that prevents most such moves from being taken. If the illegal move is detected, a subject continues the processes for the stage that generated the illegal move with the next element of the noticing order. If a subject fails to detect that a move is illegal and consequently selects it, we assume that the entire move selection process is reinitiated; i.e., subsequent processing begins with Stage I. The state defined by the illegal move is encoded as the just previous state and is therefore not considered in the subsequent Stage I process. In summary, we propose that legal and illegal moves are treated identically by the means-ends and memory processes. We assume that a move generated by any one of the three stages is evaluated by a common detection process just before it is actually made.

We have assumed that the probability of a subject rejecting an illegal move is a function of the characteristics of the resulting illegal state, i.e., when Orcs outnumber Hobbits on one side of the river. For most legal states of the problem, there are two or three moves leading to illegal states. The illegal states can be partitioned into those in which the violation (Orcs outnumbering Hobbits) occurs on the same bank as the boat rests currently (near side illegal state) and those in which the violation occurs on the opposite bank (far side illegal state). Near side illegal

states are shown on the left side of the graph in Figure 1, and the far side illegal states are shown on the right side. In addition, we assume that there are configurations where Orcs outnumber Hobbits that are easier to detect than others. An easy configuration has three Orcs and one Hobbit on either bank of the river; a hard configuration is any other combination where Orcs outnumber Hobbits (e.g., three Orcs and two Hobbits or two Orcs and one Hobbit). Combining the above two classifications, we partition the illegal states into four categories: near-easy-to-detect (NE), near-hard-to-detect (NH), far-easy (FE), and far-hard (FH). We assume that the probability of detecting an illegal state and rejecting the associated move is a function of the category of the illegal state. Let ε_d, $\underline{d}$ = NE, NH, FE, or FH, be the probability of detecting a move that leads to a state in category $\underline{d}$.

Related Research

Two other explicit process models for the Hobbits-Orcs problem have been proposed by other researchers: General Problem Solver (GPS) (Ernst and Newell, 1969), and the strategy shift model (Simon and Reed, 1976).

Greeno (1974) argues that GPS is not a viable model of human performance in sequential problem solving tasks like Hobbits-Orcs. First, he observes that GPS cannot be considered a valid simulation of human performance because it assumes internal structures that have no reasonable psychological analog; e.g., GPS has an almost unlimited STM capacity. Second, Greeno has shown that there is a very poor correspondence between GPS's performance on Hobbits-Orcs (actually Missionaries and Cannibals) and the behavior of human subjects. For a large majority of the states of the Hobbits-Orcs problem, there is an almost perfect inverse relationship between observed difficulty of various states and predicted difficulty (Greeno, 1974, Fig. 5, p. 282).

However, Greeno (1974) and Atwood (Note 1) have shown that GPS's difficulties with the Hobbits-Orcs problem are due to the procedures that it uses to set-up and manipulate its subgoals. GPS assumes that it can select effective subgoals by evaluating differences between the goal state and the state resulting from taking a possible move. This assumption is not always correct for either the water jug task or river crossing problems. Atwood and Polson (1976) show that this is also a primary source of difficulty for human subjects in the water jug task. Second, Greeno points out that the major cause of GPS's difficulties with the Hobbits-Orcs problem is the particular way that GPS sets-up and executes subgoals when a primary goal cannot be directly accomplished - the looping problem (Quinlan and Hunt, 1968; Atwood, Note 1).

The looping problem is basically a problem of goal stack management. When multiple subgoals are established, they must be executed in an order that ensures that the preconditions necessary for accomplishing subsequent goals are not destroyed when prior goals are attempted. In order to implement effective goal stack management heuristics, the programming language employed must include efficient backtracking procedures. Such procedures, however, are lacking in the IPL-V language (Newell, Tonge, Feigenbaum, Green, Kelly, and Mealy, 1964) in which GPS is implemented. Problem solving programs that were developed after GPS using more advanced programming languages have much more sophisticated heuristics for detecting and resolving conflicts among subgoals (STRIPS, Fikes, Hart, and Nilsson, 1972; NOAH, Sacerdoti, 1975). Thus, we

conclude that GPS's failure to emulate or surpass human performance on Hobbits-Orcs is due to the particular details of the implementation of the concepts underlying GPS and not due to the unreasonableness of basic ideas like means-ends analysis.

Simon and Reed (1976) have developed a model for a slightly more complex variant of the Missionaries and Cannibals problem (5 Missionaries, 5 Cannibals, and a boat that holds three travellers) that is very similar in many respects to the model proposed in this paper. Simon and Reed's basic assumption is that subjects use one of two strategies in selecting a next move: the balance strategy or the means-ends strategy. When using the balance strategy subjects attempt to keep the number of missionaries on each side of the river equal to the number of cannibals on that side. The means-ends strategy leads subjects to move the maximum possible number of travellers when going from left to right and the minimum number going right to left. Observe that either one of these strategies can be represented in our model by incorporating appropriate terms in the evaluation function. The balance strategy would be represented by a term with a high weighting for pairs, while the means-ends strategy would be a term with a high weight for the number of travellers on the right (goal) bank. Also, the particular fixed noticing order assumed by our model could be characterized as representing a mixture of the two strategies.

The strategy shift model assumes that all subjects start off using the balance strategy and on any move can shift with some probability (the strategy-shift parameter) to the means-ends strategy. At each state subjects will take the best move as defined by their current strategy with a probability specified by the use-current-strategy parameter. This latter probability increases as a function of trial number. Otherwise, subjects randomly select a next move. The anti-looping parameter is the probability that a move leading to the just visited state will be rejected; this parameter also increases over trials. The strategy shift model does not attempt to account for illegal moves, although one could incorporate an illegal move process similar to the one proposed in this paper.

Of the alternative theories of the solution processes for Hobbits-Orcs that have been discussed, the Simon and Reed strategy shift model is the most similar to the theory proposed in this paper. Both theories assume that moves are selected on the basis of local information, that subjects do not plan move sequences, and that means-ends heuristics are used in these tasks. Thus, it would seem that the two models share primary, underlying assumptions. In addition, both models reject moves that lead to previously occupied states, both attempt to select optimum moves (as defined by the current strategy), and both randomly select moves if other move selection processes do not result in the choice of a move.

Earlier, we pointed out that Thomas (1974), Greeno (1974), and Reed, et al. (1974) all concluded that subjects organize the problem into a small number of "cognitive states" that involve successful planning of forward move sequences. None of these authors develop explicit process models incorporating such an assumption. Clearly, the strategy shift model and the theory developed in this paper have no planning mechanisms, and thus the two classes of theories have very different basic assumptions concerning problem solving processes. We will attempt to demonstrate the validity of our model derived from Atwood and Polson's (1976) assumptions by showing that we can fit data from Hobbits-

Orcs and related problems, and that we can explain many of the phenomena observed by Thomas, Greeno, and Reed, et al.

METHOD

The two experiments had identical designs and very similar procedures. In both experiments different groups of subjects solved Hobbits-Orcs or one of three isomorphic problems: Elves and Men-I (Greeno, 1974), Elves and Men-II, and Silver and Gold. The second experiment was a replication of the first with slight changes in the instructions in order to increase the rates of illegal moves and errors. The primary purpose of these studies was to demonstrate the generality of Atwood and Polson's (1976) assumptions by showing that the model could fit the data from all four problems. A secondary objective was to gain a better understanding of the problem by examining the effects of very different isomorphs on performance.

Subjects

The first experiment used 100 subjects who were recruited through a newspaper advertisement and were paid $2.00 for participating in the experiment. Subjects were randomly assigned to one of the four experimental conditions; thus, 25 subjects solved each isomorph.

In the second experiment, there were 152 students from introductory psychology courses who volunteered to participate in partial fulfillment of a course requirement. They were randomly assigned to one of the four experimental conditions with 35 subjects to each isomorph. Ten subjects were not included in the analysis and were replaced because they failed to solve the problem or to make 100 legal moves within one hour. Data from two subjects were lost due to computer failures.

