COGNITIVE PSYCHOLOGY 13, 113–148 (1981)

Formal Assessment of Problem-Solving and Planning Processes in Preschool Children

DAVID KLAHR AND MITCHELL ROBINSON
Carnegie-Mellon University

While much is known about adult problem-solving, the materials, analyses, and theoretical issues from the adult literature rarely make contact with the tasks typically used to investigate children's thinking. This paper examines the behavior of 4-, 5-, and 6-year-old children attempting to solve a novel variant of the Tower of Hanoi task. Problems varied in difficulty (one to seven moves for the minimum path solution) and in goal type: tower (all objects on one peg) or flat (all pegs occupied). For each problem, children gave verbal statements of their complete solution plan. The Plan Analysis examined performance as a function of goal type and age. Better performance was observed for tower ending problems, and among older children. The Error Analysis revealed that specific error propensities were related to both age and goal type. The Strategic Analysis compared the first move profiles of 6-year-olds to those of several plausible move selection models, and a high degree of correspondence was obtained between specific models and individual children. Young children appear to have rudimentary forms of many of the problem-solving processes previously identified in adults, but they may differ in encoding and representational processes.

How do children become proficient problem solvers? In order to answer this question, we need to obtain precise and accurate descriptions of the problem-solving methods used by children at different developmental levels. We know that by the time they are adults, most people have acquired a rich repertoire of problem-solving methods such as means-ends analysis, search, evaluation, and planning (c.f., Egan & Greeno, 1974; Jeffries, Polson, Razran, & Atwood, 1977; Newell & Simon, 1972; Simon & Reed, 1976). However, with few exceptions (e.g., Baylor & Gascon, 1974; Greenfield & Schneider, 1977), the tasks typically used to study young children's reasoning processes do not address these aspects of problem-solving. For example, the overwhelming favorites in neo-

This research was supported in part by Grant BNS77-16905 from the National Science Foundation, and in part by Grant No. G-78-0035 from the National Institute of Education. We thank David Neves and Robert Siegler for helpful comments on a previous draft of the paper. Thanks also to the parents, the teachers, and particularly the children of the CMU Children's School for their enthusiastic cooperation, and to Linda Szwedko and Margaret Kinsky for their meticulous assistance in running and analyzing these studies. Model simulation and evaluation programs were run on the facilities of the CMU Computer Science Department. Preliminary results of the Plan Analysis were presented in a paper read at the 1979 Biennial Meeting of the Society for Research in Child Development (Note 1). Requests for reprints may be addressed to David Klahr, Department of Psychology, Carnegie-Mellon University, Pittsburgh, PA 15213.

0010-0285/81/010113-36$05.00/0

Piagetian investigations are class-inclusion, transitivity, and conservation. Underlying successful performance on each of these tasks are decision rules that determine which of two quantities is greater, rather than anything that could be construed as a strategy for constructing a path to a final goal (c.f., Klahr & Wallace, 1976). Nor do most of the many other Piagetian tasks designed to assess the child's knowledge of the physical or social world address issues of constructive, goal-seeking behavior. Thus, it seems fair to say that much remains to be discovered about how children solve problems.

In this paper we present a description and analysis of the performance of preschool children in a difficult problem-solving situation: a variant of the well-known Tower of Hanoi (TOH) puzzle (Simon, 1975). The TOH has well-defined initial and final states, and a set of legal operations that, when applied in the appropriate sequence, can transform the initial state into the final state. The problematical aspect derives from the fact that the sequence of operations is not immediately apparent to the young problem solver, but rather must be produced through some combination of trial and error, systematic search, testing, planning, and so forth. Our basic question concerns the nature of the young child's problem solving with respect to such processes.

TASK DESCRIPTION

The "standard" version of the TOH consists of three pegs and a "pyramid" of n disks of decreasing size from bottom to top. The disks start out on one of the pegs, and the goal is to move the entire n-disk pyramid to another peg, subject to two constraints: only one disk can be moved at a time, and at no point can a larger disk be placed above a smaller disk on any peg. The minimum number of moves for an n-disk problem is 2^n-1. The task conforms to the definition of a well defined problem (Newell & Simon, 1972), in that it contains unambiguous descriptions of an initial state, a final state, and legal moves. The difficulty lies in discovering the sequences of legal moves that transform the initial configuration into the desired one.

Related Work

We know of only two prior studies of young children's performance on this task. In both cases, young children were found to perform poorly. Piaget (1976) used 2-, 3-, and 4-disk problems with children from about 5½- to 12-years-old. He reports that most 5- and 6-year-old children "cannot move the three-disk tower even after trial and error. They do succeed in moving the two-disk tower, but only after all sorts of attempts to get around the instructions and without being conscious of the logical links." (p. 288) From this performance, Piaget concludes that "none of these subjects make a plan or even understand how they are going to move the tower" (p. 290), and later, "There is . . . a systematic primacy of the trial-and-error procedure over any attempt at deduction, and no cognizance of any correct solution arrived at by chance." (p. 291)

Byrnes and Spitz (1977) used 2-disk (3-move) and 3-disk (7-move) TOH problems in a comparison of retarded and nonretarded children. Their nonretarded subjects ranged from 6- to 11 years old. On the 2-disk problem, the 6- and 7-year-olds made errors on about one-third of the trials, and the older children were nearly perfect. Almost all of the younger children failed the 3-disk problem, and even the older children could not solve it more than half of the time. A subsequent study including subjects up to college age was recently reported by Byrnes and Spitz (1979). Group averages on 2- and 3-disk problems showed abrupt increases between the ages of 7 and 9 and again between 11- and 15-years-old.

Thus, all of these studies find that only with difficulty can children around 5- or 6-years-old solve the 2-disk problem. This is a rather curious result, because the 2-disk problem requires only that the subject remove a single obstacle, the smaller disk, and place it temporarily on an unused peg in order to move the large disk and then the small disk. It is about the most rudimentary problem that one could pose. The results are also curious because, even an infant can remove a single obstacle to achieve a desired goal (Gratch, 1975). Furthermore, casual observation of young children coping with their daily circumstances suggests that they are capable of solving "problems" in familiar environments requiring 3 or 4 "moves" (such as getting a chair to reach a cabinet containing a string to tie on a doll).

We believe that preschool children do indeed have a greater problem-solving capacity than has yet been revealed. In this study, we attempted to construct a version of the TOH that would be sensitive to such capacities. At the same time, we wanted to guard against the problem of false positive interpretations. The steps we took to increase task sensitivity included modifications of the materials themselves, presentation of partial problems, prior familiarization with the materials, and a motivating cover story. Our attempt to guard against false positive assessment consisted of requiring the child to present a plan for his entire move sequence rather than simply making moves one at a time. All of these procedures will now be described in more detail.

Children's Version of the TOH

For use with young children, we modified the task in several ways that changed its superficial appearance while maintaining its basic structure.

Materials. We use a set of nested inverted cans as shown in Fig. 1. The cans fit so loosely on the pegs that when they are stacked up it is impossible to put a smaller can on top

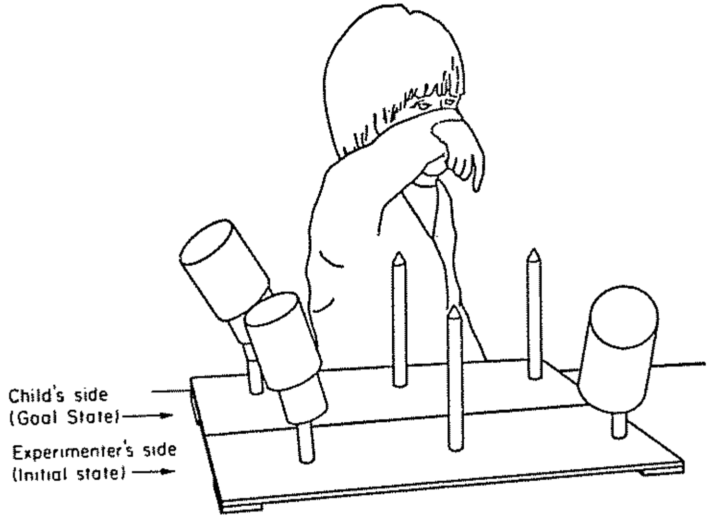

FIG. 1. Child seated in front of "Monkey Cans" working on a 1-move problem. (State 2 to State 1: see Fig. 2.)

of a larger can. Even if the child forgets the relative size constraint, the materials provide an obvious physical consequence of attempted violations: little cans simply fall off bigger cans. Furthermore, the materials are intuitively more "reasonable" in two regards. First, unlike the standard problem where small disks may obstruct larger ones, with these materials, bigger cans obstruct smaller cans, either by sitting atop them or by being on a goal peg. Second, larger cans not only sit on top of, but also partially contain the smaller cans.

Externalization of final goal. In addition to the initial configuration, the goal—or target—configuration is always physically present. We set up the child's cans in a target configuration, and the Experimenter's cans in the initial configuration. Then the child is asked to tell the Experimenter what she (the Experimenter) should do in order to get her (Experimenter's) cans to look just like the child's.[1] This procedure is used to elicit multiple-move plans: a child is asked to describe the complete *sequence* of moves necessary to solve the problem.

Cover story. Problems are presented in the context of a story in which the cans are monkeys (large Daddy, medium size Mommy, and small Baby), who jump from tree to tree (peg to peg). The child's monkeys are in some good configuration, the Experimenter's monkeys are "copycat" monkeys who want to look just like the child's monkeys (more details on the cover story are given below). The cans are redundantly classified by size, color, and family membership in order to facilitate easy reference. Children find the cover story easy to comprehend and remember, and they readily agree to consider the cans as monkeys. The remaining variations are best described after considering some of the formal properties of this task.

Formal state properties

Figure 2 shows all possible legal states and moves for these materials. It is called the "state space." The 27 unique configurations are arbitrarily

[1] All prenominal reference in this paper uses "she" for the experimenter and "he" for the subject.

