

Implicit Learning in Problem Solving: The Role of Working Memory Capacity

Paul J. Reber and Kenneth Kotovsky
Carnegie Mellon University

Participants solving the Balls and Boxes puzzle for the first time were slowed in proportion to the level of working memory (WM) reduction resulting from a concurrent secondary task. On a second and still challenging solution of the same puzzle, performance was greatly improved, and the same WM load did not impair problem-solving efficiency. Thus, the effect of WM capacity reduction was selective for the first solution of the puzzle, indicating that learning to solve the puzzle, a vital part of the first solution, is slowed by the secondary WM-loading task. Retrospective verbal reports, tests of specific puzzle knowledge, and a recognition test of potential strategies all indicated that participants were unaware of their knowledge of the puzzle, suggesting that it had been learned implicitly. Concurrent protocols collected from participants supported this conclusion and further suggested that participants were not aware of learning to solve the puzzle as this learning occurred. These results provide evidence that implicit learning depends on WM capacity and that implicit memory can play an important role in problem solving.

Starting with the original conception of problem solving as search through a problem space presented by Newell and Simon (1972), progressively more complex models and theories of problem solving have provided increasingly detailed accounts of human problem-solving behavior. These theories of problem solving universally assume a long-term knowledge store of rules and problem-solving strategies and a separate working memory (WM) that provides limited processing space for the evaluation of strategies and planning of operator application (e.g., move) sequences. The limited capacity of WM makes it possible to investigate the interaction of memory and problem-solving processes through manipulation (e.g., reduction) of WM capacity. However, recent work on the neurological basis of memory has shown that the canonical separation of memory into two components (WM and long-term memory) significantly underestimates the complexity of human memory. A number of studies of the phenomena of implicit learning and preserved learning in patients with anterograde amnesia (cf. Seger, 1994, and Squire, Knowlton, & Musen, 1993, for reviews) have provided convergent evidence that there are several long-term memory stores dependent on different

areas of the brain. This basic change in how memory is understood has only just begun to have an impact on theories of complex or "higher" cognition, such as problem solving. Furthermore, the relationship of implicit learning to WM capacity is not well understood. In this article, we present an investigation of the effect of WM capacity reduction in a problem-solving task that appears to be acquired implicitly.

An important difference between implicit and explicit memory is the fact that implicit memory does not afford awareness of stored knowledge (N. J. Cohen & Squire, 1980; P. J. Reber & Squire, 1994). Implicit learning (without awareness) has been studied in a variety of tasks including learning of artificial grammars (A. S. Reber, 1989) and sequence learning (P. J. Reber & Squire, 1994; Willingham, Nissen, & Bullemer, 1989; Nissen & Bullemer, 1987) as well as some problem-solving tasks (Berry & Broadbent, 1988; Squire & Frambach, 1990) but has drawn some criticism as to whether such learning truly operates without the conscious knowledge of the subject (Dulany, Carlson, & Dewey, 1985; Perruchet & Amorim, 1992; Shanks & St. John, 1994). A fundamental issue raised by these critiques is whether the operation of complex cognitive processes can be separate from awareness of knowledge that supports them. In some tasks known to depend on implicit memory, such as perceptual priming (see Tulving & Schacter, 1990, for a review) in which prior exposure to visual stimuli is shown to facilitate rapid identification, the knowledge acquired in an initial study (or exposure) phase could be considered to consist of small changes to simple representations and therefore not be expected to impinge on awareness in the manner of conscious, deliberate thought. To address the issue of the role of nonconscious knowledge in complex cognition, it is necessary to study implicit learning in complex domains such as problem solving.

The task we used to examine implicit learning in problem

Paul J. Reber and Kenneth Kotovsky, Department of Psychology, Carnegie Mellon University. Paul J. Reber is now at the Department of Psychiatry, University of California, San Diego.

This work was reported in Paul J. Reber's doctoral dissertation submitted in 1993 to the Department of Psychology, Carnegie Mellon University.

We would like to thank Gerry Starrett and Marsha Lovett for assistance in data collection and to thank Herb Simon, John Anderson, Tim Curran, and Rich Carlson for helpful comments on the manuscript.

Correspondence concerning this article should be addressed to Paul J. Reber, Veterans Administration Medical Center 116A, 3350 La Jolla Village Drive, San Diego, California 92161. Electronic mail may be sent via Internet to preber@ucsd.edu.

solving is the Balls and Boxes puzzle. This puzzle, previously studied in Kotovsky and Simon (1990), has been shown to be effective in examining the effect of varying WM demands across several isomorphic puzzles and was selected for study on this basis. The implicit nature of the process of learning to solve the puzzle was initially unanticipated and for this reason, Experiments 1 and 2 focus on the role of WM capacity reduction on learning to solve the puzzle and executing the learned solution, whereas Experiments 3 and 4 focus more directly on the conscious knowledge of the puzzle that participants are able to report. In addition, various tests were used across the four experiments to attempt to discover the content of participants' solution strategies beyond what they were able to report.

The Balls and Boxes puzzle is similar to a number of reasonably simple, artificial problem-solving tasks that take on the order of 5 min to 2 hr to solve (e.g., Tower of Hanoi, Anzai & Simon, 1979; "monster" isomorphs, Hayes & Simon, 1977; water jug, Luchins, 1942, and Atwood & Polson, 1976; river-crossing problems, Greeno, 1974; Chinese rings, Kotovsky & Simon, 1990) in that it consists of a set of elements (in this case, five balls and five boxes), each of which is constrained to be in one of a number (in this case, 32) of possible states (for the balls, the allowable states are in or out of an associated box; the boxes can either be open or closed; see Experiment 1 for details). From each state some set of pieces can be moved to another state, as defined by the underlying rules that describe the puzzle. The set of allowable configurations of elements (problem states) and the transitions allowed between states defines the *problem space* for the puzzle. In solving the puzzle, the participant is presented with an initial problem state and must select from among the possible *operators* (i.e., allowable moves) in order to change the state of the puzzle, thus "moving" in the problem space. In all these relatively simple puzzles, the process of operator selection and application results in moving through the problem space, and this problem-solving process continues until the participant reaches the goal state, concedes defeat, or runs out of time.

One important difference between the Balls and Boxes puzzle and the other puzzles is that it is not possible to deduce the underlying rule structure of the puzzle from the initial description and presentation of the problem. In other puzzles (e.g., Tower of Hanoi), it is theoretically possible to deduce the optimal solution strategy in advance of working on the puzzle. However, in practice, participants tend to explore the problem space for a time before deducing an appropriate strategy for solving the puzzle (e.g., Anzai & Simon, 1979). Following the research on these other puzzles, we expected participants to learn to solve the Balls and Boxes problem by initially using some set of weak, general-purpose problem-solving methods such as hill climbing (i.e., selecting moves such that each move reduces the apparent distance to the goal) and progressing to a more sophisticated strategy such as *means-ends* analysis (i.e., setting and working on "subgoals" to solve particular aspects of the problem and organizing a set of subgoals so that completion of these leads to the final goal state). The difficulty of the Balls and Boxes puzzle, which takes several

minutes to solve in spite of a small and highly constrained problem space, appears to result from the fact that certain "weak methods" (e.g., widely applicable general problem-solving heuristics such as hill climbing and search based on "backing up to look for another way") lead to nonoptimal operator choices in certain areas of the puzzle. These weak methods are useful in many other problem-solving tasks, but are not particularly appropriate for this puzzle. In order to solve the Balls and Boxes puzzle, it is necessary to abandon these weak methods in favor of an alternate strategy. Discovery of an alternate strategy through problem space exploration is known to depend on WM capacity. It may also depend on implicit memory, although previous research on this topic has focused on problems that are learned explicitly (i.e., problems where participants are aware of and able to report the acquired solution strategy).

WM and its role in complex cognition is a relatively well-studied area. Virtually all theories or models of cognitive skill and problem-solving behavior from GPS (Newell & Simon, 1972) to ACT-R (Anderson, 1993) propose a limited-capacity storage-processing area similar to short-term memory, but with an emphasis on the processing aspects of this structure rather than the storage aspects. Early views of WM as a single fixed-capacity buffer (e.g., Atkinson & Shiffrin, 1968) have more recently been extended into models in which WM is distributed among a set of specialized processing modules (e.g., Schneider & Detweiler, 1987) or characterized as the active elements of a memory network (as in the ACT-R theory, Anderson, 1993). This change is largely due to evidence that short-term storage does not always compete for WM resources with active processing (e.g., Baddeley & Hitch, 1974; Klapp, Marshburn, & Lester, 1983). When short-term storage does compete with active problem-solving processes (e.g., Carlson, Khoo, Yaure, & Schneider, 1990), the interference produced is hypothesized to be due to competition for access to a central control system (Schneider & Detweiler, 1987) or executive process (Baddeley, 1990) or due to competition in the process of spreading activation to the problem-solving rules or processes (Anderson, 1993).

In the study of problem solving, the limited capacity of WM means that a consideration of the WM load associated with operators and strategies is central to understanding problem difficulty (Carpenter, Just, & Shell, 1990; Kotovsky, Hayes, & Simon, 1985; Kotovsky & Simon, 1990). Participants' problem-solving ability is improved by reducing a problem's associated WM load either by modification of the problem (e.g., problem isomorphs, Kotovsky et al., 1985; Kotovsky & Simon, 1990) or through the acquisition of expertise in a particular problem or domain (Kotovsky & Kushmerick, 1991; Carlson et al., 1990). To examine the role of WM capacity on solving the Balls and Boxes task, we used a dual-task paradigm that required participants to solve the problem concurrently with performing a secondary task that required constant updating of information in WM and could be varied in difficulty. In Experiments 1 and 2, we assessed the effect of reducing available WM capacity in the very early stages of learning, before real expertise develops. In addition, these experiments also examined the

effect of reducing WM capacity on a second solution, at a point where significant information has already been acquired about the puzzle.

There has previously been only limited research examining the role of WM capacity and resource competition in implicit learning tasks. Most of this research has used sequence learning tasks. In those tasks, participants respond to a cue appearing in one of four (or six) locations by pressing a corresponding button. Participants are not told that the cue follows a repeating sequence of locations but demonstrate knowledge of this by improved reaction time performance (compared to when the sequence is changed or removed), which can occur in the absence of any apparent awareness of the repeating sequence (P. J. Reber & Squire, 1994; Willingham, Nissen, & Bullemer, 1989). Using this task, A. Cohen, Ivry, and Keele (1990) and Nissen and Bullemer (1987) showed that a secondary task that divided participants' attention impaired participants' implicit learning. Curran and Keele (1993) proposed that certain types of complex sequential information could not be implicitly learned under divided attention, although simpler sequences could be implicitly learned in the same situation. However, Frensch, Buchner, and Lin (1994) found that the same complex sequential information could be learned under divided attention although learning was slowed. Although these results focus mainly on the role of dividing attention in implicit learning, Frensch and Miner (1994) have additionally shown that under some experimental conditions, individual differences in WM are related to the amount of implicit learning that occurs. However, Stadler (1995) found that requiring participants to maintain information in WM during the task did not appear to affect learning, although secondary tasks that disrupted the apparent organization of the sequence did impair learning. In the experiments reported in this article, we have applied the dual-task methodology to reduce WM capacity in conjunction with performing the Balls and Boxes puzzle, a problem-solving task that appears to depend on implicit memory. The secondary task used was selected to allow us to control for the effect of simply dividing attention while solving the puzzle and clearly requires use of executive WM function (as in Baddeley, 1990) by requiring frequent updating of currently held information.

Experiment 1

In the first experiment, we investigated the effect of WM load on problem solving using a dual-task paradigm. Participants' solution strategies and knowledge of the puzzle were assessed through a retrospective verbal report.

Method

Participants


Participants were 80 undergraduates at Carnegie Mellon University (CMU) who partially fulfilled a course requirement by their participation.

Materials

The Balls and Boxes puzzle. The puzzle used for this research was first described in Kotovsky and Simon (1990) as an isomorph of the Chinese ring puzzle (Ruger, 1910). The Chinese ring puzzle is an extremely difficult puzzle (solution times range up to 10 hr) whose difficulty appears to be due to the high level of WM load associated with its operators (e.g., it is difficult to determine what is a move in the problem space, making the planning of move sequences extremely difficult). In the Balls and Boxes puzzle (called the "Lo Info" puzzle in Kotovsky & Simon, 1990), the operators used to manipulate the puzzle are much simpler and hence impose less of a WM load during move planning and execution. The inherently low WM load suggested that reducing WM capacity through an external secondary task might interfere with problem solving without making the task inordinately difficult for the participants.

The Balls and Boxes puzzle is presented on a computer screen (Mac Icx [Apple Macintosh Icx, Apple Inc., Cupertino, CA] or Micro VAX [MicroVAX II, Digital Equipment Corporation, Maynard, MA]) to participants and is usually solved in 5–10 min. Figure 1 shows the puzzle as it appears to participants together with the standard task instructions. The puzzle consists of five balls each associated with one of five boxes (initially, each ball is in its box). A ball can only be in one of two states, either in its box or out of its box. Thus there are only 10 total possible operators (i.e., each of the five balls could be moved either into or out of its box). For each location in the problem space (configuration of balls in or out) there are typically only 2 operators that yield legal moves (i.e., there are two boxes that are open indicating two balls that can be moved either in or out of those boxes). The rule that determines whether a ball can be moved (and that box is open) is as follows: The rightmost ball can always move; other balls can be moved if the ball immediately to the right is in its box and all other balls to the right are out of their boxes.

The problem space is fairly small, with only 32 states (possible configurations of the balls in or out), and furthermore, these states interconnect such that the problem space is linear. Figure 2 contains a complete description of the problem space, showing each possible state of the puzzle, numbering them by their distance from



The goal of the puzzle is to get the five balls out of the five boxes. A ball can be moved in or out of its box using the computer mouse by clicking on the ball. A ball may only be moved in or out of its box if its box top is open. For instance, right now the two balls on the right could be moved but not the three on the left. As you move balls in and out of their boxes, the box tops will open and close. The trick to the puzzle is to move the balls to get the right boxes to open up so that you can move all the balls out of their boxes.

Figure 1. Initial appearance of the puzzle with the standard task instructions. The puzzle is presented on a computer. The stated goal is to get all five balls out of their boxes (balls can only be in or out, no other position is allowed). Boxes can be moved in or out of their boxes only if the box top is open. The right two boxes are initially open, indicating that either of these moves is available but no other moves are.




