Tower of Hanoi disk-transfer task: Influences of strategy knowledge and learning on performance

Welsh, M.C.; Huizinga, M.

Published in:
Learning and Individual Differences

DOI:
10.1016/j.lindif.2005.05.002

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Learning and Individual Differences 15 (2005) 283–298

Learning and Individual Differences

www.elsevier.com/locate/lindif

Tower of Hanoi disk-transfer task: Influences of strategy knowledge and learning on performance

Marilyn C. Welsh [a,*], Mariëtte Huizinga [b]

[a]Department of Psychology, University of Northern Colorado, Greeley, CO, 80639, USA
[b]Department of Psychology, University of Amsterdam, The Netherlands

Received 22 February 2005; received in revised form 18 May 2005; accepted 18 May 2005

Abstract

Tower of Hanoi has become a popular tool in cognitive and neuropsychology to assess a set of behaviors collectively referred to as executive functions. Substantial variability in performance on the Tower of Hanoi (TOH) disk-transfer task among normally functioning young adults, and potential contributions to these individual differences, were examined. In this expanded 60-problem version of the four-disk TOH, the degree to which problem administration (blocked vs. random) and strategy knowledge influenced overall performance and changes in accuracy across problems was examined. Eighty-seven college students were randomly assigned to a Blocked Group (problems given in ascending order of move-length) and a Random Group (problems given in a random order). After administration of the TOH task, participants described their problem solving and these verbal protocols were analyzed with regard to four elements of a strategic approach to problem solving. Problem administration order demonstrated no effect on task performance or on expressed strategy knowledge; however, strategy knowledge did predict performance on the TOH. An expected decrease in performance across trials was observed in the Blocked Group, and an increase in accuracy in the Random group indicated a learning effect. Strategy knowledge did not interact with these changes in performance across the items. These results suggest that external cues do not influence performance on the TOH to the same extent as individual differences in strategy induction relatively early in the problem solving process.

Keywords: Tower of Hanoi; Executive functions; Problem solving; Strategy

* Corresponding author. Tel.: +1 970 351 1058.
E-mail address: marilyn.welsh@unco.edu (M.C. Welsh).

1041-6080/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.lindif.2005.05.002

The Tower of Hanoi (TOH) is a complex problem-solving task that has become a popular measure of the executive function construct, which has been defined as "the ability to maintain an appropriate problem solving set for attainment of a future goal" (Welsh & Pennington, 1988; p. 200). This task has demonstrated sensitivity to prefrontal lobe function and dysfunction (e.g., Fuster, 1997; Glosser & Goodglass, 1990; Goel & Grafman, 1995; Goldstein & Green, 1995; Lezak, 1995; Stuss & Benson, 1984). However, the specific executive processes recruited for successful performance, and by implication impaired in prefrontal dysfunction, is the subject of current debate. Researchers have suggested that this task taps executive processes such as planning, working memory, and inhibition (e.g., Goel & Grafman, 1995; Roberts & Pennington, 1996), and empirical evidence supporting these proposals has been converging in recent years. For example, research findings have suggested that inhibition (Welsh, Satterlee-Cartmell, & Stine; 1999), working memory (Goel, Pullara, & Grafman, 2001; Handley, Capon, & Harper, 2002; Lock, Welsh, Adams, & Kurtz, 2002; Numminen, Lehto, & Ruoppila, 2001; Welsh et al., 2002), procedural learning (Bagley, Welsh, Retzlaff, Wolf, & Bryan 2002; Davis & Klebe, 2001; Devine, Welsh, Retzlaff, Yoh, & Adams 2001; Goldberg, Saint Cyr, & Weinberger, 1990), and fluid intelligence (Devine et al., 2001; Lock et al., 2002; Numminen et al., 2001) each contribute to performance on the TOH task.

The TOH task has been termed a disk-transfer task, that is, in a series of different problems, the object is to move the disks from a start state to a goal state in the fewest number of moves as possible while following a specified set of rules. There are three rules that dictate moves that can and cannot be made: (1) only one disk may be moved at a time; (2) disks may be moved only to another peg (e.g., they cannot be placed onto the table or held in the hand while another disk is moved); and (3) a disk may never be placed on top of a disk smaller than itself. It has been assumed that these constraints force the participant to engage in planning activity in working memory, as well as the inhibition of intuitive, albeit maladaptive, moves (Pennington, 1997; Scholnick, Friedman, & Wallner Allen, 1997). For example, to successfully solve the problem it is often necessary to temporarily block the desired ending arrangement in order to free up disks needed to achieve the goal pattern (Goel & Grafman, 1995). Successful task performance therefore requires that individuals plan their moves to achieve the goal state and inhibit the tendency to focus on short-term goals.

In his seminal analysis of the TOH task, Simon (1975) suggested that the constraints of this task are conducive to the spontaneous generation of several problem solving strategies that vary in effectiveness and may explain normal individual differences in performance. The optimal strategy is referred to as goal recursion and includes the following basic elements: (1) recognizing that the first subgoal is to move the largest disk to its goal position; (2) moving the smaller disks out of the way; (3) building a "sub-pyramid" stack of these smaller disks on the "open" peg; (4) moving the largest disk to it's goal position; and (5) repeating these steps with the "next-largest" disks and progressively smaller sub-pyramid stacks until the goal state is achieved. In its most complex form, the goal recursion strategy involves the understanding that within the each major subgoal, i.e., building the subpyramid and delivering the largest disk to its final goal position, are smaller cycles of the five steps above. Simon (1975) proposed that the goal recursion strategy does not require perceptual updating with regard to the configuration of the disks on the peg, but simply the knowledge of where the person is in the stack of goals that must be achieved. In contrast, Simon (1975) suggested that the perceptual and sophisticated perceptual strategies include an appreciation of some subset of the elements comprising the recursive strategy (e.g., recognizing the first subgoal), but the person must decide where to move the disks in order to sequentially deliver each largest disk to its goal peg. Importantly, when the person uses these perceptual strategies, there is no appreciation of building successively smaller pyramids on the "other" peg (neither source peg nor goal peg). In the "production statements" set forth by Simon to describe the sophisticated perceptual strategy, the movement of the smaller disks is perceptually-driven, rather than conceptually-driven, and often the result of an educated guess between two legal moves. In contrast, the goal recursion strategy provides a "plan" for how the smaller disks should be moved out of the way to clear the largest disk's path to its goal. This "plan" derives from the rule that a smaller subpyramid must be built on the open peg before the largest disk is moved to the goal peg.

