Knowledge versus Experience in Financial Problem Solving Performance

DOUGLAS A. HERSHEY
Oklahoma State University

DAVID A. WALSH
University of Southern California

The present study examined the solution quality and information processing patterns of experts, novices, and "trained novices" as they solved a series of six complex retirement investment problems. The goal of the investigation was to better understand how prior problem solving experience and knowledge of the task influenced individuals' cognitive efforts. As expected, expert financial planners produced higher quality solutions than did novices. However, trained novices (individuals who attended a six hour educational intervention) were found to produce solutions that were twice as good as those generated by experts. A variety of analyses were conducted that focused on both the types of information individuals selected to solve the problems, and the efficiency with which they processed task information, in an effort to explain group differences in solution quality. Findings are discussed in terms of the potential benefits of offering cognitively engineered educational training programs in order to improve task-specific problem solving competency.
Keywords: Problem Solving, Expertise, Instruction, Retirement, Financial Planning, Decision

Research on expertise has witnessed a phenomenal growth in the past three decades (Ericsson & Smith, 1991; Sternberg, 1995). Early empirical examinations of expertise focused on the information search heuristics exhibited by experts and novices as they solved novel, well-structured problems (Chase & Simon, 1973; Chi et al., 1981; Gick & Holyoak, 1980). More recent research efforts, however, have employed complex, real-world tasks in an effort to better understand the relationship between knowledge and problem solving performance (Hegarty, 1991; Hershey & Farrell, 1999; Hershey et al., 1996; Lesgold & Lajoie, 1991; Voss et al., 1991; Walsh & Hershey, 1993). In fact, some researchers (Holyoak, 1991; Sternberg, 1995) have suggested that these two separate lines of work constitute different "generations" of research on expertise, each with its own goals, methods, and models.

The latter generation of research has provided us with numerous well-established findings regarding information processing differences between experts and novices. According to Glaser & Chi (1988), relative to novices, experts: (1) excel mainly in their own domain, (2) perceive large meaningful patterns in their domain, (3) solve problems quickly, (4) exhibit superior short-term and long-term memory, (5) see the deep structure in a problem, (6) spend time conducting a qualitative analysis of a

Current Psychology: Developmental • Learning • Personality • Social Winter 2000/2001, Vol. 19, No. 4, 261–291.

problem, and (7) have strong self-monitoring skills. However, the expertise literature is riddled with empirical inconsistencies and theoretical anomalies to the point that exceptions could be identified for each of the above generalizations (Holyoak, 1991).

Despite a burgeoning literature on expert/novice differences, relatively little attention has been paid to ways in which individuals can be trained to display expert-like performance. As Glaser (1989) points out, "How knowledge becomes organized and how the processes that accompany it develop with learning and experience are fundamental questions" (p. 271). Whereas it has often been suggested that it takes a minimum of ten years to develop expert problem solving abilities (see Ericsson & Crutcher, 1990 for a review), others have recognized the potential for enhancing, systematizing, or short-cutting learning experiences in ways that could lead to high levels of problem solving competence (Glaser, 1989).

The goal of the present article was to examine some of the factors that might facilitate the development of expert-like performance. The present research examines the effectiveness of training designed around a task analysis of a complex retirement investment task to improve problem solving performance. The task analysis, described later and shown graphically in Figures 1–3, was used: (a) to construct a six-hour training program focused on the core conceptual content of the problem, and (b) provide simulated experience in producing solutions to the problems. The effects of this training on the performance of novice problem solvers was assessed by examining both the quality of their solutions and their information processing patterns as compared with untrained novices and expert financial planners. Next, we outline the rationale behind the design of our research.

Two factors that are central to the development of expertise in solving problems are: (a) the acquisition of knowledge of the domain and (b) actual "production experience" at solving domain-specific problems. In many "hands-on" problem solving domains such as auto mechanics and plumbing these two factors generally accompany one another. That is, by working on cars and solving automotive problems you learn important lessons about how cars function and how they should be fixed. In other problem solving domains, however, such as legal adjudication, financial planning, and medical diagnosis, individuals are likely to follow a lengthy educational training program before facing their first real-world problem. The goal of this prolonged training period is to establish a strong knowledge network in order to ensure that the individual is equipped to deal with problems he or she may encounter. Hershey et al. (1990) have suggested that problem solving scripts develop through repetitive exposure to specific types of problems. Scripts are said to develop as the procedural steps in the problem solving process become ordered into a computationally efficient subroutine (Hershey et al., 1991). This script formation process is conceptually similar to the knowledge compilation process described by Anderson (1983).

From a training perspective there are weaknesses associated with the "prolonged education" approach described earlier. In domains such as law, finance, and medical diagnosis, a complex knowledge base is acquired over a course of years as the individual reads texts and journals and interacts with teachers and peers. This form of education leads to the piece-meal acquisition of numerous "mini-lessons" about the domain; however, these lessons are encountered and learned in an unsystematic fashion. This non-systematic experience acquired in training and work environments can result in gaps or "missing links" in the individual's mental representation of the domain. These missing links can have an adverse affect on the quality of the individual's scripts, and ultimately, a negative impact on the quality of one's solutions.

An alternative approach to training would involve providing the individual with a clear, coherent, well-designed representation of a task at the outset of training in order to establish a strong, veridical *mental model.* By introducing the individual to a *conceptual model* of the task, he or she will learn which pieces of information are important to consider while solving the problem, and the way in which those informational elements are functionally related. Norman (1983) describes a conceptual model as a high-quality, veridical model of the problem space, as would be defined by a group of experts in a given field. We believe that the level of understanding provided by this conceptual model training approach will provide problem solvers with a strong foundation from which to forge well thought-out solutions.

A research paradigm used in previous studies (Hershey et al., 1990; Walsh & Hershey, 1993; Hershey et al., 1998) challenges subjects to solve one or more hypothetical retirement financial planning problems. Retirement planning is one area in which complex problems provide cognitive challenges to large segments of the population (Gregg, 1992; Hershey, 1995). Financial planning problems are interesting for cognitive investigations for many reasons: knowing *which* variables to consider in solving a financial problem is just as important as knowing *how* to use those variables to reach a solution. Furthermore, there are many variables relevant to solving financial planning problems and those variables interact with one another in a dynamic fashion, requiring multiple issues to be considered in combination. Thus, the comprehensive and veridical nature of a problem solvers' domain-specific mental model (Gentner & Stevens, 1983) is a very important element in determining the quality of one's solutions.

In the retirement planning studies cited earlier we used a two-phase experimental procedure. In the first phase subjects specify the information they would need to solve a retirement planning problem, and in the second phase, they solve the problem using specific values for the information requested during phase one. Our primary interest is in how subjects use information to solve problems, and how their information use determines the quality of their solutions. We derive many of our dependent measures of information use from problem solving process maps (PSPMs) of a subject's performance. These process maps are graphic representations of the sequence in which subjects consider information while solving a problem.

This methodology has been useful in identifying information processing differences between expert and novice financial planners. Hershey et al. (1990) found that experts used information in a more goal directed and efficient manner than novices. Relative to novices, experts knew what information to use to solve a problem before they began, and they used fewer but more informative variables. The performance of novices, in contrast, suggested that they lacked a well-defined mental model of the domain. Relative to experts, they lacked knowledge of what information they would need to solve the problem at the outset, and they used more but less informative variables to reach solutions. In addition, the novices took longer to solve the problems and they reconsidered the same information in many recursions.

The present study follows from our earlier work on expert/novice differences in complex problem solving. The goal is to examine the extent to which the conceptual model training program described earlier can improve retirement planning performance. Most studies of expertise fail to account for performance at intermediate levels of problem solving competence, focusing only on the performance of experts and novices. As in previous studies, we again plan to examine the performance patterns of experts and novices as they solve a series of complex retirement planning problems. However, a novel contribution to the complex problem solving literature will involve the addition of a "trained novices" group to the study (hereafter referred to as trainees), who are novices trained to understand the conceptual model of the task and given some simulated experience in solving actual problems, prior to solving problems on their own. These trainees will be unique in that they will possess a great deal of knowledge about the task without the benefit of prior task-specific problem solving experiences.

Participants in all three groups (novices, trainees, and experts) will solve a series of six retirement investment problems. By contrasting the initial performance levels of the three groups we hope to gain a better understanding of how knowledge of the task affects problem solving performance. By examining changes in the performance of the three groups over trials, we hope to gain a better understanding of how experience affects complex problem solving performance.

The research reported in this article addresses a number of questions raised, but not answered, in our earlier work. We wondered if the differences we found between experts and novices solving a single retirement planning problem would remain stable when they solved a series of similar problems. Our hypothesis was that the disorganized search patterns of novices would show increased goal directedness and efficiency across trials as their mental model of the problem evolved with experience. In contrast, we predicted that experts' information use patterns would remain relatively stable, indicating their solutions were determined by a highly refined mental model.

The second question addressed in this research was how training in the conceptual structure of the task would affect both the solution quality and information use patterns of the trainees. We predicted that trainees would have a good mental model of the problem but lack a clear set of rules for reaching a solution. Thus, we predicted they would use important variables to reach a solution for their first problem, but we expected that their early solutions would be inefficient as compared to those of experts.

METHOD

Task Analysis

Conceptual Model of the Task. The object of the 401k task is to determine whether a hypothetical investor should invest money in a 401k retirement plan. 401k plans are voluntary retirement savings accounts offered by an increasingly large number of companies. In the present study, subjects were required to solve a series of six 401k investment problems.

Hershey et al. (1990) identified three higher-order conceptual issues that should be addressed when solving this type of individual retirement investment problem: (a) whether the individual has a need for additional savings during retirement, (b) whether there are surplus funds available for investment purposes, and (c) whether the characteristics of the account are suitable, given the needs of the investor. These three issues comprise the major dimensions of the 401k problem space. The complete conceptual model of the 401k task included 73 different pieces of information relevant to generating an accurate solution.