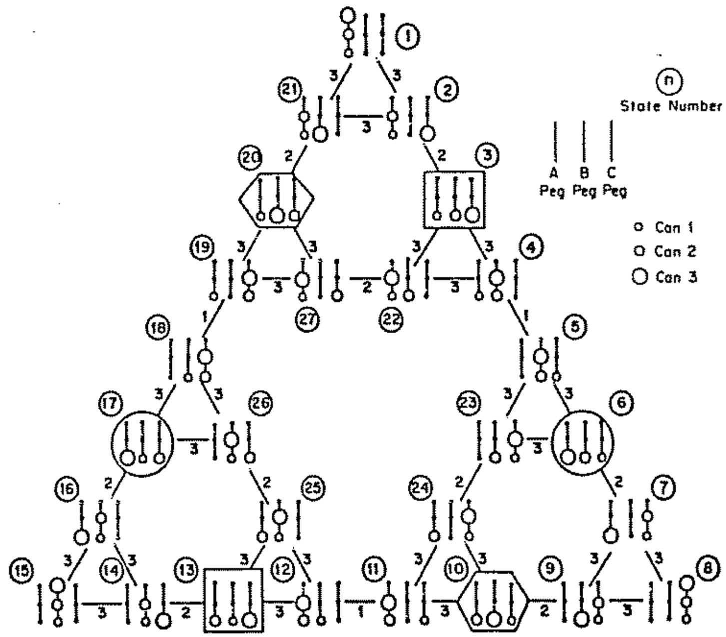

FIG. 2. State space of all legal configurations and moves for 3-can problem.

numbered in a clockwise direction, starting at the top. Each state is one move removed from its neighbors, and the can that is moved is indicated by the number on the line connecting adjacent states. The solution to a problem can be represented as a path through the state space. For example, the minimum solution path for the problem that starts with all three cans on peg A and ends with them on peg C is shown along the right-hand side of the large triangle in Fig. 2, moving from State 1 to State 8. The first move involves shifting the largest can (can 3) from peg A to peg C, producing State 2. The next move places can 2 on peg B (State 3), followed by a move of can 3 to peg B (State 4), and so on.

Reduced problems. Another task variation we introduce is to present a range of problems requiring from 1 to 7 moves for solution. (There are no two states for which the minimum path solution requires more than seven moves.) There are no dead ends in this task—any state can be reached from any other state—so that it is possible to choose from a large pool of unique problems (702 to be exact), simply by picking an arbitrary initial state and final state. Thus, although all problems involve three cans, some require only one move to solution (e.g., State 23 to State 6), some two moves (e.g., State 20 to State 1), and so on. By varying path length we can present children with problems of substantially different difficulty.

Flat ending problems. The "standard" TOH problems always end (and start) with all the cans stacked up on one peg, we call these "tower-ending" (T-end) problems. In the six states indicated by the large squares, circles, and hexagons in Fig. 2, all pegs are occupied. We call any problem that ends in one of these states "flat-ending" (F-end). Our final variation on the standard TOH is the use of F-end as well as the more commonly used T-end problems. The F-end problems provide a good test of the generality of the child's solution strategies. None of the previous formal analyses of the TOH have considered the relative difficulty of T-end and F-end problems. As we shall see, for the children in this study, F-ends were much more difficult.[2]

Subjects

All but a few of the children attending the Carnegie-Mellon University Children's School participated in the study. There were 19 children each in the "4-year" (mean 4-0; range 3-6 to 4-5) and "5-year" (mean 4-11; range 4-6 to 5-1) groups, and 13 in the "6-year" (mean 5-10; range 5-6 to 6-3) group. The educational program for the two younger groups is designed as a preschool program and for the older group as a kindergarten. The children come from predominantly, but not exclusively, white, middle-class backgrounds. There were approximately an equal number of boys and girls at each age level.

Procedure

Since our youngest subjects were only 3½ years old, we attempted to make the testing sessions as pleasant and stress-free as possible. Children were tested in a small playroom, adjacent to their regular classrooms, which was equipped with closed circuit video tape facilities. About 6 months earlier (in the preceding Fall semester) each child had spent about 2 hours in the room while engaged in activities unrelated to the problem-solving task. Then about 2 months after that, all the children had worked with a 24-item series of TOH problems in which the experimenter actually made moves, one-at-a-time, under the direction of the children. However, no direct instruction or evaluation was provided.[3] The experimenter was a 24-year-old, white female who interacted with all the children as a teacher throughout the year.

After being brought into the room, the child was again familiarized with the materials shown in Fig. 1, in the context of the following cover story.

> Once upon a time there was a blue river (Experimenter points to space between rows of pegs). On your side of the river there were three brown trees. Can you count your trees? On my side there were also. . . , etc. On your side there lived three monkeys: a big yellow daddy (present yellow can and place on peg), a medium size blue mommy, (present and place), and a little red baby. The monkeys like to jump from tree to tree [according to the rules]; they live on your side of the river. (Establish legal and illegal jumps.) On my side there are also three: a daddy, etc. . . (Introduce Experimenter's cans.) Mine are copycat monkeys. They want to look just like yours, right across the river from yours. Yours are all stacked up like so. . . [state 1], mine are like so [state 2 or 21]. Mine are very unhappy because they want to look like yours, but right now they are a little mixed up. Can you tell me what to do so mine can look like yours? How can I get my daddy across from your daddy, (etc.)? (The actual script is, of course, more elaborate. It is available upon request.)

[2] We suggest that the reader try to solve the State 1 to State 8 problem—a 7-move tower ending problem—and then State 10 to State 20—a 7-move flat ending problem. Most people find the latter harder.

[3] Some 5- and 6-year-olds had participated in pilot studies with the TOH in the previous year.

During the initial part of the familiarization phase, the child was allowed to handle the cans, but was gradually dissuaded from doing so. He was instead encouraged to tell the experimenter what she should do in order to get her cans to look like the child's.

The full set of problems is shown in Table 1. Each row of the table lists the problem number, the initial and final state numbers (corresponding to Fig. 1), the problem type (T-end or F-end), a numerical representation for the initial and final configurations and the sequence of moves for the minimum path solution. (For example, in Problem 5, the first move is to move the middle-size can from the rightmost peg to the center peg—2CB; the second move takes the largest can from the leftmost peg to the rightmost peg—3AC).

By the time the familiarization phase ended, the child had been given the first four problems (all requiring one move) and had solved them (with correction if necessary). For these problems, the experimenter actually carried out the child's instructions and moved her cans.

The remainder of the experiment was run in a "pure planning" mode: for each problem the child told the experimenter the full sequence of proposed moves, the experimenter gave supportive acknowledgment (or encouragement if the child was stuck) *but did not move the cans*, and then the next problem was presented. As each problem was being set up, the child turned away from the cans or shut his eyes until the experimenter said "ready."

In our previous work with the TOH (Klahr & Wallace, 1976), the children's suggested moves were actually executed—either a move at a time, or at the end of the full plan for each problem. However, this procedure provided the child with information about the efficacy of his initial problem-solving strategy, and frequently led to improvements over the course of the experiment. In the present study we wanted to minimize such learning, in order to obtain a more stable assessment of base-line performance.

As shown in Table 1, the full set of 40 problems consists of four problems having minimum path lengths of 1, 2, 3, and 4, and eight problems each with path lengths 5, 6, and 7. For each path length, half the problems are T-ending and half are F-ending. Problems were presented in two blocks with only T-ending problems in one block, only F-ending problems in the other. Children were randomly assigned to one or the other block order (F−T or T−F).

Within a given block, increasingly difficult problems were presented in the order listed in Table 1, until the child appeared to reach his upper limit. There were several indicators of this upper limit: (a) explicit statements of confusion or inability to continue; (b) abrupt violation of rules of the game (e.g., putting monkey in the river); (c) sudden loss of motivation; (d) consistent errors in planned moves. At this point, the session was terminated.

The time required to run through the 20 problems in a block ranged from 20 to 45 min. In some cases a block was run in two separate sessions, while in other cases—especially with the younger children—only a few problems beyond the practice phase were completed, and the session took only about 10 min. All sessions for each child were run within a period of a few days.

PLAN ANALYSIS

The raw data on which we base our analysis consist of detailed transcripts of videotape recordings of all sessions with all children. Two examples are shown in Table 2. The transcriptions include information about timing, backups, errors, corrections, restarts, and the child's own

TABLE I
Complete Problem Set With Mimimum Path Solutions

Problem	State I–G (type)	Configurations initial	goal	Minimum path moves 1st	2nd	3rd	4th	5th	6th	7th
1	23-6(F)	_/2/31	3/2/1	3CA						
2	3-4(F)	1/2/3	1/32/_	3CB						
3	8-7(F)	_/_/321	3/_/21	3CA						
4	16-15(T)	3/21/_	_/321/_	3AB						
5	27-3(F)	31/_/2	1/2/3	2CB	3AC					
6	20-1(T)	1/3/2	321/_/_	2CA	3BA					
7	10-8(T)	2/3/1	_/_/321	2AC	3BC					
8	24-6(F)	2/_/31	3/2/1	2AB	3CA					
9	19-3(F)	1/_/32	1/2/3	3CA	2CB	3AC				
10	19-1(T)	1/_/32	321/_/_	3CB	2CA	3BA				
11	11-6(F)	32/_/1	3/2/1	3AC	2AB	3CA				
12	5-8(T)	_/32/1	_/_/321	3BA	2BC	3AC				
13	5-1(T)	_/32/1	321/_/_	1CA	3BC	2BA	3CA			
14	18-3(F)	_/1/32	1/2/3	1BA	3CA	2CB	3AC			
15	18-1(T)	_/1/32	321/_/_	1BA	3CB	2CA	3BA			
16	7-3(F)	3/_/21	1/2/3	2CB	3AB	1CA	3BC			
17	23-1(T)	_/2/31	321/_/_	3CB	1CA	3BC	2BA	3CA		
18	8-3(F)	_/_/321	1/2/3	3CA	2CB	3AB	1CA	3BC		
19	1-6(F)	321/_/_	3/2/1	3AC	2AB	3CB	1AC	3BA		
20	3-8(T)	1/2/3	_/_/321	3CB	1AC	3BA	2BC	3AC		
21	11-3(F)	32/_/1	1/2/3	3AC	2AB	3CB	1CA	3BC		
22	22-8(T)	31/2/_	_/_/321	3AB	1AC	3BA	2BC	3AC		
23	6-1(T)	3/2/1	321/_/_	3AB	1CA	3BC	2BA	3CA		
24	19-6(F)	1/_/32	3/2/1	3CA	2CB	3AB	1AC	3BA		
25	16-1(T)	3/21/_	321/_/_	2BC	3AC	1BA	3CB	2CA	3BA	
26	16-3(F)	3/21/_	1/2/3	2BC	3AC	1BA	3CA	2CB	3AC	
27	18-6(F)	_/1/32	3/2/1	1BA	3CA	2CB	3AB	1AC	3BA	
28	2-8(T)	21/_/3	_/_/321	2AB	3CB	1AC	3BA	2BC	3AC	
29	7-1(T)	3/_/21	321/_/_	2CB	3AB	1CA	3BC	2BA	3CA	
30	26-6(F)	_/31/2	3/2/1	2CA	3BA	1BC	3AC	2AB	3CA	
31	14-8(T)	_/21/3	_/_/321	2BA	3CA	1BC	3AB	2AC	3BC	
32	12-3(F)	32/1/_	1/2/3	1BC	3AC	2AB	3CB	1CA	3BC	
33	15-1(T)	_/321/_	321/_/_	3BA	2BC	3AC	1BA	3CB	2CA	3BA
34	15-8(T)	_/321/_	_/_/321	3BC	2BA	3CA	1BC	3AB	2AC	3BC
35	13-3(F)	2/1/3	1/2/3	3CB	2AC	3BC	1BA	3CA	2CB	3AC
36	17-6(F)	3/1/2	3/2/1	3AB	2CA	3BA	1BC	3AC	2AB	3CA
37	20-8(T)	1/3/2	_/_/321	3BA	2CB	3AB	1AC	3BA	2BC	3AC
38	15-3(F)	_/321/_	1/2/3	3BA	2BC	3AC	1BA	3CA	2CB	3AC
39	13-1(T)	2/1/3	321/_/_	3CB	2AC	3BC	1BA	3CB	2CA	3BA
40	15-6(F)	_/321/_	3/2/1	3BC	2BA	3CA	1BC	3AC	2AB	3CA