One can also hypothesize that knowledge of the recursive rule guiding solution should reduce the demands on executive processes such as planning, working memory, and inhibition. Since Simon's pioneering work, there has been surprisingly little research on the problem solvers' knowledge of the recursive or perceptual strategy while engaged in the TOH task. There has been evidence that the recursive nature of the problem is one of the aspects that makes the task difficult for individuals to solve. For example, Welsh, Cicerello, Cuneo, and Brennan (1995) found a systematic pattern of errors and pause times that parallel each major recursive cycle within the move-path. Kotovsky, Hayes, and Simon (1985) found that, on an isomorph of the TOH task, participants tend to go through a large exploratory phase prior to "seeing" the full solution and then proceeding quickly to the goal. The verbal protocols of their participants suggested that they did not have the full solution in mind prior to the first move, and one can speculate that it was only during this exploratory phase that they discovered the recursive structure of the task. Goel et al. (2001) developed a computer model in which the TOH problems were solved via the perceptual strategy, and their results demonstrated the heavy working memory demands posed by longer-move problems. A main purpose of the current study was to contribute to our understanding of TOH solution by examining the general nature of strategy knowledge and the contribution of this knowledge to performance.

The search for the cognitive strategies underlying TOH performance is complicated by the fact that administration of this task varies considerably from laboratory to laboratory. Some variations include, among other things, the imposition of time limits, allowing either a single attempt or multiple attempts to correctly solve a single problem, and variation in orders of problem presentation. The approach to problem presentation has taken two basic forms. One form involves presenting only the original 3- or 4-disk, tower-to-tower to participants (e.g., Davis & Klebe, 2001; Goldberg et al., 1990; Squire, Cohen & Zouzounis, 1984; Welsh et al., 1995). The second form entails presenting a series of problems in order of ascending or descending order of difficulty (e.g., Borys, Spitz, & Dorans, 1982; Welsh, 1991; Welsh & Huizinga, 2001; Welsh et al., 1999). In this case, as the participant progresses through the task, the problems increase or decrease in terms of the number of moves required to solve the problems successfully. Problems with a greater number of moves presumably place more demands upon working memory since participants are required to hold and manipulate more moves in memory before they overtly implement those moves with the physical TOH. It is important to note that the most complex TOH problem for a given number of disks (e.g., 15-move problem for four disks) can be decomposed into a series of shorter problems that are embedded within the larger problem, and it is some subset of these smaller problems that is typically presented in ascending order leading up to the most complex problem. The fact that the same errors continue to be repeated across these problems suggests that participants do not recognize the embedded structure of the TOH task as presented in this manner (Welsh, 1991). The current study was designed to examine the effect of experience with embedded problems on performance within the context of a 60-problem task that provides a substantial amount of this experience.

Given the variety of administration procedures characterizing the TOH tasks, it is critical to examine the influence of these procedures on performance and on the cognitive processes involved. The first goal of the present study was to explore whether the order in which the individual TOH problems are administered affects (a) overall performance and (b) performance at differing levels of difficulty. Specifically, a blocked administration, that is, in an ascending order of difficulty (e.g., 6-move problems, 7-move problems, etc.), was compared to a random order of administration. Because the blocked order of administration capitalizes on the natural embedded structure of the TOH, it was predicted that participants in the blocked order group would exhibit better overall performance than would participants in the random order group.

A second goal of the study was to explore whether the two orders of problem administration affected the strategy knowledge expressed at the end of the large set of 60 TOH problems. Given that Simon (1975) identified the goal recursion strategy as the optimal method of problem solution on the TOH, would the order of problem presentation affect the participants' knowledge of the various aspects of this strategy? In a study by Schunn, Lovett, and Reder (2001), the results indicated that explicit awareness of changing characteristics of their task was more predictive of appropriate adaptation of strategy than were individual differences in either working memory or inductive reasoning. External cues provided by administration order of the items is one method by which such self-awareness could be influenced. Therefore, we hypothesized that administering the problems in the blocked order would facilitate the discovery of an effective strategy because as problem difficulty (i.e., number of moves) increases, this necessitates another cycle or recursion through the strategy. We expected that participants in the random order group would have less explicit awareness of the elements of an effective strategy (including aspects of goal recursion), given that the natural embedded structure is violated and the recursive nature of the task would be less obvious.

The problem-solvers' knowledge of an effective solution strategy may be conceptualized as an effect of the administration procedure experienced by the individual. Alternatively, one can view strategy knowledge as an individual difference variable that, in and of itself, contributes to performance on the TOH task. Performance on the TOH has been linked to inductive reasoning and other aspects of fluid intelligence (Devine et al, 2001; Lock et al, 2002; Welsh et al, 2002), as well as to the hypothetico-deductive reasoning characteristic of formal operational thought (Emick & Welsh, in press). These cognitive processes may contribute to the likelihood of inducing a goal-oriented strategy and applying it effectively to the task, irrespective of problem order. Therefore, a third goal of the present study was to characterize participants in terms of their level of strategy knowledge, and to explore whether there were observable group differences in TOH performance.

Finally, the issue of learning across a large set of TOH items was of interest in light of the fact that this task has been referred to as a procedural learning task by some research groups (e.g., Davis & Klebe, 2001; Goldberg et al., 1990). Typically, those studies that utilize the TOH as a procedural learning task administer a single TOH problem repeatedly, and improvements in performance have been observed. However, in those studies administering multiple problems of increasing difficulty, a robust difficulty effect is typically found; that is, there is a substantial decline in performance across trials (e.g., Welsh, 1991; Welsh & Huizinga, 2001). This administration procedure clearly obscures the learning that may also be occurring across trials. A fourth goal of this study was to explore change in performance across 60 TOH trials: a decline in performance was expected in the blocked order group, whereas an increase in performance indicative of learning was expected in the random order group. It also was of interest whether strategy knowledge, as an individual difference variable, would influence the decrement in performance over trials in the blocked order group and/or the learning over trials in the random order group. If strategy knowledge is acquired over the course of the TOH trials, it is possible that the two strategy knowledge groups (i.e., high vs. low knowledge) will not differ in the early trials of the task, but progressively diverge in performance as the task continues.

To address these four issues, a sample of young adult volunteers were administered an expanded 60-item TOH tasks in one of two administrations: blocked order of random order. It was predicted that those participants in the blocked order group would exhibit better overall performance on the TOH, as well as a greater amount of strategy knowledge. In addition, it was expected that strategy knowledge itself, will relate to overall TOH performance and may influence the degree to which one observes changes in performance over trials.