It is important to understand that each of the six 401k investments have "correct answers" when specific values are specified for each of the 73 relevant variables. In order to arrive at a correct answer for each of the six problems some definitive values must be assumed for future values of variables that cannot be fully predicted, such as inflation rates, investment returns, and personal life expectancies. Specific values for all such variables were specified for each of the 73 variables in each of the six data sets developed for this research. Using these specified values, and the logic of the task analysis, a single correct investment amount could be obtained for each hypothetical scenario.

Information related to the issue of determining the hypothetical individual's retirement need are represented in Figure 1. On a conceptual level, retirement need is determined by balancing projected retirement expenses against anticipated retirement income from various sources (e.g., social security, pensions, and personal savings). Retirement expenses are mediated by the expected rate of inflation during the retirement period and the individual's anticipated longevity. Furthermore, under certain conditions, retirement income may be reduced by federal and state taxes. Figure 1 is organized hierarchically with the most critical piece of information, the total retirement need, on the left, and lower level variables on the right. The model is arranged in such a way that lower level information can be arithmetically combined to calculate the values of the higher level nodes.

Variables related to determining the affordability of an investment are represented hierarchically in Figure 2. Again, the most critical node in the hierarchy (affordability of the investment) is on the left, and lower level information is represented on the right. The dominant branch in this hierarchy involves determining the individual's surplus income (i.e., current net income less expenses). Other factors that come into play in determining the affordability of an investment are the investor's capital assets and the tax savings that are realized by contributing to a 401k plan.

FIGURE 1
Variables related to the issue of determining an individual's retirement need. Abbreviations (in parentheses) are used in the information use density plots (Figs. 6, 7, & 8).

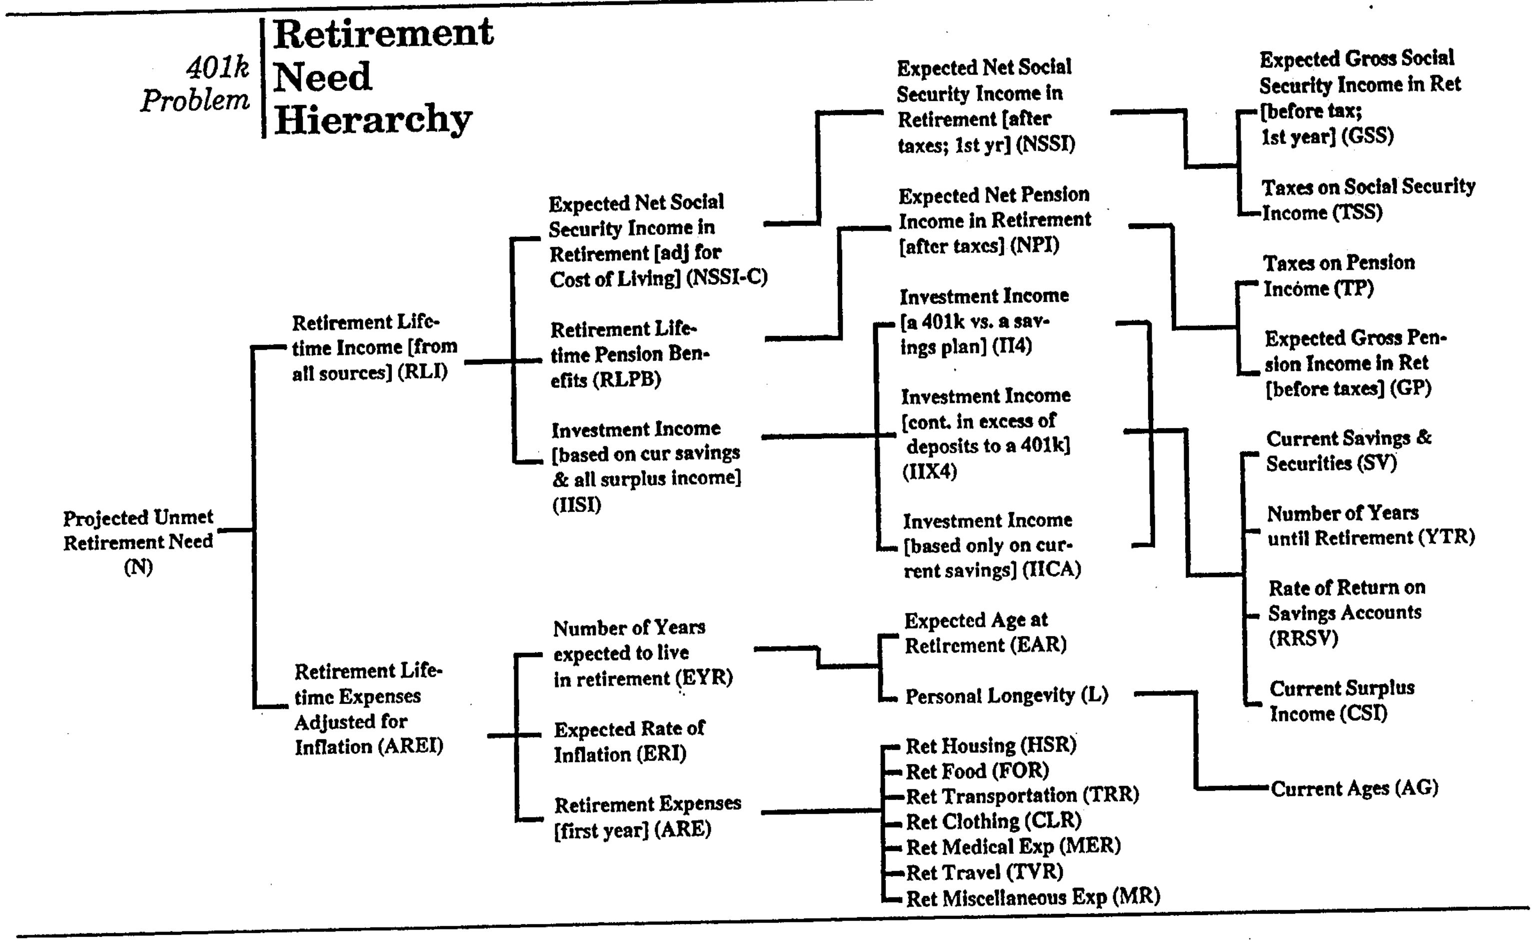

FIGURE 2
Variables related to the issue of determining the affordability of an investment. Abbreviations (in parentheses) are used in the information use density plots (Figs. 6, 7, & 8).

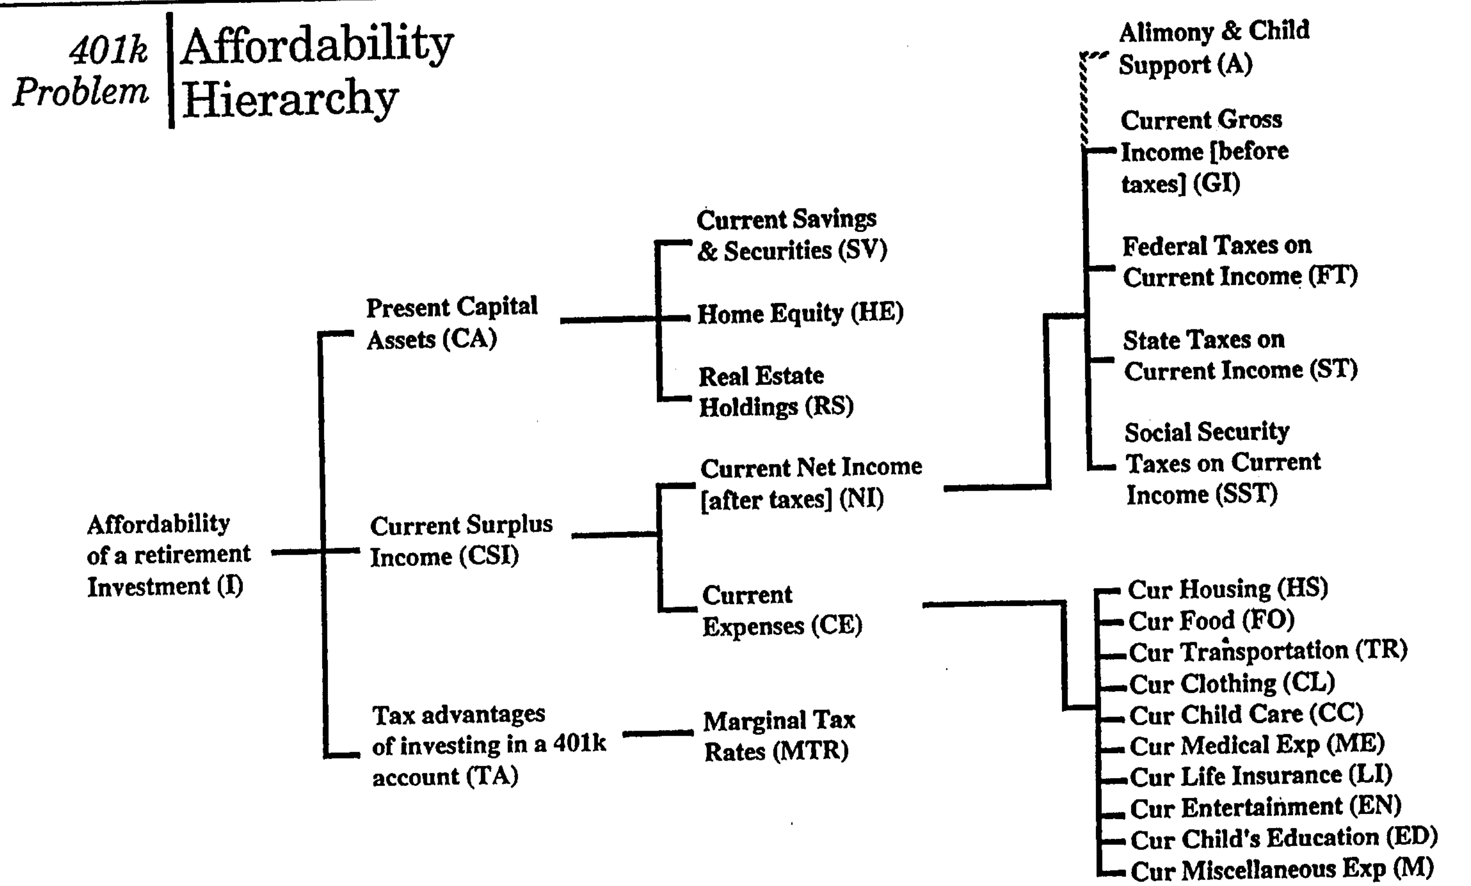

FIGURE 3
Variables related to the issue of determining the suitability of an investment in a 401k plan. Abbreviations (in parentheses) are used in the information use density plots (Figs. 6, 7, & 8).

