

INSTITUTE FOR THE
STUDY OF
INTELLECTUAL
BEHAVIOR

Further Explorations with a Process Model
for Water Jug Problems

Michael E. Atwood & Peter G. Polson

University of Colorado

Program* on Cognitive Factors in Human Learning and Memory Report No. 49

Institute for the Study of Intellectual Behavior

University of Colorado

Boulder, Colorado 80302

January, 1976

ABSTRACT

An experiment was performed to test the three stage process model for water jug problems developed in Atwood and Polson (1975). A principal assumption of this model is that move selection is a memory data-limited process. The experiment compared the performance of a control group which was required to mentally calculate and select the desired successor of the current problem state with a "move availability" group which was explicitly presented with the appropriate successors. As predicted by the model, no between group differences were observed.

*Supported by NSF Grant GB 34077X

Further Explorations with a Process Model¹
for Water Jug Problems

Michael E. Atwood & Peter G. Polson
University of Colorado

1. This research was supported by NSF Grant GB-34077X to the Institute for the Study of Intellectual Behavior, University of Colorado. This paper is a revised version of a paper read at the Psychonomic Society Meetings, November 8, 1975, Denver, Colorado.
-

In a previous paper (Atwood and Polson, 1975), we presented a model for the water jug task. Our model is a synthesis of the means-ends analysis mechanism from General Problem Solver (Ernst and Newell, 1969) with a simplified version of Atkinson and Shiffrin's (1968) buffer model of memory. It represents a specific example of an approach that has been proposed by Greeno (1973) and Newell and Simon (1972). The purpose of this paper is to provide additional support for the assumptions of our model. In particular, we want to show that alternative models assuming more complex move selection processes can not provide tenable descriptions of performance in the water jug task. We will first describe the task and present a brief outline of our model. We will then describe the experiment and its results.

The Task

The water jug problems used in this study are similar to those employed by Luchins (1942) in his experiments on rigid behavior. Our task employed three water jugs (A,B,C) of varying sizes; e.g. (8,5,3) where the A jug has a capacity of 8 units, the B jug 5 units and the C jug 3 units.

Initially, the largest jug is full and the two smaller jugs are empty. The subject's task is to determine a series of moves which would divide the initial contents of the largest jug evenly between the two larger jugs. Water is transferred until the jug the subject is pouring from is emptied or the jug being poured into is filled. The jugs are not graduated and water could not be added or deleted during the course of solving the problem.

The problem graph and legal transitions for a typical problem are shown in Figure 1. Let states S, T, R, and L in this figure be denoted as "early" states. From each of the remaining states, four transitions

are possible: one reversible backward move, two irreversible backward moves into the "early" states, and one reversible forward move. Observe that there are two unique solution paths and all problem states lie on one of these paths with the exception of state T, which connects the two paths.

Various aspects of the experimental problems are presented in Table 1. This table indicates the solution path lengths, the number of violations of the means-ends move selection heuristic, the number of "loops" (irreversible moves) to states S, T, R, and L; and the number of loops which are classified as acceptable by our means-ends heuristic.

These tasks seem well-suited for this study for several reasons. The problem space, as defined by the problem states and the rules for making transitions from these states, is small enough to be completely specified by the theorist and easily manageable for both the subject and theorist. Like problems used in concept formation studies, they permit a complete, concise description of the task, but allow the creation of a variety of alternative problems by experimental manipulation within one domain. For example, the number of solution paths, the length of the solution paths, the number of violations of the move selection heuristic, etc., can easily be manipulated.

Study of the structure of the water jug tasks shown in Figure 1 shows that there is a large number of heuristics that would make solution of this type of problem a trivial task. In particular, any sort of forward planning process would enable a subject to avoid irreversible moves. Slightly more complex planning processes involving looking two or three moves ahead would enable a subject to overcome the difficulties at those transitions which violate the move selection heuristic. However, our earlier studies have shown that water jug problems are difficult tasks for naive subjects.