Figure 2. The problem space of the Balls and Boxes puzzle. Each state is shown with the state number underneath. The problem space is linear, so that states with consecutive numbers are connected in the problem space (e.g., from State 21, it is possible to move to States 22 or 20 but no others).

the goal state (the state where all the balls are out). The linear structure of the problem space is best appreciated by noting that from any state, the only legal moves are to the states numbered one more and one less (i.e., from State 21, it is possible to move to State 20 or State 22; for State 31, the only legal move is to State 30 because there is no State 32). The legal moves for each state are indicated by the fact that moveable balls are either in or above (out of) boxes with open tops. Note that which boxes are open often changes when a move is made to another problem state (e.g., State 21 to State 20). The change in which boxes are open or closed is done by the computer program presenting the puzzle and is done according to the underlying rules (outlined in the following section).

Given the small size of the problem space, it might be surprising that it takes as long as 5–10 min to achieve a solution. One reason it takes participants some time to solve the Balls and Boxes puzzle is that a solution strategy cannot be deduced by inspection of the initial presentation of the puzzle and its instructions. Exploration of the problem space is required to discover a strategy for solving the puzzle. This search is constrained, however, by the fact that the problem space is linear. In all but two of the problem states, there is a choice between two legal moves, whereas at the ends of the problem space there is only one legal move. Whenever there are two moves, one of the moves is to a state one step closer to the goal and the other to a state one step further from the goal (though this

is not clear to the naive participant trying to solve the puzzle for the first time). Another way of viewing this is that after a move has been made, one of the two available moves “undoes” the previous move, whereas the other is a move to a new state. A move that undoes the immediately previous move will be referred to as a *reversal*. For example, if a participant moves from State 21 to State 20 (by moving the far right ball out, see Figure 2), he or she is then faced with the choice of moving the middle ball out or replacing the far right ball into its box (see Figure 2 for a diagram of how State 20 appears to the participant). If the participant chooses to replace the far right ball, thus returning to State 21, we say that a reversal has been made; the participant has reversed direction in the problem space.

The starting state presented to participants is State 21 (states are numbered by their distance from the goal). The structure of the puzzle is such that the weak methods participants usually attempt initially (i.e., hill climbing) do not successfully solve the puzzle. Because the stated goal is to remove all the balls from their boxes, there is a natural tendency to begin attempting to solve the puzzle using a strategy in which progress is measured by the number of balls out of their boxes. Much of the difficulty of the puzzle is due to the inapplicability of this approach in several areas of the puzzle. For example, progress from State 15 to State 10 appears to someone employing this heuristic to be regressive because three balls must be replaced in their boxes (see Figure 2). As is shown later, participants make many reversals in this region while trying to determine a path to the goal.

How the Balls and Boxes puzzle can be solved. Analysis of the puzzle suggests three alternative and effective solution strategies:

1. *Rules.* The rule for the puzzle that determines when it is legal to move a ball (i.e., when the ball’s box is open) is

Except for the rightmost ball, which can always move, a ball can be moved only when the ball immediately to the right is in its box and all other balls to the right are out of their boxes.

Using this rule, it is always possible to deduce the next correct move and, in theory, the entire sequence of moves necessary for a solution. For instance, because balls are always dependent on other balls to the right, the leftmost ball is the most constrained and needs to be considered first from the starting state in order to solve the puzzle in the minimum number of moves. According to the rule, to remove the leftmost ball, the second ball from the left must be in and the other three balls must be out. From the start state, the second ball from the left is already in; the next step is to remove the middle ball. From the rule, the middle ball can be moved when the second ball from the right is in and the rightmost ball is out. Therefore the correct first move is to remove the rightmost ball. This is followed by removing the middle ball and then considering how to remove the second ball from the right, and so forth. This strategy represents a means–ends approach to solving the puzzle.

2. *Avoid-reversal strategy.* Because the problem space is linear, there are always two moves (except at the end states) to choose from. If one chooses randomly at the first state and then never undoes the last move, never moving the same ball twice in a row, progress is guaranteed to either the goal state or the top of the problem space (State 31) where there is only one legal move. (Even if one chooses the first move incorrectly and ends up at State 31, after the forced reversal there, the problem is solved in 31 more moves, 41 total.) This strategy by itself does not guarantee the shortest solution because it is possible to choose the wrong move initially, but it does guarantee at least a reasonably efficient solution.

3. *Move-pattern strategy.* The sequence of moves to traverse the problem space is highly patterned. The pattern is easily seen by

listing which ball is moved in sequence moving from State 31 to the goal: 545 3 545 2 545 3 545 1 545 3 545 2 545 3 545 (spacing added for emphasis). This pattern does not overtly indicate whether balls are being moved in or out, but knowing this pattern, it is possible to correctly decide which ball should be moved next and thereby solve the puzzle.

Secondary task. Participants listened to an audiotape recording that consisted of a stream of letters with beeps mixed in at random intervals. Three different levels of memory load were generated by asking participants to remember one, two, or three of the most recent letters from the tape (e.g., participants were asked to remember either the last letter they had heard from the tape, the last two letters heard, or the last three letters). The set of letters to be remembered thus changed with each new letter while imposing a constant WM load. All three memory-load groups were also instructed to listen to the tape for the beeps. Participants were told that when they detected a beep, they were to write down the first (or "oldest") of the group of n letters they were remembering (i.e., the participants remembering one letter would write down the last letter they had heard, the participants remembering two letters would write down the letter before the last letter that they had heard, and the participants remembering three would write down the first [earliest] of the three letters they were remembering). Because of the difficulty of the secondary task in the high memory load conditions, all participants were told to be sure that they paid special attention to the tape so that "you make sure you hear every beep, and always are able to write down the correct letter." They were further told that getting the letter right every time was more important than "rushing through solving the puzzle." Participants in the control group (zero memory load) were instructed to ignore the letters and listen for the beeps. Every time they heard a beep they were to write down a "random" letter (or a letter from the tape if they desired) as long as they did not write the same letter repeatedly.

Procedure

Participants were split into four groups for the first trial ($n = 20$ each). Each group received instructions for a different level of WM load: (0) control, (1) remember the last letter, (2) remember two letters, (3) remember three letters. After solving the puzzle the first time, participants were asked to write about "how you solved the puzzle, how the puzzle worked and especially anything you could say that would help someone else solve the puzzle." After writing this, participants solved the puzzle a second time. On the second solution, each group was split ($n = 10$ each), and half the participants solved the puzzle the second time in the zero WM load-control condition, whereas the other half solved the puzzle under a WM load of two letters.

Results

Problem-Solving Performance

The increasing level of WM load led to an increased number of moves to solve the puzzle on the first trial, as can be seen in Figure 3. However, on the second trial, all groups performed similarly regardless of condition. The results from the second trial are presented in Figure 4 showing both the memory-load condition on the first trial and the memory-load condition on the second trial. A $4 \times 2 \times 2$ analysis of variance (ANOVA) with between-subjects fac-




Figure 3. Problem-solving performance, Experiment 1, Trial 1. Number of moves to solve the puzzle is measured on the y-axis. The four levels of working memory (WM) load are indicated on the x-axis. Bars indicate the standard error of the mean.

tors of first and second trial WM-load condition and a within-subjects factor of trial was performed and indicated a significant effect of first trial WM-load condition, $F(3, 68) = 3.23, p < .03$, no effect of second trial WM-load condition, $F(1, 68) = 0.73$, and no interaction, $F(3, 68) = 1.61, p > .19$. Participants were impaired on solving the puzzle on the first trial with a WM load, but on the second trial the WM load had no effect. The within-subject effect of trial was significant, $F(1, 68) = 20.15, p < .001$, with a marginal interaction with Trial 1 condition, $F(3, 68) = 2.62, p < .06$. Other interactions were not significant ($F_s < 1$). The significant trial effect reflects the fact that participants

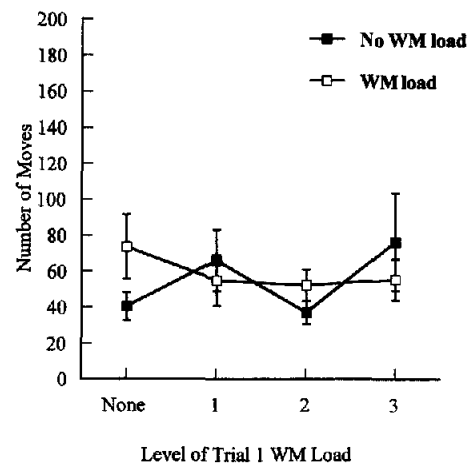


Figure 4. Problem-solving performance, Experiment 1, Trial 2. Number of moves to solve the puzzle is measured on the y-axis. Filled triangles indicate participants who solved the puzzle under a working memory (WM) load; open triangles indicate participants who solved the puzzle without a memory load. The four levels of WM load from Trial 1 are indicated on the x-axis. Bars indicate the standard error of the mean.

made fewer moves to solve the puzzle on the second trial. The marginal interaction of Trial 1 condition and Trial 2 condition reflects the fact that all participants showed speedup essentially to a common point despite having taken different numbers of moves in the different WM-load conditions on Trial 1.

A similar analysis with overall time taken to solve the puzzle as the dependent measure showed the same pattern of results: a significant effect of Trial 1 WM-load condition, $F(3, 75) = 7.32, p < .01$; no effect of Trial 2 WM-load condition, $F(1, 75) = 0.09$; and no interaction, $F(3, 75) = 1.64, p > .18$. The effect of trial was significant, $F(1, 76) = 63.2, p < .001$, with an interaction with Trial 1 condition, $F(3, 76) = 9.16, p < .01$, whereas the other interactions were not significant: WM load with Trial 2 condition, $F(1, 76) = 0.002$; WM load with Trial 1 and Trial 2 condition, $F(3, 76) = 1.42, p > .30$.

Secondary Task Performance

Performance on the secondary task, as measured by the number of correct responses to the probe (i.e., the correct letter being written down), was analyzed to determine whether participants might be trying to improve their problem-solving performance by ignoring the secondary task. On the first trial, error rate on the secondary task was not found to be predictive of number of moves to solve the puzzle within any condition (assessed by linear regression of error rate vs. number of moves, $p > .9, p > .15$, and $p > .7$ for memory load Levels 1, 2 and 3). Thus there was no speed-accuracy trade-off in performance of the secondary task. There was, however, a relationship between memory-load condition and secondary task performance as participants in the higher WM-load conditions (2 and 3 letters) performed less well on the secondary task (error rates: 2% for Level 1, 15% for Level 2, and 40% for Level 3). This was unsurprising because the secondary task was by design more difficult in these high memory-load conditions.

An analysis of the responses to the probe in the control group indicated that the majority of these participants were generally writing the last letter heard at a probe, similar to the behavior of participants in the lightest WM-load condition (1) where participants were required to remember the last letter heard. However, 4 control participants wrote the last letter heard only 3% of the time, significantly less than the other control participants who wrote the last letter an average of 61% of the time, $t(19) = 9.20, p < .001$. These 4 participants also performed very poorly on the puzzle on the first trial, taking 135, 141, 203, and 237 moves to solve. This group took an average of 179 moves to solve the puzzle, whereas the rest of the control participants took an average of 74 moves, $t(19) = 3.31, p < .01$. Although there were a few other control participants who performed in a similar range on the problem-solving task (e.g., 157, 161, and 201 moves), the group of 4 in question appeared to have solved the puzzle under different (self-imposed) memory-load conditions. These participants appear to have been generating random letters, a task shown by Baddeley (1966)

to interfere with use of WM. These participants were eliminated from the analyses reported above. However, if these participants are included in the analysis, the same pattern of results holds. Thus, the increase in solution time across WM-load condition in Trial 1 was significant for overall time to solve the puzzle, $F(3, 80) = 7.11, p < .01$, although the effect of Trial 1 condition for the number of moves to solve the puzzle is only marginal, $F(3, 80) = 2.49, p < .07$. The effect of trial is significant for both dependent measures— $F(1, 80) = 75.8, p < .001$ for overall time; $F(1, 80) = 26.5, p < .001$ for number of moves—and the interaction between Trial 1 condition and trial is significant for overall time to solve the puzzle, $F(3, 80) = 8.53, p < .001$. As before, with these participants included there is no effect of Trial 2 condition for overall time to solve the puzzle, for number of moves or on any interaction ($F_s < 1$). Thus the major conclusion, that a WM load significantly interferes with Trial 1 performance but has no effect on Trial 2 performance, is unaffected by the inclusion or exclusion of these anomalous participants.

Verbal Report

The written retrospective protocols given after one solution of the puzzle indicated that participants were not able to specify how they had solved the puzzle. To assess overall puzzle-knowledge content for each of the written protocols, each protocol was rated as to how close to a complete solution strategy the participant came. This rating was on a scale from 1, meaning the protocol had no useful information about the puzzle, to 5, meaning the participant described a complete strategy for solving the puzzle (a list of possible strategies for solving the puzzle was described earlier). No protocol earned a 5 rating. A protocol was given a rating of 1 if it contained no informative statements about the puzzle at all (e.g., containing only statements such as "it was hard to listen to the letters" and "there were five balls and five boxes"). A rating of 2 meant that the protocol contained one somewhat informative statement (e.g., "the leftmost ball was hard to get out"). For a protocol to receive a rating of 3, it had to include several somewhat informative statements (e.g., "The ones on the left were hard. In order to get these out, you had to keep taking out and replacing the ones on the right."). A rating of 4 or better was assigned when the protocol, for example, contained partial rule information to a degree that could significantly aid solution (e.g., "the leftmost ball could be removed when the right three were out, the second from the left could be removed when the right two were out") or contained a general strategy (e.g., "whenever a box opens, move that ball, otherwise move the rightmost ball"). The majority of protocols (58 of 76 or 76%) were rated as 2 or less, generally indicating little verbalizable knowledge of the puzzle. Table 1 shows the number of protocols that were rated at each level and the corresponding mean number of moves that group of participants took to solve the puzzle on the second trial.

Despite their incompleteness, the quality of the protocols were predictive of how well the participant then fared on the

Table 1
*Retrospective Protocol Ratings and Second Trial
 Performance From Experiment 1*

Protocol rating	No. of participants	Mean no. of moves on Trial 2
1.0	18	64
1.5	25	72
2.0	15	52
2.5	8	46
3.0	2	23
3.5	4	29
4.0	3	29
4.5	1	41

Note. Protocol ratings range 1–5. Higher numbers indicate a protocol with more information about how to solve the Balls and Boxes puzzle.

second trial. Participants who gave better descriptions solved the puzzle more rapidly on the second trial (regression), $F(1, 75) = 5.34, p < .03$. However, only 10 of the 76 participants gave protocols rated as 3 or higher and even the participants with poor protocols (e.g., rated as less than 3; the vast majority of participants) showed a significant increase in performance on the second puzzle, $t(65) = 3.12, p < .01$. In addition, of the 10 participants with minimum length solutions on Trial 2, only 2 gave protocols (after Trial 1) that were rated as 3 or above. The mean rating of their protocols was 1.90 (± 0.96), indicating that, on average, even these participants were able to describe very little of what they knew about the puzzle.