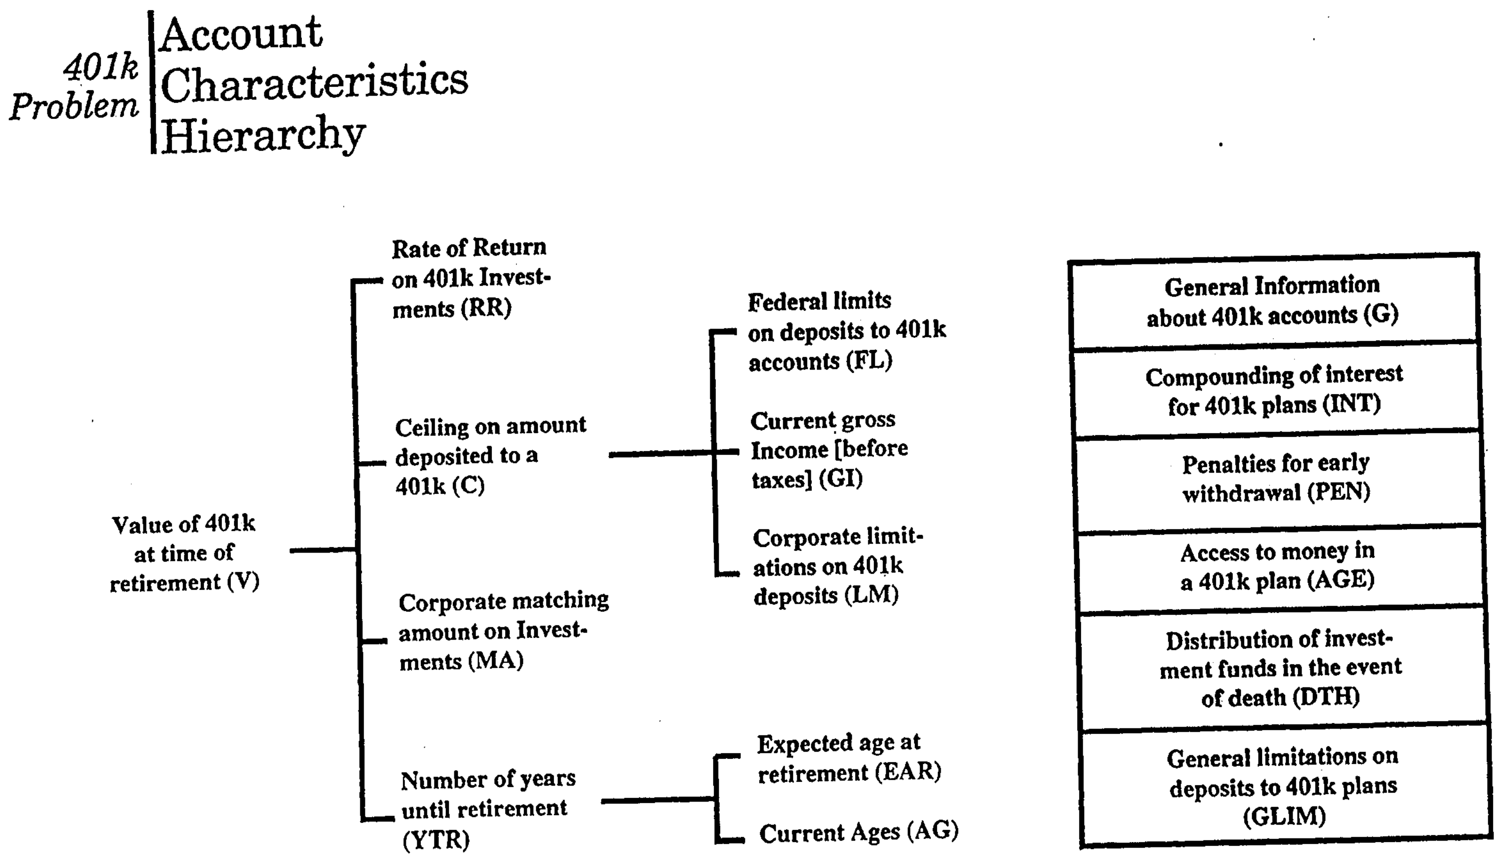

Information related to the third major issue, the characteristics of the investment vehicle, are represented in Figure 3. In making long-term retirement savings decisions, factors such as the rate of return on investments, the amount of matching contributions made by the employer, the ceiling on amounts deposited to the plan, and the number of years during which investments can be made must all be considered if the correct, deterministic solution to a specific scenario is to be achieved. Other pertinent account-related factors include the severity of penalties for early withdrawal, age-related policies governing withdrawals, and the manner in which interest is compounded.

Validation of the Model. In an effort to validate the investment model described earlier, three tax and financial planning specialists were hired as consultants to review the complete set of experimental materials. These experts were recruited from a group of financial planners working in the tax planning department of a large international accounting firm. It was from this same work group of financial planners that our sample of expert participants were drawn. In the first step of the validation process consultants reviewed the conceptual model of the problem in an effort to identify any conceptual flaws. All three consultants found the three major issues depicted in the model to be reasonable, well-organized representations of the 401k problem. Furthermore, the consultants judged the model sufficiently comprehensive to determine ideal investment values for each of the six sets of parameters we developed for the hypothetical investors.

In the second step of the validation process the consultants computed "correct" or "optimal" investment decisions from the full sets of 73 parameters for each of the six hypothetical investment scenarios. To facilitate their analysis of each scenario, all of the 73 values were presented on three large poster boards, one board devoted to each of the three conceptual issues shown in Figures 1–3. All three of the consultants provided identical solutions to five of the six problems. (An important factor leading to this high level of agreement among the three consultants was that precise values for each variable and scenario were specified that might otherwise have been estimated differently by different consultants. Examples of such variables were future inflation rates, investment returns, and life expectancy rates.) The obtained values for the five agreed upon problems were used as the optimal investment amounts. Two of the three consultants agreed upon an investment amount for the sixth problem, and this mutually agreed-upon value was used as the optimal investment. This validation procedure provided support not only for the conceptual organization of the investment model, but more importantly, it provided a criterion value against which the subject-generated solutions could be scored. Additional details regarding this validation process can be found in Hershey (1990).

Participants

Participants were recruited from two sources: The novices and trainees were selected from among undergraduate students attending a large metropolitan university, the experts were recruited from financial planners working at a large international accounting firm. Novices were fourteen individuals (7 males; 7 females) with a mean age of 19.0 years; SD = 0.88. At the time of testing they had completed a mean of 13 years of formal education. Trainees were 15 individuals (8 males; 7 females) selected from the same university population with a mean age of 18.2 years; SD = 0.80. At the time of testing they too had completed a mean of 13 years of formal education.

While a few of the student participants had taken an introductory course in economics, none had any specific training or experience in personal financial planning. In this sense, these participants could better be described as "naive," rather than "novices," since problem solving researchers often use the latter term to describe people knowledgeable about a problem but lacking extensive practice and skill in its solution (Chi et al., 1988). The students who served as novices and trainees in this experiment lacked both knowledge and experience in solving financial investment problems prior to their involvement in this study.

Expert subjects (mean age 29.2 years; SD = 4.56) were sixteen experienced financial planners (11 males; 5 females) who were employed in the tax planning department at a large international accounting firm. All but one individual held a college degree in business, accounting, or economics. As many of the experts were certified public accountants or certified financial planners, they would have also received substantive in-service training during their tenure in the field. These sixteen subjects were the most knowledgeable individuals from among a group of twenty-seven professional financial planners who were tested as part of a larger research project. Their knowledge of financial planning was assessed using a comprehensive measure of financial knowledge (described in detail later).

All subjects voluntary participated in the study. Novices received ten dollars per hour for their participation. Trainees also received ten dollars per hour for the problem solving session, and seven dollars per hour for attending two three-hour training sessions. Experts received a flat fee of thirty dollars for their participation. Attrition was low—only one subject in the training group failed to complete both training sessions, thereby making himself ineligible for the test phase of the project.

Design

A 3 (groups) x 6 (trials) mixed between-within factorial design was employed. The three levels of the grouping factor were novice, trainee, and expert. The six levels of the trials factor were the six investment problems each subject solved. A partial incomplete counterbalancing technique with random start was employed (Shaughnessy & Zechmeister, 1985) to determine the order of presentation of the problems. Five of six problems were counterbalanced across the first five trials. The sixth problem was held constant for purposes of comparison across groups.