The Model

The model to be presented in this section is based on three general assumptions. They are consistent with Newell and Simon's (1972) discussion of the interactions between memory limitations and problem solving strategies, as presented in the last chapter of their book, Human Problem Solving. The model assumes that subjects use a means-ends strategy when selecting moves, that information is stored in long term memory about the identity of each state entered during attempts to solve the problem, and that subjects tend to reject moves that lead to previously occupied states.

When the subject enters a problem state, he attempts to find a next move which he evaluates as being closer to the goal state (on the basis of a means-ends heuristic) by successively considering all successors to the current problem state. Due to short term memory limitations, we assume that the subject is not planning sequences of forward moves. Solution of the problem can be broken down into a series of episodes, each episode involving the selection of a successor to the current problem state. That is, our model assumes that a subject uses only information about the current state and its immediate successors in selecting a move.

The means-ends heuristic evaluates a potential move by computing a discrepancy measure between the resultant state and the goal state. The evaluation function used is the sum of the absolute differences between the current contents of the two larger jugs and the desired (goal) contents. An earlier study (Atwood and Polson, 1975), indicated that the evaluation function should include not only the means-ends heuristic, but also information concerning previous visits to the considered state. State evaluations are therefore altered as a function

of this frequency information. Such an assumption does not change the basic processes assumed by the model.

We assume that a subject uses the above evaluation function to classify moves as either acceptable or unacceptable. A subject compares the evaluation of the current state with the evaluation of the state resulting from the considered move; he is indifferent to small differences in the values of the evaluation functions and will ignore any difference less than or equal to some criterion. A potential move is classified as acceptable if the difference between the current state's evaluation and that of the considered state is less than or equal to this criterion. As a result of this move evaluation process, the relative values of the evaluation functions for the potential moves do not directly determine the likelihood that a subject will select a given move; potential moves are classified as either acceptable or unacceptable.

Due to short term memory limitations, we argue that a subject is not planning a sequence of moves. We assume that move generation and selection is a memory data-limited process (Norman and Bobrow, 1975). The solution of a problem involves a series of episodes, each episode being the selection of a successor to the current problem state. Our theory assumes a three stage process for move selection. The first stage is the explicit evaluation of each legal move using the means-ends heuristic and information stored in long term memory. If no move is selected during the first stage, the subject tries to find a move that will lead to a previously unvisited problem state (this is the second stage). The third stage is entered when a subject decides that he has been at all successors of the current problem state.

During Stage I, the subject randomly selects and evaluates each move. This process continues until the subject selects a move or

rejects all moves. A move that would cause the subject to return to the just previously visited problem state is never chosen in this stage.

Each move is evaluated using the means-ends heuristic discussed earlier. If the move is acceptable, the subject checks to see if the resultant state has been stored in long term memory as a previously visited state. If it is not so recognized, it is taken with a fairly high probability (α), if it is recognized, it is taken with a fairly low probability (β). As each move is considered, the resulting state, whether or not the resulting state is recognized, and the evaluation are stored in long term memory.

A subject may reject all possible moves during the first stage of the move selection procedure. During Stage II, a subject attempts to find a successor state that has not been entered during a previous episode. A subject randomly selects a successor if there is more than one such new state. A subject enters Stage III only if there are no new successors.

During Stage III a subject attempts to select an optimal move on the basis of information generated during the Stage I process. However, short term memory can only retain reliable information about r moves. If the current state has more than r successors, it is assumed that an accurate record of this information is no longer available and the subject moves randomly.

A subject decides whether or not he has been in a state during a previous episode using information stored in long term memory. It is assumed that a recognition process is employed to check a potential move. This assumption has been incorporated through an all-or-none model of memory. On entry into a state, the subject stores a representation of that state in long term memory with probability g . A

representation, once stored, remains in long term memory. Thus, "new" states are those whose representations have not been stored in long term memory or that have not been visited during previous episodes. A simplified flow diagram of the simulation model is shown in Figure 2.