Overall, the statements given by the participants indicate at best only partial knowledge of the problem. Expression of knowledge of the structure of the problem space was almost completely absent, and statements related to explicit strategies were rare and at best partial or vague.

Final Path

In marked contrast to the paucity of information available in the retrospective protocols, subjects exhibited what Kotovsky et al. (1985) referred to as *final path* behavior at the end of each solution of the puzzle. The final path refers to a sequence of rapid, error-free moves leading to completion of a problem-solving task. Such behavior, occurring just prior to the successful solution of the problem, would normally be interpreted as indicating the achievement of some crucial insight or strategy for solving the problem. For this puzzle, we defined the final path as the number of error-free (reversal-free) moves made immediately prior to reaching the goal state. The average final-path length was 18.4 (± 1.0) moves on Trial 1 and 19.0 (± 0.9) moves on Trial 2. An ANOVA on final-path length with factors of Trial 1 WM-load condition, Trial 2 WM-load condition, and trial (first or second) showed no differences in final-path length for either Trial 1 condition, Trial 2 condition, or their interaction, $F(3, 75) = 1.29, p > .25, F(1, 75) = 0.12$, and $F(3, 75) = 0.12$, respectively. In addition, there was no effect of trial, $F(1, 76) = 0.19$, nor any interactions between

trial and either Trial 1 condition, Trial 2 condition, or their interaction ($F_s < 1$). The final-path behavior was thus ubiquitous across conditions.

The constancy of the length of final path for all groups of participants suggests that the effect of the WM load is to extend the preceding exploratory phase. To test this hypothesis, a $4 \times 2 \times 2$ ANOVA with between-subjects factors of first and second trial WM-load condition and a within-subjects factor of trial was performed (as described earlier) with exploratory phase length in moves as the dependent variable. This analysis indicated a significant effect of first trial WM-load condition, $F(3, 68) = 2.83, p < .05$, with no effect of second trial WM-load condition, $F(1, 68) = 0.58$, and no interaction, $F(3, 68) = 1.38$, as was found earlier with total number of moves to solve the puzzle as the dependent variable. The within-subject effect of trial was significant, $F(1, 68) = 20.25, p < .001$, with a marginal interaction with Trial 1 condition, $F(3, 68) = 2.53, p < .07$. Other interactions were not significant ($F_s < 1$).

A post hoc analysis of the reversal-free sequences in participants' move records examined whether the final path occurred as a result of the length of reversal-free sequences progressively increasing over the course of solving the puzzle or whether it indicates a more sudden shift in knowledge or skill akin to the achievement of a sudden insight. The average number of moves between reversals was calculated for the first and second half of the exploratory phase (all moves preceding the final path are considered part of the "exploratory" phase) for each participant. For the first trial, the average consistent sequence in the first half of the exploratory phase was 4.41 (± 0.27) moves and 4.80 (± 0.28) moves for the second half (whereas the final path was 18.4 consistent moves). The difference between the first and second halves is not significant, $t(75) = 1.10, p > 0.25$. On the second trial, the average number of moves between reversals on the first half of the exploratory phase was 5.72 (± 0.40) and 4.94 (± 0.45) on the second half. On the second trial there is a slight (nonsignificant) tendency for the length of reversal-free move sequences to shorten during the exploratory phase, although the final path represents a sequence of consistent moves averaging 19.0 moves in length. The onset of the final path is therefore somewhat abrupt rather than progressive.

The sudden shift from exploration of the problem space to a rapid, error-free final path suggests that participants gained some insight into how to solve the puzzle. However, the retrospective verbal reports collected immediately afterwards gave no evidence of what this insight could have been.

Discussion

Reduction of WM capacity impaired problem solving when the puzzle was solved for the first time. This fact in isolation could imply either a deficit in learning to solve the puzzle or in executing the solution. However, the additional finding that a similar reduction of WM capacity on the second trial no longer impaired problem solving suggests

that the secondary task does not impair executing a strategy for solving the puzzle after it has been acquired (i.e., the secondary task does not create a performance deficit). In addition, the second trial performance was superior to the first trial and equivalent for the four groups of participants, regardless of the WM-load condition of the first trial. Because the four groups of participants appear to be at a similar level of knowledge before solving the puzzle a second time (and, in fact, before initiating the final-path sequence of moves on Trial 1), it may be useful to consider the process of solving the puzzle the first time as gradual acquisition of knowledge up to some criterion. When this criterion is reached, participants execute the final path on the first trial and are prepared to solve the puzzle rapidly a second time. The effect of the first trial WM load is to increase the number of moves that precede reaching this knowledge criterion, suggesting a slowed rate of learning (per move). Because the impairment effected by the secondary task on the first trial increased monotonically with the level of memory load, we can conclude that WM capacity is the critical factor affected by the secondary task rather than other factors such as simply dividing attention over two tasks (which was similar across conditions).

Participants showed the final-path behavior seen in other problem-solving studies (Kotovsky et al., 1985; Kotovsky & Simon, 1990) with a sudden transition to rapid, error-free moves leading to solution. In fact, the length of the final path (18–19 moves) indicates that participants are solving the puzzle with an error-free traversal of almost the entire distance from the start state to the goal (21 moves) on each trial. The length of the final path is not affected by the level of WM load in either trial, indicating that the longer solutions obtained by participants in the higher memory-load conditions reflect a longer exploratory phase, the phase of problem solving during which the problem is learned.

However, despite this evidence that participants seem to have acquired a thorough understanding of the puzzle, the protocols given by the participants after solving the puzzle contained, for the most part, very little information about how the puzzle works or how to solve it. The improved performance on the second trial indicates that participants acquired some significant knowledge about the puzzle when solving it the first time, but 70 of the 80 participants reverted to exploring the problem space again before solving the puzzle the second time. The return to exploratory behavior, together with the inability of many of the participants to describe a strategy for solving the puzzle suggests, for many participants, that the apparently insight-like behavior at the end of the first trial is not the result of any explicitly available knowledge. It is also important to note that the appearance of a final path indicates that the application of implicit knowledge can occur suddenly in problem solving.

The fact that a small number of participants are able to describe elements of their solution strategy and then go on to solve the puzzle very rapidly on the second trial suggests that these participants may be acquiring an explicit strategy for solving the puzzle. A similar finding is discussed in Willingham, Nissen, and Bullemer (1989) using the serial

reaction time (SRT) task. A subset of their participants also developed some explicit knowledge of the task and used this to perform extremely well on the task. However, the majority of their participants showed little or no explicit knowledge, although their problem-solving behavior demonstrated a great deal of implicit knowledge of the task. Further work on this task with patients with anterograde amnesia and matched controls (P. J. Reber & Squire, 1994) showed that explicit knowledge can be epiphenomenal to SRT task performance. We suggest a similar conclusion for the current problem-solving results, namely that although a few participants do acquire some explicit knowledge about the problem (protocol ratings of 3–4.5), the majority of participants are relying exclusively on implicitly acquired information.

One advantage of using a problem-solving task to investigate the process of implicit learning is that it may be possible to determine the content of participants' implicitly acquired knowledge. The task analysis of the puzzle indicates a small set of possible solution strategies for learning to solve the Balls and Boxes puzzle based on either the underlying rules, avoid-reversals, or move-pattern strategies (see Experiment 1, *Materials* section). In the following experiments, the content of participants' knowledge of the puzzle is assessed with further retrospective and concurrent verbal reports, a test of participants' ability to choose correct moves from problem states presented in isolation, and a recognition questionnaire for strategically descriptive statements about the puzzle. These measures are used to both confirm the finding that participants appear to have learned the puzzle implicitly and to attempt to assess the content of participants' implicit knowledge of the puzzle.

Experiment 2

The second experiment both replicates the effects of WM capacity reduction across two solution trials and assesses the content of participants' implicit and explicit knowledge of the puzzle through additional retrospective reports and a move-selection test. The move-selection test, given after completion of two solutions, measures participants' ability to choose correct moves when problem states are presented in isolation, rather than in the normal problem-solving context where information about the recently made moves is available.

Method

Participants

Participants were 30 undergraduates from CMU who received partial course credit for participation.

Materials

Secondary task. Participants in the WM-load condition listened to an audiotape containing letters read at a rate of one letter every 3 s and were required to always remember the last two letters heard. The set of letters to be remembered changed every time a

new letter was presented on the tape, keeping a relatively constant load on WM. Participants were periodically interrupted and required to say out loud the letter before the last letter heard to ensure that they were performing this task. If they did not respond correctly, they were reminded by the experimenter to pay attention to the tape (this happened infrequently).

Procedure

Participants were asked to solve the Balls and Boxes puzzle two times. As in Experiment 1, the WM load (ML) group ($n = 16$) additionally performed the secondary task while solving the puzzle. A control group ($n = 14$) did not perform the secondary task, but were interrupted in their problem solving at the same rate (as the secondary task probes) and asked to name a letter appearing on the computer screen. All participants solved the puzzle twice with the second solution occurring under the same conditions as the first.

Between the two trials, participants were asked to verbally describe the strategies they had used to solve the puzzle and to attempt to describe how the puzzle worked. They were then further prompted to reveal any knowledge they had of the puzzle by being asked to give any information that they felt would help others solve the puzzle. Responses were tape recorded.

After both trials were completed, participants were given a *move-selection* test to determine whether a depicted state of the puzzle could cue the correct choice of a move. Participants were presented with a randomly selected isolated state from the puzzle (i.e., no information was provided about the moves leading up to that position). Participants were then asked which of the two available moves they would select as the next move. An incorrect choice moved one step further from the goal (in the problem space), and a correct choice moved one step closer to the goal. Each participant was presented with every state in the problem space in a random order. It should be noted that the move-selection test was given after Trial 2, in contrast to the protocols obtained after a single solution.

Results

Problem-Solving Performance

Problem-solving performance was measured by the number of moves and overall time taken to solve the puzzle. The mean performance on each measure is presented in Table 2 for both groups. On the first trial, the participants solving the puzzle under a memory load were much slower to solve the puzzle than control participants. On the second trial, both ML and control groups solved the puzzle at a similar rate, considerably faster than either group on the first trial. A 2×2 ANOVA of trial by condition, with number of

moves as a repeated measure, indicated a significant effect of trial, $F(1, 28) = 32.73, p < .001$, and a significant interaction between trial and group, $F(1, 28) = 6.22, p < .02$. A similar analysis examining overall time to solve the puzzle found a similar pattern of results, with a significant effect of trial, $F(1, 28) = 24.45, p < .001$, although the interaction between trial and group is just marginal, $F(1, 28) = 4.19, p < .051$. This overall pattern of performance is identical to the result obtained in Experiment 1; the secondary task interfered with Trial 1 performance but had no effect on Trial 2 performance, and there was considerable improvement in performance between Trial 1 and Trial 2.

Figure 5, which shows some sample participants' progress toward solution, illustrates two important aspects of participants' problem-solving behavior. These graphs were selected to present a variety of behavior and overall times to solve the puzzle. Each of these graphs shows, for a single participant, distance from the goal on the y-axis (minimum moves to solution) with number of moves along the x-axis (left to right). Because the problem space is linear, at every move each participant either moved one state closer to the goal or one state further away.

Two prominent features are illustrated in these progress graphs. First, several *barrier areas* that contribute to the difficulty of the puzzle are apparent. In these areas (States 31–26 and States 15–10 in particular), participants must replace a number of balls to make progress (e.g., going from 4 balls out to 1 ball out from States 15 to 10). Participants make a large number of reversals in these areas of the problem space indicating that they may initially be focused on getting as many balls out as possible (a logical hill-climbing strategy based on the stated goal of removing all five balls). Figures 5B and 5E show participants who encountered difficulty around State 15, several times coming down to this state and electing to reverse direction. Figures 5A, 5C, and 5D show participants who encountered difficulty in the barrier area near the start state when it was approached from the top of the problem space. Figure 5F provides an example of the sudden onset of a maximal length final path.

The final-path behavior of participants is similar to that found in Experiment 1. The overall average length of final path was 16.2 (± 1.9) moves on the first trial; 18.0 (± 2.8) moves for the control group and 14.6 (± 2.6) moves for the ML group. Similarly, on the second trial, the mean length of the final path was 18.5 (± 1.6) moves overall; 17.3 (± 3.0) for the control group; and 19.5 (± 1.6) for the ML group.

Table 2
Problem-Solving Performance, Experiment 2

Group	Trial 1		Trial 2	
Memory load ($n = 16$)	161 moves (29)	339 s (78)	41 moves (5.9)	64 s (12)
Control ($n = 14$)	94 moves (12)	192 s (39)	49 moves (8.6)	82 s (23)

Note. Numbers in parentheses in the Group column indicate number of participants in each group. Numbers in parentheses in the Trials 1 and 2 columns represent one standard error of the mean.