Educational Intervention

Members of the training group attended two three-hour training sessions designed to provide them with a basic understanding of financial planning for retirement and specific training in solving the 401k problems used in this research. During these sessions, Ss were introduced to the conceptual model of the 401k problem and shown two complete hypothetical instances of how task information could be combined to determine whether a retirement investment should be made.

Materials

Scenarios. Six different hypothetical scenarios were created which were designed to represent a range of plausible retirement planning situations. Personal and situational characteristics of the hypothetical investor were varied across scenarios (e.g., age, marital status, whether the person owned or rented housing, and whether or not they had children). Furthermore, the problems were designed to vary in (a) the extent to which there was a need for additional money in retirement, and (b) the affordability of a contribution to a 401k plan. Differing values on these two dimensions limited the amount of the contribution that could (or should) be made on behalf of the hypothetical investor.

Information Cards and Display Apparatus. Based on the task analysis of the 401k problem, 73 variables (shown in Figures 1, 2, and 3) were created for each of the six hypothetical cases. The variable name and the value for each of the parameters were printed on 4 x 6 index cards. The variable name was also printed on the reverse side of each card (the rationale for this is described in the procedure section later). Because it was anticipated that some Ss would request numerous parameters, a large information board was used that was designed to hold as many as 48 cards. The information board approach used in this study is identical to the procedure used in Hershey et al. (1990).

It should be noted that important differences exist in the informational content of variables shown on the left and right sides of Figures 1–3. Variables on the far right of each figure represent "raw data" parameters for each of the six hypothetical scenarios, such as Current Age, Expenses, Tax rates, etc. (see Figure 1). Variables to the left represent aggregations of this raw data into higher level values such as Retirement Life-Time Income, Retirement Life-Time Expenses, and Projected Unmet Retirement Need (see Figure 1). These left most variables can be thought of as the output of calculations and aggregations that experts solving real world problems might obtain either with the assistance of financial calculators or computer support. We computed these values, for each scenario, using computer assistance. These high-level, computed variables were available to all subjects, but only upon their specific request.

Procedure

All subjects were tested on an individual basis. Prior to testing, participants were told that they would be asked to solve six retirement investment problems, and they were given a brief, general description of a 401k savings plan. Participants were asked to imagine that they were serving as financial consultants for the six hypothetical individuals, and it was their job to recommend whether each individual should contribute funds to a 401k plan.

The experimental task was conducted in two phases. During the first phase of each trial, subjects were asked to identify the specific pieces of information they would need to make a sound investment decision. During the second phase of the task, subjects were provided with index cards containing the information they had requested, and then asked to determine how much money (if any) the hypothetical investor should contribute to the plan in the upcoming year.

The cards the subjects requested were placed on the information board with only the name of the variable showing. The parameter for the variable was printed on the back of the card. Subjects were allowed to remove only one card from the board at a time, and the card had to be returned to the board (with only the variable name showing) before subsequent cards were viewed. During the information search phase of the task, subjects could look at any or all of the cards they had requested, and they could look at individual cards as often as they wished. If more information was requested during the information search phase of the task, additional cards were provided. Virtually all of the subjects' requests for information were anticipated based on the thorough prior task analysis of the problem. A pencil, paper, and a simple calculator were placed on the table and subjects were told that they could make notes or use the calculator as needed. Each of the six problems were solved using the same two-stage procedure—requesting information during phase one, and using that information to make an investment decision in phase two. At no time during the test session were subjects given feedback regarding the quality of the solutions they generated.

After the final problem was solved, subjects completed a questionnaire that assessed self-perceptions of task performance, and measured their knowledge of financial and retirement planning issues. The later measure served as an objective index of domain-specific knowledge. Once subjects had completed the questionnaire they were debriefed and paid for their participation.

Knowledge Assessment

A 32 item financial knowledge and retirement planning test adapted from Hershey et al. (1990) was administered in order to assess group differences in domain-specific knowledge. The test focused on three different sets of issues: (a) knowledge of general financial trends, (b) knowledge of 401k accounts, and (c) knowledge of issues related to financial planning for retirement. The psychometric properties of the scale were evaluated using an extended sample of 87 individuals who were participants in a larger study of complex problem solving.[1] Alpha coefficients were .75 for the full scale, .72 for the financial knowledge subscale, .62 for the account knowledge subscale, and .59 for the retirement knowledge subscale—all of which are of suitable magnitude for a research instrument of this type (Nunally, 1978). The mean corrected item-total correlation for the full-scale measure was .32. Finally, a factor analysis provided empirical support for the three conceptual knowledge categories engineered into the test. The above analyses indicated that the psychometric properties of the test were sound and in line with the conceptually driven structure of the measure. These findings serve to pave the way for an examination of group differences in domain-specific knowledge.

It is important to empirically demonstrate that the three groups differ in terms of their knowledge of the task before examining between-group differences in problem solving performance. Planned comparisons of the knowledge test scores between pairwise combinations of groups revealed that experts' scores ($M = 74\%$) were significantly higher than the scores of trainees ($M = 66\%$) and novices ($M = 39\%$), $t\ (26) = 2.69$, $p < .05$, and $t\ (28) = 12.75$, $p < .01$, respectively. Furthermore, the 27 percent disparity in the scores of trainees and novices was significantly different, $t\ (26) = 8.14$, $p < .01$, a finding that provides empirical support for the efficacy of the training program.

RESULTS

The results are presented in three parts. In the first section group differences in the quality of subjects' investment decisions are reported. Section two focuses on group differences in the selection of task information. The third section focuses on group differences in how subjects processed task information.

FIGURE 4

Mean decision quality values (in dollars) for the three groups with standard error bars shown. These values are six problem averages of the absolute value of the deviation from the optimal solution. Lower values indicate more accurate solutions. The quality of the trainees' decisions were twice as good as those made by experts', and ten times better than those made by novices.

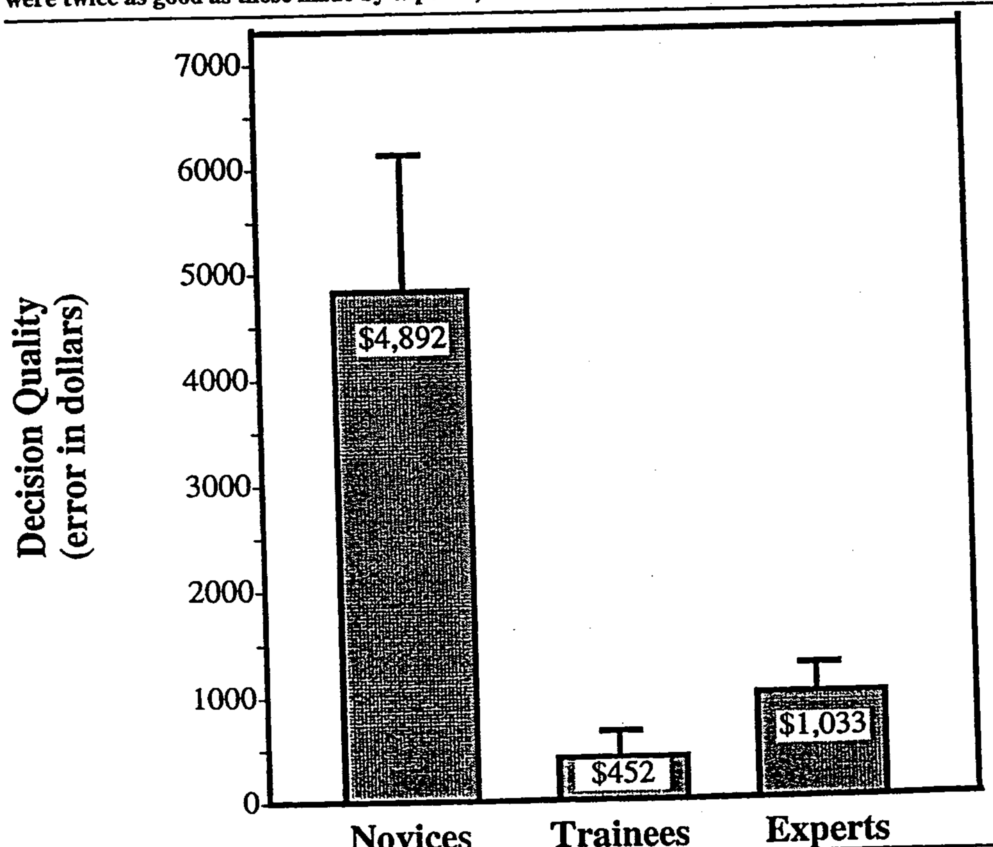

Quality of the Investment Decisions

Optimal investment amounts for each of the six problems were established on an *a priori* basis when designing the parameters for each scenario. Furthermore, these solutions were reviewed and validated by a three-person panel of expert consultants, as described previously. The six optimal investment values, which ranged from $0 to $7,000, were used as a "gold standard" against which subjects' solutions were compared. An index of the overall quality of each subject's performance was created by aggregating across the six problems the absolute value of the deviation from each optimal investment amount. Absolute values of the deviations were taken because otherwise, overinvestment errors, which carry a negative sign, would cancel out underinvestment errors, which carry a positive sign. Thus, the absolute value approach helps to ensure a truer aggregate error value.