Problems

The structure of all three isomorphs is identical to that of the Hobbits-Orcs problem presented earlier and in Figure 1. The problem statements for the four different isomorphs are presented below:

Hobbits and Orcs. Once upon a time, in the last days of Middle Earth, three Hobbits and three Orcs set out on a journey together. They were sent by the great wizard Gandalf to find one of the lost palantiri, or oracle stones.

In the course of their journey, they come to a river. On the bank is a small rowboat. All six travellers need to cross the river but the boat will hold only two of them at a time.

The Orcs are fierce and wicked creatures, who will try to kill the Hobbits if they get the opportunity. The Hobbits are normally gentle creatures, but are very good fighters if provoked.

The Orcs know this, and will not try to attack the Hobbits unless the Orcs outnumer the Hobbits. That is, the Hobbits will be safe as long as there are at least as many Hobbits as Orcs on either side of the river.

Elves and Men-I. The same text as Hobbits and Orcs except that the word "elves" replaces the word "Hobbits" and the word "men" replaces "Orcs". The last two paragraphs are changed to read:

Although the men and elves are friendly, the elves are very nervous among big people. The elves will be so uncomfortable if they are ever outnumbered by the men, that they will simply disappear and not return.

The men know this, so both groups want to be sure that there are always at least as many elves as men on either side of the river.

Elves and Men-II. The first two paragraphs are the same as Elves and Men-I. The text continues as follows:

The crossing is complicated by the need to consider the elves' customs. The elves on the bank will either talk among themselves or mingle with the men.

If they join the men, courtesy requires that each elf pair up with exactly one man. There must never be anyone left unpaired - man or elf.

If the elves choose to talk with each other, a superstition demands that there be at least three elves in the group.

The customs allow the elves to talk among themselves even if there are men nearby.

These customs are believed so strongly by the elves that, if they are violated, the elves will disappear and not return. So whenever the boat is loaded, the travellers must ensure that those who remain on the bank can comply with the rules of pairing elves with men or of having three elves in a group.

Silver and Gold. Once upon a time, in a far away place, there lived a monk who was the guardian of the temple of the three silver and three gold magic talismans. The monk has been ordered to deliver the talismans to the dedication of a new temple. He must cross the enchanted forest of Rangimali to get from one temple to the other.

The talismans must be carried through the forest in a special box. All six of the talismans need to be taken across the forest, but the box will hold only two of them at a time.

The three silver talismans are the sacred symbols of the goddess Silverina. The three gold talismans are those of the god Goldmund.

Silverina will protect any temple which holds one or more of her talismans. She becomes very upset if, in any temple under her protection, there are more talismans dedicated to Goldmnund than to her. If the monk should accidently allow this to happen, Silverina will strike him dead. Goldmund, on the other hand, is an easygoing god, who doesn't mind if his talismans are outnumbered.

To overcome the hazards of the forest, the monk must always carry at least one talisman with him whenever he crosses the forest. Either a silver or a gold talisman or both can protect the monk, but he cannot cross the forest alone.

In the remainder of the paper, experimental conditions will be referred to by the title of the isomorph solved by subjects in that condition or an abbreviation: Hobbits and Orcs (HO), Elves and Men-I (EM-I), Elves and Men-II (EM-II), and Silver and Gold (SG).

Apparatus

The execution of these experiments was controlled by a Xerox Sigma 3 computer. The problems were presented to the subjects on a IV Phase System CRT Display Terminal. The subject responded by pressing buttons mounted on a box that was located in front of the display terminal. The three rightmost of the five buttons were labelled (left to right) SEE IT, DO IT and ERASE, respectively, for all conditions. The two buttons at the left side were labelled HOBBIT and ORC for the HO group, ELF and MAN for both EM groups, and SILVER and GOLD for the SG isomorph. Presentation of the instructions and the problem and data recording were performed by a program written in FORTRAN IV.

From one to six subjects were run concurrently under the control of the CLIPR/RBM Operating System. The procecure was subject paced, and an independent sequence of events was presented to each subject. Each pair of terminals was in a small room off a large common room.

Procedure

The subject was given general instructions about what to expect from the experiment and then taken into an experimental room. Detailed instructions concerning both the problem and the method of responding were presented on the CRT, which the subject read at his own pace, paging forward and backward by pressing the buttons labelled DO IT and ERASE, respectively. The instructions for the two groups differed only in that subjects in Experiment 1 read a single screen review (12 lines of text) of the instructions just before they began the problem. In both experiments, once a subject started the problem, he could not return to the instructions. There was a separate button on the table beside the display that enabled the subject to call the experimenter if he had any questions.

Insert Figure 2 about here

A sample display from the Hobbits and Orcs isomorph is shown in Figure 2. The Elves and Men conditions had E's instead of H's and M's instead of O's. The Silver and Gold condition had TEMPLE replacing BANK, FOREST instead of RIVER, and S's and G's in place of H's and O's. For half of the subjects the two lines representing the travellers (the H's and O's) were reversed.

The subject entered a move by pressing one of the two leftmost buttons once for each creature he wanted to load into the boat. No changes appeared on the display. When he had filled the boat to his satisfaction, he pressed the SEE IT button. The boat and the "from" bank were then changed to reflect his move choice, and the message "DO IT OR ERASE" appeared at the bottom of the screen. He then pressed DO IT if he wanted to complete the move. The boat was erased from that side of the display, reappeared on the other side, and was

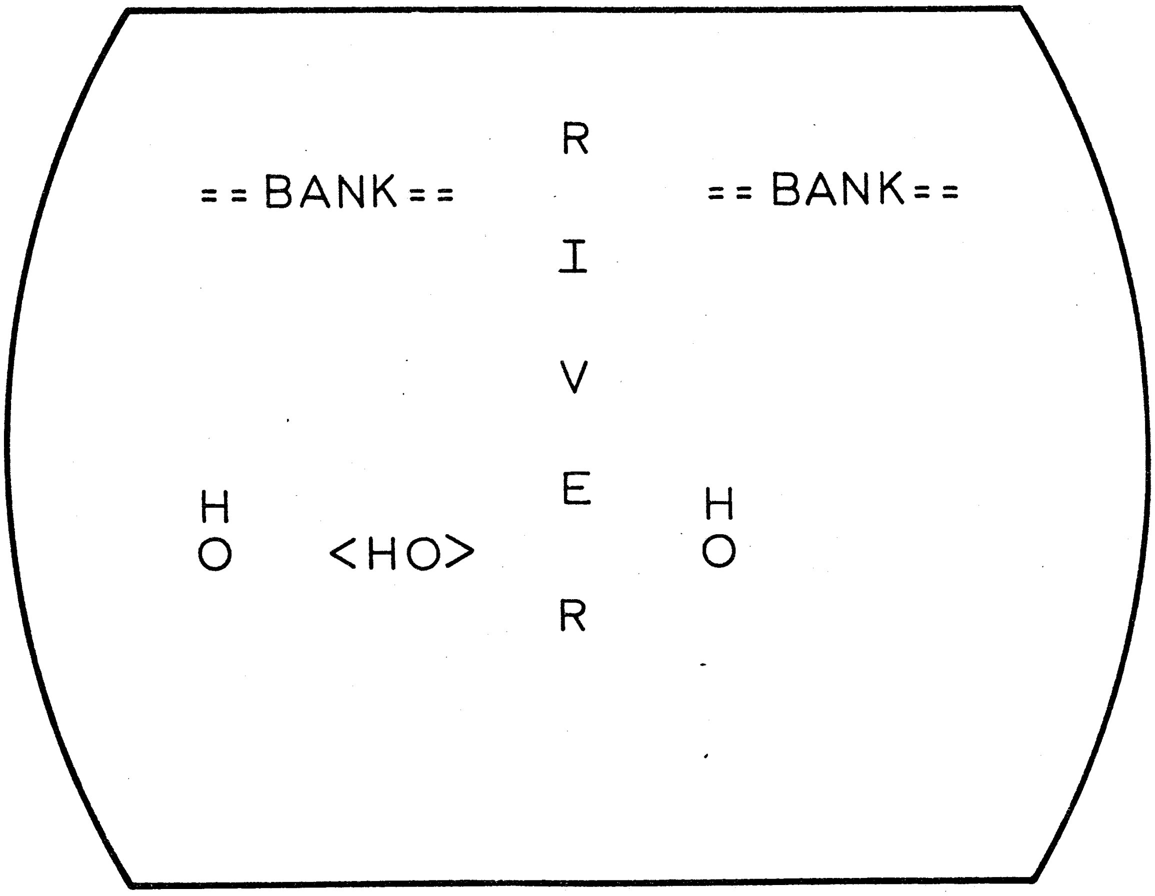

Figure 2. The display as seen by the subject at state 9 of the Hobbits-Orcs problem. The subject has loaded the boat to make the move to state 8.

emptied. This process took about five seconds. The subject could press the ERASE button at any time before he pressed DO IT. That caused the message "ERASED" to appear at the bottom of the screen and, if the boat had been loaded, caused all the creatures to be returned to the bank.