1. Method

1.1. Participants

A total of 87 undergraduate students enrolled at a mid-sized university participated in the study. Upon volunteering for the study, the participants were questioned as to known diagnosis of any learning disabilities or head injuries. If the participant had a history of head injury and/or had a diagnosed learning disorder, he or she was excluded from the study. Six participants were excluded from the analyses for the following reasons: three participants (two females, one male) were excluded due to self-reports of past head injury, and three participants (three females) were dropped from the analyses due to extremely low scores (well below 2 standard deviations from the sample mean) and behaviors indicating a failure to understand the task. Therefore, the sample of participants included 81 students, 60 females and 21 males, with a mean age of 18.32 years (SD = .70). These students were compensated with course credit for their participation.

1.2. Apparatus

The Tower of Hanoi (Simon, 1975) consists of a flat board (40 × 15 × 2 cm) on which three vertical wooden pegs of equal diameter (1 cm) and equal height (14.5 cm) are spaced equidistantly (12.5 cm). Three or four wooden disks of graduated size (13.5, 11, 8.5, and 6 cm diameter), each have one hole (1.3 cm in diameter) drilled through the center so that they fit onto any of the three pegs. A set of 17 × 22 cm cards in a three-ring binder displays the goal states of the individual items that are presented to the participant.

The TOH requires that an initial start configuration of disks across the three vertical pegs be transformed into a specific goal configuration of these objects, in the minimum number of moves. Disks must be moved according to a set of specified rules that constrain the manner in which these objects may be moved from peg to peg. These rules include the following: (a) only one disk may be moved at a time; (b) a disk may not be placed on the table or held in the hand while another disk is being moved; and (c) a larger disk may not be placed on tope of a smaller disk.

1.3. Procedure

Participants were randomly assigned to either the blocked-order group or the random-order group by a coin toss. The task consisted of 30 tower-ending problems (e.g., all odd-numbered items on

Fig. 1) and 30 flat-ending problems (e.g., all even-numbered items on Fig.1). The 60 items were all possible items derived from two possible different end-states (tower or flat), 10 different move-lengths per item (6-through 15-moves), and three different starting pegs (Welsh & Huizinga, 2001). The blocked-order of administration involved presenting problems in an ascending order from 6-move problems through 15-move problems. The random-order of administration was developed by randomly selecting problems in an alternating sequence of tower-ending and flat-ending items. Participants were administered the TOH individually after an explanation of the three rules (see above) and two three-disk practice problems (one flat-ending and one tower-ending). For each item, the tester set up the start state on the TOH apparatus, and the participant was presented with a card that exhibited the goal state. Both this card and the tester indicated the number of moves required to achieve the goal state. Participants had to reach the goal in the designated number of moves on the first attempt, and there was no time limit imposed. Scoring involved awarding one point for each correct solution (i.e., transforming the start state to the goal state in the required number of moves).

At the end of the session, participants were questioned as to how they solved the problems. The experimenter set up a 4-disk, 15-move, tower-ending problem and asked the participant what should have been done to solve the problem. The experimenter then recorded the participant's response verbatim. Based on these written records of the participants' verbal protocol of their problem solving process, two trained, independent researchers assessed which, if any, of the four strategy components were articulated in the response given by the participant. The four components included: (a) move the largest disk to its goal first; (b) move the smaller disks out of the way; (c) build a "mini-tower" on the "open" peg; and (d) repeat the process to completion. The inter-rater agreement for the assessment of presence of these four elements was 95%, and all disagreements were resolved by mutual consensus prior to data analysis. A total recursive strategy score was calculated by assigning one point for each component present in the response.

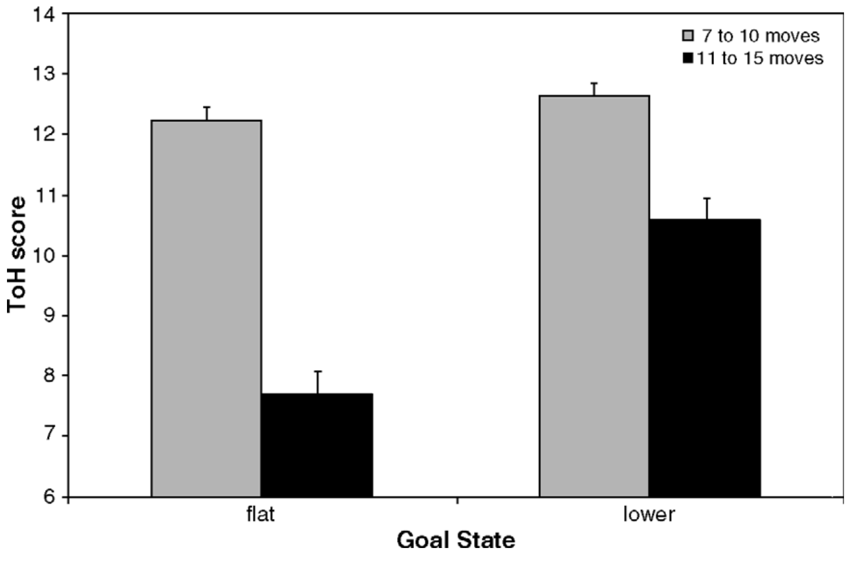

Fig. 1. Mean TOH score (number of correctly solved items) and standard errors as function of Goal State and Move Length.

2. Results

2.1. Task administration effect on performance

The prediction examined was that accuracy on the TOH task would be positively impacted by the blocked administration procedure in which 60 problems of 7- through 15-moves were presented in ascending order. In contrast, the random order of problem administration was expected to result in less accurate performance. The descriptive statistics for TOH performance in the two administration groups are presented in Table 1.

Three independent variables were the focus of this analysis: Administration Group (blocked order vs. random order), Move Length (7 to 10 moves vs. 11 to 15 moves), and Goal State (problems with a tower-ending goal vs. problems with a flat-ending goal). A 2 (Administration Group) × 2 (Move Length) × 2 (Goal State) mixed model ANOVA identified significant main effects of Move Length, MSE = 867.11, $F(1,79) = 177.51$, $p < .05$ and of Goal State, MSE = 221.10, $F(1,79) = 69.83$, $p < .05$. There was no main effect for Administration Group, MSE = 25.09, $F(1,79) = 1.44$, $p > .05$, nor did this factor interact with any other independent variable in the analysis. For both administration groups, the shorter move-length problems and the tower-ending problems were solved more accurately. The Move Length × Goal State interaction was significant, MSE = 124.18, $F(1,79) = 35.64$, $p < .05$. As seen in Fig. 1, there was a greater difference between the shorter move-length problems and the longer move-length problems for the flat-ending problems than for the tower-ending problems. Dependent *t*-tests at each Move Length level confirm that there is no significant difference between goal states at the shorter move-length level ($t(80) = -1.76$, $p < .05$); whereas, there was a significant difference between goal states at the longer move-length level ($t(80) = -8.62$, $p < .05$).