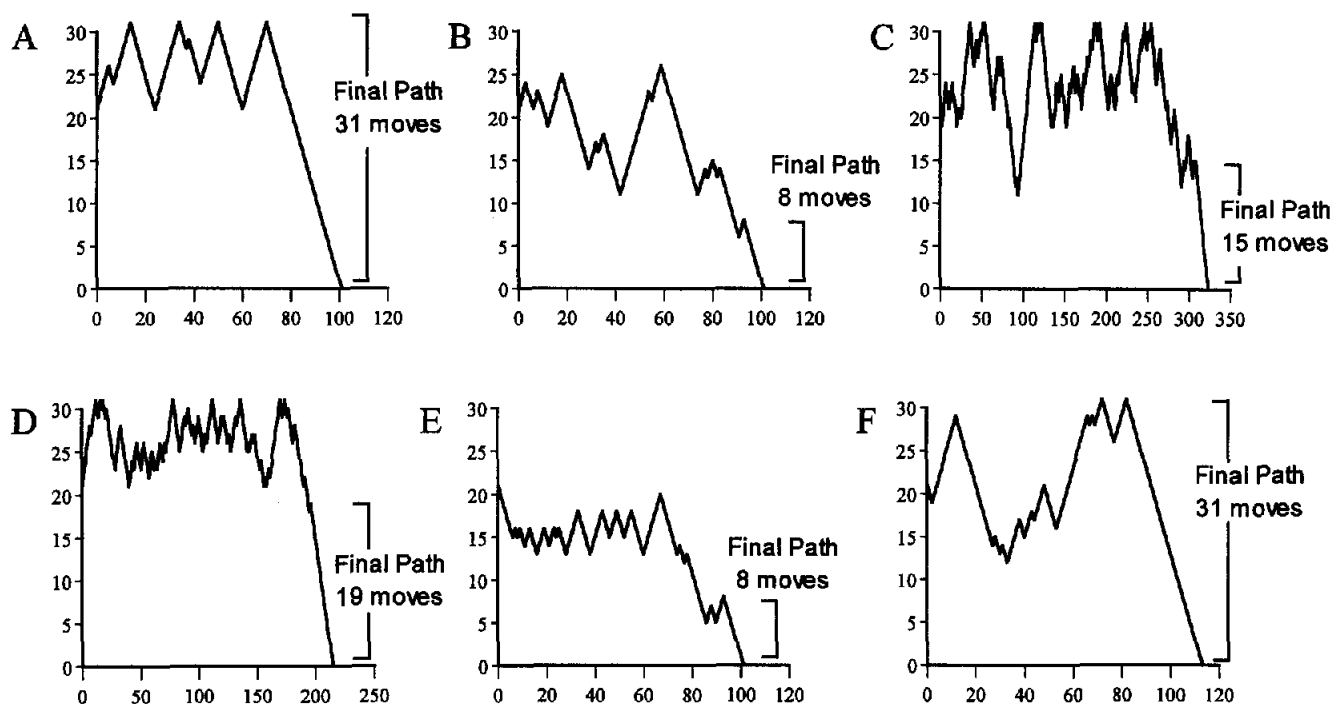


Figure 5. Sample move records (with final path), Experiment 2. Distance from the goal is measured on the y-axis, and move number (for a particular participant) is measured on the x-axis. When the distance reaches zero, the puzzle has been solved. Changes in slope indicate reversals in direction in the problem space.

Verbal Report

Unfortunately, it was impossible to fully analyze the details of the participants' verbal responses due to poor recording conditions and the fact that participants used many gestures in describing the puzzle (making it difficult to reconstruct statements such as "then I moved this one . . ."). However, participants were generally unable to say almost anything of substance about how the puzzle was solved. This stands in contrast to the fact that they had just completed solving the puzzle once, traversed a lengthy and error-free final path, and in almost every case were about to solve the puzzle a second time even more rapidly than the first. A representative example follows (experimenter in italics):

Now I want to ask you about the puzzle you just solved, how it worked, what you did.

No idea.

No idea?

No idea. It was very painful. (Laughs)

You did get it, right?

Yeah, but it was basically luck . . . that I got it.

You had no idea what you were doing?

Not really.

Suppose somebody else was going to do the puzzle who had never seen it before and you had to give them some hints, tell them how to solve it.

Well, let's see . . . I don't know what to say . . . but, I guess, (garbled word) the puzzle . . . The good part was that there

usually wasn't any more than like, one or two choices. I think there was one choice. Was there any more than one choice? I don't know. But I had (garbled word). Which is why I kept ending up back where I started from, which was frustrating. I would tell them, I would tell them, good luck. That's all.

This participant insisted he had no idea how he solved the puzzle, but eventually managed to recall that the choices at each state were somewhat limited. This participant solved the puzzle in 21 moves (the minimum number of moves) on the second trial, immediately after giving this fairly uninformative description of his process of solving the puzzle. The second trial performance of this participant was superior to most of the other participants, but the verbal protocol given by this participant was typical in its lack of information about how the puzzle was solved.

Move-Selection Test

In spite of the marked improvement in performance on the second solution trial, participants were only correct on a mean of 55% ($\pm 3.3\%$) of the trials on the move-selection test (immediately following the second solution), a rate not significantly different from chance (50%). There was no difference in performance on the move-selection test between the two groups, $t(29) = 0.70, p > .40$. Participants' inability to select the correct move from a problem state presented in isolation (i.e., in the absence of any cues about

what the previous move had been) indicates that at least some information about previously made moves is necessary for participants to be able to choose the next move correctly.

Discussion

As in Experiment 1, participants' problem-solving performance was impaired by a WM load the first time they solved the puzzle, but the WM load did not create a performance impairment the second time the puzzle was solved. Once again, the lack of an effect of WM capacity reduction on performance of the second solution indicates that the secondary task does not create a performance deficit. The effect of WM capacity reduction on Trial 1, together with the similar performance of both groups on Trial 2, again suggests that the participants solving the puzzle under a WM load are taking longer than the control participants to achieve a similar level of knowledge before the second trial. We therefore conclude that the effect of the secondary task is to slow learning on the first trial.

The location of the barrier areas in the exploratory phase indicated that participants initially approached the problem using a hill-climbing strategy (trying to get as many balls out as possible) but at some point converged on an effective strategy that led to a long final path. However, the strategy used to accomplish this final path was not reportable verbally even though participants were asked to do so immediately after it occurred.

Using an indirect measure, the move-selection test, it was possible to determine that participants were not learning to derive the next move based on the current state of the puzzle. Therefore, it appears that participants need information about the moves leading up to a state in order to correctly choose the next move (as they did effectively in the two preceding solutions). This finding argues against the rules strategy (discussed earlier) and is consistent with both the avoid-reversal and move-pattern strategies that are based on using contextual information (previous move or moves leading to the current state). The lack of reportable knowledge relevant to the avoid-reversal or move-pattern strategies in either this experiment or Experiment 1 suggests that whichever of these strategies is implemented by the participants, it is learned and represented implicitly (i.e., without awareness).

Although the fact that participants are unable to report their task knowledge is considered a hallmark of implicit learning (cf. Seger, 1994), research on implicit learning that relies solely on retrospective reports has been challenged. Several researchers have claimed that learning does not ever proceed without awareness and that the apparent lack of explicit knowledge is due to insensitivity or inaccuracy in the methods by which explicit knowledge is assessed (e.g., Perruchet & Amorim, 1992; Shanks & St. John, 1994). Retrospective verbal report has been particularly questioned as a method for accurately assessing explicit knowledge. To address this concern, in the next experiment a recognition-

oriented measure is used to attempt to measure participants' explicit strategic knowledge of the puzzle. Recognition is generally more sensitive to explicit knowledge than recall, and recognition-oriented questions have been successful in identifying when explicit knowledge has been acquired in sequence-learning tasks (e.g., P. J. Reber & Squire, 1994; Willingham, Greeley, & Bardone, 1993). By asking participants whether they recognize important strategic statements about solving the puzzle, we should be able to discover the content of participants' explicit strategic knowledge of the puzzle, if any is available.

Experiment 3

Participants' knowledge of puzzle-solution strategies was assessed after two solutions of the puzzle by asking them to indicate which of a set of statements would be helpful to someone else who had to solve the puzzle. By including statements that describe effective solution strategies or partial strategies (e.g., never undo a move, or move the right-most ball every other time), we can determine if participants can recognize these strategies as being useful in solving the puzzle and also whether participants may have been utilizing these strategies in spite of the fact that these strategies have not been reported in their retrospective protocols.

Participants' ability to use features of the current problem state to select moves was also retested using a move-selection test that extends the technique used in Experiment 2. After solving the puzzle twice (without any secondary task), participants were presented with different starting positions from which to work toward a solution. However, unlike Experiment 2, where participants chose a single move, participants were required to make a series of moves in an attempt to solve the puzzle. This allowed us to both replicate the result of Experiment 2, that participants have difficulty choosing the first move, and additionally to test whether participants could later correct for having made a first-move error by using the context provided by the moves they have selected. As before, if participants have learned the rule strategy, they should be able to choose the first move correctly. In addition, after an occasional erroneous first choice, the rule strategy could be used to correct this error by deducing the move that reverses direction in the problem space in order to move toward the goal. If, on the other hand, participants are solving the puzzle by learning either the move-pattern strategy or the avoid-reversal strategy but cannot determine the first move to make from the presentation of an isolated state (because no contextual information is available about the previous moves), then they should be essentially at chance on choosing the first move from a randomly selected problem-space position and then make few or no reversals in subsequent moves, showing no sensitivity to whether the direction of movement in the problem space is toward or away from the goal.

Method

Results

Participants

Fifty-one CMU undergraduates received partial course credit for participating in this study. Three participants were eliminated because the puzzle reminded them of other puzzles they had previously solved that were isomorphic to this puzzle.

Materials

The puzzle was presented on a Macintosh (Apple Macintosh IIcx, Apple Computers, Cupertino, CA) personal computer in the standard five-balls-and-five-boxes format. Instructions were presented on the screen of the computer. Participants used the computer mouse to indicate when they were ready to solve the puzzle and also to manipulate the balls in order to solve it. The instructions given to participants were the same as in Experiment 1.

The written questionnaire consisted of 25 strategy-oriented statements about the puzzle. These statements were mainly derived from our task analysis of the puzzle (specifically, the possible solution strategies) and also from participants' statements in the verbal retrospective protocols from Experiment 2. Participants were asked to judge how helpful each of these statements would be to somebody else who had to solve the puzzle. All statements were rated on a scale from 1–7, where 1 indicated the statement would not be very helpful and 7 indicated that the statement would be very helpful. For the last 15 statements, participants were asked to decide whether the statement was true or false in addition to rating how helpful it was (the first 10 statements cannot be properly judged to be true or false). The statements are listed in Appendix A (together with the strategies consistent with each statement).

In the move-selection test, participants were instructed to attempt to solve the puzzle from a randomly chosen starting state. Each participant received either 8 or 9 such trials in the posttest. Each trial started with the presentation of a different starting state, and participants chose moves until either the puzzle was solved or 10 moves were made. State 0 (the goal) and State 31 (the top of the problem space) were never given as starting states. The 29 potential starting positions were broken into four groups, each containing 7 or 8 states, and each participant received one group of starting states. The standard initial state, State 21, was given to every participant as the final trial (in addition to 7 or 8 other states).

Procedure

Group A ($n = 27$). Participants solved the Balls and Boxes puzzle twice, followed by the move-selection posttest and then the written questionnaire about the puzzle.

Group B ($n = 24$). To provide a baseline for interpreting the results from the postsolution questionnaire, participants completed the written questionnaire about the puzzle before solving it. Participants were shown a diagram of the starting state together with the task instructions and told to imagine they had just solved the puzzle and to provide ratings as best they could. Participants then solved the puzzle twice and then completed the written questionnaire a second time.

Problem-Solving Performance

Group A. Participants solved the puzzle the first time in an average of 116 moves ($SE = 19$) and showed considerable improvement to an average of 49.4 ($SE = 7.4$) moves on the second trial. This improvement was significant, $t(26) = 3.65$, $p < .01$.

Group B. The mean number of moves to solve the puzzle was 87.6 ($SE = 20.2$) on Trial 1 and 28.7 ($SE = 2.7$) on Trial 2. Improvement from Trial 1 to Trial 2 was significant, $t(23) = 3.03$, $p < .01$.

The problem-solving performance of Group B was numerically superior, although a 2×2 (Groups \times Trials; within-subject) ANOVA indicated that the group effect is not reliable, $F(1, 49) = 2.90$, $p < .10$, although the effect of trial is reliable, $F(1, 49) = 22.40$, $p < .001$, with no interaction between group and trial, $F(1, 49) < 1$. The tendency for participants in Group B to solve the puzzle more rapidly might reflect some slight advantage of seeing the written questionnaire before solving the puzzle.

Strategy Recognition

Groups of statements consistent with each of the three solution strategies—move-pattern, rule, or avoid-reversal—were constructed. A statement could be consistent with more than one strategy (e.g., "Move the rightmost ball first" is important for both the move-pattern and avoid-reversal strategies. Statements that provided information contradictory to a strategy (e.g., "Move the second ball from the right first") are included in the relevant strategy statement groups after inversion of the assigned rating. By inverting the ratings given to contradictory statements, higher aggregate scores always reflect more knowledge shown by participants about the groups of statements. An "Other" group was constructed of all the statements not consistent with any of the three strategies. Appendix A lists the statements and the strategies that are consistent with each. For each participant, a score was calculated based on aggregating the ratings given to all statements consistent with each of the three possible strategies (with ratings of contradictory statements being inverted such that the inverted rating equals 8 minus the rating, so that a rating of 1 is inverted to be 7 and vice versa). Each aggregate score reflects the perceived helpfulness of a group of statements with a higher score indicating that participants felt that the statements would be more helpful to another hypothetical participant. These scores were calculated for Group A from ratings given after solving the puzzle and for Group B both before and after solving the puzzle. The ratings given by Group B from ratings given before solving the puzzle provide a baseline for assessing how the aggregate ratings changed with experience in the puzzle. The mean rating scores are presented in Figure 6.

The main comparison of interest is between the baseline ratings (Group B before solving the puzzle) and the ratings provided by Group A after solving the puzzle. The average rating given for statements consistent with using the move-

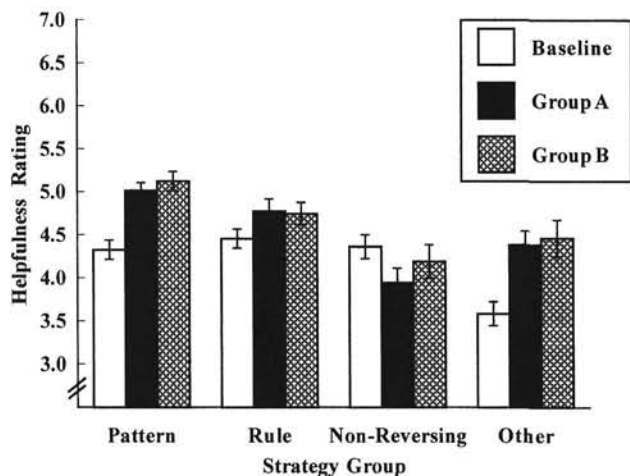


Figure 6. Mean ratings given to statements grouped by strategy, Experiment 3. Higher ratings indicate greater perceived helpfulness for a group of statements. The baseline group reflects ratings provided by Group B participants before solving the puzzle. Ratings for Group A and Group B participants were provided after two solutions of the puzzle.

pattern strategy was significantly higher for Group A than in the baseline condition, $t(49) = 4.71, p < .01$. The mean rating increase for rule strategy statements was only marginal, $t(49) = 1.80, p < .08$, and the mean rating for statements relevant to the avoid-reversal strategy exhibited a trend toward decreasing, $t(49) = -1.92, p < .07$.