The Stage I process defines the primary interaction between the means-ends heuristic and the initial strategy for using information stored in long term memory. The process is a satisficing strategy and it enables the subject to reduce his reliance on his limited short term memory capacity. A major source of difficulty in a MOVE problem is the number of unacceptable moves on the solution path; i.e. the number of moves that violate the means-ends heuristic. As was pointed out by the Gestalt psychologists (e.g. Luchins, 1942), required divergences from a subject's initial strategy are a fundamental source of difficulty. Stage II is a mechanism which enables subjects to handle this detour problem in MOVE problems. It is a strategy which makes use only of the information stored in long term memory.

Introduction

A primary assumption of this model is that the move selection process is memory data-limited. Short term memory limitations prevent the subject from consistently selecting optimal moves. The results reported in Atwood and Polson (1975) indicate that the processes assumed by the model are capable of explaining human behavior in this situation. Another possible interpretation of our earlier results is that a subject is attempting to use more complex strategies but fails to employ them consistently because of memory and resource limitations. We required a subject to mentally calculate the consequences of each possible move as well as select a next move. We assume that move calculation uses a significant portion of the subject's total resources.

In the study to be described, a subject in the move availability condition was given each possible move and its consequences. We assume that providing this information would dramatically reduce the resources necessary to select a move. The present experiment was designed to test whether reducing the subject's processing load would enable the subject to utilize more efficient problem solving processes (e.g. forward planning or strategies requiring global information about the structure of the problem graph). We should observe dramatic improvements in performance in the move availability condition if a subject is able to consistently employ more complex strategies. The model presented in this paper predicts little or no facilitation of performance.

METHOD

There were two experimental conditions. In the control condition, subjects were kept informed of their current problem state and required to mentally calculate and select the desired successor of the current problem state. In the move availability condition, an external memory display was added to the display used in the control condition. This additional display presented all legal successors to the current problem state and the move necessary to reach each successor. In this condition the subject's task was to select the desired successor state and input the corresponding move. Such a display considerably reduces the processing load required in move selection. Sample displays for each condition are given in Figure 3.

Three experimental problems were used. The graphs of all problems used in this study have the same basic form as the problem shown in Figure 1. The problems differed in the length of the two solution paths, the evaluations of each of the possible states, and the number of violations of the means-ends move selection heuristic. These



problems are summarized in Table 1.

Subjects

Subjects were recruited through a newspaper advertisement and were paid \$2.00 for participating in the experiment. Subjects were randomly assigned to experimental conditions and problems in order of their arrival to the laboratory. Forty subjects were assigned to each condition and problem, with the exception of Problems 2 and 3 of the move availability condition which involved 41 subjects. Six subjects were lost due to apparatus failures.

Apparatus

The execution of these experiments was controlled by a Xerox Sigma 3 computer. The problems were presented to the subject on a IV Phase System CRT Display Terminal. The subject responded by pressing buttons mounted in a row 6 cm. from the far edge of a 19 cm. x 31 cm. wooden box that was located in front of the display terminal. Three buttons were labelled "A", "B", and "C"; two buttons were labelled "DO IT" and "ERASE". The buttons were approximately 4.5 cm. on center. Presentation of the problems 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 procedure 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. The display terminals and button boxes were on a 1.2 x .75 meter table. The tables were positioned in each experimental room so that subjects faced different walls.

Procedure

On entering the experimental room, the subject was given a set of written instructions. They contained a detailed description of the task, information about the content of the CRT display, and a description of how to respond. After having any questions answered by the experimenter, the subject solved a practice problem (jug sizes of 6, 4, and 1). Subjects were then given the experimental problem. The problem was terminated if subjects failed to solve within 100 moves (not counting erased or illegal moves).

A common procedure was used for all problems. Each problem began with the appropriate information, as shown in Figure 3, being presented on the CRT display. The line labelled "JUG SIZE" specified the sizes of the three jugs for the current problem; the "GOAL" line specified the respective goal state. The "CONTENTS" line changed appropriately after each move.