TABLE 2
Two Protocols and Plan Encodings

Protocol	move	result	
Problem 25: 3/21/ _ (initial) 321/_/_ (goal)			
		3/21/ _	(initial)
What you do is you put the daddy (3) . . .			
What you do is you move the daddy (3) over this tree (points to C),			
and move and move the baby . . .			
and then you move the mommy (2),			
wait; where could you move the mommy (2) to ?,			
well first move the mommy (2) on this tree (points to C)	2BC	3/ 1/ 2	
then put the daddy (3) on that tree (points to C)	3AC	_/ 1/32	
and put the baby (1) over there (points to A).	1BA	1/ _/32	
Then how would the mother? . .			
and after you put the baby (1) over here (points to A)			
you could put the daddy (3) (points to B)	3CB	1/ 3/ 2	
then you could put the mommy (2) over the baby,	2CA	21/ 3/ _	
and the daddy over the mommy.	3BA	321/ _/ _	
Problem 29: 3/ _/21 (initial) 321/_/_ (goal)			
		3/ _/21	(initial)
Oh, that. O.K. That's easy.			
Just take the yellow one (3) and put it on there (B).			
Take the (pointing to 2(C)) . . . and take . . . and take, take the ba . . .			
No, take the blue one (2), put it on there (B),	2CB	3/ 2/ 1	
and then, then take the yellow (3)			
and put it on the blue (points toward C, then to B),	3AB	_/32/ 1	
and then take the red (1) one and put it on here (A).	1CA	1/32/ _	
And then take the blue (2) one			
and . . . no, and then . . . and then put the yellow (3) one here (C),	3BC	1/ 2/ 3	
and then put the blue one (2) on the red one,	2BA	21/ _/ 3	
and then put the yellow one on the blue one.	3CA	321/ _/ _	

Can 3	Yellow	Daddy
Can 2	Blue	Mommy
Can 1	Red	Baby

_/ _/ _
A B C
Pegs

rhetorical questions. For the plan analysis, we use only the final planned move sequences, encoding them as shown on the righthand side of Table 2. (Recall that no cans are actually moved during these protocols, so all the "results" except the initial and final ones are imagined rather than real.) Both of these protocols would be scored as perfect 6-move plans.[4]

In determining the maximum planning level for each child we use a very strict criterion: at a given problem length every plan must be the minimum path solution. (For all but problems 35 and 36, the minimum path solution is unique.) For example, to be classified as passing the 5-move planning level on T-ending problems, a child would have to produce the minimum path plan for problems 17, 20, 22 and 23.[5] Note that we do allow the utterance of illegal or nonminimum path moves, as long as they are recognized by the child and self-corrected, ultimately producing the correct solution path.

Tower-ending and Flat-ending problems were analyzed separately. The proportion of subjects in each age group producing correct plans for all problems of a given length is shown in Fig. 3a for T-ending problems and Fig. 3b for F-ending problems. Note that the abscissa in Fig. 3 is not overall proportion correct, but rather a much more severe measure: the proportion of subjects with perfect plans on all problems of a given length. For example, 9 of the 13 (69%) 6-year-olds were correct on all four of the 5-move problems, while only 3 of the 19 5-year-olds (16%) and 2 of the 19 4-year-olds (11%) produced four flawless 5-move plans.

[4] There is some question about the extent to which this is a plan in the sense used by Newell and Simon (Chap. 10, 1972), since it is at exactly the same level of detail as the solution itself. Further ambiguity arises from the likelihood that some of these utterances appear to be produced concurrently with the solution process, while others are probably immediate retrospective reports on problems first solved silently. For purposes of this analysis, we will use the term "plans" rather than the more cumbersome—albeit more accurate—"verbalization of proposed move sequences." (For a detailed examination of the effects of verbal protocols on the problem-solving process, see Ericsson & Simon, 1980).

[5] The probability of a random responder being misclassified is exceedingly slim. If we consider only legal moves, for all problems (except 33, 34, 38, and 40), the probability of the correct first move is ⅓. For all moves after the first, the correct move probability is ½. (We exclude the option of moving the can just moved—if we include it the probability of any move being correct drops to ⅓). Thus for a 5-move problem the probability of the correct solution is $\frac{1}{3} * (\frac{1}{2})^1 = 1/48$, and for all four of the 5-move problems (assuming independence of solution paths) it is $(1/48)^4 = 19 * 10^{-8}$. Similar computations show that this probability varies from approximately $7 * 10^{-3}$ for the first two 3-move problems to $7 * 10^{-10}$ for the four 7-move problems. One might argue that this is a gross overestimate of the power of our test, since it assumes intermove and interproblem independence. For adults, or for children actually making moves, it is reasonable to assume a much more tightly constrained solution path (c.f., Anzai & Simon, 1979). However, for children in *planning* mode, both illegal and impossible moves can be generated and the number of potential alternatives at each branch can be reasonably argued to be at least six (see error analysis section below), rather than the conservative two or three used here.

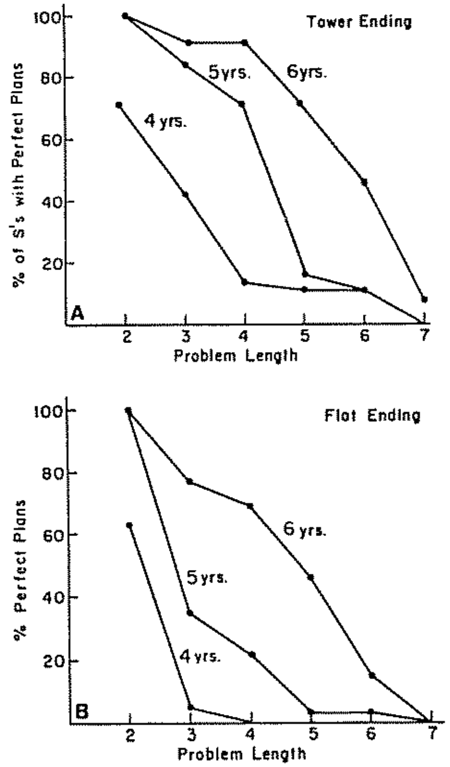

FIG. 3. Proportion of children producing perfect plans. (a) T-end problems; (b) F-end problems.

What is striking—given results of previous studies with children on this task—is the absolute level of performance. On the T-ending problems, over two-thirds of the 5-year-olds and nearly all of the 6-year-olds consistently gave perfect 4-move plans, and over half of the 6-year-olds gave perfect 6-move plans. Almost half of the 4-year-olds could do the 3-move problems. Recall that these plans are verbal descriptions of transformations of hypothetical future states. Furthermore, all intermediate states are different from, but highly confusable with, the two physically present states (i.e., the initial and final configurations).

A different picture emerges with F-ending problems. One-third of the youngest children could not do anything beyond a 1-move plan, and barely one-third of the 5-year-olds could reliably do the 3-move problems. Although the 6-year-olds did much better than the two younger groups, their scores were also substantially below their T-ending levels.

Individual subjects were assigned a planning level according to the maximum problem length for which all of their plans were perfect. Table 3

TABLE 3
Number of Children at Each Planning Level

		Planning level on T-end problems							Planning level on F-end problems						
Group	*N*	1	2	3	4	5	6	7	1	2	3	4	5	6	7
4-year-olds	19	6	4	6	1	0	2	0	7	11	1	0	0	0	0
5-year-olds	19	0	3	3	10	1	2	0	0	13	2	3	0	1	0
6-year-olds	13	0	1	0	3	2	6	1	0	3	1	3	4	2	0
Total	51	6	8	9	14	3	10	1	7	27	4	6	4	3	0

shows the number of children at each planning level for T-end and F-end problems. These planning level assignments form a Guttman scale, so that any child classified at a given level could solve all problems at or below that level. Since we were unable to detect any effect of presentation order, the results in Table 3, and all subsequent analyses, are collapsed across order.

The effects of goal type and age on planning level are very strong. Planning level scores for tower ending problems were greater than those for flat ending problems for 35 of the 51 children. Relative difficulty was reversed for only 4 of the remaining 16, and 12 children had equal planning levels on T- and F-end problems.

The effect of age on planning level was analyzed separately for tower and flat ending problems. A median test (Siegel, 1956) revealed a significant effect of age for both types of problems (tower ending $\chi^2_2 = 15.42$, flat ending $\chi^2_2 = 17.89$, $p < .001$).

Between-age group comparisons were made using Cochran's (1954) method for decomposing the chi-squared table. For both tower ending and flat ending problems, the distribution of the scores of the 4- and 5-year-olds were not reliably different with respect to the median $\chi^2_1 = .13$, $p > .5$, and $\chi^2_1 = 2.95$, $p < .10$, respectively. However, together they were both different from 6-year-olds $\chi^2_1 = 15.29$, $\chi^2_1 = 14.93$, respectively, both $p < .001$. A separate comparison of the scores of 5- and 6-year-olds supported this conclusion $\chi^2_1 = 4.26$, $p < .05$. The distribution of planning scores shift from below the median to above it with age, mostly among the 6-year-olds.[6] Some children classified at level 1 were in fact below that level, since the 1-move problems were treated as practice and often solved with the Experimenter's assistance. Thus, these children might justifiably be viewed as simply unable to do the task. All 2-move problems

[6] Dividing the planning scores alternatively into groups below the median and equal to or above it retains the overall age effect, but the grouping changes, with 5- and 6-year-olds better than the 4-year-olds, but not reliably different from each other. Thus, there is support for an overall age effect, but the position of the 5-year-olds is ambiguous.

require only that the two moves be performed in the correct order, but there are no obstacles in the way of either move. The first problems requiring removal of an obstacle—and hence the generation of the subgoal to remove that obstacle—are the 3-move problems. The column sums in Table 3 indicate that on T-end problems, only 8 of the 45 children who could solve 2-move problems could not get past this hurdle, while for F-end problems 27 of 44 were stymied at this point.