Figure 4 shows the mean quality of the solutions for each of the three groups as represented by the summed deviations from the optimal decisions. As expected, novices generated the poorest investment decisions, with an average error of $4,892 per trial. Expert's deviations averaged $1,033 per trial, and trainees, surprisingly, made the best overall decisions with an average investment error of $452 per problem. Planned comparisons between the three groups revealed that the solutions of both experts and trainees were superior to those of novices, $t\,(28) = 3.07$, $p < .01$, and $t\,(26) = 3.35$, $p < .01$, respectively. A third comparison indicated that trainees solutions were significantly better than those of experts, $t\,(28) = 2.05$, $p = .05$.[2]

It was anticipated that experts and trainees' superior knowledge of the domain would allow them to make better decisions than those made by novices, which was indeed the case. The unexpected outcome, however, was that trainees produced higher quality decisions than experts. The relative quality of the groups' solutions can be better understood by examining (a) the specific types of information subjects selected to solve the problems, and (b) the way in which that information was processed. The following two sections of the article focus on these two issues.

Selection of Task Information

Three different techniques were used to characterize the types of information subjects selected to solve the problems: the calculation of hierarchy scores, the development of group information use density plots (IUDPs), and the computation of group-based similarity scores. Each of these measures of information selection are described separately in the following paragraphs.

The hierarchy score is a measure of the hierarchical level of information a subject considers when solving a problem. "Levels" are defined in terms of the various levels of task information shown in the three information hierarchies presented in Figures 1, 2, and 3. An individual's hierarchy score on a given trial is the average (mean) level of information considered during the problem solving process. Small values on this variable indicate the use of information from the upper levels of the task hierarchies (i.e., nodes on the left of the hierarchies such as Projected Unmet Retirement Need in

Figure 1), whereas large values indicate that lower-level information was selected from the bottom levels of the task hierarchies (e.g., Current Transportation Expenses in Figure 2). Scores on this variable could range from a low value of 1.0, to a high value of 5.0.

Figure 5 shows the mean hierarchy scores for each of the three groups plotted across the six trials. A visual inspection of the figure indicates that trainees used higher level information than both experts and novices to solve each of the six problems. A 3 (groups) x 6 (trials) analysis of variance (ANOVA) using the hierarchy scores as the dependent measure revealed significant main effects for both groups and trials, F (2, 41) = 76.61, $p < .01$, $MSe = .57$, and F (5, 205) = 10.17, $p < .01$, $MSe = .06$, respectively). Also identified was a significant group by trials interaction, F (10, 205) = 5.05, $p < .01$, $MSe = .06$, that was due to the marked decrease in trainees' scores over trails relative to those of experts and novices. Unlike the other two groups, which used a consistent level of information across trials, trainees used increasingly higher-level information on successive trials (information that represented computed aggregation of low-level information into more useful, high-level information).

Information use density plots (IUDPs) were also created to depict subjects' information use profiles. The IUDPs shown in Figures 6, 7, and 8 are a graphic representation of the problem space (see Walsh & Hershey, 1993 for additional details on this

FIGURE 5

Mean hierarchy scores for each of three groups across the six problems (N = Novices, T = Trainees, and E = Experts). Relative to trainees, novices and experts used low level information to solve the problems. Additionally, trainees were found to eliminate their consideration of various pieces of level information as they gained experience with the task (as indicated by the negative slope of their scores over trials), a strategy not found in the performance of the other two groups.

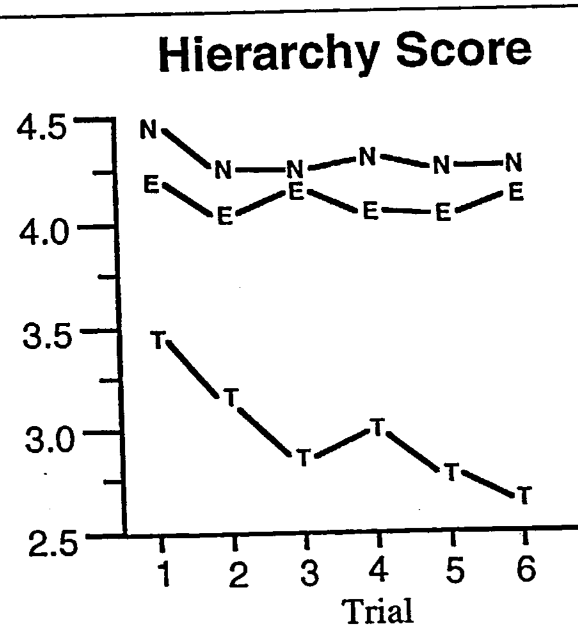

technique). Abbreviations in these plots correspond to those found in the conceptual model of the problem (see Figures 1, 2, and 3). Figures 6, 7, and 8 indicate the types of information sets the novice, expert, and trained groups used to solve the sixth problem, respectively. These three IUDPs are group aggregates of the information subjects selected to solve the problem presented on the sixth trial.[3]

Nodes on the density plots are differentially shaded to indicate the percentage of subjects within a group who selected particular pieces of information. Nodes with dark shading indicate cues that a large percentage of subjects within a group selected. Nodes with light shading are those variables that few members of a group considered. Nodes that are unshaded are those variables that no one selected. The bar at the base of each figure indicates the level of shading associated with different percentages of subjects who requested the various pieces of information.

One of the most striking features of Figures 6, 7, and 8 is that only a small subset of the total pool of 73 pieces of information were considered. Novices (Figure 6) tended to select lower-level cues found near the bottom of the task hierarchies. Experts (Figure 7) showed a similar pattern of information selection. A small percentage of experts selected higher-level cues, as indicated by the light shading in the 'I' and 'V' nodes at the top of the middle and bottom hierarchies. However, as a group, experts considered roughly the same number of variables as novices, and those variables tended to be those toward the middle and bottom of the information hierarchies. The density plot for the trainees (Figure 8) is clearly different from the other two plots. Relative to experts and novices, trainees used about 33 percent fewer variables from among the set of 73 available, however, the majority of variables they selected were from the top of the task hierarchies. This finding suggests that trainees' mental models of the problem led them to select a few, highly relevant pieces of information that enabled them to generate relatively high quality solutions.

In addition to the analysis of hierarchy scores and IUDPs, a third measure of information selection was developed in order to quantify the degree of overlap in information selected by the three groups. This measure of information selection is based on a novel application of the Pearson correlation coefficient. In this context, r values can be thought of as *similarity coefficients*, in that they represent the extent to which information selected by one group is similar to information selected by another. On a conceptual level, values of these coefficients address the question—to what extent do the information use density plots overlap (or correlate) with one another? The similarity coefficients were computed from two vectors of percentage scores (i.e., the percentage of individuals within a group who selected particular pieces of information) for pairwise combinations of groups. These vectors are different from those typically used to compute correlation coefficients in that the rows represent each of the 73 variables of the problem space—not subjects, as is typically the case when computing r. A high r value indicates that two groups used similar sets of information to solve a problem, whereas a low or near-zero correlation indicates low between-group agreement on variables selected.

Novice and expert groups selected the most similar sets of information to solve the sixth problem ($r = .79$), and members of both of these groups selected cues that were

FIGURE 6

The information use density plot for the novice group on the sixth problem. The structure of the plot corresponds to the conceptual model of the problem shown in Figs. 1, 2, and 3. Different levels of shading found in the node areas correspond to the proportion of subjects in the group who selected a particular variable (percentage ranges are found in the shaded bar at bottom of figure). Note that novices tended to ignore information found near the top of the information hierarchies. Light shading across the nodes activated indicate a lack of consensus among members of this group regarding the most important variables to consider.