Although the ratings provided by Group B after solving the puzzle could potentially have been contaminated by exposure to the statements before solving the puzzle, there were no differences in the postsolution ratings provided by Group B and Group A in any category, $t(49) < 1.05$. Similarly, for Group B, ratings of statements consistent with a move-pattern strategy increased significantly, $t(46) = 5.02, p < .01$. Ratings for statements relevant to a rule strategy increased marginally, $t(46) = 1.72, p < .10$, and although the ratings for the avoid-reversal statements decreased numerically, this effect was not reliable, $t(46) = 0.71$.

The decrease in ratings given to statements consistent with the avoid-reversals strategy provides evidence that participants are not explicitly using this strategy to solve the puzzle. The increase in ratings of perceived helpfulness for the statements consistent with the move-pattern strategy suggests the acquisition of some explicit knowledge about the puzzle. However, the statements about the move pattern are also generally good descriptions of the solution (e.g., "Move the balls on the right more frequently than balls on the left," "Put balls on the right back in, but not balls on the left.") The retrospective protocols of Experiments 1 and 2 suggest that knowledge of how to solve the puzzle is implicit, but do not preclude the fact that participants remember aspects of the problem-solving episode and may be able recognize statements that accurately describe the solution path.

Ratings given to statements that were not consistent with any of the three strategies also increased relative to the baseline ratings, $t(49) = 3.79, p < .01$ for Group A, and $t(46) = 3.47, p < .01$ for Group B. However, the ambiguous nature of these statements (e.g., "Move as quickly as you can") and the inclusion of general weak heuristics (e.g., "Use trial and error") make it difficult to conclude that this change in rating reflects acquisition of useful knowledge about the puzzle and might simply reflect a tendency for the helpfulness ratings given after solving the puzzle to increase.

Move-Selection Posttest

Each participant encountered either 8 or 9 trials in this posttest depending on which of the four groups the participant was randomly assigned to (three groups required 8 trials, one required 9 trials in order to encompass all 29 starting positions). There was no difference in performance across the four groups as assessed by a one-way ANOVA on average percentage correct for each participant, $F(3, 23) = 0.95$. A total of 222 trials were administered. Participants chose the first move correctly 116 times, a rate of 52%. Because there are always two moves to choose from, this represents essentially chance performance (as in Experiment 1) and is far too low to account for final-path performance.

The moves selected after the first move were examined to determine if participants were able to correct for an erroneous first move. Second move-selection performance was compared in cases where the first move was correct with those for which it was incorrect. After choosing the correct first move (toward the goal), participants made a reversal on the second move 13% of the time (88% of the second moves chosen were in the same direction as the first). When the first move was chosen incorrectly, participants immediately made a reversal 12% of the time. Participants were thus equally likely to reverse on the second move when the first move is correct as when the first move is incorrect and therefore exhibit no directionally sensitive error correcting on the second move.

To examine whether participants were correcting errors at any point during the nine moves following the first, each of the 176 reversals was classified as either (a) occurring at the top of the problem space (State 31) where a reversal is forced: 45 (25%); (b) correctly changing direction in the problem space: 62 (35%); or (c) incorrectly changing direction in the problem space: 69 (39%). Participants therefore had no tendency to correct their erroneous direction in the problem space as they made erroneous reversals (c) slightly more frequently than correct reversals (b).

In general, participants made very few reversals during the posttest. On 119 of the 222 trials (54%), participants made no reversals at all. Participants made one reversal in 62 trials (28%) and two or more reversals in 41 trials (19%). Overall, participants made reversals on 8% of their moves. This rate was roughly similar to the reversal rate of 12% seen on participants' Trial 2 solution performance. These

results indicate that participants generally made the first move selection essentially at chance and then chose nonreversing moves without being sensitive to the direction of their movement within the problem space (i.e., to correct for a first-move error). These results also replicate our finding in Experiment 2 where we found that participants could not use their knowledge of the puzzle to correctly select the first move from a problem state presented in isolation. This result is extended further to show that once this first move is made, participants are able to use their knowledge of the puzzle to the extent that they choose nonreversing, although not necessarily correct, moves.

Discussion

Participants' improved performance on the second problem-solving trial demonstrated that they acquired significant knowledge of the puzzle from their experience on the first trial. Participants who had been preexposed to the recognition questionnaire (Group B) tended to solve the puzzle slightly more quickly on the first trial, suggesting that seeing the list of potentially useful strategic statements aided them. However, even if it were reliable, this result would not indicate that other participants (e.g., Group A) were acquiring an explicit solution strategy, merely that solution time could be reduced by giving the participants some explicit knowledge. Certainly, it would be possible to teach participants how to solve the puzzle explicitly (e.g., by instructing them to use the avoid-reversal strategy. This fact is independent of the fact that participants who learn the puzzle without instruction appear to do so implicitly.

Two tests were used to further assess participants' knowledge of the puzzle and problem-solving episode: a strategy statement questionnaire and a move-selection test. Participants' responses to the strategy statement questionnaire were evaluated by comparing "helpfulness" ratings given after solving the puzzle to ratings given by participants who had not yet solved the puzzle. Statements consistent with a move-pattern solution strategy were rated as likely to be helpful, and these ratings increased over the baseline ratings for these statements. Statements consistent with conscious use of an avoid-reversal strategy were rated as not being helpful and significantly decreased relative to the baseline ratings. Statements consistent with a rule-based strategy increased numerically, but not reliably. It is important to note, however, that the move-pattern related statements that were rated by participants as being helpful also include many piecemeal descriptions of the solution path (e.g., "Move balls on the right more frequently than balls on the left"). Therefore, higher ratings given to these statements might reflect explicit memory for the end of the problem-solving episode rather than explicit knowledge of a comprehensive strategy used to solve the puzzle. It is also worth noting that the increase in ratings given to these statements, although reliable, is not dramatic. The improvement in performance from Trial 1 to Trial 2 in Experiments 1, 2, and 3 is sizeable, and virtually all participants demonstrate a long final path to solution. The small degree of knowledge

expressed on the recognition test does not appear to be sufficient to account for the knowledge of the puzzle exhibited by participants through performance. In Experiment 4, we attempt to further investigate the issue of conscious strategic knowledge by assessing participants' puzzle knowledge as they solve the puzzle using a concurrent protocol design.

The result of the move-selection test replicated the result of Experiment 2 that participants were not able to consistently select the correct first move from a problem state presented in isolation. Of the three candidate strategies, only the rule-based strategy would permit successful performance on this test. Using the rule strategy, it is possible to deduce the correct next move from any problem state. The other two strategies both require some knowledge of the previous move or moves to determine the next move. Because use of a rule-based strategy would allow good performance on this test, which was not exhibited by participants, it appears that participants did not learn this strategy. In addition, participants' tended to avoid reversing on this test after making the first move (whether or not reversing would be correct), which is consistent with using either a move-pattern or avoid-reversal strategy. Because this test is based on performance rather than attempting to report conscious knowledge, the poor first move-selection performance (and lack of subsequent correction) suggests that participants are not learning the rule strategy either implicitly or explicitly: It appears that they simply do not know (at any level) this strategy. In contrast, the strategy statement questionnaire assessed explicit knowledge of the various strategies and indicated that participants are not consciously using the avoid-reversal strategy, although it is possible that they could be using this strategy implicitly (i.e., without being aware of it). Participants therefore appear to be learning to solve the puzzle based on either an implicitly acquired avoid-reversal strategy or a move-pattern strategy. The move-pattern strategy, if it is the source of participants' problem knowledge, would seem to be implicitly represented based on the retrospective protocol results. However, the strategy statement questionnaire suggests the possibility that participants might have some explicit knowledge of the move pattern.

In Experiment 4, an attempt is made to further disambiguate among these hypotheses by assessing participants' development of a problem-solving strategy while exploring the problem space. To do this, participants are instructed to give a "think-aloud" or concurrent protocol while solving the puzzle. If participants are developing some explicit strategic knowledge (e.g., of the move-pattern strategy), they should be able to report some information about the development of this strategy as they solve the puzzle.

Experiment 4

In other implicit learning tasks in which learning occurs over a number of short trials, it has not been possible to assess participants' awareness of their learning as the task is performed. Because the Balls and Boxes puzzle differs in

that it is solved over several minutes and a sizeable number of moves, it provides an opportunity to assess participants' conscious thoughts on-line using the technique of concurrent verbal-protocol analysis. Another important issue in the study of implicit learning that is addressed in this experiment is whether participants are aware of knowledge that is being acquired implicitly as that knowledge is acquired or whether this knowledge is simply difficult to report later (e.g., Shanks & St. John, 1994).

The concurrent protocol technique has previously been used impressively to determine the sequence of strategies employed to solve the Tower of Hanoi puzzle (Anzai & Simon, 1979) as well as many other problem-solving tasks (Ericsson & Simon, 1993). By collecting as much information as possible about participants' thought processes while solving the puzzle, we anticipated either being similarly able to determine the strategy used by participants while solving the Balls and Boxes puzzle or to demonstrate that participants were unaware of their knowledge of the puzzle, even as it was being acquired.

A small group of participants were videotaped under the instructions to "teach aloud" while solving the balls and boxes puzzle. Instructions to teach aloud differ from standard "think aloud" instructions (Ericsson & Simon, 1993) in that participants are asked not only to report the contents of WM but to also give some insight into the strategies and approach that go into choosing operators (balls to move). The advantage of this approach is to get an idea of why participants are choosing the moves they choose at or very near the time of choice. The disadvantage is that because this is not a normal part of problem solving, it may affect problem solving. Because the focus of this experiment is to gain any possible insight into the process of learning to solve this puzzle, we decided the possibility of obtaining additional information about how participants made decisions was worth the potential risk of modifying participants' solution strategies.

Method

Five CMU undergraduates (P1-P5) participated in the experiment and received partial course credit. Participants were informed that they would be videotaped doing some problem solving and were instructed to attempt to give reasons for (justify) each move

as well as to report the contents of WM. Participants were first given a warm-up task consisting of mental multiplication while talking aloud (as recommended in Ericsson & Simon, 1993, to orient them to the protocol method). Participants were then asked to solve the puzzle while talking aloud. Participants were prompted by the experimenter to verbally describe the reasons for choosing their moves in cases where a move was made without an accompanying verbal report by the participant. After one solution, participants were asked how they had solved the puzzle and what advice they could give to another person who had to solve the puzzle. Participants were then asked to solve the puzzle a second time while still talking aloud.

Results

The protocols were analyzed in an attempt to determine what types of rules or justifications participants gave for making the moves that they chose. If the transition to final path is the result of participants acquiring a new strategy for the puzzle, we should be able to identify the acquisition and use of that strategy during problem solving and also determine whether the strategies used by participants change across the two trials. To determine what strategies were employed by these participants, we attempted to determine the rationale behind as many of each participant's moves as possible. In this way we could identify whether new rules appeared as participants gained experience with the puzzle.

Three general heuristics were commonly used by each of the participants to guide move selection (*A, D, B*): Avoid-reversing (i.e., avoid undoing the previous move); identify a *Difficult* ball to focus on; and "Back up and look for a different way." The use of these heuristics at the beginning and end of each problem-solving trial is shown in Table 3 for each of the 5 participants. These heuristics and statements such as, "I'm going to move this one to see what happens," account for the majority of justifications given by participants.

There were also a few specific hypotheses and rules that participants (P1-P5) generated and used during problem solution (all hypotheses listed next were generated on the first trial unless otherwise stated):

P1 first hypothesized that moving Ball 5 causes boxes to open and close. Later in Trial 1, he stated that Balls 4 and

Table 3
Participant Performance and Reported Heuristic Use in Experiment 4

Participant	Trial 1: Moves	Trial 1: Early			Trial 1: Late			Trial 2: Moves	Trial 2: Early			Trial 2: Late		
		A	D	B	A	D	B		A	D	B	A	D	B
1	146	+	+	+		+		37	+	+			+	+
2	148	+	+	+		+		47		+			+	
3	29	+	+		+	+		39	+	+		+		
4	33	+	+	+		+	+	128	+	+	+		+	+
5	67	+	+	+		+		97		+			+	

Note. The three move-selection heuristics are designated *A, D, and B*: Avoid undoing the previous move; identify a "Difficult" ball; Back up and look for another way, early (first half) and late (second half) in each of the two problem-solving episodes. The "moves" column indicates how many moves were taken to solve the puzzle on the first and second trials. Use of a heuristic is noted by a plus sign (+). A blank space indicates that a heuristic was not used by that participant during the indicated time period.

5 act as “triggers” and subsequently that Ball 5 is “a key piece.” These statements are all essentially true. P1 further elaborated that opening Box 3 depends on Ball 4 being in and that Box 4 is open when Ball 5 is in (both true). These rules represent some progress towards deducing a means–ends strategy for solving the puzzle. However, no rules were suggested for either Ball–Box 1 or 2, which P1 recognized to be the most difficult pieces to remove. Therefore, it seems that his description of the means–ends strategy falls short of his successful problem-solving behavior. For example, in spite of the repeated prompting to explain his strategic approach, P1 does not mention using these rules to do any move sequence planning on the second trial.

P2 stated the configuration required to open Box 2 when it occurred. A few moves later, when attempting to replace Ball 2, this configuration is recalled. However, on the second trial, when it was necessary to remove Ball 2, P2 could not recall the rule (but managed to open the box in spite of this). P2 also tried to examine every possible configuration in order to generate a rule for removing Ball 1 (which had been identified as difficult) but lost track of the configurations tested. Later, when Box 1 opened, P2 stated the required configuration. Remembering the required configurations for opening these two boxes would have allowed a means–ends strategy for solving the puzzle, but only if rules were known for achieving these configurations. This participant never stated any rules describing the movement of Balls 3–5. In addition, the only planning of any move sequences observed occurred during the beginning of Trial 1, when P2 would plan to remove both Balls 4 and 5 at State 29 (when both boxes are open). This plan was accomplished but resulted in revisiting State 31 repeatedly (which is counterproductive).

P3 generated no hypotheses during the first trial (other than mentioning the heuristics described earlier). On Trial 2 P3 suggested that Ball 4 moving is linked to Box 2 opening. This is partially true, as Ball 4 needs to be out (and Ball 3 in, which was not mentioned) as well as Ball 5 in order for Box 2 to open. Later the ability to move Ball 3 is similarly linked to Ball 5 moving (Ball 5 must be out and Ball 4 in to move Ball 3). These statements occurred during a fairly rapid second trial during which no planning of move sequences was described.

P4 mentioned no relevant hypotheses other than using the three heuristics described earlier.