All 4-move problems start with the move of a can directly to its ultimate goal peg. Thus, 4-move problems should be negligibly harder than 3-move problems, and we should find very few children classified at level 3 (i.e., able to consistently solve 3-move, but not 4-move problems). Although this occurred in an earlier study (Klahr, 1978), it was evidenced here only by the 6-year-old group. For T-end problems, there are no level 3 planners in the oldest group, whereas ten of the thirty eight 4- and 5-year-olds were so classified. For the F-ending problems there is a slight dip at level 3 for the two older groups.

An abrupt decline occurs between level 6 and level 7 (10 out of 11 fail here) for T-ending problems. Recall that in all previous studies, young children almost never solved the 3-disk (7-move) problem. Neither did our subjects. However, if the problem is presented with the first move "already made"—reducing it to a 6-move problem—then about 20% of the preschoolers can solve it perfectly four out of four trials.

Summary: Plan Analysis

The aggregate analysis of plans has produced two major findings—one relating to absolute performance, the other to the effects of goal configuration. First, it is clear that by the time they are ready for First Grade, many children can produce plans up to six moves in length, even in an impoverished and arbitrary task domain. Second, the relative difficulty of formally equivalent problems (same materials, rules, state space, and path length) depends upon the form of the goal configuration. For F-end problems, it is not immediately apparent in which order the final configuration will be attained: in Problem 38 or 40 which can will reach its ultimate destination first? However, for T-end problems (e.g., 37 or 39), it is clear that the smallest can will have to reach the goal peg first, then the middle size can, and so on. Thus, one explanation for the differential difficulty of F-end and T-end problems is the extent to which the external materials facilitate the generation of subgoals. Another is that children's strategies interact with the goal configuration to make one type easier than the other, without any intrinsic difficulty being attached to either type of ending.

In order to get a better understanding of how such subgoals might be generated and used, we need to go beyond the aggregate analysis presented thus far. Although it has yielded a picture of what children can *do* with these problems, it has not addressed the issue of what they *know*, i.e., of what problem solving strategies they use. In the next section, we will describe some procedures we have used to generate and evaluate a set of models for children's strategies on these problems.

FIRST MOVE ANALYSIS

Perfect plans reveal those cases in which a child can access and execute an appropriate strategy. However, in order to characterize children's inadequate strategies, we need to search for regularities in all of their plans, including the incorrect ones. Production of perfect plans makes severe demands on the child's ability to maintain a mental representation for a changing configuration. A child with a correct strategy might produce a flawed plan if, in the course of execution, he forgot the location of some of the cans or neglected to describe some of the moves he knew were necessary. Since successive moves within each plan are increasingly susceptible to such failures of memory or production, we use the child's first move on each problem in our attempt to infer solution strategies.

Note that while the first move analysis ignores all subsequent moves on a given problem, we can view the first move of every n-move problem as the second move of some other $n + 1$-move problem. Thus, the first move analysis does provide information about the complete planning process. This will become particularly clear in the section on strategic analysis.

We used two levels of aggregation for the first move analysis. For *error analysis*, we aggregated across subjects within each age group, focusing on the types of nonoptimal moves and the circumstances under which they occur (see Augmented Problem Space, below). The error analysis suggests some of the strategic variations that the children might be using. For the *strategic analysis*, we focused on individual subjects, attempting to account for each subject's pattern of first moves over the entire problem set.

Error Analysis

Augmented problem space. The state space in Fig. 2 shows only legal moves. For the purpose of an error analysis we augment this space to include not just legal moves, but any other type of move that a subject might want to make. If there were no constraints at all, then at any state, the subject might produce a move composed of any of the 27 = (3 cans × 3 "from" pegs × 3 "to" pegs) combinations. Eliminating impossible moves (moving a can from a peg other than the one it is on) reduces the space to nine possible moves, but three of these are stationary (from a peg to itself). Thus, for the augmented problem space, we consider six possible moves (3 cans × 2 "to" pegs) from each state. For all but the three tower

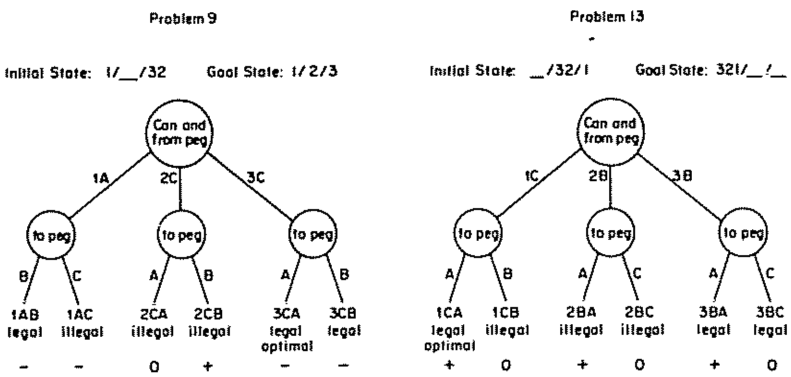

FIG. 4. Two examples of the augmented problem space with goal gradients.

states (1, 8, and 15) three of the six moves are legal; of these three legal moves, only one would be optimal. Two examples are shown in Fig. 4.

Our augmented problem space includes one additional characterization of each potential move: the extent to which the move changes the number of cans on their respective goal pegs. This *goal gradient* can be positive, negative, or zero. For example, for the augmented problem space around the initial state for problem 9, shown in Fig. 4a, the move 2CB would increase the number of cans on their goal pegs from two to three, so the goal gradient is positive (although the move is illegal). The optimal move, 3CA, has a negative gradient, since it reduces the number of cans correctly placed from two to one. For problem 13, there are three moves that might produce a positive goal gradient (1CA, 2BA, and 3BA). One is illegal (2BA) and one is optimal (1CA).

For all problems requiring more than two moves (problems 9 through 40), we scored all first moves into the categories shown in Table 4. Optimal, legal-nonoptimal, and illegal moves have already been described. "Other" moves is a residual category including partial moves (e.g., 3A?), conglomerate moves ("both of them to A"), and ambiguities in the protocols. It is sufficiently noisy to warrant excluding it from further analysis.

The main entries in Table 4 show the total number of first moves in each category by age and problem type. For example, of the 187 instances in which 6-year-olds generated first moves for T-end problems, 96 moves were optimal, 74 were legal, but nonoptimal, 8 were illegal, and 9 were "other." Since the number of problems presented to each child was determined by his performance, there are fewer responses for the younger children than for the older. The last two rows in Table 4 indicate, for example, that if each of the 19 4-year-olds had been presented with all 16 of the T-end problems beyond problem 8, there would have been 304

TABLE 4[a]

Number of Occurrences of Optimal, Nonoptimal, and Illegal First Moves, by Age and Problem Type

Type	Expected[b] frequency		T-end				F-end				
		Group	4-year-olds	5-year-olds	6-year-olds	Σ T Moves	4-year-olds	5-year-olds	6-year-olds	Σ F Moves	Σ moves
		n	19	19	13	51	19	19	13	51	51
Optimal	1/6		35(28)[a]	83(45)	96(51)	214(43)	32(34)	74(42)	66(42)	172(40)	386(42)
Legal, nonoptimal	2/6		39(31)	57(31)	74(40)	170(34)	35(37)	81(46)	88(55)	204(47)	374(40)
Illegal	3/6		35(28)	23(13)	8(4)	66(13)	24(25)	19(11)	5(3)	48(11)	114(12)
Other			16(13)	21(11)	9(5)	46(9)	4(4)	4(2)	0	8(2)	54(6)
Total first moves			125	184	187	496	95	178	159	432	928
Possible first moves			304	304	208	816	304	304	208	816	1632

[a] Percent of total first moves shown in parenthesis
[b] Relative frequency of move type in augmented problem space

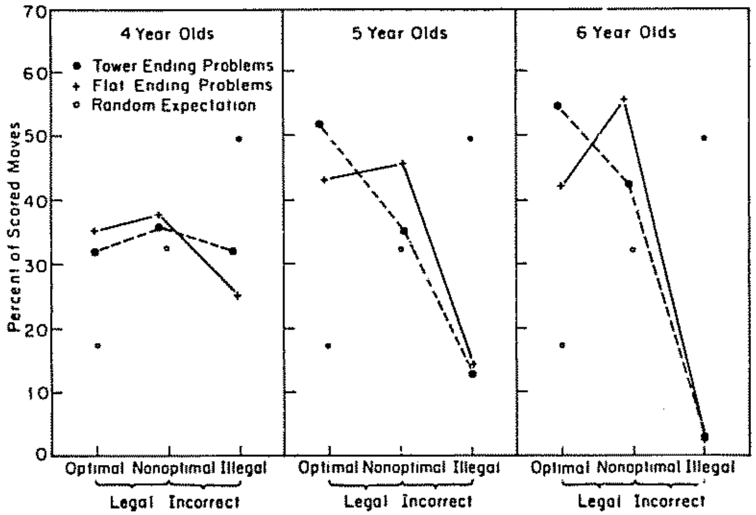

FIG. 5. Percentage of move types by age (both problem types).

responses, instead of only 125, and if the 13 6-year-olds had completed the full set of 16 F-end problems, there would have been 208 rather than 159 responses.

Entries from Table 4 are plotted in Fig. 5 as a percentage of all unambiguous moves (optimal, nonoptimal, and illegal) made by subjects in each age group. Figure 5 also shows what the expected frequency for each of the three move types would be if subjects randomly chose one of the six alternatives from the augmented problem space. For the two oldest groups, F-end problems turn out to be harder, even on the first move, with respect to the relative frequency of optimal (i.e., correct) moves. However, problem type does not affect the relative frequency of illegal moves.

The oldest children rarely choose illegal moves, and they choose optimal and nonoptimal moves about equally. The 5-year-olds show a slightly greater tendency to choose illegally.