If the subject made a nonsense error (three in the boat, etc.), an appropriate message appeared on the bottom of the screen immediately after he pressed the button that resulted in the error, and that move was restarted. If he made an illegal move choice, the message did not appear until he had pressed DO IT. For the HO condition the message "THAT WAS AN ILLEGAL MOVE. THE ORCS WILL ATTACK THE HOBBITS." appeared for 3 seconds. For the other conditions the message was similar, but appropriate to the cover story. After the illegal move message appeared, the display returned to the beginning of the move.

The subject worked on the problem until he solved it, he made 100 legal moves without solving, or one hour elapsed. Only two subjects ever made 100 legal moves within the one hour time limit without solving. They were both in the SG group of Experiment 1.

RESULTS

We edited all protocols to eliminate extraneous moves that we felt were due to one aspect of our procedure. A subject had to make a sequence of button presses terminated by pressing SEE IT before a move was displayed on the screen. The purpose of this restriction was to compel a subject to select a move and then rapidly execute the response sequence to enter the chosen move. The latency data suggest that we were successful. All variations in latencies as a

function of problem state were confined to the time between the presentation of the results of the previous move and the first button press; all other inter-response times were between one and two seconds.

We eliminated from the analyses all ERASE responses when no travellers or talismans had been selected and long latency (> 15 sec) errors in which SEE IT was the first button pressed. Another class of response sequences that we eliminated were erased moves that were re-entered immediately. In such cases, only the last move in a sequence of re-entries, whether taken or erased, was included in the analysis. We also ignored the second move in a sequence if the subject made the same nonsense error twice in succession by pressing exactly the same sequence of buttons.

Because our procedure permitted a subject to see a move before he took it, he was able to ascertain whether the move was illegal before taking it. Thus, we decided, after the fact, to recode erased moves into illegal moves if the move erased was illegal and the subject had pressed the SEE IT button before the ERASE button.

Applying the transformations described above to the data from both experiments, we eliminated 66% of the erased moves and 35% of the errors. The percentages of errors and erased moves eliminated were constant across isomorphs. Transforming erased illegal moves into illegal moves increased the number of illegal moves by 14%. The analyses reported in the following section were done on both the transformed and untransformed protocols. The transformations had no effects on the overall results.

The edited data from the two experiments were combined for the purpose of analysis. Separate analyses were performed for legal moves (both forward and backward), illegal moves (moves leading to states in which Orcs outnumber Hobbits on one bank), errors (e.g., three travellers in the boat), and erased moves. The means and standard deviations are presented in Table 1. The design was a 4 X 2 factorial with the first factor being isomorphs and the second being replications. The replications factor confounds two variables: paid vs. volunteer subjects and modifications of the instructions.

Insert Table 1 about here

There were no significant main effects or interactions in the analysis of the legal move data (all F's < 1.0). For illegal moves, there were significant differences across isomorphs, $F(3,232) = 4.32$, $p < .01$, and replications, $F(1,232) = 6.56$, $p < .02$. The interaction was not significant. The same pattern was observed for errors: isomorphs, $F(3,232) = 4.25$, $p < .01$; replications, $F(1,232) = 6.42$, $p < .02$. The analysis of the erased moves showed no significant differences or interactions although the p values were approximately .10 for both main effects. Given that all interactions were small, (all p values > .20) we decided to combine the data from both experiments in all further analyses.

We next investigated whether there were different patterns of legal and illegal moves for the four isomorphs. We wanted to examine possible differences in the profiles of legal and illegal moves plotted as a function of problem state for the different isomorphs. A group (isomorphs) by repeated measures (problem states) (Grant, 1956) analysis of variance was employed in these analyses. For legal moves we reduced each subject's protocol to a 15 item vector with the elements being the number of times he entered each of the legal states of the problem. For illegal moves, a six item vector was generated for each

subject with each element being the number of illegal moves made while in states 2, 5, 7, 8, 9, and 10 of the problem (See Figure 1). These states accounted for 67% of the illegal moves. States 1, 3, 6, 11, and 13 were not included because no illegal moves are possible from these states. State 4 was eliminated because some subjects did not enter this state and thus had no chance to make illegal moves from this state. States 12 and 14 were disregarded because out of 240 subjects, only two illegal moves were made from these states. The mean number of entries per state for each of the four isomorphs is shown in Figure 3, and the mean illegal moves are shown in Figure 4.

Insert Figures 3 and 4 about here

The analysis of variance for the legal move data strongly confirms the impression given by Figure 3. There were no differences in overall performance among the four isomorphs nor were there differences in the profiles; the F's for between groups and the groups by repeated measures interaction were both less than 1. The differences as a function of problem state shown in Figure 3 are highly significant, $F(14, 3304) = 79.15$, $p < .0001$.

The illegal move analysis showed highly significant differences between isomorphs in both total number, $F(3,236) = 7.25$, $p < .001$, and profiles, $F(15,1180) = 3.31$, $p < .001$. The differences as a function of problem state are also highly significant, $F(5,1180) = 94.93$, $p < .0001$. Examination of Figure 4 suggests that the profile differences are due primarily to variations in illegal move rates at states 5:(0H,1O,L) and 8:(2H,2O,R), the two states with highest illegal move rates. As shown in Figure 4, the relative numbers of illegal moves made in these two states seem to differ widely for the various isomorphs. For HO and EM-1, many more illegal moves are made at state 5. This pattern is reversed