P5 made two true observations about the puzzle, that when Balls 4 and 5 are in, both boxes are open and that moving Ball 5 out opens Box 1 (although Balls 2, 3, and 4 need to be in a specific configuration for this to occur). P5 then began to consider a series of hypotheses based on the idea that moving and then immediately undoing that move (e.g., move Ball 4 out then immediately back in) resulted in a different problem state (i.e., the current state depends not only on the current configuration of the balls, but also the specific sequence of recently made moves). This is untrue in that in the Balls and Boxes puzzle, the current configuration of the balls completely determines the problem state. However, P5 was able to develop an idiosyncratic rule for removing Ball 3 that required a reversal. This rule was used

to plan a successful move sequence for removing this ball. No rules were stated for moving Balls 1 or 2 (which were noted to be difficult) on either trial. Late in the second trial, however, the idiosyncratic rule was apparently forgotten and was replaced by a variant that is incorrect. Use of this new, incorrect rule led to a number of extra moves being made before solution.

Extended summaries of the problem-solving episodes for all 5 participants are presented in Appendix B.

The improvement in performance across the two problem-solving trials was not as pronounced as in the previous three experiments. Only 2 of the 5 participants showed a decrease in the number of moves required to solve the puzzle (P1 and P2). In addition, the final-path length exhibited by these 5 participants was somewhat shorter than previously seen, 7.0 (± 1.1) moves for Trial 1 and 13.0 (± 4.7) moves for Trial 2. Although there are only a small number of participants in this experiment, it is possible that generating the concurrent protocol interfered with participants' learning to solve the puzzle or expression of this knowledge on the second trial.

Discussion

Every participant initially used the heuristic *A* (avoiding reversing) at some point early in solving the puzzle. However, further use of this heuristic at the end of the first problem-solving trial or during the second trial was fairly rare. Only 1 participant used this heuristic to explain his move selection at the end of either trial. It should also be noted that no participant made any statement indicating use of this heuristic as a general principle (e.g., “I’m going to try to avoid undoing any moves”) but only invoked the heuristic in specific instances. Because of the linearity of the problem space, if this heuristic was applied consistently in selecting moves, participants would solve the puzzle very quickly with at most one reversal. Because this heuristic is only apparently used early in solving the puzzle, when many reversals are being made (in spite of the occasional use of this heuristic), we can conclude that participants are not consciously relying on this strategy when they begin the final path to successful problem solution.

The second commonly used heuristic, *D* (identify a difficult ball), represents a step towards development of a means–ends–based strategy for solving the puzzle. A ball identified as one that is difficult to remove can be used as a subgoal that presents a specific subproblem to be solved en route to solving the puzzle. However, for a means–ends strategy to be effective, it is necessary to be able to solve the subproblems that are identified (i.e., identify the ball to be removed and then execute the moves required to open the box that contains the ball of interest). Although some of the participants appeared to begin to determine some of the dependencies between the balls and boxes necessary to solve subproblems identified by using this heuristic, none of the participants made statements indicating sufficient knowledge of these dependencies to successfully use a

means–ends strategy. In addition, use of a means–ends strategy often requires planning a sequence of moves for solving a subgoal. Only P5 did any move–sequence planning based on an acquired rule (briefly using an idiosyncratic rule that required 2 extra moves and a reversal to accomplish opening Box 3). This participant, however, did not plan any move sequences for removing either Ball 1 or 2, both of which he had noted to be difficult to remove.

The use of heuristic *D* throughout both problem-solving trials by virtually all of the participants suggests that although participants attempt to use a means–ends strategy, they meet with little or no success. This conclusion is consistent with the results of the move–selection tests used in Experiments 2 and 3, which indicated that participants were not able to implement the means–ends analysis necessary to choose the correct first move from a problem state presented in isolation.

The third commonly used heuristic, “back up and look for another way” (*B*), is a generally useful search heuristic for a large and branching problem-space problem but is actually counterproductive for this puzzle because the search space is so constrained (linear). This heuristic leads participants to make reversals that can interfere with solving the problem (e.g., if no reversals are made voluntarily, the problem can be solved relatively quickly). The frequent use of this heuristic reinforces our previous conclusion that participants are not solving the puzzle by intentionally avoiding reversals. It also indicates that participants are not developing an accurate model of the problem space (i.e., they are unaware of its linear character).

Overall, the protocols are consistent with the retrospective reports, especially in that several participants indicated that they knew how to solve the puzzle, but could not describe it well. Although some participants did seem to be making progress towards developing a means–ends, rule-based strategy, they were not successful enough in this for it to account for their final path. The concurrent protocols indicated a set of general problem-solving heuristics that participants used when initially exploring the puzzle. Over the course of solution, however, there appeared to be a marked reduction in the use of these heuristics without the development of a reportable strategy. Thus, participants did not seem to be aware of the acquisition of a problem-solving strategy while solving the puzzle. Therefore, we conclude that participants are learning an implicit strategy for solving the puzzle by acquiring the pattern of moves required to solve the puzzle or gradually learning to avoid making reversals without being aware of this tendency.

The relatively poor learning exhibited by participants in this experiment raises the possibility that participants were not learning the same information as in previous experiments. The concurrent protocol may interfere with either learning a strategy for solving the puzzle or with performance of this strategy. One intriguing idea is that the instructions to provide a detailed concurrent protocol encouraged participants to perform more explicit processing than is normally attempted (i.e., in previous experiments) and prevented the implicitly acquired information from being expressed. A. S. Reber (1976) has shown that implicit

learning of artificial grammars is impaired when instructions to explicitly determine the underlying rule structure are given. Further work is required to determine whether this instructional effect is reliable for the Balls and Boxes puzzle.

It is worth noting that the reportable content of participants' puzzle knowledge is very different in this experiment than that found in Anzai and Simon's (1979) study examining concurrent verbal reports obtained from a participant solving the Tower of Hanoi (TOH) puzzle. In their reports, the strategy used by the participant each time the puzzle was solved was clear enough to track changes in solution strategy over the first three times their participant solved the puzzle. The ability of that participant to report a problem-solving strategy for the TOH puzzle suggests that knowledge of that puzzle acquired during problem solving is explicit (in contrast to the implicitly acquired strategy for the Balls and Boxes puzzle) and is concurrently reportable during the problem-solving episode.

An important question raised by the contrast between the results of this experiment and the findings of Anzai and Simon (1979) is, What are the critical characteristics of the two puzzles that cause them to differentially afford implicit or explicit learning? To definitively answer this question, it is necessary to systematically vary aspects of the puzzles to attempt to encourage different types of learning, implicit and explicit. However, a critical difference between the puzzles is the fact that for the Balls and Boxes puzzle, participants do not initially know what the problem state will look like after a move is made (they are not given any specific information about how the boxes might open and close). In the TOH, it is at least theoretically possible to plan long sequences of moves immediately and even potentially deduce an effective problem-solving strategy before any exploration of the problem space is done. The lack of deducible information about subsequent problem states in the Balls and Boxes puzzle may, for example, simply make explicit deduction of a rule-based strategy extremely difficult and therefore afford implicit learning of the move-pattern strategy. This hypothesis could be tested by constructing an isomorph of the TOH that does not lend itself to planning move sequences (e.g., by not informing participants of the move constraints in advance) and examining whether it is possible to learn to solve this new puzzle using the highly patterned sequence of moves necessary to solve the TOH. Alternately, the issue could be explored by constructing an isomorph of the Balls and Boxes puzzle that provides enough structural information to predict what moves are available after the current move is made and thereby enable planning of move sequences. Kotovsky and Simon (1990) report one such isomorph, their “all-info” puzzle, that could have afforded planning move sequences. Although they did not specifically investigate the degree of explicit (or implicit) strategy use, they did find much longer move latencies for this puzzle than the more standard Balls and Boxes puzzle, suggestive of planning and attendant use of explicit problem-solving strategies.

General Discussion

The first experiment showed that an external WM load interfered with learning to solve a novel problem and that the impairment was proportional to the degree of WM load. In addition, after a single solution, the same WM load no longer interfered with problem solving, suggesting that WM capacity reduction does not affect problem-solving performance. Second trial performance was improved over first trial performance and was unrelated to the level of memory load on Trial 1, indicating that all participants had acquired similar levels of knowledge about the puzzle (although different groups took different numbers of moves to reach this point). Therefore, we conclude that a reduction in WM capacity selectively interferes with the rate of learning to solve the Balls and Boxes puzzle.

Because participants were not maximally efficient on their second solution, it would appear that there should still be some learning occurring during this second solution, but the WM load had no effect on this trial. There are three hypotheses for this lack of an effect: (a) By the second trial, there is much less to learn, and therefore the slight impairment is not detectable in a sample size of 76; (b) there is a threshold in the WM capacity required for learning below which resource competition does not occur; or (c) on the first trial, the WM load interferes with learning some specific information associated with the first solution, but further learning and improvement is not affected by the WM load. Our results do not distinguish between these hypotheses, but the similarity in second trial performance indicates that any undetected effect of second-trial WM load is fairly small.

In all four of the experiments presented earlier, we attempted to determine the content of participants' knowledge of the puzzle by asking them to tell us how it was solved. Virtually no participants were able to give us a coherent verbal description of the solution strategy in spite of almost universal improvement in performance across two solutions. This result implies that participants are learning the puzzle implicitly and do not have conscious access to the knowledge that leads to the improvement in performance. However, many researchers view claims of implicit learning with skepticism, and some go as far as to suggest that evidence for implicit learning is unconvincing. In their review of the field, Shanks and St. John (1994) have suggested that all claims for implicit learning should meet criteria of "information" and "sensitivity." These criteria are meant to ensure that tests of explicit knowledge test the same knowledge that leads to performance improvement and that these tests are at least as sensitive as tests of implicit knowledge. We have attempted to meet these criteria by using a recognition-oriented questionnaire and the move-selection test, as well as asking participants to verbally report their solution strategy both retrospectively and concurrently. In our paradigm, as in others, there is no clear threshold indicating when these criteria are sufficiently well met to justify accepting the hypothesis that there is null awareness (as is discussed in Merikle, 1994, in a reply to Shanks & St. John, 1994). However, the verbal protocol

findings are consistent with our other findings and hold across experiments. These results provide a distinct contrast to the elaborate concurrent verbal protocols generated by solution of a similar puzzle, the Tower of Hanoi (Anzai & Simon, 1979). The lack of reportable strategic knowledge indicates that the Balls and Boxes puzzle is being solved using different processes than those typically associated with problem solving.

The combination of examining the role of WM in a problem-solving task that appears to be supported by implicit learning is a novel one. However, there is previous research in each of the three pairwise combinations of these three elements: WM and problem solving, implicit learning and problem solving, and WM and implicit learning. Each of these areas can be considered separately to provide a context for interpreting our results.

Working Memory Capacity and Problem Solving

There is no doubt that WM plays a crucial role in problem solving. In fact, we are defining WM as a resource-limited processing construct for executing the computations necessary to execute problem-solving behavior. Increasing the WM load imposed by use of the problem operators results in increasing the difficulty of a problem (Kotovsky et al., 1985; Kotovsky & Simon, 1990). The approach used here, reducing WM capacity through a secondary task, has not been widely used with complex problem solving (Carlson et al., 1990, being a notable exception). Secondary tasks have been previously studied with simple reasoning tasks (Klapp et al., 1983), and no interference was observed between maintaining a list in short-term memory (STM) and performing the task. Halford, Bain, and Mayberry (1984) found that when the problem-solving task was more difficult (algebra problems), interference between problem solving and maintaining an STM list did result. Given that the Balls and Boxes task appears to be supported by implicit learning, it is not obvious how to categorize its difficulty with respect to these previous results. In addition, our secondary task procedure, which required constant updating of a list of 1, 2, or 3 letters (rather than maintaining a long, near-span list as in Halford et al., 1984; Klapp et al., 1983), requires executive function as well as STM use (see Baddeley, 1990) and would therefore be expected to more effectively reduce available WM. In the highest level of WM load, the three-back condition, when a new letter is heard, participants need to update their representation of the three most recently heard letters, and this is a more difficult task than updating the representation of the two most recently heard letters required by the two-back condition (and, of course, both are more difficult than the one-back condition). The increasing length of average solution time across the levels of WM load found in Experiment 1 indicates that increasing the demands of the WM task increases the interference with problem solving. This result indicates that the interference is related to reduction in WM capacity rather than simply dividing attention between the two tasks.

The interference between the secondary task and

problem-solving performance on the first trial followed by no interference on the second trial has a certain structural similarity to the pattern of the WM load effects on problem-solving performance found by Carlson et al. (1990). They showed that interference in problem solving from a secondary task only occurred when participants were new to the problem-solving domain (digital circuit troubleshooting). Once participants acquired some expertise in the domain, interference from the secondary task was reduced. One difference is that in the Carlson et al. study, participants trained for many hours to achieve expertise in their much more complex domain. In the experiments described here, participants acquire sufficient expertise to solve the Balls and Boxes puzzle without interference from a secondary task after a single solution lasting about 5 min. Although the problem-solving tasks are very different in terms of difficulty, the dramatic difference in time required to achieve an ability to perform the task without being slowed by a WM load suggests the expertise acquired in the two tasks may be of a different nature. The obvious candidate for describing this difference is the fact that participants in the experiments reported here were unable to report their acquired knowledge, whereas participants in the Carlson et al. study were clearly acquiring conscious, explicit strategies. Further exploration of the nature of the expertise acquired in tasks like the Balls and Boxes task and other problems that afford more conscious strategies are required to determine if there are fundamental differences in the nature of this expertise (other than reportability).

It is also worth noting that the Carlson et al. (1990) study examined the effect of reducing WM capacity on problem-solving performance after different amounts of training and not the effect of WM capacity reduction on learning. Sweller (1988) has suggested that reduction of WM capacity (through use of more complex problem-solving heuristics) may have an effect on learning, specifically, that WM capacity reduction may impair participants' ability to acquire domain-specific schemas. Our results are consistent with Sweller in that learning to solve the puzzle is impaired by the WM capacity reduction. Given that the Balls and Boxes puzzle is learned implicitly, it is not immediately obvious whether the knowledge acquired in solving the puzzle might be usefully considered a schema. One intriguing possibility is that schema acquisition may be supported by implicit learning under some conditions. If so, it may be useful to examine conditions under which acquired schemas are or are not reportable.

Our results indicate that a sufficiently difficult secondary task can have a large effect on problem solving. In addition, we have been able to show that this effect is specifically on learning and not performance of problem-solving skill for the Balls and Boxes problem. Carlson et al. (1990) showed that for a more complex problem-solving task, WM capacity reduction impairs problem-solving performance (until expertise is developed). For the Balls and Boxes problem, a single solution affords the acquisition of sufficient knowledge of the problem so that subsequent solutions can be executed without an effect of WM capacity reduction.