The subject entered a move by first pressing one of the buttons labelled "A", "B", or "C" to indicate the jug he wished to pour from. The subject then pressed a second button to indicate the jug he wished to pour into. The "POUR ? INTO ?" line was changed appropriately to display the responses made by the subject after he had entered both responses. The subject could erase an incorrectly input or undesirable move by pressing the "ERASE" button. Pressing "DO IT" caused the move to be made and the "CONTENTS" line to be changed appropriately. The message "ILLEGAL MOVE TRY AGAIN" was displayed for two seconds in the event the subject entered an illegal move (pouring from an empty jug or into a full jug). At the end of this time, this message was erased and the message "POUR ? INTO ?" reappeared.

RESULTS

The means and standard deviations for all conditions and problems are presented in Table 2. One randomly selected subject from each of Problems 2 and 3 of the move availability condition was dropped to simplify the data analysis. A two-factor analysis of variance for moves to solution indicated no difference between the control and move availability conditions, $F(1,234)=3.3$, $p>.05$; differences in problem difficulty, $F(2,234)=22.4$, $p<.001$; and no interaction between problems and experimental conditions, $F<1$.

The mean number of times subjects entered each state of the experimental problems was reduced to a vector. These data were analyzed with a groups by repeated measures analysis of variance (Grant, 1956). Since the three problems contained different numbers of problem states, a separate analysis was performed for each problem. The only significant difference was obtained on the conditions by repeated measures interaction for the (8,5,3) problem. Differences between the control and move availability conditions failed to reach significance (at the .05 level) for all problems.

Evaluation of the Model

All theoretical predictions were obtained by simulation. A FORTRAN IV program simulated the processes assumed by the model. Each run involved the performance of 250 subjects on a given problem with fixed values of the parameters $\alpha, \beta, \underline{s}$ and $\underline{r}$.

In order to account for the overall performance of subjects, the structure of the problems requires that we accurately simulate the number of times subjects return to the start state ($\underline{s}$). Post experimental interviews of subjects in an earlier experiment (Atwood and

Polson, 1975) suggested that at least some returns to $\underline{S}$ were made for the purpose of starting the problem over, a "desperation" move. We have incorporated a similar process in Stage III of our model. It is assumed that a subject in this stage will elect to return to the start state (if this is a legal transition) with a relatively small probability before attempting to select an optimal move. In the simulations presented, this probability has been set to .10. The structure of our model indicates that Stage III is a plausible location for this process. Our model uses two heuristics, means-ends analysis and the use of information stored in long term memory to detect moves to previously unvisited problem states. Stage I uses both of these heuristics, while Stage II uses only the second. The processes in both Stage I and Stage II "degrade gracefully" (Norman and Bobrow, 1975). That is, Stage I functions accurately even when no acceptable move is found; it simply rejects all unacceptable moves. Stage II functions in a similar manner. In Stage III, however, neither heuristic is appropriate if short term memory capacity has been exceeded (which is likely in water jug tasks). We postulate, therefore, that subjects choose, with some probability, to return to the start state ($\underline{S}$) rather than rely on possibly inappropriate heuristics. Recall that Stage III is reached only if both heuristics have been unsuccessfully applied. The processes underlying a subject's decision to return to the start state are undoubtedly complex; we have attempted to incorporate a simplified, plausible version of these processes.

The means-ends heuristic evaluates each problem state as a function of the number of times each state is recalled as having been visited in previous episodes. The addition of frequency information to the evaluation function was suggested by the results of our earlier investigations (Atwood and Polson, 1975). Notice that in all

experimental problems, the evaluation (excluding frequency information) of state R is very good relative to other problem states. Our model consistently predicted too many returns to this state. Since state k lies at the top of one solution path, over-predicting subjects' visits to this state seriously affected the model's ability to accurately simulate subjects' behavior.

State R is the most frequently visited state in all problems in both the observed and predicted results. While this tends to substantiate the means-ends evaluation function used by our model, our consistent over-prediction of the visits to state R suggest that transitions into this state become less attractive to a subject as he makes repeated attempts at solving the problem. That is, a state's evaluation is increased as a function of the number of remembered previous visits. (Recall that our evaluation function returns a value of zero for the goal state and increasing a state's evaluation makes the corresponding transition less desirable.)