The youngest children are only slightly affected by problem type. This could simply result from the fact that the 4-year-olds were exposed to many fewer of the harder problems—and thus to less structural variation—than their older classmates. Although they were well above chance in choosing optimal moves on the easy problems that they did get, the 4-year-olds were much more likely to make illegal moves than older children. Apparently, many of these 4-year-olds did not fully understand either the constraints of the problem, or the reporting regimen. By far, the most frequent illegal move for the 4-year-olds was to attempt to move a can directly to its goal peg when it was covered up by another can. This error could be due to a lack of understanding of the move constraint, or it could be a confusion between reporting a goal (e.g., ultimately getting the can to its final destination) and a move ("move it there now").

In the error analysis so far, we have considered only the most global structural property; the type of ending configuration. We can get additional information about why subjects make errors by looking at some of the structural relationships between current and final configurations. One type of common error occurs on states having a configuration that is neither a tower nor a flat. All such configurations have the general form $(xy/\_/z)$, and the common error is to move can x to the empty peg. (Of course, on some problems, this is the optimal move: thus on Problem 26, 2BC is not an error, whereas on Problem 9, 3CB is.) With T-end problems, subjects are likely to make this "flat error" slightly more than other kinds of nonoptimal moves, with F-end problems they make it five times more frequently than other errors. These ratios are essentially unchanged across the three age groups. If the flat error were simply the result of a local heuristic to move to an unoccupied peg, then it would be expected to occur with equal frequency for both types of final configuration. However, the overwhelming preference for flat errors on F-end problems suggests that children tend to use a global evaluation of the shape of the final configuration. Since, at the global level, the flat error always produces a configuration approximating the goal (even though different in detail), it is a very tempting move.

We suggested earlier that the T-end problems might be easier because the subgoal order was implicit in the final configuration. Some additional support for this conjecture comes from an analysis of all problems with a current configuration of a 2-can stack not on the goal peg and a T-end. When they err on these problems, subjects prefer, by a two to one margin, to move the smaller, rather than the larger, of the two cans directly to the goal peg. Although either move would have a positive gradient, it appears that the T-end facilitates the creation of a subgoal ordering according to can size, and that subjects are affected by this subgoal stack.

Finally, what can we say about the effect of the goal gradient? The most revealing indication of its role comes from an analysis of nonoptimal moves: those with positive gradients are chosen almost 80% of the time.

Three features of this error analysis suggest that we might be able to understand even more about the children's behavior if we analyzed their first moves from a different perspective. First, we have treated effects of problem type, gradient, etc., at a very aggregate level, and with only a minimum of sensitivity to particular relations that exist between initial and final states in different problems. Second, we have had to posit subjects' use of goals, strategies, and heuristics, but we have not been at all explicit about such cognitive processes. Finally, we have been completely insensitive to individual differences beyond the occasional age aggregations. In the next section we shall attempt to remedy these deficiencies.

Strategic Analysis

On every move, the child must decide which can to move and where to move it. A stable set of rules for making these decisions constitutes a *strategy* which produces, over the entire problem set, a characteristic "profile" of moves. (For the reasons given earlier, in this paper we will examine only first move profiles.) The basic idea underlying the strategic analysis is that, given a set of plausible strategies, we can compare each subjects first move profile with the profiles generated by each strategy over the same problem set. For each subject, the strategy producing the best matching profile is then taken as the strategy used by that subject. This sort of profile matching has been used to assess children's rule induction strategies by Klahr and Wallace (1970), and it has been applied to characterizing children's knowledge in a wide range of conceptual domains by Siegler (1976, 1978), under the rubric "rule assessment methodology."

We assume that well before they encounter the TOH problems, young children have at least the following two general principles for transforming their environment from one state to another.

P1: If you want object X to be in location B, and it is currently in location A, then try to move it from A to B.

P2: If you want to move object X from A to B, and object Y is in the way, then remove object Y.

In addition, we assume that children have an even more general principle that guides all of their goal oriented behavior:

P3: If the thing you are doing is too hard, then do some part of it that is easier.

The question of interest in this investigation is how children assimilate the TOH situation to these very general principles. The strategic analysis consists of describing a series of move selection models that might result from this assimilation process and comparing their predictions with subjects' behavior.

The move selection process can be decomposed into a series of decisions corresponding to the TOH instantiations of the general principles:

(a) *Subgoal selection:* which of the cans not on its goal peg should be attended to next?
(b) *Obstructor detection and removal:* if the subgoal move can't be made, then which of the obstacles should be moved, and to where?

(c) *Effort determination:* how much effort should be allocated to achieving the current subgoal?

We will explicate each of these in turn, and then incorporate their various forms in a series of move selection models. Close examination of each subject's protocol suggested that different aspects of these decisions were combined in quite distinct ways by different children. The 9 models to be described represent our attempt to explicate particular combinations of decision processes that might account for the behaviors we observed.

Subgoal selection. In the more difficult of our problems, where none of the cans are initially on their goal pegs, there are six (3!) possible preference orders that might be used to select subgoals. As indicated in the Error Analysis, subjects appear to use a smallest to largest (1, 2, 3) order most frequently, but not exclusively. It can be shown that this is equivalent to a heuristic of always attempting to solve the most difficult problem first (since can 1 is the most constrained) and that it would produce an optimal solution for a means-end strategy. Another way to select subgoals derives from a forward search strategy. First see which subgoal can be achieved in one move. If there is one, select it; if not, see what can be done in two moves, and so on.

Obstructor detection and removal. If the subgoal can be attained directly, then there is no problem. But if it is blocked by one or two cans either on top of it (on the "from" peg) or on its goal peg (the "to" peg), then a new subgoal must be created to move one of those cans. The move selection models embody important differences in the way this obstructor detection occurs, ranging from no detection at all to a recursive procedure.

Effort determination (depth of search). Each move in the TOH problem is selected with respect to a sequence of moves computed to achieve a specified subgoal. In our problem set, the length of such sequences vary from 1, in which the subgoal is achieved directly, to 4, in which three preparatory moves must be made first. For example, in a tower to tower problem, in order to achieve the first subgoal of moving can 1, three moves are required to remove the 2-can stack from the top of can 1. If the number of required moves exceeds the depth to which the child is willing or able to search, then he cannot compute the first move of a subgoal-achieving sequence. He can then do one of two things. He can apply a default rule, or he can select a new subgoal to work on. Our models differ in their depth of search capacity, and in what they do when it is exceeded.

Rather than list the formal rules used by each model, or the computer programs used to simulate model performance, we will simply give a brief description of each model and a list of its first moves.[7] Table 5 shows those

[7] The models are implemented in MACLISP, a variant of LISP. They are available upon request. People with access to the ARPANET can address inquiries to Dr. Klahr, CMUA, Pittsburgh, PA 15213.

TABLE 5
First Move Selections by All Models on Criterial Set (Moves Not Shown are Identical to Model 10 Moves)

Problem number	Configurations		Models								
	Initial	Goal	1	2	3	4	5	6	7	8	9
9	1/_/32	1/2/3	2CB	3CB							3CA
10	1/_/32	321/_/_	2CA				3CA				3CB
11	32/_/1	3/2/1	2AB	3AB							3AC
12	_/32/1	_/_321	2BC				3BC				3BA
13	_/32/1	321/_/_									1CA
14	_/1/32	1/2/3									1BA
15	_/1/32	321/_/_									1BA
16	3/_/21	1/2/3	1CA		3AB				3AB		2CB
17	_/2/31	321/_/_	1CA	3CA			2BA				3CB
18	_/_/321	1/2/3	1CA		3CB	3CB			3CB		3CA
19	321/_/_	3/2/1	1AC	3AB	3AB	3AB			3AB		3AC
20	1/2/3	_/_/321	1AC	1AC							3CB
21	32/_/1	1/2/3	1CA	1CA	3AB	3AB			3AB		3AC
22	31/2/_	_/_/321	1AC	3AC			2BC				3AB
23	3/2/1	321/_/_	1CA	1CA							3AB
24	1/_/32	3/2/1	1AC	1AC	3CB	3CB			3CB		3CA
25	3/21/_	321/_/_	1BA		3AC		3AC	3AC	3AC		2BC
26	3/21/_	1/2/3	1BA		3AC		3AC	3AC	3AC		2BC
27	_/1/32	3/2/1	1BC	1BC	3CA	3CA	3CA	3CA	3CA	3CA	3CB
28	21/_/3	_/_/321	1AC		3CB		3CB	3CB	3CB		2AB
29	3/_/21	321/_/_	1CA		3AB		3AB	3AB	3AB		2CB
30	_/31/2	3/2/1	1BC	3BA		3BA	3BA	3BA	3BA		2CA
31	_/21/3	_/_/321	1BC		3CA		3CA	3CA	3CA		2BA
32	32/1/_	1/2/3	1BA	1BA	3AC	3AC	3AC	3AC	3AC	3AC	3AB
33	_/321/_	321/_/_	1BA		3BC	3BC		3BC	3BC	3BC	3BA
34	_/321/_	_/_/321	1BC	3BA	3BA	3BA		3BA	3BA	3BA	3BC
35	2/1/3	1/2/3	1BA	1BA	2AB	2AB	2AB	2AB	3CA	2AB	3CB
36	3/1/2	3/2/1	1BC	1BC	2CB	2CB	2CB	2CB	3AC	2CB	3AB
37	1/3/2	_/_/321	1AC	1AC	2CA	2CA	3BC	3BC	3BC	3BC	3BA
38	_/321/_	1/2/3	1BA		3BC	3BC	3BC	3BC	3BC	3BC	3BA
39	2/1/3	321/_/_	1BA	1BA	2AB	2AB	3CA	3CA	3CA	3CA	3CB
40	_/321/_	3/2/1	1BC	3BA	3BA	3BA	3BA	3BA	3BA	3BA	3BC

first moves for all models on the 3, 4, 5, 6, and 7 move problems. For the sake of clarity, Table 5 actually contains—for models 1 to 8—only moves that differ from those made by model 9. Reference to Table 5 should resolve any ambiguities that may exist in the descriptions to follow.

Model 1 ignores the problem constraints: it makes direct moves to the goal peg, regardless of any obstructors. This was typical of many of the youngest subjects, and reflects the extent to which they totally ignored the constraints on moves. *Subgoal selection:* choose the smallest can which is not on its goal peg. *Obstructor detention:* none. *Search depth:* no search.

Model 2 is highly susceptible to the immediate salience of the physical configuration. The top obstructing can on the "from" peg is simply placed on an empty peg, without any consideration of subsequent effects. *Subgoal selection:* choose the smallest can which is not on its goal peg. *Obstructor detection:* detects only those cans on top of the can to be moved; selects the topmost can on the "from" peg and moves it to an empty peg. *Search depth:* no search.