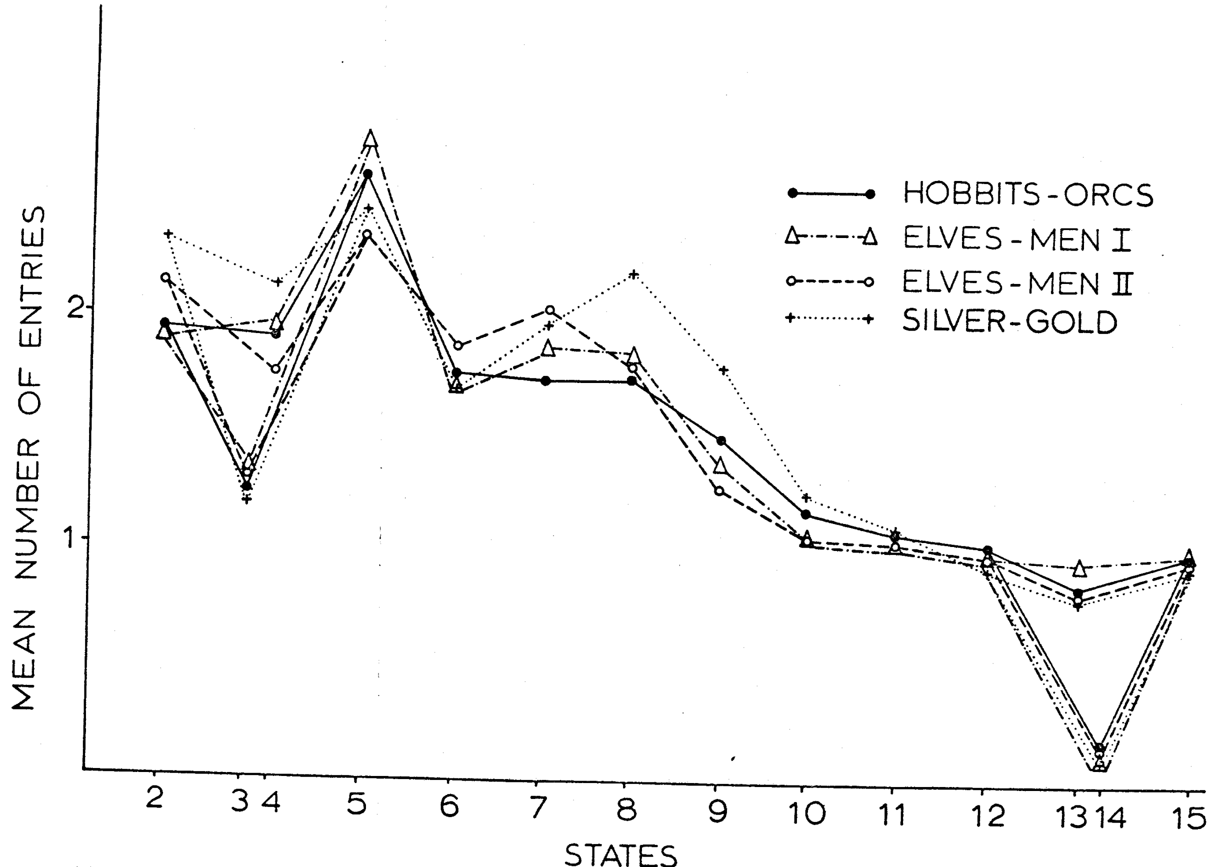

Figure 3. Mean observed visits to legal states for all four isomorphs.

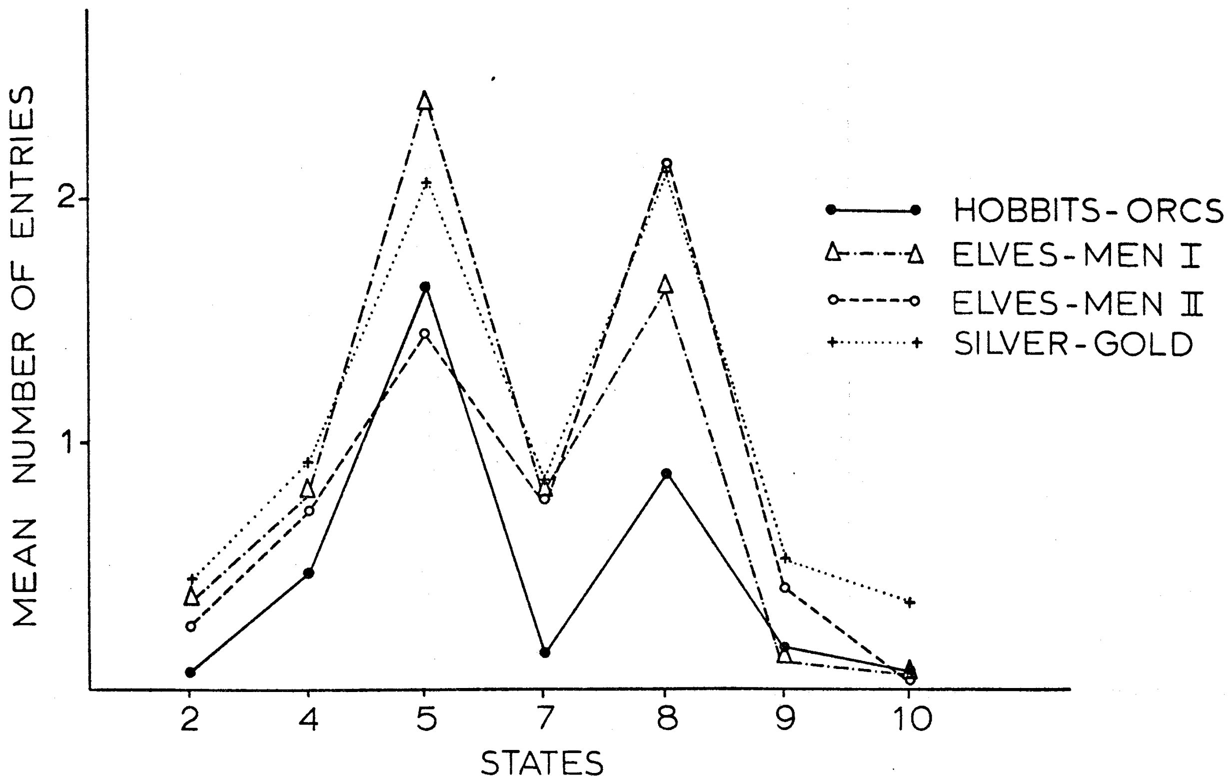

Figure 4. Mean observed moves to illegal states for all four isomorphs.

for EM-II, and approximately equal numbers of illegal moves are made in these two states on the SG isomorph.

Since the differences between the isomorphs appear to be localized in the illegal moves and other errors, we decided to investigate these more closely. In particular, we wanted to determine if a significant percentage of the errors or illegal moves were due to subjects not encoding the rules as they read them, but rather "learning by doing". In that case, we hypothesize, subjects should make many more errors early in the problem than toward the end. We divided each subject's error and illegal move data into Vincent quartiles and applied the Chi-square test of stationarity developed by Suppes and Ginsberg (1963).

For the illegal moves we found large differences across quartiles (HO: $\chi^2(3) = 52.64$; EM-I: $\chi^2(3) = 31.13$; EM-II: $\chi^2(3) = 34.80$; SG: $\chi^2(3) = 30.08$; all $\underline{p} < .001$). However, this was entirely due to a decrease in the illegal move rates for the fourth quartile. Considering only the first three quartiles, we get no differences for three of the isomorphs (HO: $\chi^2(2) = 1.65$; EM-I: $\chi^2(2) = .069$; SG: $\chi^2(2) = 2.28$). EM-I fails the stationarity test ($\chi^2(2) = 6.02$, $\underline{p} < .05$), but the pattern is not one of a decrease in illegal moves; the percentage of moves that are illegal in each of the first three quartiles for this group, are: 21%, 30%, 26%, respectively.

For the errors other than illegal moves there was evidence that subjects learned not to make these errors as they worked on the problem. Because some of the cell frequencies were too low, it was not possible to test each type of error separately, but a stationarity test was done on the three types of errors combined, for all four quartiles and for just the first three quartiles. In both cases, for all the isomporhs except HO, there was a significant decrease over

quartiles (4 quartiles - HO: $\chi^2(3) = 5.69$, $\underline{p} < .20$; EM-I: $\chi^2(3) = 26.83$; EM-II: $\chi^2(3) = 22.29$; SG: $\chi^2(3) = 33.83$; all $\underline{p} < .001$. First three quartiles - HO: $\chi^2(2) = 3.91$, $\underline{p} < .20$; EM-I: $\chi^2(2) = 15.20$, $\underline{p} < .001$; EM-II: $\chi^2(2) = 10.31$, $\underline{p} < .01$; SG: $\chi^2(2) = 28.14$, $\underline{p} < .001$). In the three isomorphs for which the decrease was significant, at least half of the errors were made in the first quartile. There appears to be a floor effect operating in the HO isomorph. Approximately the same number of errors are made in the last half of the problem by subjects in all four conditions; those in the HO group, however, made many fewer errors on the first half of the problem.

Simulation Results

The model is realized as a FORTRAN program. Each run involved the simulation of the performance of 250 subjects using fixed parameter values. A set of best fitting parameters for the Hobbits and Orcs condition was found by a coarse grid search followed by local relaxation. In fitting the other isomorphs we distinguish between two types of parameters: fixed, or common, and variable. Those parameters that are part of the means-ends and memory processes we called fixed. We assume that the values of these parameters are dictated either by the structure of the task or by the strategies the problem-solver brings to the task and should not vary across isomorphs. The fixed parameters are: α^i, the probability of taking the i^{th} new move in the noticing order in Stage I if it is acceptable; β, the probability of taking an acceptable recognized move in Stage I; $\underline{s}$, the probability of storing a representation of a state in LTM; and, $\underline{r}$, the number of moves that can be accurately retained in STM. We kept these parameters constant over all four isomorphs. The variable parameters are those that we believed are affected by the cover story. They include the illegal move detection parameters and the evaluation function. These parameters were manipulated to fit the different isomorphs. The values of all parameters are given in Table 2.