2.2. Analysis of strategy knowledge

The following analyses were conducted to examine the nature of the strategy knowledge expressed by the participants in the post-task interview, as well as the effect of the blocked vs. random order of problem administration on the amount of strategy knowledge expressed. The a priori prediction was that the blocked order of administration would result in greater strategy knowledge. In addition, strategy knowledge was analyzed as an individual difference variable to explore the differences in TOH performance between participants with relatively greater strategy knowledge vs. those with relatively less knowledge.

A 2 (Administration Group) × 4 (Strategy Element) mixed model ANOVA was conducted to examine differences in knowledge for the two administration groups, and for the four types of strategy elements: (1) move the largest disk to its goal first; (2) move the smaller disks out of the way; (3) build a “mini-tower” on the “open” peg; and (4) repeat the process to completion. The main effect of Administration

Table 1
Means and SDs (between brackets) of several TOH scores across administration groups

Administration group	Total TOH overall score	Total TOH score on tower-ending items	Total TOH score on flat-ending items	TOH score on 7- to 10-move items	TOH score on 11- to 15-move items
Blocked	44.3 (8.4)	24.1 (4.3)	20.2 (4.8)	25.0 (3.6)	19.3 (5.7)
Random	42.1 (8.3)	22.3 (4.5)	19.7 (4.6)	24.8 (3.5)	17.3 (5.7)

Group was nonsignificant, MSE=.09, $F(1,79)=.24$, $p>.05$, and Administration Group did not significantly interact with the Strategy Element variable. Therefore, there was no effect of task administration type, random vs. blocked, on overall recursive strategy knowledge or on knowledge of the specific components of the strategy. There was a significant main effect of Strategy Element, MSE=2.38, $F(3,237)=15.83$, $p<.05$, indicating that there were differences in the degree to which the four elements were expressed in the verbal protocols of the participants. Given that the mean of the dichotomous score (1 or 0) for each strategy element reflects the percentage of participants expressing that element in their protocols, each of the first two elements were referred to by 80% of the participants, the fourth element was mentioned by 62%, and the third element was described by only 45% of the participants. Within subjects contrasts indicated that response rate to the first two elements did not differ significantly; however, each differed significantly from element three and four, which also were significantly different from each other.

In order to explore the association between strategy knowledge and TOH performance, each participant was assigned to one of two groups, based on the mean strategy knowledge score ($M=2.5$ points of 4 possible points) for the entire sample (i.e., collapsed across administration group). Participants with a strategy score of 2.5 and above were assigned to the High Strategy group ($N=35$) and participants with a score less than 2.5 points were assigned to the Low Strategy group ($N=46$). A 2 (Strategy Group) × 2 (Move Length) × (Goal State) mixed model ANOVA examined the differences between the two strategy knowledge groups on overall TOH performance, as well as the degree to which these differences might be specific to TOH problems of different move-lengths or goal state configurations. The main effect of Strategy Group was significant, MSE=161.13, $F(1,79)=10.27$, $p<.05$, reflecting a higher overall TOH score in the High Strategy group ($M=45.61$, SD=6.97) than in the Low Strategy group ($M=39.91$, SD=9.03). Strategy group assignment did not interact with the type of TOH problem, with regard to move-length or goal state. Consistent with the main effect of Strategy Group, strategy knowledge significantly correlated with TOH accuracy, $r(80)=.45$, $p<.05$, when knowledge was analyzed as a continuous variable (i.e., recursive strategy score from 0 to 4).

In light of the different rates at which participants in the post-test interview expressed the four strategic elements, it was of interest to examine whether these elements also differed with regard to prediction of the total TOH score. Point biserial correlations were calculated between scores (1 or 0) on each of the four recursive strategy elements and the total TOH score, as well as among the elements themselves. Significant correlations were found between TOH total score and scores for Point One ($r(.79)=.29$, $p<.05$: "move the largest disk to its goal first"), Point Three ($r(79)=.41$, $p<.05$: "build a "mini-tower" on the "open" peg"), and Point Four ($r(79)=.37$, $p<.05$: "repeat the process to completion"). Moreover, intercorrelations among the scores on the elements of recursive strategy were found, most notably an association between the score on Point Three and Point Four ($r(79)=.50$, $p<.05$). To identify which of these four points were most predictive of TOH performance, a stepwise multiple regression was conducted and the only variable to enter the equation was Point Three (standardized beta=.41). The model was significant, $F(1,79)=15.80$, $p<.05$, $R^2=.17$.

2.3. Changes in performance over trials

The 60 trials of this expanded version of the TOH were divided into quartiles of 15 problems each and these four sections included different problems for the two administration groups. For the Blocked Group, the first quartile included 6- to 8-move problems, the second quartile included 8- to 10-move problems, the third quartile included 11- to 13-move problems and the fourth quartile included 13- to 15-move problems. For the Random Group, the average move-length for the problems in each quartile were 9.8 moves (first quartile), 11.2 moves (second quartile), 10.3 moves (third quartile), and 9.9 moves (fourth quartile). The a priori prediction was that performance would decline in the Blocked Group as the move-length of problems increased, and that performance would increase in the Random Group if learning was occurring across trials.

The degree to which administration group and time across task (i.e., the quartiles) affected performance was analyzed by means of a 2 (Administration Group) × 4 (Time) Mixed Model ANOVA. There was a significant effect of Time, MSE = 30.899, $F(3,237) = 8.291$, $p < .05$; however, more importantly the predicted Time × Administration Group interaction was significant, MSE = 210.97, $F(3,237) = 56.61$, $p < .05$. As Fig. 2 illustrates, the performance of the Blocked Group indicated an increase from quartile one to two, and then the expected decrease in the following two quartiles. Within-subjects contrasts demonstrated that there was a significant change in performance within each pair of adjacent quartiles (e.g., between quartiles 1 and 2, 2 and 3, etc.). In contrast, performance of the Random Group exhibited an increase in accuracy across problems in the Random Group. Comparing adjacent quartiles, within-subjects contrasts found that the only significant difference in performance was between quartile 2 and 3 ($F(1,38) = 30.37$, $p < .01$).