Following Weber's law, we assume that frequency information is incorporated through a negatively accelerated function; as a consequence, we have adopted a logarithmic representation of frequency. We have assumed this constraint in our attempts to incorporate frequency information in our evaluation function. Due to the nature of the tasks used in this study, however, we do not require that frequency information affect the evaluation function in a strictly logarithmic manner.

In water jug tasks, there are generally four possible transitions from each problem state, and two of these transitions are irreversible moves into the "early" states. A subject who is relatively close to the goal state, therefore, would be returned to the initial states of the problem as a result of such a move. Since the consequences of such

transitions are clearly undesirable, we postulate that frequency information is interpreted as a function of its logarithmic representation. In the current simulations, we have assumed that frequency information is not utilized in evaluating states which have been visited relatively infrequently; in the simulations presented, frequency information does not affect a state's evaluation until after the third visit to that state. After this limit is reached, evaluations are increased by 3.5^a , where a is the natural logarithm of the difference between the number of recalled visits and this limit.

Since we have assumed that state representations of previously visited states are stored probabilistically, we are using the number of recalled visits rather than the number of actual visits to the problem state.

The exact form of the function we use to incorporate frequency information into our evaluation function is not crucial to our theory. The basic assumptions of our theory are unaffected by this addition. Our earlier results indicate that some form of frequency information should be incorporated into the evaluation function. Our primary purpose in introducing this information is to demonstrate that such an addition improves the overall performance of our simulation. The discrepancy in predicted visits to state R was most serious in the more difficult (16,10,3) problem. There were 9.32 observed visits to this state in the control condition. Our model, without frequency assumptions, predicted 12.78 visits; this was reduced to 9.83 visits by incorporating frequency information.

Theoretical predictions were compared with observed results on the basis of means and variances of moves to solution, distributions of solvers by solution path, and distributions of solvers and non-solvers (see Table 2). Significant differences were observed on the solution

by path distributions for the (12,7,4) problem in both experimental conditions.

There are two differences between the parameters used to simulate performance in the control and move availability conditions. For the control condition, short term memory was assumed to be capable of maintaining reliable information about at most 3 moves; this parameter ($\underline{r}$) was increased to 4 for the move availability condition. This parameter change is consistent with the experimental manipulation performed. In addition, we decreased the value of β (the probability of making a transition to a previously visited, acceptable state) from .13 (for the control condition) to .12 (for the move availability condition). Within the context of our model, this reduction implies that either a subject is less likely to move to a previously visited state, or (which is more likely) that β is confounded with the probability of recognizing a state as "old" and, consequently, a subject in the move availability condition was aided in this regard by being explicitly presented with all successors of the current problem state.

Discussion

We conclude that the model presented in this paper gives an adequate overall description of the data. The model correctly predicts the relative difficulties of different problems and adequately describes the overall behavior of subjects in the process of solving a given problem. The predicted moves to solution for all problems across both conditions are consistent with the observed results.

This experiment has provided a strong test of the current model. This model was developed in the context of naive subjects operating under severe memory data-limits. If subjects were capable of more efficient strategies, such as forward planning, such strategies should

become more evident as the subject's processing load is reduced and performance should be substantially improved. This does not appear to be the case. The results indicate that the assumptions of the model are plausible and deserve further consideration.

Recall that a primary assumption of this theory is that move selection is a memory data-limited process. As such, move selection uses a great deal of the subject's available resources and possibly prevents the consistent utilization of strategies which are more efficient than those assumed by our theory. In the move availability condition, the subject's processing load was considerably reduced. Our results, however, indicate that the additional resources were used to more effectively utilize the processes assumed by our theory. This suggests that these processes may represent the naive subject's best guess as to how such problems should be approached. More efficient strategies are apparently not observed not because of the resource limitations on move selection, but because they develop only after considerable attempts at problems in a given task domain.