Insert Table 2 about here

The predicted means and standard deviations for legal and illegal moves for the different isomorphs are shown in parentheses in Table 1 next to the observed values. Figure 5 compares the predicted and observed results for visits to legal states; Figure 6 gives the same information for moves to illegal states. t tests were used to evaluate differences between predicted and observed means and F tests were used for the variances. The largest difference between the observed and predicted means was for the EM-I illegal moves ($t(308)=1.81$, $p > .05$). The other comparisons all produced t values of less than 1.0. The variance comparisons showed that the model's legal move variances were consistently high (all $p < .01$). We believe that this is due to our decision to discard the data of subjects who neither solved the problem nor made 100 legal moves within our one hour limit. In a given simulation run, approximately 5% of the predicted legal moves to solution were greater than 50; of our 240 subjects, this only occurred 5 times. For illegal moves, the variance comparisons yielded the following F ratios - HO: $F(60,250) = 1.14$, $p > .25$; EM-I: $F(60,250) = 4.12$, $p < .01$; EM-II: $F(60,250) = 1.78$, $p < .01$; SG: $F(60, 250) = 1.05$, $p > .25$. The distribution of number of legal moves to solution obtained from the model was compared to the data via the Kolmogorov-Smirnov two-sample test; no significant differences were observed.

Insert Figures 5 and 6 about here

We also evaluated the correlation between the predicted and observed state-to-state transitions. For every state we compared the number of times subjects made the move to each successor of that state with the model's predictions. Considering legal move transitions only, the correlations for the HO isomorph and both EM groups were .98; for the SG isomorph the value of the correlation was .97. Taking into account both legal and illegal move transitions, the correlation for the EM-I group was .94; the other three isomorphs all produced correlations of .97.

As Figures 5 and 6 show, qualitatively the overall fit is quite good. There are minor instances of misfit in the legal moves, but they vary from isomorph to isomorph, and thus could be attributed to sampling error. In fact, we conclude that the fit to the legal moves is excellent.

As Figure 6 indicates the model is able to account for the illegal move data quite well, except for one point that it consistently underestimates: the illegal moves made at state 5: (0H,1O,L). There are two possible illegal moves at this state, but the one that leads to 21:(1H,2O,R) is greatly preferred to 20:(2H,1O,R). For some isomorphs this error is made more than 1.5 times per subject. Even if we set the probability of detecting this illegal move to 0.0, the model would not make any illegal move that frequently, simply because the memory processes would cause the move to be rejected by most subjects after it had been made once. We do not have any conjectures about what causes subjects to be so seduced by this move.

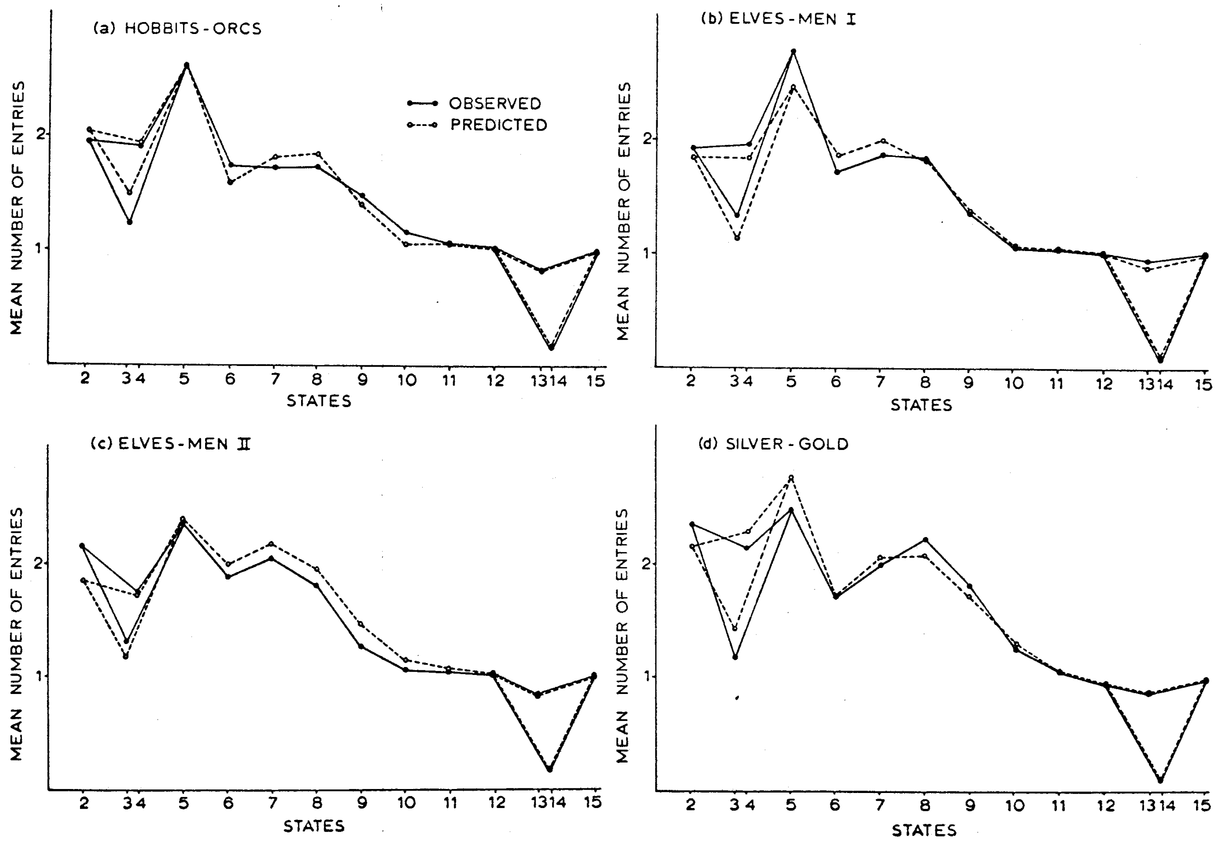

Figure 5. Observed and predicted mean visits to legal states for all four isomorphs. Predictions were computed by simulation using the parameter values given in Table 2.