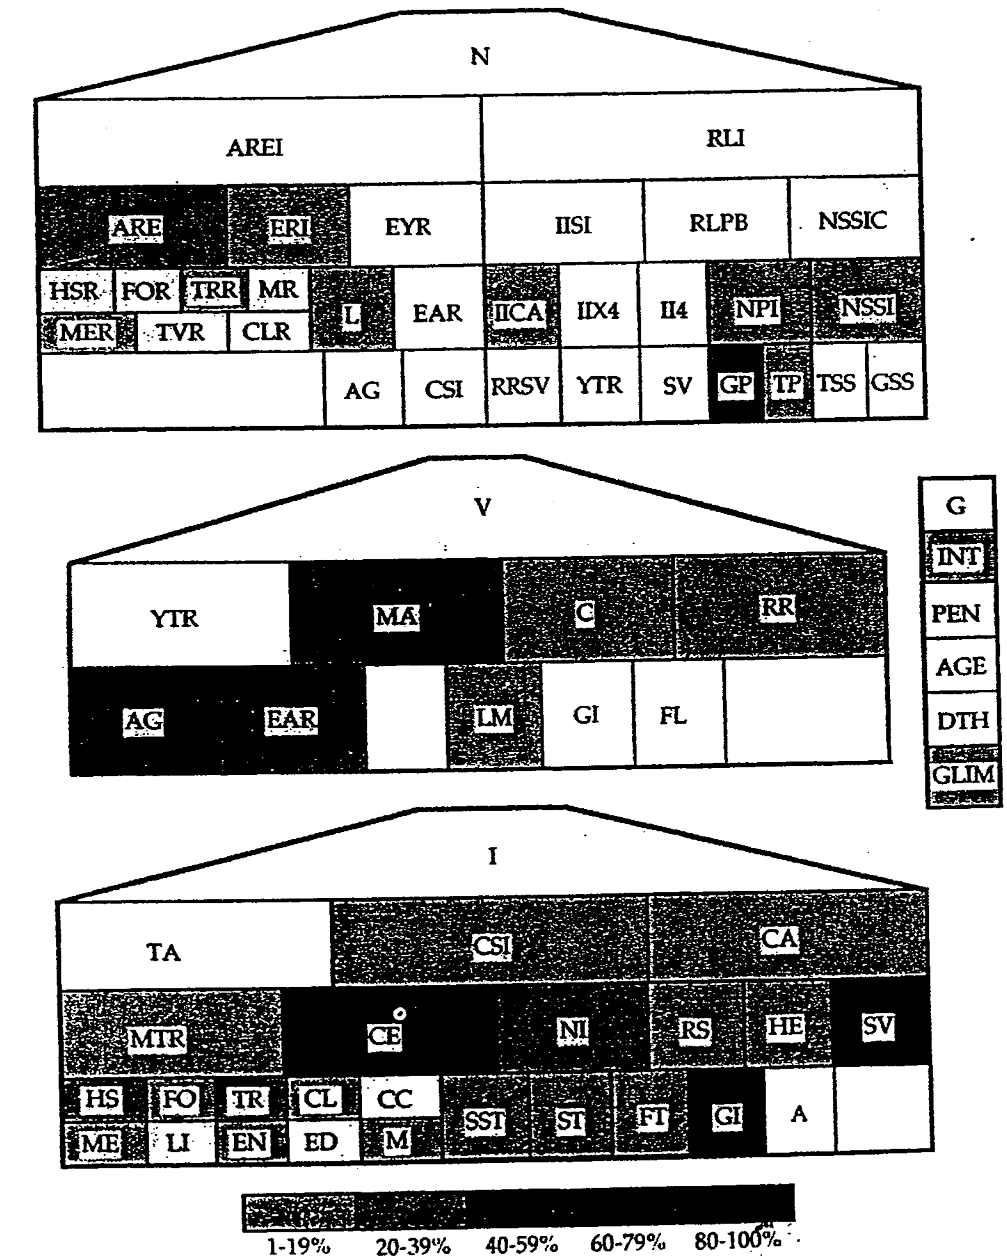

FIGURE 7

The information use density plot for the expert group on the sixth problem. The structure of the plot corresponds to the conceptual model of the problem shown in Figs. 1, 2, and 3. Different levels of shading found in the node areas correspond to the proportion of subjects in the group who selected a particular variable (percentage ranges are found in the shaded bar at bottom of figure). Experts appeared to consider more higher level variables than novices (cf., Fig. 6), a finding which can also be seen by an inspection of the hierarchy scores in Figure 5.

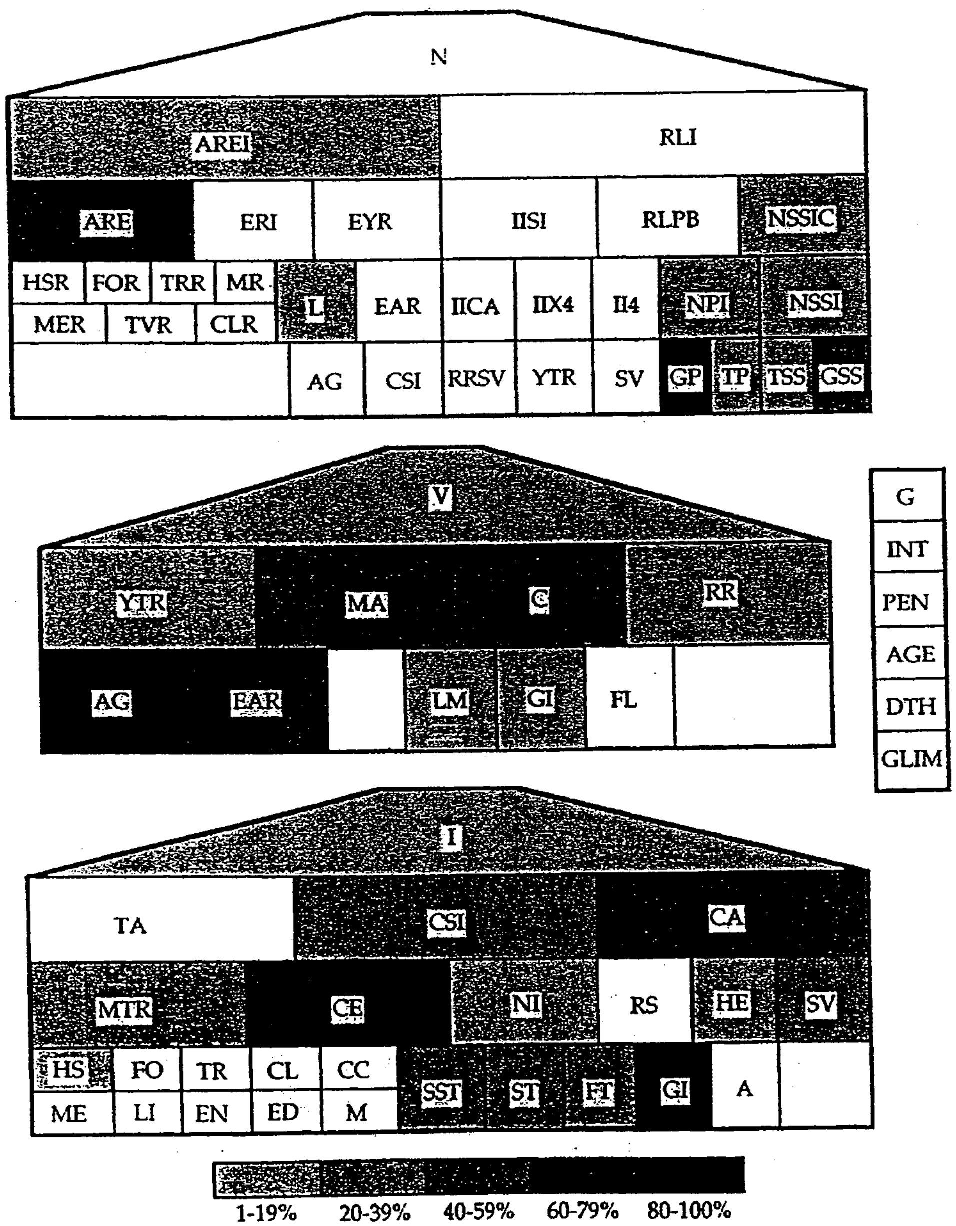

FIGURE 8

The information use density plot for the trained group on the sixth problem. The structure of the plot corresponds to the conceptual model of the problem shown in Figs. 1, 2, and 3. Different levels of shading found in the node areas correspond to the proportion of subjects in the group who selected a particular variable (percentage ranges are found in the shaded bar at bottom of figure). Note that by the sixth problem trainees had focused their search on the most critical (higher-level) information contained in the hierarchies. Moreover, the relatively dark shading in this figure suggests that there was a high level of within-group agreement regarding the variables that should be selected.

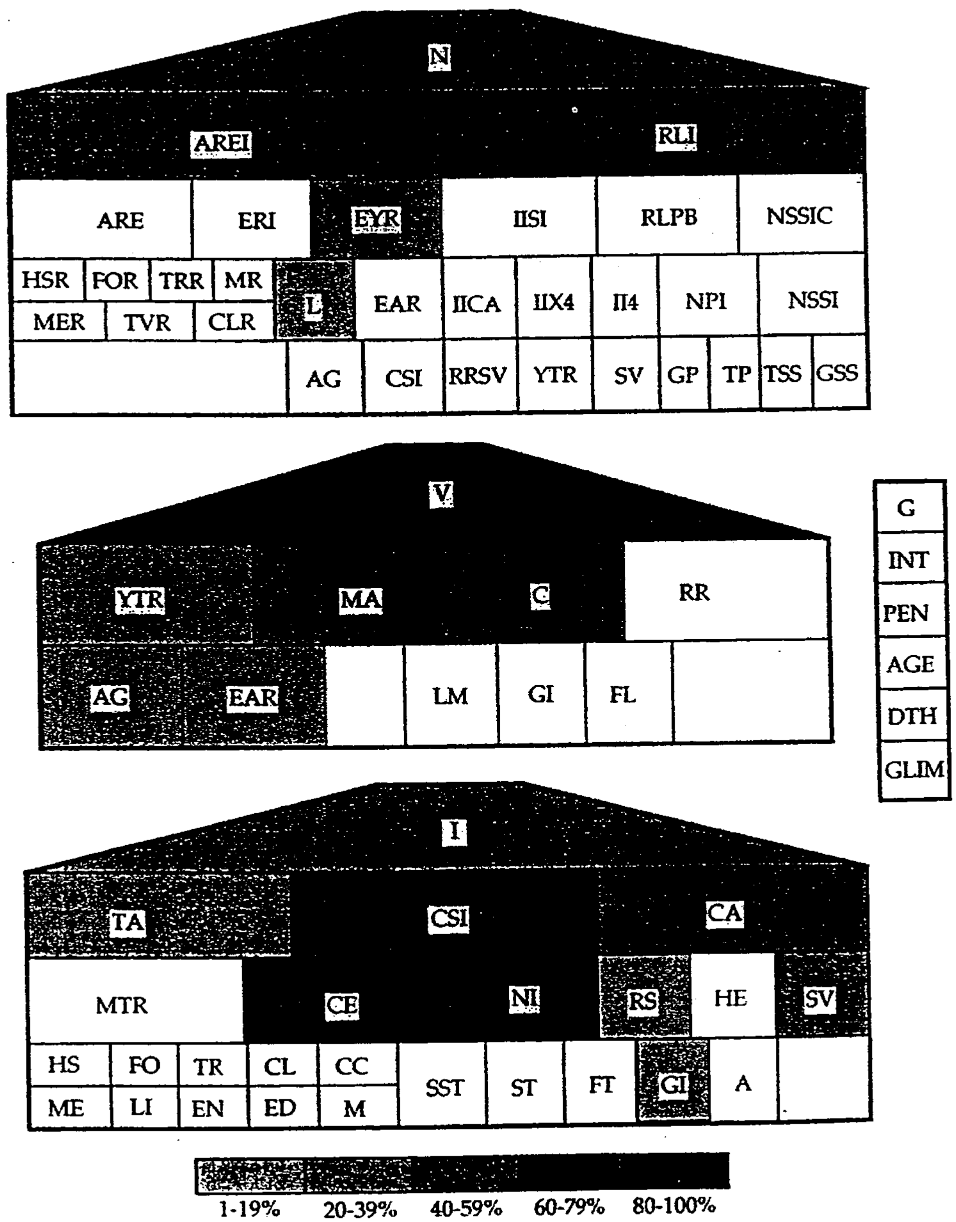

dissimilar to those selected by trainees (novices with trainees, $r = .27$; experts with trainees, $r = .33$). Tests of differences between dependent correlations revealed that the novice/expert correlation value was significantly different from the novice/trainee value, $t\ (70) = 7.46$, $p < .05$, and the novice/expert value was significantly different from the expert/trainee value, $t\ (70) = 6.45$, $p < .05$. The novice/trainee value was not found to be different from the expert/trainee value, $t\ (70) = 1.06$, *ns*. It is worth mentioning, however, that all three of these coefficients were statistically different from zero, which indicates that all three groups employed a small, similar core set of information.