Implicit Learning in Problem Solving

Although the findings presented here appear to be the first demonstration of an extended problem-solving episode supported principally by implicitly acquired information, Berry and Broadbent (1988) have previously demonstrated learning without awareness in a problem-solving context using their dynamic control system paradigm. These tasks typically consist of a number of trials on which participants choose an operator and receive immediate feedback on whether the stated goal for the task has been achieved. Over time, participants learn to select operators so as to achieve the stated goal more frequently. Squire and Frambach (1990) showed that amnesic patients exhibit normal learning of these control system tasks (over the initial learning trials), providing additional support that these tasks can be learned implicitly. These findings indicate that cognitive skill learning can be supported by implicit memory. Our results additionally suggest that implicit learning can play a vital role in more traditional problem-solving contexts where participants explore a problem space over an extended period of time before achieving the originally stated goal (i.e., most of the learning occurs before any explicit feedback is given to the participant).

The approach used by Squire and Frambach (1990) of examining the performance of amnesic patients on a task hypothesized to be learned implicitly is often an effective one. Given the findings presented here, it may be useful to attempt this approach with the Balls and Boxes puzzle. However, the impaired declarative memory problems of patients with amnesia can necessitate changes to the presentation of the task. Patients with amnesia are much more likely than healthy participants to forget the task goal or instructions (i.e., this information is typically learned explicitly by control participants). An example of the consequences of this difficulty is provided by research on the TOH problem. Initially the TOH problem was reported to be learned normally by patients with amnesia (N. J. Cohen, 1984; N. J. Cohen, Eichenbaum, Deacedo, & Corkin, 1985). However subsequent attempts were unable to replicate this finding using a different group of patients (Butters, Wolfe, Martone, Granholm, & Cermak, 1985). It has been suggested that the initial observation of normal learning by patients with amnesia depended on the frequent use of prompts and cuing by the experimenter during learning (Gabrieli, Keane, & Corkin, 1987). Examining the performance of patients with amnesia on extended problem-solving tasks likely depends on first developing a more accurate model of what aspects of problem solving can be acquired implicitly.

To gain additional understanding of how implicit learning operates in a problem-solving context, we have attempted to determine the actual procedure used by participants to solve the puzzle. Because we have a reasonably simple puzzle, there are a limited number of ways to solve it: (a) use the rules that describe the opening-closing of boxes based on current ball position to deduce the next move; (b) avoid reversing (undoing the previous move); and (c) follow the pattern within the sequence of moves from start state to

solution. The results of the move-selection tests given in Experiments 2 and 3 show that participants need information about the previous move or moves to choose the next move correctly, therefore effectively eliminating the possibility that participants are acquiring and using a strategy based on the underlying puzzle rules (Solution a; e.g., means-ends analysis). With regard to reversal avoidance (Solution b), we found in Experiment 3 that participants rated statements indicating a strategy of avoiding reversals to be unhelpful. In fact, statements supporting this strategy were perceived as more helpful before solving the puzzle than after it. In Experiment 4, participants used this heuristic, occasionally avoiding reversals when initially exploring the problem space, but appeared to abandon this heuristic as they proceeded rather than adopting it as a solution strategy. Therefore, we conclude that participants did not appear to be consciously avoiding reversals (Solution b) as a mechanism for solving the puzzle. In addition, no verbal description of the pattern of moves to solve the puzzle was provided in either the retrospective protocols (Experiments 1 and 2) or concurrent protocols (Experiment 4). This result suggests that participants are not consciously acquiring knowledge of the move pattern (Solution c) in order to solve the puzzle. Thus it appears that participants are implementing a strategy based on Solution b or c without being aware of the strategy. Because no strategy was reported by participants, and these two strategies cannot be distinguished by performance on the tasks used here, we cannot conclude exactly which of these two strategies (or some combination of the two) was used by participants.

The performance of participants on the "final path" of a solution of the Balls and Boxes puzzle provides a contrast to previous characterizations of implicit learning. Previous findings with implicit learning tasks have suggested a gradual learning rate (e.g., Reed & Johnson, 1994; Willingham et al., 1989; as in learning a habit, Knowlton, Squire, & Gluck, 1994; Mishkin, Malamut, & Bachevalier, 1984), and this is consistent with the extended exploratory phase that precedes the final path. The fact that the exploratory phase is lengthened as WM capacity is reduced suggests that this gradual learning process is affected by WM capacity (i.e., the learning rate is slowed when WM is reduced). However, the sudden onset of the final path described in Experiments 1 and 2 has the character of performance driven by a sudden insight into a solution strategy. One possibility for the sudden onset of the final path is that the process of solving the puzzle using the implicitly acquired knowledge operates in competition with the explicit, heuristic search observed in the concurrent protocols of Experiment 4. When enough knowledge is acquired implicitly, the unproductive heuristic search could be abandoned in favor of using the accumulated implicit knowledge to solve the puzzle easily (generating a final path). The return to exploratory behavior on the second trial could also result from this competition because the participants, unaware of their knowledge of the puzzle, once again attempt heuristic search for a solution strategy. Competition arising from attempting explicit problem-solving strategies might also have led to the poorer learning performance of participants in Experiment 4, who were

encouraged repeatedly to explicitly account for their move selections. The hypothesis that the implicit knowledge arises suddenly and in a complete form at the end of the first trial and is not immediately available at the beginning of the second trial, is less consistent with previous observations of implicit learning.

The development of evidence from several sources for the conclusion that implicit learning can play a crucial role in problem solving has consequences for general models of problem solving and cognition. Anderson's ACT-R theory (1993) provides a framework for making predictions about the ability to verbalize and awareness during the acquisition of expertise but does not directly account for our findings. In ACT-R, knowledge is divided into two categories: declarative (explicit and reportable) knowledge (e.g., facts), and procedural production rules (which are not reportable). Declarative knowledge provides the primary control over which production rules execute and thus, which procedures are effected. Active declarative memory elements (those in WM, by definition) are matched to relevant production rules, which are then executed. Procedures for solving a problem are initially acquired declaratively (explicitly) and production rules are subsequently derived from the declarative knowledge. The development of expertise could be accompanied by the forgetting (decay) of the original declarative knowledge. However, we found no declarative knowledge of the solution strategy immediately after a single problem solution, with no time for forgetting. One way to reconcile our findings with the ACT-R theory is to suggest an alternate learning mechanism (implicit) that leads to direct acquisition of production rules for solving the puzzle without ever having acquired the initial declarative knowledge. The regulation of these production rules would be based on declarative knowledge of the current state (perceptual information) and active elements representing previous moves or states (information shown to be required for effective solution in the move-selection tests of Experiments 2 and 3). The effect of the WM load on the first solution indicates that this postulated implicit learning mechanism shows a graded dependency on WM. However, the process of using (and possibly further tuning) these rules on a second solution of the puzzle does not depend on WM in the conditions used here.

Working Memory Capacity and Implicit Learning

Our results indicate that reducing WM capacity slows learning without affecting performance for the Balls and Boxes puzzle. Previous studies of the effect of a secondary task while performing an implicit learning task have shown effects without being able to separate the impact of the secondary task on learning and performance. One particular implicit learning task, sequence learning, has been extensively studied in the context of dual-task paradigms that produce an effective reduction in WM capacity. In this task, participants respond to a cue appearing in one of four (usually) locations by pressing a corresponding button. The participants are not told that the cue follows a repeating

sequence of locations, but exhibit knowledge of the repeating sequence by faster reaction times to the cue when it follows the sequence compared with a condition in which the cue follows a random sequence of locations. This task is thought to be learned implicitly because (a) patients with amnesia show normal learning of the task with no corresponding awareness of the repeating sequence (P. J. Reber & Squire, 1994); and (b) among healthy participants (typically college undergraduates), although some participants acquire explicit knowledge of the repeating sequence, some participants exhibit significant knowledge of the sequence (through their decreasing reaction times) without exhibiting corresponding explicit knowledge (e.g., Willingham et al., 1989).

The original investigation of this paradigm by Nissen and Bullemer (1987) examined the effect of a secondary task on performance, and a number of subsequent studies have continued to use this approach. The most commonly used secondary task is a tone-counting task in which participants have to maintain a running count while performing the sequence-learning task. The effect of this secondary task is typically presented as dividing attention and appears to impair learning (A. Cohen, Ivry, & Keele, 1990; Nissen & Bullemer, 1987). Curran and Keele (1993) subsequently suggested that certain types of sequences, those with more unique structure, could be learned while performing a secondary task, whereas other, more complex sequences could not (although Reed & Johnson, 1994, observed learning of this type of repeating sequence under dual-task conditions over extended training). The effect of dividing attention is certainly similar to reducing WM capacity (in some models, e.g., ACT-R, these are essentially identical), although increasing levels of impairment for correspondingly larger WM capacity reduction has not been demonstrated. Frensch and Miner (1994) have shown a relationship between individual differences in WM capacity (measured by digit span) and learning ability in the sequence-learning task. However, Stadler (1995) suggests that manipulations that disrupt the organization of the sequence have a much larger effect on sequence learning than do WM capacity reductions. Stadler found that when participants were required to maintain lists of seven letters while performing the sequence-learning task, there was no effect on learning the sequences. The paradigm used by Stadler of preloading STM is similar to that of Klapp et al. (1983) and was also used by Carlson et al. (1990). No effect of a STM preload was observed in any of these three studies (Carlson et al. found an effect of WM load only when the secondary task required frequent updating, as did the secondary task used here with the Balls and Boxes puzzle). It may be that, as in problem solving, WM capacity is only of consequence for sequence learning when participants are operating to maximally use their WM (i.e., interference only occurs when both the primary and secondary tasks are sufficiently difficult).

The exact impact of WM capacity reduction on sequence learning is not yet understood. It appears that there is some effect on performance when participants are required to concurrently perform a secondary task, and this may be analogous to the deficits in problem solving observed on

Trial 1 in Experiments 1 and 2 here. Nissen and Bullemer (1987) examined the performance of participants trained with a dual task and then transferred to sequence learning without the dual task. The transfer task appeared to show no evidence of learning, suggesting that the dual task interfered with learning the repeating sequence. This result is consistent with our findings. It is worth noting that implicit learning occurs separately in a number of different areas in the brain, (Squire, 1992) and thus it is premature to conclude that sequence learning and the Balls and Boxes puzzle are identically affected by WM capacity reduction (e.g., the two tasks may depend on separate brain systems). It may also be useful to directly compare the effect of WM capacity reduction on learning and performance on sequence learning by examining the effect of a secondary task both early and late (separately) in sequence learning. Bringing these research areas together will provide some insight into whether our finding that a WM capacity reduction selectively impairs learning and not performance applies generally to implicit learning tasks. In addition, for both the sequence-learning task and the Balls and Boxes puzzle, it is necessary to determine the content of participants' acquired task knowledge in order to determine what operations required by learning, performance, or both depend on WM capacity.

There is one published report in which implicit learning and WM were studied in a problem-solving context. Hayes and Broadbent (1988) examined the effect of a secondary task on the dynamic control task shown by Berry and Broadbent (1988) and Squire and Frambach (1990) to be learned implicitly. Hayes and Broadbent reported that the imposition of a secondary task impaired explicit learning in this task but had no effect on implicit learning. Their result is very different from either the studies of sequence learning or our results reported here and suggests the possibility that different implicit learning paradigms may depend differentially on WM capacity (but it may also be important to note that their results have proven difficult to replicate, Green & Shanks, 1993; Sanderson, 1989).

Conclusions

By demonstrating that implicit learning can play a vital role in problem solving, we have opened up the possibility of examining how different memory systems participate in the problem-solving process. Previous studies of implicit memory have generally focused on simpler, more perceptual tasks such as priming and sequence learning (although the learning of artificial grammars may require complex processing; see A. S. Reber, 1989). Because implicit learning occurs automatically and without any conscious intention, the fact that it can play a role in learning at the level of operator selection and strategy use suggests that it may be an important part of the general problem-solving process. Of course, most problem-solving tasks, which afford the development of explicit, reportable strategies, are not supported primarily by implicit learning. However, there may be aspects of these tasks that are supported by implicit

learning (e.g., application of frequently used or patterned sequences of operators). Tasks such as the Balls and Boxes puzzle, which appear to be particularly suited to implicit learning, can provide information about what types of information can be acquired implicitly.

We have also documented some of the operating characteristics of implicit learning in problem solving by showing a graded dependency on WM when a task is learned and also that later performance is immune to WM capacity reduction. The similarity between our findings and the reduction of WM dependence during the acquisition of expertise might even suggest that some of the same mechanisms are involved. If so, a better understanding of the operating characteristics of implicit learning may directly impact our understanding of the acquisition of expertise in a cognitive skill.

We acknowledge that many interesting questions, particularly ones about the content of the solution strategy and the nature of expertise in this puzzle, have not been fully addressed here. Because our research brings together issues of implicit learning, problem solving, and the role that WM plays in each, it is inevitable that some interesting issues have not been fully explored. Further investigation into what kinds of strategies can be acquired implicitly, how these strategies are represented, and how they interact with explicit strategies will provide more information about the operation of implicit learning in problem solving. A more thorough understanding the roles of implicit and explicit learning in problem solving and the interaction between the learning types will augment our ability to understand general problem solving in more complex domains.