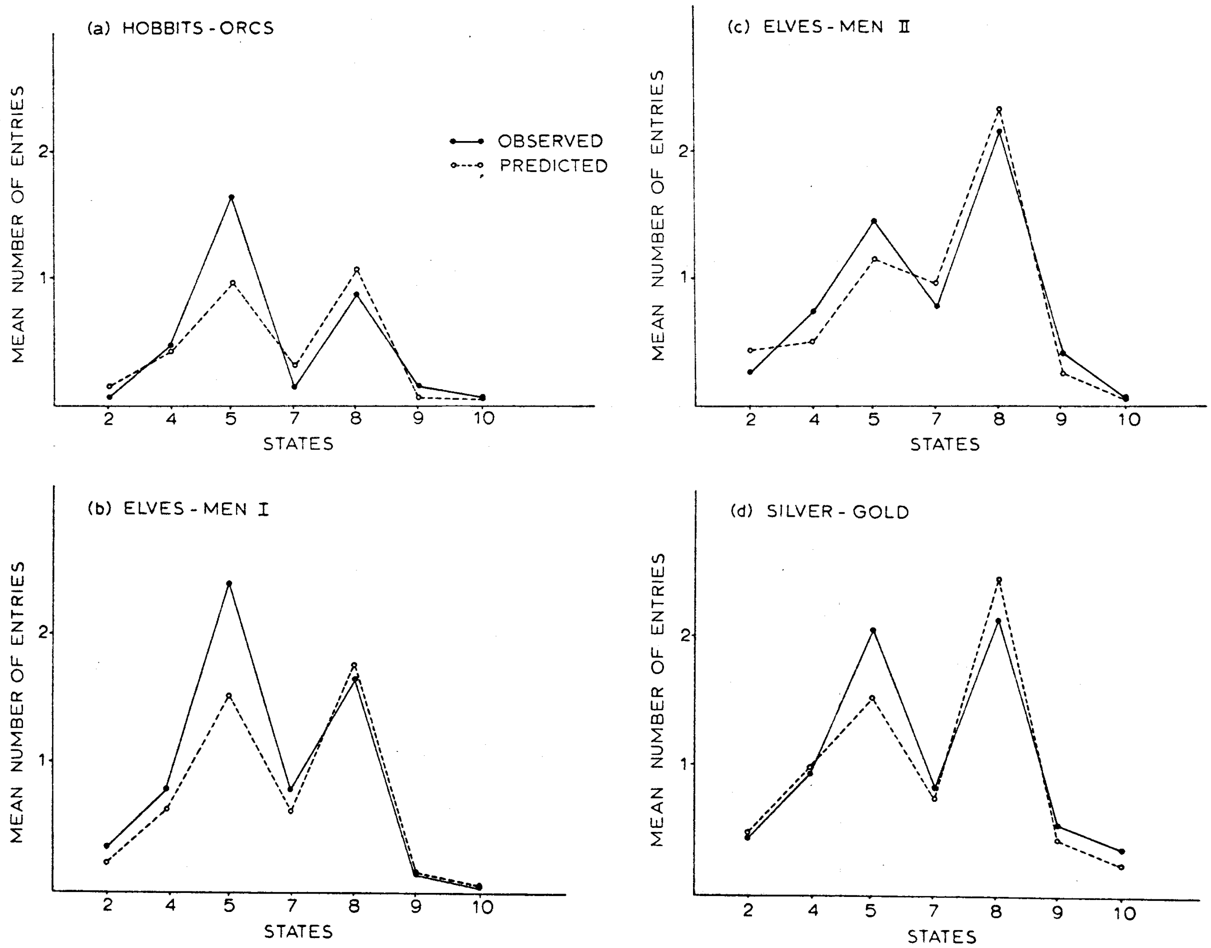

Figure 6. Observed and predicted mean moves to illegal states for all four isomorphs. Predictions were computed by simulation using the parameter values given in Table 2.

Evaluation Functions

Subjects in the various isomorphs showed very different patterns of illegal moves (See Figure 4). The mechanisms in the model that lead to distinct patterns are the evaluation function and the illegal move filter. We assume that the evaluation function results from the subjects' comprehension of the cover story. Although we use mathematical formulae to represent the subject's evaluation of considered positions, we do not believe he actually performs those calculations on the moves he analyzes, but rather that his evaluation of the relative goodness of various moves is homomorphic with the values produced by the expression incorporated into the model. The actual numbers generated by our functions are meaningless; what is important is the value of the evaluation of a state relative to that of its successors, especially in relation to the acceptability criteria. Any combination of function and acceptability criterion that produced similar patterns of acceptable and unacceptable moves would be equivalent to the particular functions we have chosen.

We found that a very wide range of evaluation functions produced similar patterns of legal moves. In fact, an acceptable fit to the legal moves can be produced by almost any function of the form:

$$\underline{e}_i = aH+bO+cP \qquad (2)$$

where H is the number of Hobbits, elves, or silver talismans on the right bank, O is the number of Orcs, men, or gold talismans on the right, P is the number of pairs on the right, and a, b, and c are weighting factors, subject to the constraints that $a,b > 0$, $c \geq 0$, and $a \geq b$. We consider the insensitivity of the model to changes in the evaluation function strong evidence for the

processes we postulate, since subjects' performance was similarly insensitive to our manipulations of the cover story.

The evaluation function used to fit the HO isomorph is:

$$\underline{e}_i = 3H+O+2P \qquad (3)$$

with the addition of a rule that made the transition from state 5: (0H,1O,L) to state 6: (0H,3O,R) unacceptable. (A constant was subtracted from the evaluation of state 6 so that

$$\underline{e}_6 < \underline{e}_5 + \delta_L.) \qquad (4)$$

This evaluation function embodies a strategy of getting Hobbits over to the right before Orcs, while trying to keep as many pairs as possible on the right bank.

The unacceptability of the transition from state 5 to 6 is related to the phenomenon Thomas (1974) and Greeno (1974) call "not trusting the Orcs." There is something about the configuration of three Orcs alone that subjects find very undesirable. We found in post-experimental interviews that some subjects considered this move as leading to a "dead-end," but they were unable to explain coherently why this was so. When this transition is acceptable, not enough visits to states 3:(0H,2O,R) and 4:(1H,1O,R) are made; with it unacceptable, subjects make too many such visits, but the fit is noticeably better. This seems to be an indication that individual differences are not as unimportant as our model has assumed them to be.

We doctored the evaluation of this one state rather than incorporate an evaluation function which had the requisite dip at that point for two reasons: 1) we felt that the unacceptability of state 6 was due to a different process than the ones which caused other moves to be unacceptable, and 2) such an evaluation function would of necessity be highly non-linear, which we felt was psychologically implausible.

We used the same evaluation function to fit EM-I as we used for HO, except that the move to (0H,3O,R) kept its original acceptable evaluation. (Perhaps Greeno and Thomas's original conjecture was correct; subjects in HO "did not trust the Orcs," and the EM-I isomorph eliminated this obstacle. However, the structure of the problem is such that this difference is only weakly reflected in the legal move data.) The pattern of illegal moves is very similar for the HO and EM-I isomorphs, except that the absolute number of errors is higher for EM-I. Again the best fit seems to be somewhere between having state 6 acceptable and unacceptable, but in this case closer to acceptable.

The evaluation function used to fit EM-II is

$$e_i = 2H+O+3P. \quad (5)$$

The cover story for this isomorph stresses the need to keep the men and elves paired; thus this evaluation weights pairs more than the previous ones do.

The SG isomorph was fit with an evaluation function of

$$e_i = 2H+2O+P \quad (6)$$

with a constant subtracted from the evaluations of state 6:(0H,3O,R) and state 11:(3H,0O,L) to make them both unacceptable. With this function silver talismans (H's) are not given preference over gold talismans (O's), and the strategy of maximizing the number of pairs is present to a lesser extent. The motivation for making state 6:(0H,3O,R) unacceptable was the same as for the HO group. The mirror image of this state in the lower half of the graph is also unacceptable to this group. A possible explanation for this fact is that

the SG group is confusing states 6 and 11. Thus, state 11 is perceived as both an old state and an undesirable state. In post-experimental interviews, subjects in the SG group tended to exhibit more confusions about which talismans had which properties than subjects in the other groups. If the subjects are less able to distinguish between silver and gold talismans than between the other types of creatures, this would account for the state confusions.

Associated with each of these evaluations functions are values for δ_L and δ_R, the acceptability criteria. For all four isomorphs, $\delta_L = 0$, i.e., a move is acceptable from any state with the boat on the left so long as it leads to a better evaluated state. δ_R depends on the range of the evaluation function. We could have kept δ_R constant by normalizing all the evaluation functions, but we felt they were easier to interpret when they only took on integer values. The evaluation function for HO and both EM isomorphs ranges from 0 to 18; for these isomorphs δ_R=5. For SG the function's range is from 0 to 15; the associated δ_R=4.

Legal Moves

There were two situations where the model had difficulty matching the legal move data. One was in visits to state 1:(0H,1O,R). Subjects almost never made the move into state 1. While this was an undesirable move for the model, it was still taken more frequently by the model than by subjects; this is because in Stage II any new move will be taken eventually, and in Stage III any move can be taken. State 1 is clearly a dead-end; the traveller that was just taken over must return with the boat. Apparently most subjects were able to look ahead at this level. Evaluation of a single successor of the currently considered state seems to be an upper limit to subjects' abilities to look ahead. Consider state 5:(0H,10,L). A subject who is capable of better look-ahead should reject the move which takes him to state 3:(0H,20,R), since its only two successors return him either to his current state or the start state. However, this transition was frequently taken. Because we did not want to incorporate a look-ahead process just to account for one transition, the model was modified to never make the move to state 1.

The second point of mismatch was the number of times subjects went from state 4:(1H,10,R) back to the start state. The phenomenon that we think accounts for this disparity is subjects' failure to encode the need for a round trip. Several subjects moved a Hobbit and an Orc on their first move, did the same on their second move, and then called the experimenter to complain that "the computer wasn't working right", because it had taken their travellers from the right bank instead of the left. We assume this move sequence occurred many more times than subjects reported it. We incorporated this phenomenon by arbitrarily having the model make the move from 4:(1H,10,R) to 2:(0H,00,L) 10% of the time it was possible, before even considering any other moves.