The information selection patterns described earlier help to explain the observed group differences in decision quality. The hierarchy score and density plot data lend support to the notion that the high quality decisions made by trainees were in part due to the fact that they used high-level information. Further support for the relationship between information selection and decision quality can be found in the information use patterns displayed by experts and novices. Experts selected moderately low-level information to solve the problems, and accordingly made poorer investment decisions than trainees. Novices, who as a group selected the lowest-level information, were found to make the poorest investment decisions.

The information selection data provide insights into how the groups' problem solving performance differed. In the following section of the article data are presented that describe differences in how the groups processed task information.

Processing of Task Information

Six problem solving process maps (PSPMs; Hershey et al., 1990) were created for each participant in order to trace their step-by-step information search behavior. These maps, which like the IUDPs are based on the structure of the conceptual model of the problem, use arrows to represent the sequential nature of the information search process. In creating PSPMs for each problem the variables selected by a subject are highlighted on the map and then interconnecting arrows are drawn between nodes. The resulting graphic provides a "snapshot" representation of a subject's complete information search process. These PSPMs were then used as a source of raw data for a number of the different process-oriented analyses reported in the following text, and summarized in Table 1.

Panel "A" in Figure 9 shows the number of primary variables subjects requested to solve each of the six problems. Primary variables are those pieces of information that were requested prior to the beginning of each trial. A 3 (groups) x 6 (trials) mixed model ANOVA was computed using the number of primary variables requested as the dependent measure. This analysis revealed a main effect for group, $F\ (2, 41) = 7.15$, $p < .01$, $MSe = 32.61$, and a significant group by trials interaction, $F\ (10, 205) = 9.14$, $p < .01$, $MSe = 4.15$. The main effect of trials failed to obtain significance, $F\ (5, 205) = 1.76$, *ns*. As can be seen in Figure 9a, trainees requested many more primary variables than experts and novices to solve the first problem. By the third trial, however, trainees and experts requested a similar number of primary variables. Although the expert and novice groups requested the same number of information cards on the first trial,

FIGURE 9

Three dependent measures that indicate the number of pieces of information subjects used to solve the problems plotted as a function of trials. Panel (a) shows the mean number of primary variables selected, (b) shows the mean number of secondary variables requested, and (c) shows the mean number of unique variables activated. N = Novices, T = Trainees, E = Experts.

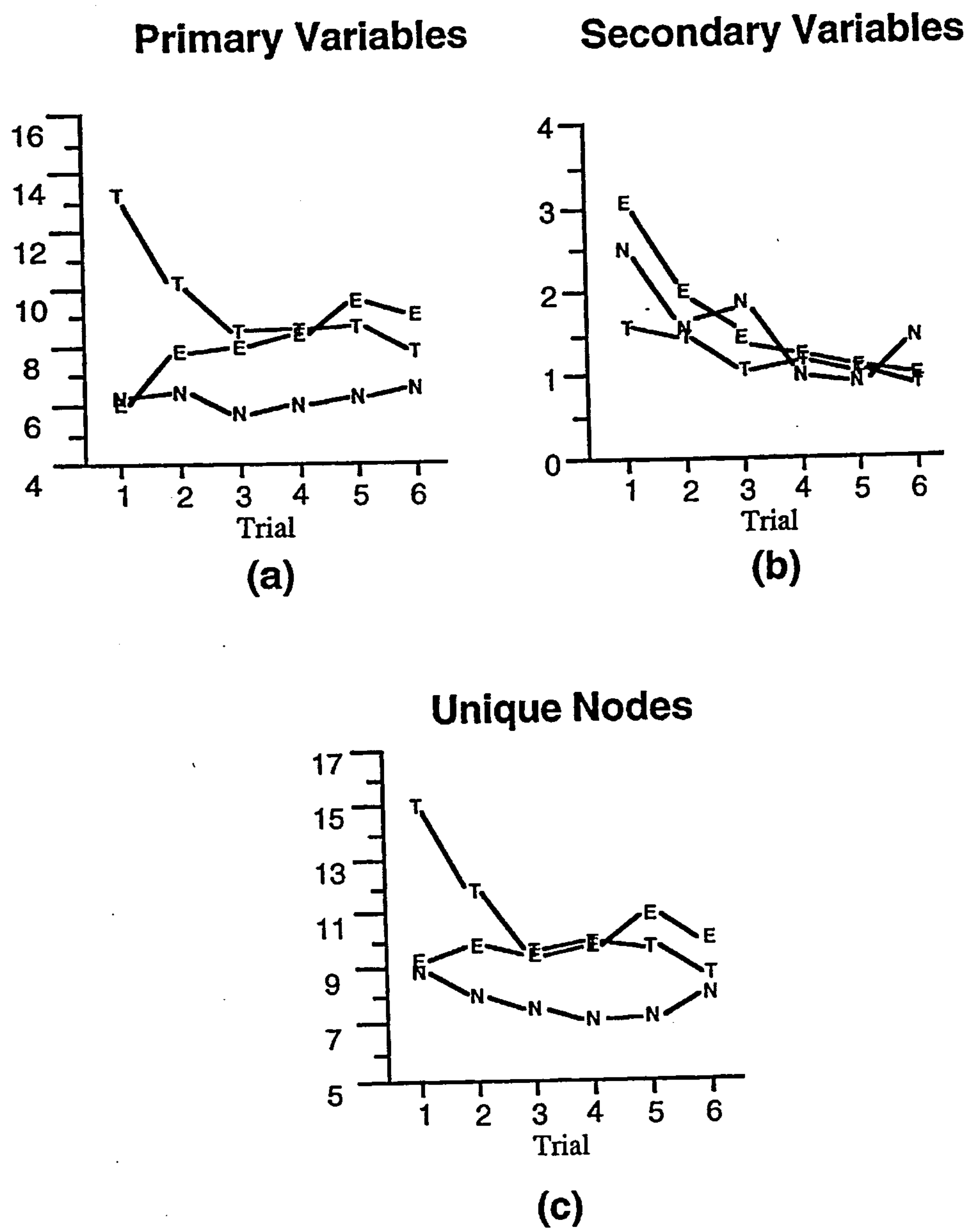

experts increased the amount of information they requested across trials, whereas the novices' requests remained stable and trainees' requests dropped.

Although experts and novices requested fewer primary variables than trainees before beginning the first problem, Figure 9b shows that experts and novices compensated by requesting additional information (secondary variables) on the first trial while solving the problem. However, by the fourth trial members of all three groups were requesting only one secondary variable per problem, on average. A 3 (groups) x 6 (trials) ANOVA revealed a significant main effect for trials on the number of secondary variables selected, $F\ (5, 205) = 3.67, p < .01, MSe = 3.09$, with the trend showing a decreasing number of variables selected from trials one through six. The main effect of group and the group by trials interaction was not significant, $F\ (2, 41) < 1$, *ns*., and $F\ (10, 205) < 1$, *ns*., respectively. It is interesting to note that the substantial drop across trials in the number of primary variables requested by trainees was not offset by an increase in the number of secondary variables selected.

While no subject failed to use a secondary variable they asked for, many subjects ignored one or more of the primary variables they had requested. Figure 9c shows the final tally of variables subjects actually used to solve each problem—in other words, the number of unique variables they considered in their effort to reach a solution. This measure ignores the primary variables that were requested but not used and the repetitive use of variables already considered (i.e., recursions). The pattern of results for unique variables is similar to the pattern for primary variables. On the first problem trainees considered about 15 unique pieces of information whereas experts and novices considered nine. By the fifth problem, however, experts considered more unique variables than trainees and novices, a reversal resulting from a gradual increase across trials in the number of unique variables considered by experts and a substantial drop for the trainees. A 3 (groups) x 6 (trials) ANOVA was computed based on the number of unique variables considered. This analysis indicated a significant main effect for groups, $F\ (2, 41) = 3.98, p < .05, MSe = 51.34$, and trials, $F\ (5, 205) = 5.69, p < .01, MSe = 5.32$, as well as a significant group by trials interaction, $F\ (10, 205) = 5.34, p < .01, MSe = 5.32$. This two-way interaction is due to the sharp reduction in the number of unique variables viewed over trials by trainees, combined with a gradual increase over trials in the number of unique variables considered by experts.

Panel "A" in Figure 10 shows the mean number of recursions subjects made over the six trials. A recursion is counted each time a subject reconsidered a previously viewed information card. A 3 (groups) x 6 (trials) ANOVA was computed using the number of recursions as the dependent measure. This analysis revealed a significant main effect of trials, $F\ (5, 205) = 7.98, p < .01, MSe = 2.38$, due to a general decrease in the number of recursions made by all groups across the six problems. The main effect of groups failed to obtain significance, $F\ (2, 41) < 1$, *ns*. However, a significant two-way interaction was identified, $F\ (10, 205) = 2.53, p < .01, MSe = 2.38$, in which novices averaged nearly four recursions on the first problem, as compared to about two and one for experts and trainees, respectively. By the second problem, novices had dramatically reduced their recursions to the level of the trainees, as had experts by the third problem. One of the more striking features of this figure is the consistently low level of recursions made by trainees across the six problems.