To examine whether strategy knowledge influenced the change in performance over trials, each of the administration order groups were analyzed separately. For the Blocked Group only, a 2 (Strategy Group) × 4 (Time) Mixed Model ANOVA identified main effects of Time (MSE = 180.54, $F(3,114) = 49.85$, $p < .05$) and of Strategy Knowledge (MSE = 75.14, $F(1,38) = 4.88$, $p < .05$). There was no interaction between the Time and Strategy factors; thus, both the high and low strategy groups demonstrated a similar decrease in accuracy across the quartiles and the High Strategy Group maintained superiority in performance over the Low Strategy Group at each quartile. For the Random Group only, a 2 (Strategy Group) × 4 (Time) Mixed Model ANOVA identified main effects of Time (MSE = 64.94,

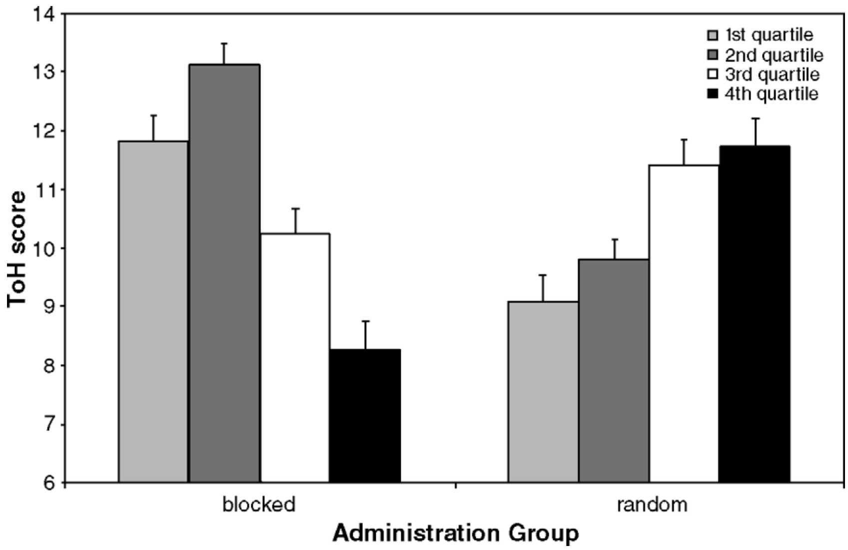

Fig. 2. Mean TOH score (number of correctly solved items) and standard errors as function of Administration Group and Time across the task.

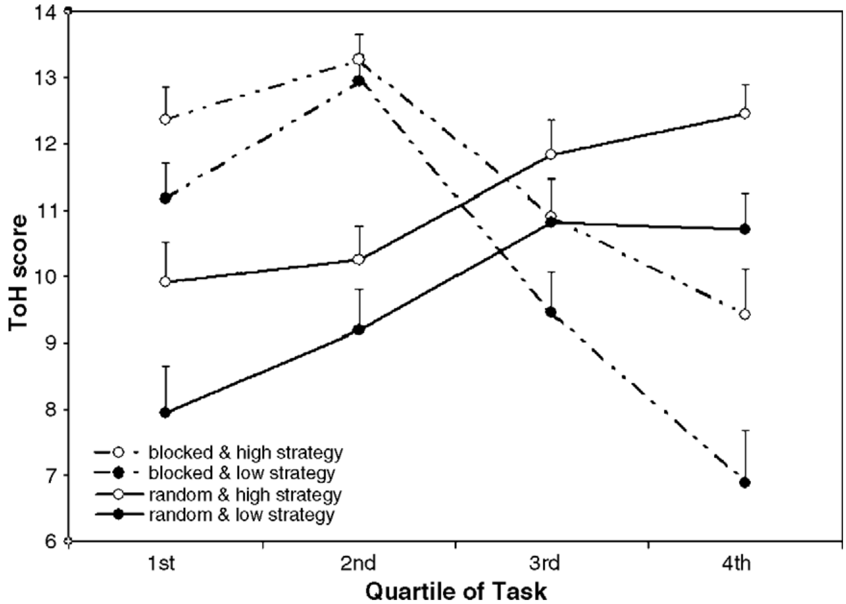

Fig. 3. Mean TOH score (number of correctly solved items) and standard errors as function of Goal State and Move Length in two administration groups and two strategy knowledge levels as a function of Time across the task.

$F(3,117)=17.31$, $p<.05$) and of Strategy Knowledge (MSE=84.02, $F(1,39)=5.43$, $p<.05$). Again, there was no interaction between the Time and Strategy variables. As was the case for the Blocked Group, both the high and low strategy groups exhibited the same type of learning curve across the quartiles with the High Strategy Group demonstrating a higher accuracy score at each quartile. These performance trends over time can be seen on Fig. 3.

3. Discussion

The purpose of this study was to gain a greater understanding of the mechanisms underlying performance on a structured problem-solving task, the Tower of Hanoi (TOH). By manipulating the order in which the sixty problems was administered, as well as examining the nature of the participants' strategy knowledge, the degree to which an appreciation of the inherent embedded structure of the TOH problem influences problem solving success was explored. Taken together, the results suggest that the putative external cues provided for the participant through the order of problem administration does not aid in the recognition of this problem structure. Instead, it appears that strategy knowledge is an important individual difference variable that contributes to problem solving performance relatively early in the task.

It was predicted that the Blocked Order of administration would be related to greater problem solving success and greater recursive strategy knowledge because the recursive structure of the TOH problems would be more salient to the participants. Given that the earlier problems in the blocked order are simply smaller problems embedded within the larger 15-move tower-to-tower sequence, the participant can see the recursive strategy as it continues to "build" cycle by cycle. In contrast, the Random Order completely breaks down this natural embedded structure, presenting the problems in a completely random fashion. If participants are sensitive to cues provided by the embedded structure and use these to induce an effective strategy, then performance and/or strategy knowledge should have been significantly better in the Blocked Order group. This was not the case, as the Blocked Group participants did no better in overall performance, performance on longer-move or flat-ending goal problems, or in expressed strategy knowledge. As was found in previous studies in which the TOH problems were administered in ascending order, neither children (Welsh, 1991) nor adults (Welsh et al., 1995) appear to be sensitive to the embedded structure of the task. In both administration modes, strategy knowledge appeared to be acquired through extended "transactions" with the task, similar to that observed by Best (1990) in a different problem solving context.