Illegal Moves

The values of the illegal move detection parameters used to fit each isomorph are shown in Table 2. We believe that the failure to detect an illegal move occurs when the subject exceeds his resource limits. In such circumstances the process that determines whether a move results in an illegal state does not consistently return the correct answer; in fact, the

test may be skipped entirely if the subject is operating right at his resource limit. The amount of resources required to choose moves, evaluate them, and decide if they are legal should be related to how difficult the cover story is to encode and how complex a representation it produces. The means for errors and illegal moves seem to indicate that the order of difficulty for the four isomorphs is: HO < EM-I $\leq$ EM-II < SG. The detection parameters generally reproduce this ordering. The one exception is EM-II, which has a detection parameter pattern different from any of the other isomorphs. Recall that this isomorph is the only one for which illegal states are not defined as occurring when one group outnumbers the other. A move in EM-II is illegal if either group outnumbers the other, except when all three elves are together. The distinction between near side and far side violations does not apply in this version, as any illegal move will violate the rule on both sides of the river. Similarly, the distinction between easy and hard to detect illegal moves does not apply. As a result, in EM-II, we have effectively only one error parameter.

The essential point of the illegal move process is that illegal moves are considered in exactly the same fashion as legal moves. This implies that frequently taken illegal moves should be both acceptable and high in the noticing order. The data show exactly this pattern. While we are unable to account for the actual frequency with which illegal moves are made at state 5, the model does predict that a large number of illegal moves will be made from this state. The illegal move detection parameters describe the outcome of the legality checking process. Certain tests will give erroneous results more frequently than others; moves that fail those tests are more likely to be taken.

DISCUSSION

One striking result of this study was the extreme consistency of subjects' patterns of legal moves, even with problems having very different cover stories. In addition, our legal move means and patterns are very similar to those obtained by others in their experiments with this problem (Thomas, 1974; Greeno, 1974; Reed et al., 1974; Reed and Abramson, 1976). Perhaps even more surprising was the fact that the model produced essentially this same pattern of legal moves over a wide range of parameters and evaluation functions. We consider this a strong test of the model. In particular, our assumption that the illegal move detection process interacts only minimally with the rest of the move selection process is borne out; we obtained very different illegal move patterns for the four isomorphs, but this had almost no effect on the pattern of legal moves.

The invariance of legal move performance was at first puzzling. We do not believe that the four cover stories caused all the subjects to infer the same problem representation; the stories are too dissimilar. Simon and Hayes (1976) have shown that different cover stories can produce different representations that lead to large variations in performance. However, in this problem, performance is limited by the structure of the graph. Subjects' progress through the problem is memory-driven. If a subject tries to avoid backtracking, i.e., not take the move that returns him to where he just came from, often his only remaining legal move choice is the move that takes him closer to the goal. The model uses just such a strategy: in Stage I, the move to the just visited state is never taken, and a move to any other "old" state is taken only with a low probability (β). In Stage II the subject is actively seeking new moves;

he will only enter the third stage if all the successors of the current state have been visited before. We also found in fitting the model that the most sensitive parameters were s, the probability of storing a move in LTM, and β.

Other investigations of Hobbits-Orcs (Thomas, 1974; Greeno, 1974; Reed and Abramson, 1976) and its homomorph, Jealous Husbands, (Reed et al., 1974) have produced a pattern of results similar to ours: a mean of 18-20 legal moves to solution for naive subjects, states 5:(0H,1O,L) and 8:(2H,2O,R) as the hardest states, and the most illegal moves made from those two states. However, these authors conclude that subjects attempting to solve this problem progress through a series of "cognitive states", each of which encompasses several problem states. Thomas is led to postulate such a construct because the pattern of moves from certain states does not have the Markov property, i.e., the probability of choosing a particular successor is not independent of the number of times the state has been previously visited. Thomas points out that the non-Markovian nature of the move choices indicates that memory for previous attempts enters into the consideration of a move; our model incorporates just such an assumption. He was led by this result to conclude that the move selection process had to involve consideration of multi-move sequences.

Greeno (1974) postulates look-ahead to account for the difficulty of states 5 and 8. If these states are the final moves of a sequence of moves, then subjects begin planning a new sequence at these states and are likely to make many errors. Within a pre-planned sequence, fewer errors should be made. We are able to account for the difficulty of states 5 and 8 without assuming multi-move look-ahead. The difficulty of these states is as much due to the structure of the problem graph as it is due to the strategies of the problem solver. State 8 is the only state that requires the subject to violate a general means-ends strategy; he must take two creatures back to his starting bank. This causes state 9: (1H,1O,L) to be evaluated as unacceptable in most evaluation functions of the form of (2), since there are two fewer travellers and one less pair on the right bank.

State 5:(0H,1O,L) is difficult because it is the only state from which there are two moves other than a return to the just previously visited state, but only one of these leads closer to the goal. The model shows that this structural property can account for most of the difficulty at this state. In two isomorphs we also added the constraint that the correct forward move be unacceptable. The fact that the state 5 was one of the most difficult states of the problem is further evidence that the structure of the graph is the main determinant of performance on this problem.

The strategy-shift model proposed by Simon and Reed (1976) for the 5 Missionaries-Cannibals problem shares many basic assumptions with the model developed in this paper. Unfortunately, Hobbits-Orcs is not an adequate problem for differentiating between the two theories. In the majority of cases, both the balance strategy and the means-ends strategy would choose the same move. Moreover, as we have shown above, behavior on this problem is memory driven, and the two models have similar memory assumptions.

It seems relevant to point out that what Simon and Reed call the "balance strategy" is highly confounded with a strategy of "get Hobbits over first" in both the 5 Missionaries-Cannibals problem and Hobbits-Orcs. The only legal states that include Hobbits on the right bank are "paired" states until all the Hobbits are on the right. Once the Hobbits are all across, a stragegy of "now

get the Orcs over would be equivalent in its move choices to the means-ends strategy. A model using this single strategy (get Hobbits over first, then get Orcs over) throughout the problem would produce move patterns similar to the strategy shift model. Our model incorporates attributes of the "balance", the "Hobbits first", and the "means-ends" strategies. For example, the evaluation function contains terms for all three. Similarly the noticing order operates to first consider moves that keep the largest number of travellers on the right bank, but also perfers HO to HH.

We consider the model presented in this paper to be a non-trivial generalization of the Atwood-Polson (1976) model for water jug tasks. The two tasks are different in many ways. In water jugs the unacceptability of the correct move is a major source of difficulty. In this problem, performance is much more memory driven. The evaluation function serves a very different purpose in the two tasks. In the water jug problem subjects used the evaluation function to help them find a "better" state. In Hobbits-Orcs for half the moves (i.e., when the boat is on the right bank), all successors lead to "worse" states. The evaluation function is used to choose states that are not "too bad". For the other half of the moves, any choice will improve the current position; for those the evaluation function does little to help the subject eliminate alternatives.

Our initial conjecture about this problem was that its structure might encourage subjects to use very different processes than the Atwood-Polson model assumes. The fact that the model agrees so well with observed behavior contradicts this hypothesis and gives us more evidence that the processes we postulate in this model represent the general strategies subjects use when confronted with a completely unfamiliar problem.

There are three major differences between the original Atwood-Polson model and the model presented here. The first is the noticing order. The water jug model assumed moves were considered in a random order. Our fixed noticing order reflects the fact that it is possible to assign a value to moves as well as to states in this problem. The subject prefers to try the better moves first. By establishing a fixed order for considering moves, a person is able to minimize one aspect of his memory load. In the water jug task, subjects would also lessen their memory requirements by using a fixed noticing order, but there is no one order that is suggested by the problem description or by the means-ends strategy. Noticing orders for this problem are presumably idiosyncratic and were approximated by using a random noticing order. Interacting with the fixed noticing order is the idea that α^i, the probability of taking the i^{th} move if it is acceptable and unrecognized, is greater for moves higher in the noticing order.