Model 3 and *Model 4* both use some means-ends analysis, but they are limited with respect to their ability to respond to the difficulties they detect. Each model is sensitive to obstructors on both the "from" and "to" pegs; they differ only in which peg they prefer to clear first. If obstructors exist on both pegs, Model 3 decides to remove the "to" peg obstructor, while Model 4 focuses on the "from" peg. They both select the top can on their chosen peg. The next question is: where to place the chosen obstructor? Both models have the generalized concept of "other." For any two specified pegs, the models can designate the third—as yet unspecified—peg. They produce a tentative target peg for the obstructor which is the "other" of the "from" and "to" peg. For example, on Problem 25 (3/21/_ → 321/_/_), Model 3 generates the goal 1BA; it selects can 3 as the primary obstructor, and decides to move it to C—the "other" of A and B. Model 4, on the other hand selects can 2 as the obstructor, and chooses to move 2BC. Before making the move however, both models check to see if it is itself blocked. If it is, then instead of moving the obstructor to the illegal peg, it is moved to the "other" of the obstructor's source peg and the illegal peg. Thus in problem 39 (2/1/3 → 321/_/_), both models would initially choose 2AC, but they would notice its illegality, and instead produce 2AB. *Subgoal selection:* the smallest can not on its goal peg is selected first. *Obstructor detection:* both the "from" and the "to" obstructors are detected; Model 3 prefers the "to" can, and Model 4 prefers the "from" can. Move the obstructor to the "other" peg if legal, otherwise to the "other" of "other" and obstructor peg. *Search depth:* one level to check legality of obstructor removal.

Model 5 uses a limited, breadth-first search. First it looks for a direct legal move that will get a can on its goal peg (using the 1, 2, 3 subgoal ordering). If no such move can be made, then it looks for a 2-move sequence that might get a can to its goal peg. On problem 9 (1/_/32 → 1/2/3), Model 5 discovers that although can 2 cannot be moved directly to its goal peg, a 2 move sequence (3CA, 2CB) will succeed, so it chooses 3CA as the first move. On problem 10 (1/_/32 → 321/_/_) it discovers that it can move can 3 directly to its goal peg in a single move, and it chooses to do so. *Subgoal selection:* By move length, and within move length, by can size. *Obstructor removal:* Not relevant. *Search depth:* up to two moves.

Model 6 uses a limited depth-first search. It is similar in structure to Model 5, but the priority ordering on can selection and search depth is reversed. Model 6 chooses a can according to the "smallest first" preference order, and then seeks to achieve its goal in one or two moves, before moving on to the next can. On problem 17 (_/ 2/31 → 321/_/_) it looks first for a 1-move sequence that will get can 1 on peg A. Failing that, it retains the can 1 goal, and now looks for, and finds, a 2-move sequence (3CB, 1CA). *Subgoal selection:* By can size, and within can size by move length. *Obstructor removal:* same as Model 5. *Search depth:* two moves.

The effects of the difference between Model 5 and Model 6 are shown in Table 6. On problem 33 the models select different moves, while on problem 18, they choose the same move, but for different reasons.

Model 7 is sensitive, but faint hearted. It can detect all the obstructors to its preferred subgoal. If there are none, it makes the direct move. If there are any, it simply decides to move the largest can (can 3) to a peg that won't block the "from" peg of the current subgoal (or the "to" peg either, if possible). On problem 21 (32/_/1 → 1/2/3), it initially wants to move 1CA, but since that move is blocked, it chooses 3AB. On problem 28, (21/_/3 → _/_/321), its initial goal is 1AC, but cans 2 and 3 block that move, so it chooses 3CB. On problem 35 (2/1/3 → 1/2/3), having detected an obstructor to the 1BA subgoal, it decides to move can 3. Forced to choose between 3CB, which would block the "from" peg, and 3CA, which would block the "to" peg, it chooses the latter. *Subgoal selection:* smallest first. *Obstructor removal:* detects all, but chooses to move three to nonobstructing peg. *Search depth:* one level.

Model 8 is similar to Model 6, but it has an additional level of search. That is, it does a 3-move depth-first search, using the standard subgoal order. On problem 34, a tower-to-tower problem, it cannot find a 3-move sequence that would get can 1 to the goal peg. It then focuses on can 2, discovers the sequence 3BA, 2BC, and chooses 3BA as the first move. Note that problems 18, 19, 21, and 24 present similar difficulty with the can-1 subgoal, and in each case, Model 8 focuses instead on the can-2 subgoal. *Subgoal selection:* Same as Model 6. *Obstructor removal:* Not relevant. *Search depth:* 3.

Model 9 is taken from Simon's (1975) "sophisticated perceptual

TABLE 6
Trace of Models 5 and 6 on Problems 33 and 18

Problem 33 _/321/_ → 321/_/_			
Model 5		*Model 6*	
Desired sequence length: 1		Current goal: 1 to A	
current goal: 1 to A	fail	desired sequence length: 1	fail
current goal: 2 to A	fail	desired sequence length: 2	fail
current goal: 3 to A	succeed	Current goal: 2 to A	
		desired sequence length: 1	fail
		desired sequence length: 2	3BC, 2B
select: 3BA		select: 3BC	

Problem 18 _/_/321 → 1/2/3			
Model 5		*Model 6*	
Desired sequence length: 1		Current goal: 1 to A	
current goal: 1 to A	fail	desired sequence length: 1	fail
current goal: 2 to B	fail	desired sequence length: 2	fail
Desired sequence length: 2		Current goal: 2 to B	
current goal: 1 to A	fail	desired sequence length: 1	fail
current goal: 2 to B	3CA, 2CB	desired sequence length: 2	3CA, 2C
select: 3CA		select: 3CA	

strategy." It recursively attends to the minimum obstructor by moving it to the "other" of the immediately preceding subgoal. On problem 35 the sequence would be: Goal 1BA, blocked by 2. Move 2 to "other" of BA; subgoal 2AC. Blocked by 3. Move 3 to "other" of AC: subgoal 3CB. Choose 3CB. This strategy produces the minimum path solution (and thus the "correct" first move) on all problems but 27 and 32, where it adds one extra move.

Evaluation of Models

Since all models perform nearly identically on the 1- and 2-move problems (1–8), these easy problems provide little diagnostic power. All of the analysis, therefore is based on the *criterial set* of the last 32 problems (9–40). The basic measure we use is percentage of first moves correctly predicted.

Model–model performance similarity. The previous discussion treated model similarity in terms of subgoal ordering, obstructor detection, and depth of search. In order to assess the similarity of the *performance* of each model, we computed the number of times each model chose a move identical to every other model. Table 7 shows the percentage of identical first moves selected by all pairs of models among Models 3 to 9 on the criterial set and on selected subsets of problems.[8] The first column in Table 7 is based on the entire set of 32 criterial problems. The three pairs of models that are structurally quite similar (3–4, 5–6, and 6–8), are also fairly similar in performance (78, 81 and 81% identical moves, respectively). But the highest similarity measure comes from a pair of models that do not, on the surface, have much in common: 3–7 (84%). Furthermore, 3–6 and 4–8 choose the same first moves on 75 and 78% of the problems, respectively.[9]

This "unexpected" similarity exemplifies the point made by Simon (1975): quite different strategic variations can produce functional equivalence. Simon's analysis of the TOH focused on equivalent *optimal* strategies, while our analysis of children's performance has discovered some functionally nearly-equivalent nonoptimal strategies.

[8] As might be expected, Models 1 and 2 matched the other models less than 50% of the time.

[9] Hit rates above 70% are extremely "significant." A model that randomly generated first moves in the augmented problem space would have an expected hit rate with any other model of 17% (1/6), and one that generated only legal moves would have an expected rate of 33%. But the probability that randomly generated legal moves would achieve a hit rate with any given model of at least 72% (23 hits on the 32 criterial problems) is

$$\sum_{i=23}^{32}\left[\binom{32}{i}\left(\frac{1}{3}\right)^{i}\left(\frac{2}{3}\right)^{32-i}\right] < 10^{-5}.$$

TABLE 7
Model–Model Similarity (Models 3–9). Percent Identical First Moves for Model Pairs on Criterial Problem Set and Subsets

Model pair	All (32)	Flats (16)	Towers (16)	Easy (16)	Hard (16)
3–4	78	81	75	93	63
3–5	56	62	50	43	69
4–5	46	68	24	50	42
3–6	75	62	88	68	82
4–6	65	68	62	75	55
5–6	81	100	62	75	87
3–7	84	81	87	100	68
4–7	68	75	61	93	43
5–7	59	56	62	43	75
6–7	78	56	100	68	88
3–8	62	62	63	68	56
4–8	78	68	87	75	81
5–8	62	87	38	75	50
6–8	81	87	74	100	62
7–8	59	43	75	68	50
3–9	37	25	49	68	06
4–9	53	31	75	75	31
5–9	43	50	36	75	11
6–9	50	50	50	100	0
7–9	34	18	50	68	0
8–9	68	62	75	100	37

The degree of similarity depends on the problem set. A glance at Table 5 shows that there is little difference among models on the easiest problems. In order to assess the effect of problem type on model similarity more formally, we split the criterial set in two different ways: Flats vs Towers, and easy (3, 4, and 5 moves) vs hard (6 and 7 moves). Model similarity was then computed on each of these subsets. The results are shown in the last four columns of Table 7.

For some model pairs, functional equivalence is extremely sensitive to problem subset, while for others it is fairly stable across subsets. For example, the 65% similarity between Models 4 and 6 over all problems does not change much for Flats, Towers, easy or hard problems, whereas Models 6 and 9 choose identical first moves on easy problems and completely different moves on hard problems.

Since the longer problems generally provide more opportunity for the models to respond to obstructors, they tend to provide the most diagnostic power. All but 5 of the 21 pairs among Models 3 to 9 have higher similarity measures for the set of shorter problems than for the longer ones, and Models 6, 8, and 9 are functionally identical for the shorter set.

Finally, it is worth noting that none of the nonoptimal models comes very close to being correct (i.e., matching Model 9) most of the time. This is particularly evident in the last column of Table 7, which shows model–model fits on the harder problems.

Matching Individual Subjects

First move profiles for all criterial problems attempted by the 6-year-olds were compared with each model's profile for the corresponding set. Table 8 shows the "hit rates"—the percentage of identical first moves—for each of the 117 subject–model pairs. For example, on 55% of the 29 problems he attempted, Subject 1 and Model 4 selected the same first move. Model 4 also matched 70% of Subject 2's 24 first moves and 85% of Subject 4's. Overall, individual model–subject hit rates range from 6 to 90%. For individual models, mean hit rates across subjects range from less than 17% for Model 1 to almost 70% for Model 8.