With regard to strategy knowledge, it also was not affected by the manipulated variable of problem order. However, there was clear variability in the degree to which the four elements of the goal-oriented strategy were expressed in the post-test interviews, such that a "high strategy knowledge" and "low strategy knowledge" group could be identified. The high knowledge group scored significantly better on the TOH than the low knowledge group and degree of knowledge was positively correlated with the total score. Given that this was an individual difference variable, and was based on information gathered *after* the 60-problem task was over, one cannot conclude that greater strategy knowledge per se *caused* the superior performance. For example, it is entirely possible that one group of participants was more prone to inducing the strategy over the course of the task and the strategy was only fully realized toward the end of the large set of problems, right before the interview session. The analysis of change in performance over the four quartiles of the task provided a window through which one could examine this possibility. If indeed the strategy was being induced by some participants gradually across the 60-problem task, one might expect relatively equal performances of the low and high strategy groups on the first quartile (in both the Blocked and Random Orders) with divergence between the two groups in the later quartiles. That is, the increasing strategy knowledge would be reflected in relatively less decrement with more difficult problems in the Blocked Order or relatively greater learning in the Random Order. Such an interaction was not revealed by the analyses, and instead, the high strategy group showed superior performance compared to the low strategy group at every quartile of the task. Thus, something about this high strategy group, greater knowledge or perhaps another factor such as self-monitoring (Schunn et al., 2001), was related to better performance in the first fifteen problems of the task. Given the novelty of the TOH task, it is unlikely that these participants possessed knowledge of this strategy prior to the administration of these problems; however, it is reasonable to hypothesize that these participants were more effective at inducing the rule-governed strategy very early in the problem solving process.

What cognitive processes might underlie this ability to quickly discover a successful strategy for solving a novel problem-solving task such as the TOH? Research in our lab has identified associations between performance on a 22-problem version of this task, the Tower of Hanoi—Revised (TOH-R), and measures of fluid intelligence, particularly inductive reasoning (Devine et al., 2001; Lock et al., 2002; Welsh et al., 2002). Additionally, the related process of formal operational thinking may facilitate the generation of effective problem solving strategies, and indeed, performance on the TOH-R has been found to correlation with formal operational reasoning (Emick & Welsh, in press). Across several studies utilizing the TOH-R, the individual differences in performance in normal college-student samples are extremely consistent and undeniable (e.g., Welsh & Huizinga, 2001). It appears that a subset of young adults is more likely to induce a goal-oriented strategy, and this occurs over the course of the first dozen or so problems. One might speculate that there are more basic cognitive mechanisms that contribute to the ability to induce a rule-based strategy from patterns observed across these problems, and two candidate cognitive processes are working memory and inhibition. Patterns across problems are recognized only if one integrates information across time (Hambrick & Engle, 2003), and this presumably occurs in working memory. Various researchers have identified the important contributions of memory processes, such as working memory (Goel et al., 2001), and activation and priming (Altmann & Trafton, 2002) to performance on different versions of the TOH task. By "lesioning" the working memory representations of relevant subgoals, the researchers could simulate impaired performance on the longer-move problems, such as that observed in frontal damaged adults (Altman & Trafton, 2002; Goel et al., 2001). Similarly, inhibitory processes must be engaged by the person so as to resist being pulled in the direction of reasonable disk moves that are nevertheless inconsistent with the recursive strategy, and it is this demand that Goel and Grafman suggested was central to performance in an earlier study (Goel & Grafman, 1995). At this point, there is some controversy in the empirical literature regarding the contribution of working memory and inhibitory processes to TOH performance. Whereas, a structural equation model proposed by Miyake, Freidman, Emerson, Witzki, and Howerter (2000) identified the inhibition factor as the best predictor of TOH performance, a structural model developed by Welsh et al. (2002) found working memory to be the stronger predictor. The differences in the measures of working memory, inhibition, and most importantly, the TOH, may contribute to the explanation for these incongruent results.

In addition to these explicit cognitive processes as underlying the induction of the strategy knowledge and performance on the TOH, many researchers have pointed to implicit learning as a factor in improved performance on the task. Typically this procedural learning is observed by continuous improvement in performance over repeated administrations of the same TOH problem. The classic study described by Cohen and colleagues (Cohen, 1984; Cohen, Eichenbaum, DeAcedo, & Corkin, 1985) demonstrated that amnesic patients exhibited this improved performance over repeated presentations, with no explicit memory of having done the task previously. Such evidence bolstered the claim that the TOH could be learned through implicit processes. Interestingly, in studies that employ the TOH as a procedural learning task, optimal performance on each problem (i.e., solving a 7-move problem in seven moves) is not required of the participants. In early administrations of the problem, a person might eventually solve a 15-move item in over 50 moves and procedural learning is indicated by the steady decline in number of moves, though perfect performance may never be achieved. Therefore, this type of TOH paradigm may not be appropriate for examining the discovery of an explicit strategy, such as goal recursion (Simon, 1975), which should lead to perfect performance. In our own lab we have found significant correlations between performance on a procedural learning task, Mirror Tracing (Gabrieli, 1998), and TOH-R performance (Bagley et al., 2002), with one study indicating somewhat equal contributions of both explicit learning and implicit learning to the explanation of variance in TOH-R scores (Devine et al., 2001). When a form of the recursive strategy was taught explicitly to participants before beginning the TOH-R task, these participants performed significantly better than a control group and, as expected, expressed greater knowledge of the recursive strategy (Lock et al., 2002). However, the "teaching group" performed no better than participants who were allowed to explore the task on their own without explicit strategy teaching, even though the "exploration group" exhibited less strategy knowledge. Equal performance, albeit with unequal explicit strategy knowledge, does suggest that at least a part of the performance on the TOH-R is related to implicit learning processes. The degree to which implicit versus explicit learning processes are involved in performance on various versions of the TOH continues to be a matter of debate. For example, Winter, Broman, Rose, and Reber (2001) found evidence for problem solving on the TOH via declarative (i.e., explicit) strategies in the two cases of amnesia they studied, and concluded that the TOH may not be a pure measure of implicit learning, as has been suggested in the literature.

It is of interest to consider not only *how* participants may induce move patterns and generate strategy knowledge, but also *what* types of strategy knowledge is induced. Of the four identified components of the strategy, the average number expressed in the post-test interviews was between two and three, and only 32% of participants stated all four elements. Those participants who understood the complete strategy had mean scores that were about three points higher on all performance measures. Interestingly, the 44% sample who expressed element three of the strategy exhibited scores that were indistinguishable from the group that recognized all four of the elements. The multiple regression analysis supports the conclusion that this component (i.e., building the smaller tower of disks on the "other peg" as the largest disk not currently on the goal peg is moved to the goal), expressed by the smallest number of participants, was nevertheless the most crucial element for success. Most participants appeared to understand that each major subgoal involved delivering the current largest disk to its goal as soon as possible, and to accomplish this, the smaller disks must be "moved out of the way". It was the notion that all smaller disks should be stacked in a pyramid configuration and *where* these disks should be stacked that posed the problem for most participants, and it is this lack of understanding that typically leads to the common errors one observes on the TOH task (Welsh, 1991; Welsh et al., 1995). Importantly, it is the understanding of the concept of "moving a smaller pyramid to the other peg" that distinguishes the recursive strategy from the sophisticated perceptual strategy (Simon, 1975) and, therefore, the best performers in this study appear to have a more fully realized recursive strategy. Both Altmann and Trafton (2002) and Welsh et al. (1995) have found the most frequent errors, as well as increases in pause times to occur at the beginning of each major recursive cycle, immediately after the largest disk has arrived at its goal destination. Such errors could be avoided by understanding the concept of the building a smaller pyramid characteristic, or, as has been suggested by other researchers, by an ability to "look ahead" several moves (Borys et al., 1982). The "look ahead" approach implies perception of the configuration of disks, which is closer to the sophisticated perceptual strategy as described by Simon (1975). Researchers disagree as to whether the task requires planning in the form of "look ahead," also known as "depth of search" or, alternatively, the ability to inhibit moves directly to the goal in order to make the counter-intuitive move (Goel & Grafman, 1995). In fact, it may be a combination of anticipating the results of future moves in working memory and inhibiting direct, but maladaptive, moves that allows the person to gradually appreciate the structure of the task and discover a systematic and effective strategy, such as goal recursion.