The second difference between the two models is the flavor of the evaluation function. The function used for water jug tasks was the sum of the absolute differences between the current jugs' contents and the desired (goal) jugs' contents. We tend to believe that subjects mentally performed operations similar to this in constructing their evaluations of states. The evaluation functions used for Hobbits-Orcs are much more artificial. We doubt that subjects multiply the number of each type of configuration on the right by a constant and sum the resulting products. Rather we assume the weights represent the relative importance to the subject of strategies such as "get the Hobbits across", "get the Orcs across", and "keep the Hobbits and Orcs paired". We have simply condensed a complex set of perceptual strategies into a mathematical function.

Thirdly, the Hobbits-Orcs model tries to account for subjects' illegal moves. Illegal moves are made much more frequently in the Hobbits-Orcs protocols than they are in the water jug problems; in some isomorphs, as many as one move in four was illegal. In fact, we consider the illegal moves of the water jug task to be more like the nonsense errors in Hobbits-Orcs. They may well represent wrong button presses, desparation moves, or testing the computer. There is no mechanism in the model that accounts for this apparently random behavior. Illegal move patterns in Hobbits-Orcs, on the other hand, are far from random. The model evaluates and chooses illegal moves in exactly the same manner as legal moves. All moves are examined by an illegal move detection process before they are taken. We readily admit that we do not have a process model for this illegal move filter. The model describes the output of this process - it is more likely to return the correct result if the violation occurs on the same bank as the boat rests or if the disparity between Hobbits and Orcs is large.

The three differences between the models for Hobbits-Orcs and for water jugs; the evaluation function, the noticing order, and the illegal move detection filter, are related to the evaluation of moves and states. We believe that all these processes are actually components of a common framework, which we call the evaluation structure. This structure is induced by the subject when he reads the cover story. This task specific information combines with the means-ends and memory processes to make up the move selection procedure.

We feel that the results presented in the preceding section justify our claim that our theory provides an excellent account of subjects' performance in Hobbits-Orcs and related problems, and that we have been able to give a much more rigorous explanation of the phenomena described by other investigators. Our successes in a radically different task environment strongly support our assertion that the assumptions proposed by Atwood and Polson do in fact describe the processes used by subjects to solve MOVE problems.

Simon and Hayes (1976) report that the details of the cover story (isomorph) in the Tower of Hanoi problem have very large effects on performance, and they concluded that the form of an isomorph directly determines the subject's internal representation of the problem. We found for Hobbits-Orcs that the form of the cover story had no effect on the number and pattern of legal moves, but only changed the rates and patterns of illegal moves. We were able to describe the illegal moves for each isomorph by changes in the evaluation function and the parameters of the illegal move filter. We argued that manipulations of the structure of the cover story for the Hobbits-Orcs problem have small effects on performance because of the structure of the problem graph, and that subjects seem to primarily use the "choose new state" strategy in Stage II to solve the problem. We conclude that subjects use the same process to solve the various isomorphs, but that changes in cover story modify two components of the evaluation structure.

Reference Note

1. Atwood, M. E. A theoretical analysis of behavior in a sequential problem solving task. Unpublished doctoral dissertation. University of Colorado, 1976.

REFERENCES

Atwood, M.E., and Polson, P. G. A process model for water jug problems. Cognitive Psychology, 1976, 8, 191-216.

Bower, G.H. Cognitive psychology: an introduction. In W.K. Estes (Ed.) Handbook of learning and congitive processes. Volume I. Hillsdale, N.J.:Erlbaum, 1975.

Ernst, G.W., and Newell, A. GPS: a case study in generality and problem solving. New York: Academic Press, 1969.

Fikes, R., Hart, P., and Nilsson, N. Learning and executing generalized robot plans. Artificial Intelligence, 1972, 3, 251-288.

Grant, D.A. Analysis of variance tests in the analysis and comparison of curves. Psychological Bulletin, 1956, 53, 141-154.

Greeno, J.G. Hobbits and orcs: acquisition of a sequential concept. Cognitive Psychology, 1974, 6, 270-292.

Hayes, J.R., and Simon, H.A. Understanding written problem instructions. In L. W. Gregg (Ed.) Knowledge and cognition. Potomac, Md.: Erlbaum, 1974.

Newell, A., and Simon, H. A. Human problem solving. Englewood Cliffs, N. J.: Prentice-Hall, 1972.

Newell, A., Tonge, F. M., Feigenbaum, E. A., Green, B. F., Kelly, H. A., and Mealy, G. Information processing language V manual (second edition). Englewood Cliffs, N. J.: Prentice-Hall, 1964.

Norman, D. A., and Bobrow, D. G. On data-limited and resource-limited processes. Cognitive Psychology, 1975, 7, 44-64.

Quinlan, J. R., and Hunt, E. B. The FORTRAN deductive system. Behavioral Science, 1969, 14, 74-79.

References Continued

Reed, S. K., and Abramson, A. Effect of the problem space on subgoal facilitation. Journal of Educational Psychology, 1976, 68, 243-246.

Reed, S. K., Ernst, G. W., and Banerji, R. The role of analogy in transfer between similar problem states. Cognitive Psychology, 1974, 6, 436-450.

Sacerdoti, E. D. A structure for plans and behavior. Technical Note 109, Stanford Research Institute, Menlo Park, California, August, 1975.

Simon, H. A., and Hayes, J. R. The understanding process: problem isomorphs. Cognitive Psychology, 1976, 8, 165-190.

Simon, H. A., and Reed, S. K. Modelling strategy shifts in a problem-solving task. Cognitive Psychology, 1976, 8, 86-97.

Suppes, P., and Ginsberg, R. A fundamental property of all-or-none models, binomial distribution of responses prior to conditioning with application to concept formation in children. Psychological Review, 1963, 70, 139-161.

Thomas, J. C., Jr. An analysis of behavior in the hobbits-orcs problem. Cognitive Psychology, 1974, 6, 257-269.

Tolkien, J. R. R. The hobbit. New York: MacMillan, 1937.

TABLE 1

Observed and Predicated Means and Standard Deviations for Legal and Illegal Moves and Observed Means and Standard Deviations for Erased Moves and Errors for the Combined Data

	Isomorph							
	HO		EM-I		EM-II		SG	
Legal Moves								
Mean	18.63	(18.84)[a]	18.97	(18.38)[a]	18.67	(18.88)[a]	20.27	(20.46)[a]
Standard Deviation	10.61	(14.80)[a]	11.05	(13.85)[a]	9.35	(13.42)[a]	11.77	(17.27)[a]
Illegal Moves								
Mean	3.37	(3.01)[a]	6.20	(4.44)[a]	5.83	(5.70)[a]	7.35	(6.82)[a]
Standard Deviation	2.78	(2.62)[a]	7.26	(3.60)[a]	5.98	(4.51)[a]	6.40	(6.28)[a]
Erased Moves								
Mean	0.78		1.08		0.97		1.48	
Standard Deviation	1.32		1.43		1.28		2.17	
Errors								
Mean	0.47		1.03		1.02		1.65	
Standard Deivation	0.77		1.68		1.72		2.69	

[a]Predictions computed by simulation. Parameter values are given in Table 2.

TABLE 2

Parameter Values Used in the Simulation Program

Common Parameters

α = .70, β = .15, _s_ = .85, _r_ = 2, _N_ = 250

Illegal Move Detection Parameters

Isomorph	ϵ_{NE}	ϵ_{NH}	ϵ_{FE}	ϵ_{FH}
HO	.95	.95	.90	.75
EM-I	.85	.85	.85	.60
EM-II	.70	.70	.70	.70
SG	.80	.80	.80	.50

Evaluation Functions

Isomorph		
HO	3H+O+2P*	H = Hobbits on right bank, O = Orcs on right, P = pairs on right.
EM-I	3H+O+2P	H = elves on right, O = men on right, P = pairs on right
EM-II	2H+O+3P	H = elves on right, O = men on right, P = pairs on right
SG	2H+2O+P**	H = silver talismans on right, O = gold talismans on right, P = pairs on right

* with state 6 made unacceptable

** with states 6 and 11 made unacceptable