— — — — —

REFERENCES

- Atkinson, R.C. & Shiffrin, R.M. Human memory: A proposed system and its control processes. In K.W. Spence & J.T. Spence (Eds.) The Psychology of Learning and Motivation: Advances in Research and Theory, Vol. II. New York: Academic Press, 1968.
- Atwood, M.E. & Polson, P.G. A process model for water jug problems. Program on Cognitive Factors in Human Learning and Memory Report No. 43. Institute for the Study of Intellectual Behavior, University of Colorado, Boulder, Colorado. August, 1975.
- Ernst, G.W. & Newell, A. GPS: A Case Study in Generality and Problem Solving. New York: Academic Press, 1969.
- Grant, D.A. Analysis of variance tests in the analysis and comparison of curves. Psychological Bulletin, 1956, 53, 141-154.
- Greeno, J.G. The structure of memory and the process of solving problems. In R.L. Solso (Ed.) Contemporary Issues in Cognitive Psychology. Washington, D.C.: Winston, 1973.
- Luchins, A.S. Mechanization in problem solving. Psychological Monographs, 1942, 54, No. 6 (Whole No. 248).
- Newell, A. & Simon, H.A. Human Problem Solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Norman, D.A. & Bobrow, D.G. On data-limited and resource-limited processes. Cognitive Psychology, 1975, 7, 44-64.

TABLE 1

Descriptions of the Experimental Problems

Problems	8,5,3	12,7,4	16,10,3
Right Path			
Length	7	7	13
Number of Violations	3	3	5
Number of possible loops to states <u>S</u> , <u>T</u> , <u>R</u> , and <u>L</u>	2,3,2,3	2,3,2,3	4,5,4,5
Number of acceptable looping moves to <u>S</u> , <u>T</u> , <u>R</u> , and <u>L</u>	0,1,2,1	0,1,2,1	1,1,3,1
Left Path			
Length	8	14	14
Number of Violations	2	3	3
Number of possible loops to states <u>S</u> , <u>T</u> , <u>R</u> , and <u>L</u>	3,3,3,3	6,6,6,6	6,6,6,6
Number of acceptable looping moves to <u>S</u> , <u>T</u> , <u>R</u> , and <u>L</u>	1,0,3,1	1,0,5,2	0,3,5,1

TABLE 2
Experimental Results

	^a Control	^b Move Availability
Problem (8,5,3)		
Mean	27.85(27.42)	20.58(24.60)
Standard Deviation	23.76(26.52)	18.37(21.39)
Problem (12,7,4)		
Mean	36.13(32.98)	34.78(30.09)
Standard Deviation	28.78(30.81)	27.89(26.12)
Problem (16,10,3)		
Mean	58.72(54.55)	48.10(49.76)
Standard Deviation	31.13(33.84)	31.66(31.65)

(a) Predictions computed by simulation; $\alpha=.60, \beta=.13, s=.90, r=3$
(b) Predictions computed by simulation; $\alpha=.60, \beta=.12, s=.90, r=4$

Mean Moves ANOVA

Source	SS	DF	MS	F
Conditions	2541.5	1.	2541.5	3.3
Problems	34561.0	2.	17280.5	22.4 (p<.001)
Interaction	864.0	2.	432.0	0.6
Residual	180920.3	234.	773.2	




Figure 1. Graph of the possible states and legal moves for the 8, 5, 3 problem. The numbers in parentheses are the current contents of the jars A, B, and C respectively. S:, T:, etc., are the state labels. State evaluations are listed beside the contents.




Figure 2. A flow diagram of the model for water jug problems. The symbols, α , β , $(1-\alpha)$, and $(1-\beta)$ represent the probabilities of taking the respective paths in the flow chart.

Control Condition Display

JUG SIZE	8	5	3
GOAL	4	4	0
JUG NAME	A	B	C
CONTENTS	2	5	1

POUR ? INTO ?

Move Availability Display

JUG SIZE	8	5	3
GOAL	4	4	0
JUG NAME	A	B	C
CONTENTS	2	5	1

POUR	RESULT		
B INTO C	2	3	3
C INTO A	3	5	0
B INTO A	7	0	1
A INTO C	0	5	3

POUR ? INTO ?

Figure 3. Sample Experimental Displays.