References

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
 Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
 Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2, pp. 89-195). New York: Academic Press.
 Atwood, M. E., & Polson, P. G. (1976). A process model for water jug problems. *Cognitive Psychology*, 8, 191-216.
 Baddeley, A. (1966). The capacity for generating information by randomization. *Quarterly Journal of Experimental Psychology*, 18, 119-129.
 Baddeley, A. (1990). *Human memory: Theory and practice*. Hillsdale, NJ: Erlbaum.
 Baddeley, A., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 8, pp. 47-90). New York: Academic Press.
 Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology*, 79, 251-272.
 Butters, N., Wolfe, J., Martone, M., Granholm, E., & Cermak, L. S. (1985). Memory disorders associated with Huntington's disease: Verbal recall, verbal recognition and procedural memory. *Neuropsychologia*, 23, 729-743.
 Carlson, R. A., Khoo, B. H., Yaure, R. G., & Schneider, W. (1990). Acquisition of a problem-solving skill: Levels of orga-

nization and use of working memory. *Journal of Experimental Psychology: General*, 119, 193-214.
 Carpenter, P. A., Just, M. A., & Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices test. *Psychological Review*, 97, 404-431.
 Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 17-30.
 Cohen, N. J. (1984). Preserved learning capacity in amnesia: Evidence for multiple memory systems. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 83-103). New York: Guilford.
 Cohen, N. J., Eichenbaum, H., Deacedo, B. S., & Corkin, S. (1985). Different memory systems underlying acquisition of procedural and declarative knowledge. In D. S. Olton, E. Gamzu, & S. Corkin (Eds.), *Annals of the New York Academy of Sciences—Vol. 444. Memory dysfunctions: An integration of animal and human research from preclinical and clinical perspectives* (pp. 54-71). New York: New York Academy of Sciences.
 Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, 210, 207-209.
 Curran, T., & Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 189-202.
 Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1985). A case on syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General*, 113, 541-555.
 Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (2nd ed). Cambridge, MA: MIT Press.
 Frensch, P. A., Buchner, A., & Lin, J. (1994). Implicit learning of unique and ambiguous serial transitions in the presence and absence of a distractor task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 567-584.
 Frensch, P. A., & Miner, C. S. (1994). Effects of presentation rate and individual differences in short-term memory capacity on an indirect measure of serial learning. *Memory and Cognition*, 22, 95-110.
 Gabrieli, J. D. E., Keane, M. M., & Corkin, S. (1987). Acquisition of problem-solving skills in global amnesia. *Society for Neuroscience Abstracts*, 13, 1455.
 Green, R., & Shanks, D. (1993). On the existence of independent explicit and implicit learning systems: An examination of some evidence. *Memory & Cognition*, 21, 304-317.
 Greeno, J. G. (1974). Hobbits and orcs: Acquisition of a sequential concept. *Cognitive Psychology*, 6, 270-292.
 Halford, G. S., Bain, J. D., & Mayberry, M. T. (1984). Does a concurrent memory load interfere with reasoning? *Current Psychological Research and Reviews*, 3, 14-23.
 Hayes, N. A. & Broadbent, D. (1988). Two modes of learning for interactive tasks. *Cognition*, 28, 249-276.
 Hayes, J. R., & Simon, H. A. (1977). Psychological differences among problem isomorphs. In N. J. Castellan, D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive theory*. Hillsdale, NJ: Erlbaum.
 Klapp, S. T., Marshburn, E. A., & Lester, P. T. (1983). Short-term memory does not include "working memory" of information processing: The demise of a common assumption. *Journal of Experimental Psychology: General*, 112, 240-264.
 Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, 1, 106-120.

- Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard: Evidence from the tower of Hanoi. *Cognitive Psychology*, *17*, 248–294.
- Kotovsky, K., & Kushmerick, N. (1991). Processing constraints and problem difficulty: A model. *Proceedings of the 13th Annual Cognitive Science Society* (pp. 790–795). Hillsdale, NJ: Erlbaum.
- Kotovsky, K., & Simon, H. A. (1990). Why are some problems really hard: Explorations in the problem space of difficulty. *Cognitive Psychology*, *22*, 143–183.
- Luchins, A. S. (1942). Mechanization in problem solving. *Psychological Monographs*, *54*(6, Whole No. 248).
- Merkle, P. M. (1994). On the futility of attempting to demonstrate null awareness. *Biological and Brain Sciences*, *17*, 412.
- Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of learning and memory* (pp. 65–77). New York: Guilford Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1–32.
- Perruchet, P., & Amorim, M. (1992). Conscious knowledge and changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 785–800.
- Reber, A. S. (1976). Implicit learning and synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory*, *2*, 88–94.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219–235.
- Reber, P. J., & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning & Memory*, *2*, 217–229.
- Reed, J., & Johnson, P. (1994). Assessing implicit learning with indirect tests: Determining what is learned about sequence structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 585–594.
- Ruger, H. A. (1910). The psychology of efficiency. *Archives of Psychology*, *2*(15).
- Sanderson, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 729–747.
- Schneider, W., & Detweiler, M. (1987). A connectionist/control architecture for working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 53–119). New York: Academic Press.
- Seger, C. A. (1994). Implicit learning. *Psychological Bulletin*, *115*, 163–296.
- Shanks, D. R., & St. John, M. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367–447.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis of findings with rats, monkeys, and humans. *Psychological Review*, *99*, 195–231.
- Squire, L. R., & Frambach, M. (1990). Cognitive skill learning in amnesia. *Psychobiology*, *18*, 109–117.
- Squire, L. R., Knowlton, B. J., & Musen, G. (1993). The structure and organization of memory. *Annual Review of Psychology*, *44*, 453–495.
- Stadler, M. A. (1995). Role of attention in implicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 674–685.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, *12*, 257–285.
- Tulving, E., & Schacter, D. S. (1990). Priming and human memory systems. *Science*, *247*, 301–306.
- Willingham, D. B., Greeley, T., & Bardone, A. M. (1993). Dissociation in a serial response time task using a recognition measure: Comment on Perruchet and Amorim (1992). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 1424–1430.
- Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 318–327.

Appendix A
Strategy Recognition Questionnaire Statements

Statement	Relevant strategy			
	Pattern	Rule	Nonreversing	Other
1. Remember the pattern of moves that opens each box.	+			
2. Use trial and error.				+
3. If you get stuck, back up and look for a different way.				+
4. Move as quickly as you can.				+
5. Don't move the same ball twice in a row.			+	
6. If you get lost, put all the balls back in.				+
7. Sometimes it's useful to stop thinking and just move.				+
8. Never undo the last move you made.			+	
9. Work left to right.		+		
10. Remember the combinations of balls that open the boxes.		+		
11. Take the rightmost ball out first.	+		+	
12. Move the balls on the left more frequently than the balls on the right.	-			
13. When a box opens, always move that ball.			+	
14. Never put all the balls back in.				-
15. If all the balls but the leftmost are out, you have to put them all back in.		+		
16. It's important to occasionally put balls back in their boxes.	+	+	+	
17. Take the second ball from the right out first.	-		-	
18. Move the rightmost ball every other time.	+			
19. Move the balls on the right more frequently than the balls on the left.	+			
20. The leftmost ball is the hardest to get out.		+		
21. Move the middle ball every other time.	-			
22. There are usually two moves.			+	
23. Put balls on the right back in, but not the balls on left.	+			
24. If nothing opens or closes after you move a ball, move the rightmost ball next.		+		
25. The second ball from the left is the hardest to get out.		-		

Note. Statements 1–10 were rated for helpfulness from 1–7. Statements 11–25 were judged to be true or false as well as rated for helpfulness. The strategies for which each statement are potentially relevant are listed in the rightmost 4 columns for the three potential types of applicable problem solution strategies (move pattern, rule use, and avoiding reversing), with the final column being for statements not applicable to the other three. A plus sign (+) indicates that a high rating is consistent with use of this strategy. A minus sign (-) indicates that a low rating is consistent with use of this strategy.

Appendix B
Concurrent Protocol Summaries From Experiment 4

The summary for each participant is presented in tabular form following the "path" the participant traversed through the problem space broken down into several segments. Each segment contains

the state the participant was in when that segment began and ended as well as a list of reversals made in that segment. For example, a segment state list of "7-5-6-5-7" indicates that the participant

(Appendixes continue)

began that segment on State 7, moved to State 6 then State 5, reversed to return to State 6, reversed again to return to State 5, reversed again to State 6, and continued to State 7 (see Figure 2 for diagrams of each problem state). The rightmost column indicates statements made by the participant during this segment of the protocol. Statements that simply report the participants' activity are not included (i.e., "I'm going to move this one now.").

Certain frequently invoked rules are listed with code numbers to facilitate summarizing participants' comments.

Rules

R1. Don't undo previous move (e.g., have only one option, the other move).

- R1a. Test this hypothesis by undoing and redoing a move.
- R2(n). Subgoal to get difficult ball, *n*, out, especially after encountering a state where the other four balls are out. Indicates participant is focused on removing this ball.
- R2a(n). Don't replace ball, *n*, later should that become an option.
- R2b. Clearly expresses happiness when this subgoal is satisfied.
- R3. Back up and look for another way when stuck, especially after putting several balls back in.

Other specific statements that represent hypotheses about how the puzzle works are indicated by "hyp" in Table B1.

Table B1
Concurrent Protocol Summaries From Experiment 4

Move numbers	States visited	Statements
Participant 1		
Trial 1		
1-22	21-11-21	R1 (5 times), R2 (2), tries to start over
22-35	21-27-25-27-31	R2b (2), R1a, selects "new" move, R3
35-42	31-30-31-30-31-30-31-30	Moves Ball 5 repeatedly in frustration
42-61	30-27-30-29-30-29-30-29-27-28-27-31	R2 (1), becomes frustrated, moves Ball 4 repeatedly
61-77	31-21	Explicitly starting over, hyp: moving Ball 5 causes boxes to open and close (true)
77-88	21-22-21-20-23-18	hyp: Balls 4 and 5 act as "triggers"
88-111	18-19-18-19-14-15-14-15-12-13-12-13-12-13-10	Increasing confidence that Ball 5 is "a key piece"; hyp: if Ball 4 is in, then 3 opens
111-121	10-20	R2 (2), hyp: Ball 4 is open when Ball 5 is in
121-129	20-19-20-19-24	Frustrated at revisiting start state (21), moves Ball 3 repeatedly
129-146	24-7-6-7-0	Indicated familiarity; hyp: something opens when Ball 5 moves out, R2 (2), R2 (3)
Trial 2		
1-6	21-18-21	Remembers first move, appears to intentionally start over when asked to explain moves
6-19	21-13-14-10	R1, predicts either Box 1 or Box 2 will open at State 16 (true), R2 (2)
19-37	10-14-0	R3, R2a (2), R2 (3), R2b
Participant 2		
Trial 1		
1-14	21-25-24-25-24-31	Prefers to move a ball out rather than in, R1, R1a, trying to find "pattern"
14-20	31-29-31-30-31	R2 (1), R1
20-32	31-29-31-29-[attempts illegal moves]-31	Each time at State 29, with Box 4 and Box 5 open, plans to remove both and moves Ball 4 first, gets very stuck, tries illegal moves
32-45	31-29-30-21	R3, discovers that moving Ball 5 before Ball 4 at State 29 opens Box 3, breaks out, R2 (1), states conditions for Box 2 to be open
45-59	21-30-25	Plans to remove both Ball 4 and Ball 5 at start state (like State 29 above), R2 (1), restates "rule" to open Box 2
59-69	25-31-30-31-29	R2 (1)
69-83	29-30-25-30-29-31	R2 (1), tries to examine every possibility-configuration to remove Ball 1
83-95	31-24-25-21	Loses track of configurations tested, decides to replace Ball 2, remembers rule
95-108	21-13-12	New first move, R2b (1), states configuration for Box 1 to be open, thinks Ball 1 out may have changed rule for Ball 2
108-124	12-15-9-11-9-12	R2 (2), doesn't use rule stated before
124-148	12-5-7-6-8-4-5-4-0	R2 (3), R2a (2), restates rule to remove 2
Trial 2		
1-6	21-22-19-21	Remembers first move, R2 (1)
6-19	21-13-15-12	R2b (1), R2a (1)
19-27	12-14-11-14-11-14	R2a (1)
27-47	14-7-8-6-7-6-7-0	Can't remember rule for Ball 2 right before getting Box 2 open, R2 (3)

Table B1 (continued)

Move numbers	States visited	Statements
Participant 3		
Trial 1		
1-14	21-7	R1, R2 (2)
14-16	7-8-7	R1a
16-22	7-5-6-5-7	R1, R2 (3)
22-29	7-0	
Trial 2		
1-12	21-23-22-26-21	R1, R2 (1 and 3)
12-26	21-14-15-9	Links Ball 4 moving to Box 2 opening, R2 (2), R1a, R1
26-31	9-10-9-6	Sees State 10 as similar to start state, appears to choose to move Ball 4 because linked to Ball 2
31-39	6-7-0	R1a, links Ball 3 to Ball 5 moving
Participant 4		
Trial 1		
1-12	21-14-15-11	R1, anticipates hitting a dead end at State 16, R1, R1a
12-18	11-14-11	Decides to start over at beginning, R3, R2 (2) at State 14, pauses, and doesn't seem to notice reversal
18-22	11-9-11	R3
22-33	11-0	R2 (3), R2 (4)
Trial 2		
1-20	21-31-21	Tries different move to start, recognizes Ball 2 as hard one from last time, R1, R2 (1)
20-40	21-15-16-12-21	Tries new starting move, R1a, R2 (2), R3 from State 12
40-62	21-10-21	Wonders if puzzle has changed, R2 (2), analogy to start state at 10, thinks State 12 is new (third visit)
62-66	21-19-21	R3
66-76	21-31	R2 (1)
76-98	31-29-31-22-31	"I think I skipped a step," R3
98-128	31-0	R2 (1), R2a (1), R2b (1)
Participant 5		
Trial 1		
1-8	21-25-21	Considers combinations, R2 (3 and 5)
8-22	21-14-16-11	New first move, R1, R2 (4) to replace, R1, R1a, R1, R2 (2), hyp: when Balls 4 and 5 are in, they both are open
22-33	11-17-16-17-16-18	Tries "new" option, R3, R1a
33-45	18-13-5-13-14-12	R2 (2), R2 (3) to replace, R2 (2), tries varying order of moving Balls 4 and 5 out, hyp: moving Ball 5 out opens Box 1
45-52	12-13-12-13-9	Surprised when Box 3 opens at State 12, backs up to see it again, makes analogy to start state at State 10
52-67	9-11-5-6-0	hyp: Moving Ball 4 out then in and then Ball 5 out opens Box 3, uses this rule at State 5 to solve puzzle
Trial 2		
1-12	21-20-31	Tries different starting move, states configuration when Box 2 opens, R2 (1), remembers Box 1 related to Ball 5, R2 (1)
12-20	31-29-31-30-31-29	Tests idea that reversing changes state, hyp: Ball 4 in then out then Move 5 opens Box 3 (wrong)
20-33	29-30-26-29-24	Trying different move orders
33-38	24-25-21	R2 (1)
38-49	21-24-22-25-24-26	Tries reversal to change state again, move Ball 2 out then in and then move Ball 5
49-60	26-21-22-18-19	Uses rule: Move Ball 4 in then out, then remove 5 and surprised when Box 2 opens, then uses rule moving 4 out first to open Box 3
60-72	19-10-11-9	Replaces Ball 5 to open 4, R2b (1), R2 (2)
72-82	9-11-10-11-5	Uses rule to open Box 3 to replace it at State 11, tries moving Ball 5 twice and then Ball 4, surprised when Box 2 opens at State 8, hyp: Move Ball 4 twice then Ball 5 twice
82-97	5-7-5-7-5-6-0	Tries to use "rule" to open Box 3, but uses it backwards twice, then uses it forward to solve puzzle

Note. R = rule.