The analysis reveals that quite distinct, but identifiable, strategies were used by the children. The weakest child (Subject 6) made only nine responses, and was best fit by Model 1 (the most simple minded), with a hit rate of 77%. Subject 11 made 13 responses, and was best fit by Model 2, at the 76% level. Again, Model 2 is intended to characterize children who were just a little beyond minimal levels of performance. At the other extreme, Subject 7 gave optimal responses on nearly all of the problems; he was fit by Model 9 on 90% of his first moves.

This leaves us with 10 subjects who generated error profiles of potentially high diagnosticity. For these 10, one is best fit by Model 4, two by Model 6, and seven by Model 8 (ties with Model 8 are resolved in its favor, for parsimony). Thus, three models account for the 10 subjects with a mean hit rate of 78%; nine out of the ten fits are greater than 70%.[10]

Since the last half of the criterial set has higher diagnosticity than the first, we repeated the analysis on just the last 16 problems (25–40). The general pattern is pretty much the same, except that Subject 10 is now best fit by Model 4 rather than Model 9 and Subject 14 becomes equally well fit by Models 3, 6, and 7. The mean hit rate for the 10 midrange subjects drops to 70%, and only 5 of the 10 fits are better than 70% (although 9 of the 10 are better than 60%).

In order to assess the robustness of these fits, we split the criterial set into subsets containing just odd or even numbered problems from the

[10] Recall that the probability that a *particular* model will match a random responder above the 70% level is $p_0 < 10^{-4}$. The probability that at least *one* of the 10 models will have a 70% hit rate with a subject who is responding randomly is $p_1 = 1-(1-p_0)^{10} < 10^{-3}$.

TABLE 8
Model–Subject Hit Rates on Criterial Set[a]

Subject	Response	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8	Model 9
1	29				55	55	65		*79*	65
2	24		54	62	70	54	75	66	*79*	66
4	28			82	*85*		67	78	67	53
5	32			71	75	62	*81*	65	75	56
6	9	77								
7	32				56		53		71	*90*
8	32			50	65		59	50	*78*	75
9	32		50	53	68		53	50	*71*	50
10	32			56	71	53	71	53	*84*	68
11	13		76	53	61	53	69	53	69	69
12	27			59	74	51	62	51	*74*	62
13	32				59	50	59		*81*	75
14	21		57	52	52		*66*	52	61	52

[a] Hit rates < 50% not shown
[b] Preferred model for each subject in italic.

criterial set, and computed the hit rates on those two subsets. We used the same procedure for resolving ties as we used with the full set. For all 13 subjects, the best-fitting model based on either the odd or even subset (or both) was identical to the best-fitting model based on the full criterial set; in all but four cases, both the odd and even subsets selected identical models. In other words, the split-half analysis produced only four out of a potential 26 anomalies.

In summary, an analysis of each child's first move profile has revealed substantial differences in the ability of these models to account for children's strategies. Three of the models, although plausible, were eliminated. Model 5 never quite reached a 70% hit rate for any subject, and Models 3 and 7, when they did exceed 70%, were always dominated by another model. Models 1 and 2 gave excellent fits for the weaker subjects. Model 9 fit the best subject very well, and Models 4, 6, and 8 accounted for the 10 middle range subjects.

Strategic Analysis: Discussion

What can we say about how children make the three key decisions listed earlier? The models finally selected indicate that all of the children utilize the 1, 2, 3 subgoal ordering. This could result from either the formal structure of the problem, or the integration of social convention with the semantics of the cover story: "baby first, then mother. . . , etc."

Children differ, however, in their response to obstacles. The Model 4 children treat single obstructors appropriately, by removing them to a neutral peg. When confronted with two such obstructors, they focus on the more immediate and salient one, that is, the can on top of the peg containing the can they want to move. Having chosen *which* obstructor to remove, the Model 4 children manage to keep track of *why* they want to move it—i.e., of the original subgoal—long enough to compute the "other" of the two pegs involved in that subgoal. But if they discover that the obstructor elimination move is itself blocked, they do not relinquish their selected can. Instead they move it to one of the pegs that was involved in the initial subgoal formulation.

Model 6 systematically proceeds through the subgoal order, seeking first 1-move, and then a 2-move sequence for the current subgoal. Detected obstructors do not generate subgoals for their removal, instead, that path of exploration is simply abandoned.

Model 8 children search for sequences up to three moves long that might achieve their initial subgoal. Only when thwarted at this depth of search do they replace their first subgoal with an easier one.

Recall that these best fitting models are relatively weak compared to a "mature" strategy (see Table 7.) On the criterial set, Models 1, 2, 4, and 6 make optimal choices on only 9, 43, 53, and 50% of the moves, respectively. Model 8 is correct 68% of the time. For the hardest 16 problems, these models are—for reasons described earlier—even worse: Model 4 is correct on only five problems, Model 6 on none, and Model 8 on six. Thus, the strategic analysis has clarified the sources of *non*optimal performance. Instead of simply reporting that beyond a certain level of difficulty children make errors, we have been able to specify the major processes that generate their behavior.

By eliminating several plausible models, the strategic analysis has enabled us to characterize not only what children do know about how to solve these problems, but also what they do not know. The dominance of Model 4 over Model 3 indicates a focus on an immediate rather than a remote obstructor, while the rejection of Model 5 reveals the absence of limited breadth-first search. Perhaps the most interesting comparison is between Model 8, which accounted for 7 of 11 children beyond Model 2, and Model 9—the "expert" model—which accounted for only one. Model 8 uses a 3-move depth-first search, which, while requiring a systematic consideration of move sequences, is simpler than the recursive procedure used by Model 9. The two strategies are only discriminable on those problems where can 1 can not be moved to its goal peg in three or fewer moves (e.g., problems 33–40; see Table 1).

While the depth-first search has the flavor of a rudimentary form of formal operations, it does not have the subtle requirements of subgoal management imposed by a recursive procedure. As evidenced by performance on this problem, such abilities appear to be just beyond the capacity of almost all of our sample of pretty bright 6-year-olds.

DISCUSSION

The results of this study provide clear evidence that by the time children are ready to enter First Grade, they have acquired the rudiments of a nontrivial range of general problem solving methods. Furthermore, they can apply these methods to a novel task. This finding raises two opposing questions: one concerned with why our subjects did so well; the other with why they did not do better.

As for the first: why have other investigators of this problem concluded that young children are capable of no more than trial and error? While there are many procedural differences between this and previous studies, we believe that the most important is our use of very fine-grained levels of differential problem difficulty. Recall that the Plan Analysis indicated that our children were no more successful with the standard 3-disk (7-move) problem, than were Piaget's or Byrnes and Spitz's subjects. However, a substantial number of them could solve up to 6-move problems. Careful examination of Table 5 reveals why this is so. Model 8 is nearly perfect until it encounters 7-move problems; Model 6 starts to err at 6-move problems, and Model 4 begins to weaken after 5 moves. Thus, the use of problems whose solution requirements lay between the standard 2-disk and 3-disk problems revealed some previously undetected problem-solving abilities.

It is likely that the externalization of the goal configuration also helped, principally by making it unnecessary to maintain an internal representation of the goal, and thus simplifying the difference detection process. The net effect of the rest of the task modifications (cover story, familiar environment and experimenter, interesting objects, etc.) was to maintain the children's attention long enough to have them make serious attempts to solve the many problems necessary for the profile matching procedure.

Recall that although the 6-year-olds did very well up through 5-move problems, the majority of 4-year-old children could not produce perfect plans beyond the 2-move problems (Fig. 3), and even their first moves were as likely to be illegal as legal (Fig. 5). One might conclude from this that the processes we are studying develop very rapidly between the ages of 4 and 6 years. However, such a conclusion is a bit puzzling when contrasted with the results from investigations of infants' search behavior (Gratch, 1975; Harris, 1975; Piaget, 1954). By the age of 12 months, most children have no trouble setting aside an obstacle in order to reach a desired object if it is visible. And by 18 months, most can use an object as a means to an end, such as reaching a toy on a pillow by pulling the pillow. Thus, the second major question raised by this study is: If children can solve what we have characterized as a 2-move problem at 18 months, why do they fail to solve our 3-move problems when they are 4 years old?

It is tempting to attribute these discrepancies to "decalage"—Piaget's name for unexpected failure of immediate transfer. For example, the difference between infant search and the poor performance of our youngest subjects on a task requiring verbal solutions could be attributed to *vertical decalage*, i.e., a situation in which "action is more advanced than verbal thought" (Ginsberg & Opper, 1969, p. 109). Indeed, in a previous study, we allowed children to move the cans as they solved problems, and the youngest children's performance was somewhat better than in the present study (Klahr, 1978). Of course, the TOH and the infant search tasks differ in many ways other than the verbal–nonverbal distinction; the performance differences may be yet another example of Piaget's *horizontal decalage*. That is, this may be a situation in which "Task contents . . . differ in the extent to which they resist and inhibit the application of cognitive structures" (Flavell, 1963, p. 23). However, the decalage label still leaves open the question as to the nature of the underlying difficulty.

We may begin to answer this question by distinguishing between two intertwined aspects of problem solving: strategies and representations. Thus far, we have focused entirely on the former; in these concluding comments, we offer some speculations about the latter. Throughout the Strategic Analysis, we assumed that the children's encodings were isomorphic to the external display: cans, pegs, positional relations (above and below), size relations, etc. The explanatory power lay entirely in strategic variations operating on uniform and veridical encodings. While this may be a reasonable approach for the 6-year-olds on whom we used it, it is probably not appropriate for the youngest children.

Indeed, in reviews of the infant search literature, one finds an emphasis not so much on strategies as on the encoding processes with which the child constructs an internal representation of the environment. Gratch (1975) provides an abstract characterization of the issue:

> When the actor imposes a point of view upon events, then the events take on coherence in terms of the perspective. When events are not assimilated into a framework, awareness of events is fragmented, partial, episodic. However, there are many perspectives that can be taken on events, each affording a different and limited ordering. To know reality, the actor must develop a higher level of perspective, must look at events from multiple perspectives which themselves are lodged in a larger perspective that permits the various looks to be ordered. (pp. 51–52).

Harris is more direct: "Despite the apparent sophistication of the neonate's perceptual apparatus, object displacement and disappearance are not encoded in an adult fashion during infancy" (1975, p. 332).