A limitation of this study is the fact that it is difficult to determine the precise nature of the covert strategy guiding the participants' problem solving, and the degree to which the participants who expressed knowledge of all four elements fully realized the goal recursion strategy. Goel et al. (2001) state that the Simon's goal recursion strategy can be applied only in the classic tower-to-tower configuration of the TOH problem and not in intermediate states of the move path. This is certainly the case if one views the recursive strategy as a fixed algorithm, such as that programmed into computer models of TOH solution, such as ACT-R (Anderson & Lebiere, 1998). In our view, the basic elements of the recursive strategy (e.g., get largest disk to final goal, stack smaller subpyramid of disks on other peg, etc.) can be applied successfully to intermediate problems, and even to flat-to-flat configurations, such as those included in the 60-item task. Given that perceptual information regarding the current and final goal configurations are referenced and monitored by the person, in Simon's framework our participants might be more appropriately classified as using the sophisticated perceptual strategy. This may be particularly true in the case of the more difficult flat-to-flat type of problem. It is true that the recursive strategy is optimally suited to solving a tower-ending problem, in which the largest disks are moved successively to a single goal peg. However, the same procedures can solve a flat-ending problem, which is an intermediate state on the way to the final tower-ending configuration. It is precisely the ability to "see" the flat-ending goal as an intermediate state of, or embedded in, the larger problem, that poses a serious difficulty for most participants. Anecdotally, participants who appear to have an understanding of the recursive strategy often behave as if they are caught by surprise when the current flat-ending goal is achieved a bit earlier than they expected. That is, they seem to be focused on the ultimate end of the goal recursion strategy: a stacked tower of disks on one peg. They begin executing the strategy and its rules in the deductive manner that Simon describes; that is, without constant reference to the current configuration of disks. Therefore, those participants who appear to have a fully realized goal recursion strategy may err by continuing past the intermediate goal on to the final tower configuration, unless they allow perceptual information (i.e., the intermediate goal on the card in front of them) to influence their responses. The current findings demonstrate that these flat-ending goal problems are clearly more difficult for these participants, perhaps because the goal recursion strategy that works so well in the tower-ending case does not apply as transparently in the flat-ending case.

It is not argued here that our participants have a full realization of the goal recursion algorithm as might be represented in a computer model. Instead, it is suggested that participants use the basic elements of the sophisticated perceptual strategy or the goal recursion strategy (i.e., if they have the concept of the "other peg") as a guide to achieve the subgoals in their appropriate order. For those using the sophisticated perceptual strategy, they may engage working memory process to "look ahead" in order to determine which would be considered the "other peg" in the more difficult flat-to-flat problems. A subset of participants was able to induce the rule of the "other peg" and then deductively apply this rule to all of the problems. Those participants who went through this hypothetico-deductive process were presumably the ones to express their explicit understanding of the "mini-pyramid" and where to place it. Recall that it was this third element of the strategy that was most predictive of performance. Also consistent with this proposal is the conclusion drawn by both Anderson and Douglass (2001) and Altmann and Trafton (2002) that TOH performance does not follow the "perfect memory goal stack" mechanism, as is found in the their computer program known as the ACT-R cognitive architecture (Anderson & Lebiere, 1998). Instead, both research groups suggested that performance appears to depend on the priming, activation, and retrieval of subgoals from memory.

The general disk-transfer task paradigm that is reflected in the many variants of the TOH, and its distant cousin the Tower of London (Shallice, 1982), have become popular neuropsychological measures of executive functions presumed to be mediated by the prefrontal cortex (e.g., Stuss & Benson, 1984). Although the two types of tasks are often treated as isomorphic due to commonalities in their general structure (i.e., a start state and a sequence of moves to achieve the goal state), the TOL does not involve the inherent recursive structure that characterizes the TOH task. Therefore, it has been suggested that the TOH may tap inductive reasoning and the TOL may tap a more pure form of planning (Goel & Grafman, 1995); a hypothesis that was recently confirmed via structural equation modeling (Welsh et al., 2002). Before these tasks can be used effectively within a larger neuropsychological assessment battery, the construct validity of each must be examined further. Most importantly, the question of discriminant validity, or the differential executive functions that may be tapped by each task, should be explored (e.g., Welsh et al., 1999). The current study indicates that inductive reasoning as it applies to discovering the recursive strategy is an important cognitive mediator to success on an expanded version of the TOH and contributes to the individual differences in performance observed among normal, young adults. These results are convergent with the notion that there is significant overlap between frontally-mediated executive functions, such as working memory and inhibition, and the cognitive construct of fluid intelligence (Pennington, Bennetto, McAleer, & Roberts, 1996). This perspective may substantially influence the development of future measures of executive functions in the field of neuropsychological assessment.

Acknowledgement

The authors wish to acknowledge Kasey Beatty, Sarah Devine, and Molly McGarraugh for their assistance with the data collection.

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39–83.

Anderson, J. R., & Douglass, S. (2001). Tower of Hanoi: Evidence of the cost of goal retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(6), 1331–1346.

Anderson, J. R., & Lebiere, C. (Eds.). (1998). *The Atomic Components of Thought*. Hillsdale, NJ: Erlbaum.

Bagley, A., Welsh, M. C., Retzlaff, P., Wolf, C., & Bryan, E. (2002). Towers of Hanoi and London: Contribution of procedural learning and inhibition. *Journal of the International Neuropsychological Society*, *8*(2), 229.

Best, J. B. (1990). Knowledge acquisition and strategic action in "mastermind" problems. *Memory and Cognition*, *18*(1), 54–64.