Thus, in integrating the literature on this fundamental aspect of knowing about the world, i.e., the development of the object concept, both reviewers have focused on the centrality of the development of an appropriate representational capacity. The impact on performance of developmental changes in general encoding and attentional processes has been emphasized by Baron (1978) and Klahr and Wallace (1970, 1976). Encoding has been identified as the source of developmental differences in learning (Siegler, 1976), and even with adults, changes in external forms of isomorphic problems produce substantial differences in performance (Simon & Hayes, 1976). We believe that one possible source of difficulty for our youngest subjects was the creation of an internal representation upon which their general problem-solving methods could effectively operate.

General problem-solving methods manifest themselves in rudimentary form by the end of Piaget's "sensory-motor period." They may emerge from the interaction of an "innate kernal" of regularity detectors (Klahr & Wallace, 1976, Chap. 8) and primitive encodings of sensory–motor activity. While much remains to be learned about the developmental trajectory of problem solving methods, we know even less about the development of encoding processes. In future investigations of problem solving by very young children, it will be necessary to provide a more

balanced treatment of representational and strategic variation. We will need to direct our attention to the conditions under which task environments are encoded such that they can be appropriately operated on by the rapidly emerging problem-solving processes.

REFERENCES

Anzai, Y., & Simon, H. A. The theory of learning by doing. *Psychological Review*, 1979, 86, 124–140.

Baron, J. Intelligence and general strategies. In G. Underwood (Ed.), *Strategies of information processing*. New York: Academic Press, 1978. Pp. 403–450.

Baylor, G. W., & Gascon, J. An information processing theory of aspects of the development of weight seriation in children. *Cognitive Psychology*, 1974, 6, 1–40. ✓

Byrnes, M. M., & Spitz, H. H. Performance of retarded adolescents and nonretarded children on the Tower of Hanoi problem. *American Journal of Mental Deficiency*, 1977, 81, 561–569.

Byrnes, M. M., & Spitz, H. H. Developmental progression of performance on the Tower of Hanoi Problem. *Bulletin of the Psychonomic Society*, 1979, 14, 379–381.

Cochran, W. G. Some methods for strengthening the common chi-square tests. *Biometrics*, 1954, 10, 417–451.

Egan, D. E., & Greeno, J. G. Theory of rule induction: Knowledge acquired in concept learning, serial pattern learning, and problem solving. In L. W. Gregg (Ed.), *Knowledge and cognition*. Potomac, MD: Erlbaum, 1974. Pp. 43–103. ✓

Ericsson, K. A., & Simon, H. Verbal reports as data. *Psychological Review*, 1980, 87(3), 215–251.

Flavell, J. H. *The developmental psychology of Jean Piaget*. Princeton: Van Nostrand, 1963.

Ginsberg, H., & Opper, S. *Piaget's theory of intellectual development*. Englewood Cliffs, NJ: Prentice-Hall, 1969.

Gratch, G. Recent studies based on Piaget's view of object concept development. In L. B. Cohen & P. Salapatek (Eds.), *Infant perception: From sensation to cognition*. New York: Academic Press, 1975. Pp. 51–102.

Greenfield, P. M., & Schneider, L. Building a tree structure: The development of hierarchical complexity and interrupted strategies in children's construction activity. *Developmental Psychology*, 1977, 13, 299–313.

Harris, P. L. Development of search and object permanence during infancy. *Psychological Bulletin*, 1975, 82, 332–344.

Jeffries, R., Polson, P. G., Razran, L., & Atwood, M. E. A process model for missionaries—Cannibals and other river-crossing problems. *Cognitive Psychology*, 1977, 9, 412–440. ✓

Klahr, D. Goal formation, planning, and learning by pre-school problem solvers, or: 'My socks are in the dryer'. In R. S. Siegler (Ed.), *Children's thinking: What develops?*. Hillsdale, NJ: Erlbaum, 1978. Pp. 181–212.

Klahr, D., & Wallace, J. G. The development of serial completion strategies: An information processing analysis. *British Journal of Psychology*, 1970, 61, 243–257.

Klahr, D., & Wallace, J. G. *Cognitive development: An information processing view*. Hillsdale, NJ: Erlbaum, 1976.

Newell, A., & Simon, H. A. *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall, 1972.

Piaget, J. *The construction of reality in the child*. New York: Basic Books, 1954.

Piaget, J. *The grasp of consciousness*. Cambridge, MA: Harvard University Press, 1976.

Siegel, S. *Non-parametric statistics for the behavioral sciences.* New York: McGraw-Hill, 1956.

Siegler, R. S. Three aspects of cognitive development. *Cognitive Psychology*, 1976, 8, 481–520.

Siegler, R. S. The origins of scientific reasoning. In R. S. Siegler (Ed.), *Children's thinking: What develops?* Hillsdale, NJ: Erlbaum, 1978. Pp. 109–149.

Simon, H. A. The functional equivalence of problem-solving skills. *Cognitive Psychology*, 1975, 7, 268–288.

Simon, H. A., & Hayes, J. R. Understanding complex task instructions. In D. Klahr (Ed.), *Cognition and instruction.* Hillsdale, NJ: Erlbaum, 1976. Pp. 269–287.

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

REFERENCE NOTE

1. Klahr, D. *Problem solving and planning by preschool children.* Paper presented at the Biennial Meeting of the Society for Research in Child Development, March 1979.

(Accepted June 3, 1980)

Statement of ownership, management and circulation required by the Act of October 23, 1962, Section 4369, Title 39, United States Code: of

COGNITIVE PSYCHOLOGY

Published quarterly by Academic Press, Inc., 111 Fifth Avenue, New York, N.Y. 10003. Number of issues published annually: 4. Editor: E. Hunt, Dept. of Psychology, Univ. of Washington, Seattle, Wash. 98195.
Owned by Academic Press, Inc., 111 Fifth Avenue, New York, N.Y. 10003. Known bondholders, mortgagees, and other security holders owning or holding 1 percent or more of total amount of bonds, mortgages or other securities: None.
Paragraphs 2 and 3 include, in cases where the stockholder or security holder appears upon the books of the company as trustee or in any other fiduciary relation, the name of the person or corporation for whom such trustee is acting, also the statements in the two paragraphs show the affiant's full knowledge and belief as to the circumstances and conditions under which stockholders and security holders who do not appear upon the books of the company as trustees, hold stock and securities in a capacity other than that of a bona fide owner. Names and addresses of individuals who are stockholders of a corporation which itself is a stockholder or holder of bonds, mortgages or other securities of the publishing corporation have been included in paragraphs 2 and 3 when the interests of such individuals are equivalent to 1 percent or more of the total amount of the stock or securities of the publishing corporation.
Total no. copies printed: average no. copies each issue during preceding 12 months: 2185, single issue nearest to filing date: 2100. Paid circulation (a) to term subscribers by mail, carrier delivery, or by other means: average no. copies each issue during preceding 12 months: 1684; single issue nearest to filing date: 1668: (b) Sales through agents, news dealers, or otherwise: average no. copies each issue during preceding 12 months: 0; single issue nearest to filing date: 0. Free distribution by mail, carrier delivery, or by other means: average no. copies each issue during preceding 12 months: 56; single issue nearest to filing date: 56. Total no. of copies distributed: average no. copies each issue during preceding 12 months: 1740; single issue nearest to filing date: 1724.

(Signed) Roselle Coviello, Senior Vice President

NOTICE TO CONTRIBUTORS:

The publishers wish to call your attention to the following instructions for preparing manuscripts for *Cognitive Psychology:* Format and style of manuscript should conform to the conventions specified in the "Publication Manual" of the American Psychological Association (1200 Seventeenth Street, N.W., Washington, D. C. 20036; 1974 Revision, $3.00), with the exceptions listed below. Please note that it is the responsibility of the author that manuscripts for *Cognitive Psychology* conform to the requirements of this journal.

INFORMATION FOR AUTHORS

Cognitive Psychology publishes original empirical, theoretical, and tutorial papers, methodological articles, and critical reviews dealing with memory, language processing, perception, problem solving, and thinking. This journal emphasizes work on the organization of human information processing. Papers dealing with relevant problems in such related areas as linguistics, artificial intelligence, and neurophysiology also are welcomed provided that they are of direct interest to cognitive psychologists and are written so as to be accessible to them and to make their relevance clear. There are no maximum or minimum length restrictions for journal articles.

All manuscripts should be submitted to the Editor: **Dr. Earl Hunt, Department of Psychology, University of Washington, Seattle, Washington 98195.**

Original papers only will be considered. Manuscripts are accepted for review with the understanding that the same work has not been and will not be published nor is presently submitted elsewhere, and that all persons listed as authors have given their approval for the submission of the paper; further, that any person cited as a source of personal communications has approved such citation. Written authorization may be required at the Editor's discretion. Articles and any other material published in *Cognitive Psychology* represent the opinions of the author(s) and should not be construed to reflect the opinions of the Editor(s) and the Publisher.

Authors submitting a manuscript do so on the understanding that if it is accepted for publication, copyright in the article, including the right to reproduce the article in all forms and media, shall be assigned exclusively to the Publisher. The Publisher will not refuse any reasonable request by the author for permission to reproduce any of his or her contributions to the journal.

A manuscript submitted for publication is judged by three main criteria: (a) appropriateness of the subject matter for this journal; (b) significance of its contribution to knowledge; and (c) clarity and conciseness of writing. No changes in a manuscript may be made once it has been accepted and is in press.

Form. Type at least double-spaced throughout, including tables, footnotes, references, and figure captions, with 1 inch margins on all sides. Submit four complete copies to expedite processing. One copy remains the property of the journal. Each copy *must* include all figures and tables.

Number the pages consecutively. *Page 1* should contain the article title, author(s) names(s), and affiliation; at the bottom of the page type a short title, not exceeding 35 characters and spaces, and the name and complete mailing address (including zip code) of the person to whom proofs should be sent. The address for correspondence should be given as a footnote to the appropriate author's name. *Page 2* should contain a short abstract, approximately 100 to 150 words in length.

Headings. The organization of the paper must be clearly indicated by appropriate headings and subheadings.

Abbreviations. Do not use final periods with units of measure that are abbreviated (cm, sec, kg, etc.) in text of in tables, except for "in." (inch).

Symbols. Underline letters that represent mathematical symbols; these will be set in *italic* type.

Equations. Number displayed equations consecutively, with the number placed in parentheses to the extreme right of the equation. Refer to numbered equations as Equation 1 or say "the first equation." Punctuate equations to conform to their place in the syntax of the sentence.