Borys, S. V., Spitz, H. H., & Dorans, B. A. (1982). Tower of Hanoi performance of retarded young adults and nonretarded children as a function of solution length and goal state. *Journal of Experimental Child Psychology*, *33*(1), 87–110.

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. Oton, E. Gamzu, & S. Corkin (Eds.), *Memory Dysfunctions: An Integration of Animal and Human Research from Preclinical and Clinical Perspectives* (pp. 54–71). New York: Annals of the New York Academy of Sciences.

Davis, H. P., & Klebe, K. J. (2001). A longitudinal study of the performance of the elderly and young on the Tower of Hanoi puzzle and Rey recall. *Brain and Cognition*, *46*(1-2), 95–99.

Devine, S., Welsh, M. C., Retzlaff, P., Yoh, M., & Adams, C. (2001). Explicit and implicit cognitive processes underkying Tower of Hanoi performance. *Journal of the International Neuropsychological Society*, *7*(2), 250.

Emick, J., & Welsh, M. C. (in press). Association between Formal Operational Thought and Executive Function as Measured by the Tower of Hanoi-Revised. To appear in *Learning and Individual Differences*.

Fuster, J. M. (1997). The prefrontal cortex. (3rd ed.). New York: Raven Press.

Gabrieli, J. D. E. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, *49*, 87–115.

Glosser, G., & Goodglass, H. (1990). Disorders in executive control functions among aphasic and other brain-damaged patients. *Journal of Clinical and Experimental Neuropsychology*, *12*(4), 485–501.

Goel, V., & Grafman, J. (1995). Are the frontal lobes implicated in "planning" functions? Interpreting data from the Tower of Hanoi. *Neuropsychologia*, *33*(5), 623–642.

Goel, V., Pullara, S. D., & Grafman, J. (2001). A computational model of frontal lobe dysfunction: Working memory and the Tower of Hanoi task. *Cognitive Science*, *25*(2), 287–313.

Goldberg, T. E., Saint Cyr, J. A., & Weinberger, D. R. (1990). Assessment of procedural learning and problem solving in schizophrenic patients by Tower of Hanoi type tasks. *Journal of Neuropsychiatry and Clinical Neurosciences*, *2*(2), 165–173.

Goldstein, F. C., & Green, R. C. (1995). Assessment of problem solving and executive functions. In R. L. Mapou, & J. Spector (Eds.), *Clinical neuropsychological assessment: A cognitive approach. Critical issues in neuropsychology* (pp. 49–81). New York: Plenum Press.

Hambrick, D. Z., & Engle, R. W. (2003). The role of working memory in problem solving. In J. E. Davidson (Ed.), *The psychology of problem solving* (pp. 176–206). New York: Cambridge University Press.

Handley, S. J., Capon, A., & Harper, C. (2002). Conditional reasoning and the Tower of Hanoi: The role of spatial and verbal working memory. *British Journal of Psychology*, *93*, 501–518.

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? Evidence from Tower of Hanoi. *Cognitive Psychology*, *17*(2), 248–294.

Lezak, M. D. (1995). *Neuropsychological assessment* (3rd ed.). London: Oxford University Press.

Lock, C., Welsh, M. C., Adams, C., & Kurtz, A. (2002). Tower of Hanoi: Influence of strategy instruction and extended practice on performance. *Journal of the International Neuropsychological Society*, *8*, 229.

Miyake, A., Freidman, N. P., Emerson, M. J., Witzki, A. H., & Howerter, A. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, *41*, 49–100.

Numminen, H., Lehto, J. E., & Ruoppila, I. (2001). Tower of Hanoi and working memory in adult persons with intellectual disability. *Research in Developmental Disabilities*, *22*(5), 373–387.

Pennington, B. F. (1997). Dimensions of executive functions in normal and abnormal development. In N. A. Krasnegor, & G. R. Lyon (Eds.), *Development of the prefrontal cortex: Evolution, neurobiology, and behavior* (pp. 265–281). Baltimore: Paul H. Brookes.

Pennington, B. F., Bennetto, L., McAleer, O., & Roberts Jr., R. J. (1996). Executive functions and working memory: Theoretical and measurement issues. In G. R. Lyon, & N. A. Krasnegor (Eds.), *Attention, memory, and executive function* (pp. 327–348). Baltimore: Paul H. Brookes.

Roberts Jr., R. J., & Pennington, B. F. (1996). An interactive framework for examining prefrontal cognitive processes. *Developmental Neuropsychology*, *12*(1), 105–126.

Scholnick, E. K., Friedman, S. L., & Wallner Allen, K. E. (1997). What do they really measure? A comparative analysis of planning tasks. In S. L. Friedman, & E. K. Scholnick (Eds.), *The developmental psychology of planning: Why, how, and when do we plan?* Mahwah: Lawrence Earlbaum.

Schunn, C. D., Lovett, M. C., & Reder, L. M. (2001). Awareness and working memory in strategy adaptivity. *Memory and Cognition*, *29*(2), 254–266.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London (Biology)*, *298*, 199–209.

Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, *7*(2), 268–288.

Stuss, D. T., & Benson, D. F. (1984). Neuropsychological studies of the frontal lobes. *Psychological Bulletin*, *95*(1), 3–28.

Welsh, M. C. (1991). Rule-guided behavior and self-monitoring on the Tower of Hanoi disk-transfer task. *Cognitive Development*, *6*(1), 59–76.

Welsh, M., Cicerello, A., Cuneo, K., & Brennan, M. (1995). Error and temporal patterns in Tower of Hanoi performance: Cognitive mechanisms and individual differences. *Journal of General Psychology*, *122*(1), 69–81.

Welsh, M. C., & Huizinga, M. (2001). The development and preliminary validation of the Tower of Hanoi—Revised. *Assessment*, *8*(2), 167–176.

Welsh, M. C., Huizinga, M., Granrud, M. A., Cooney, J., Adams, C., & Van der Molen, M. W. (2002). A structural equation model of executive function in normal young adults. *Journal of the International Neuropsychological Society*, *8*(2), 264.

Welsh, M. C., & Pennington, B. F. (1988). Assessing frontal lobe functioning in children: Views from developmental psychology. *Developmental Neuropsychology*, *4*(3), 199–230.

Welsh, M. C., Satterlee Cartmell, T., & Stine, M. (1999). Towers of Hanoi and London: Contribution of working memory and inhibition to performance. *Brain and Cognition*, *41*(2), 231–242.

Winter, W. E., Broman, M., Rose, A. L., & Reber, A. S. (2001). The assessment of cognitive procedural learning in amnesia: Why the Tower of Hanoi has fallen down. *Brain and Cognition*, *45*(1), 79–96.