FIGURE 10

Mean scores for three markers indicative of processing efficiency: (a) the number of recursions, (b) the total number of steps it took to generate a solution, and (c) the amount of time it took to reach a solution. N = Novices, T = Trainees, E = Experts.

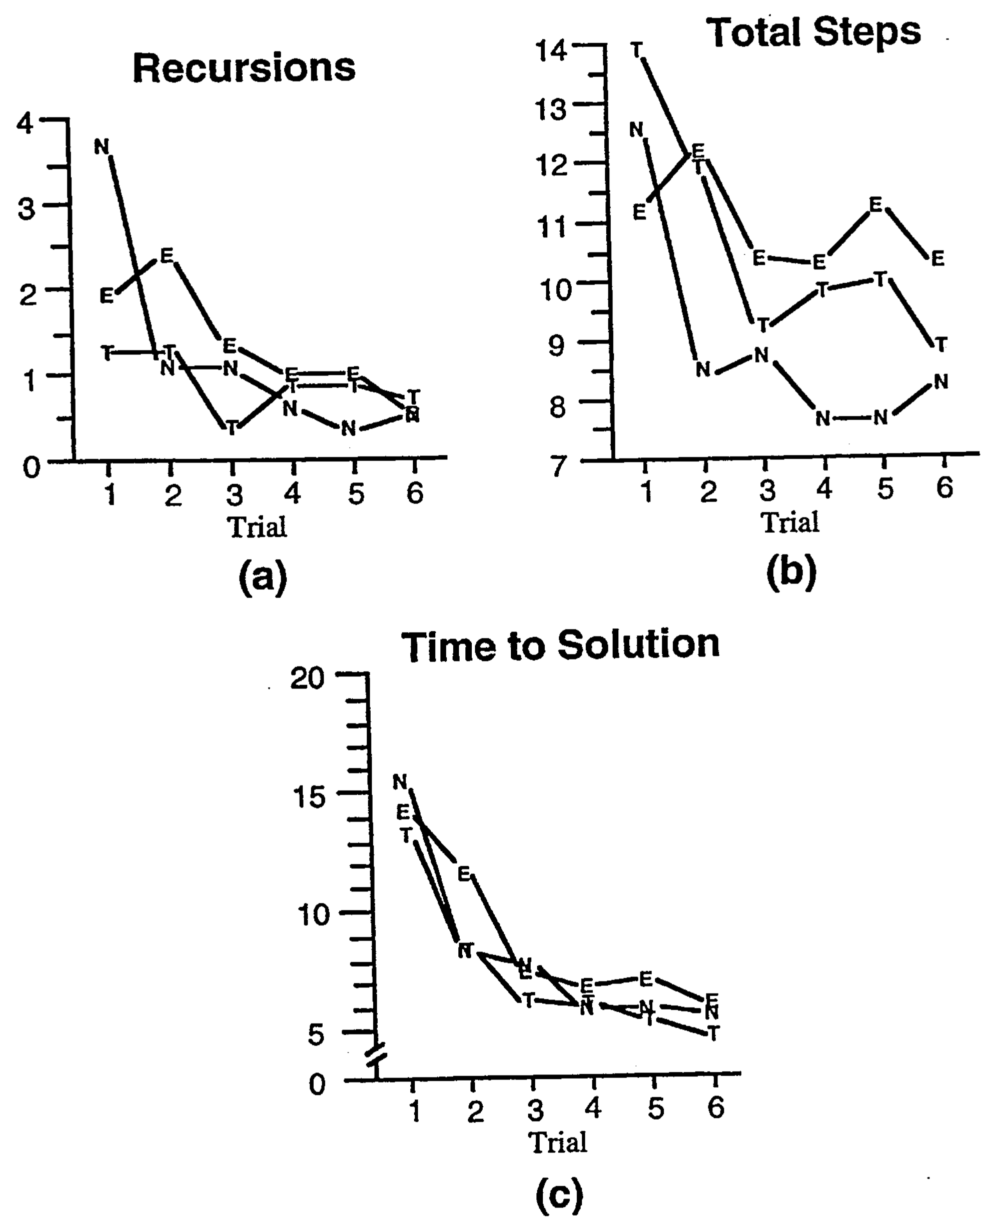

The total steps variable represented in Figure 10b combines subjects' use of unique variables and recursions to provide a composite picture of their problem solving efficiency. The clear picture shown by this variable is the stability of the experts performance across trials. They used about eleven total steps to solve each of the six problems. In contrast, trainees and novices used about five fewer steps to solve the sixth problem than they used to solve the first. This was largely due to trainees' reduction over trials in the consideration of lower level information, and novices' reduction over trials in the number of recursions generated. A 3 (groups) x 6 (trials) ANOVA using total steps as the dependent measure indicated a significant main effect of trials, F (5, 205) = 7.48, $p < .01$, $MSe = 10.25$. This effect was based on a decrease in the number of total steps involved in solving the later problems. The main effect of groups was not significant, F (2, 41) = 1.64, *ns*., however the group by trials interaction effect indicated a noteworthy trend, F (10, 205) = 1.78, $p = .06$, $MSe = 10.25$.

One process marker not derived from the PSPMs, the amount of time it took participants to generate solutions to the six problems, is depicted in Figure 10c. A 3 (groups) x 6 (trials) ANOVA indicated that the three groups did not differ in the amount of time it took to reach a solution, F (2, 41) < 1, *ns*., and each group showed significant decreases in time on task over trials, F (5, 205) = 23.15, $p < .01$, MSe = 21.16. This analysis failed to reveal a group by trials interaction, F (10, 205) < 1, *ns*.

When interpreting the earlier pattern of group performance differences we considered whether differences in motivation and/or effort could have influenced subjects' problem solving behaviors. Toward this end, we compared mean scores across groups for the following self-report questions from the post-experimental questionnaire: (a) How interesting was it for you to work on the problems? (1 = not interesting; 7 = very interesting); and (b) To what extent do you believe you carried out a thorough consideration of the problems, taking into account most of the details? (1 = poor consideration; 7 = thorough consideration). Separate univariate ANOVAs failed to reveal group differences for either of these questions, F (2, 41) < 1, both tests. Failure to reject the null hypothesis for the interest and thoroughness questions suggests that the observed information processing differences were not due to differences in motivation or effort. However, the three groups were found to differ in terms of their answer to a question regarding how challenging they found the set of problems (1 = not challenging; 7 = very challenging), F (2, 41) = 6.97, $p < .01$, $MSe = 1.11$. Novices reported the problems to be the most challenging ($M = 5.64$, $SD = .93$), trainees found the problems to be intermediately challenging ($M = 4.57$, $SD = 1.22$), and experts found the task the least challenging ($M = 4.25$, $SD = 1.00$). Thus, the perceived cognitive complexity of the task covaried with one's experience with the problem. Those with the least problem-specific knowledge and experience found the task to be most challenging, whereas those with the most knowledge and experience found the task to be least challenging.

DISCUSSION

A consistent but not surprising finding in the literature on complex problem solving is that experts generate better solutions to problems than novices. In fact, Ericsson &

TABLE 1
***F* Values and *p* Levels from Mixed Model ANOVAS for Information Processing Measures**

Variable	F value	p level
Primary Variables Requested		
Groups	7.15	< .01
Trials	1.76	ns
Groups x trials	9.14	< .01
Secondary Variables Requested		
Groups	< 1	ns
Trials	3.67	< .01
Groups x trials	< 1	ns
Unique Variables Used		
Groups	3.98	< .05
Trials	5.69	< .01
Groups x trials	5.34	< .01
Number of Recursions		
Groups	< 1	ns
Trials	7.98	< .01
Groups x trials	2.53	< .01
Total Steps to Solution		
Groups	1.64	ns
Trials	7.48	< .01
Groups x trials	1.78	.06
Time on Task		
Groups	< 1	ns
Trials	23.15	< .01
Groups x trials	< 1	ns

Smith (1991) suggest that outstanding problem solving performance is a requisite skill for one to be considered an expert. However, from an information processing perspective, it is interesting to ask what it is that enables experts to solve problems quickly, efficiently, and with a high degree of accuracy. In a historical overview of research on complex problem solving, Figenbaum echoed the words of Francis Bacon (1597) by proclaiming that "Knowledge is Power" (1989, p. 173). Figenbaum further suggests that we can look to the *Knowledge Principle* to better understand what differentiates levels of competence in complex problem solving situations. In short, the knowledge principle suggests that an individual "exhibits intelligent understanding and action at a high level of competence primarily because of the specific knowledge that...[he or she] can bring to bear [on the problem]: the concepts, representations, facts, heuristics, models, and methods of its domain of endeavor" (1989, p. 179). The multidimensional character of the knowledge principle suggests that we must look beyond simple domain-specific factual knowledge as the sole determinant of the quality of the subjects' performance. We must also closely examine both the characteristics of the individuals' information search activities and the quality of the information they consider in order to paint a complete profile of problem solving competence.

The finding that experts' investment decisions were nearly five times better than those produced by novices was not an unexpected outcome. What was an unexpected finding, however, was that trainees' solutions were, on average, twice as good as those produced by experts. In order to understand how this result could have come about, we first look to differences in the specific types of information considered by the three groups of problem solvers.

Trainees produced the most accurate solutions to the six problems by considering high level information in a very efficient and goal directed fashion. Their mental model of the problem mirrored the conceptual model of the 401k task, which enabled them to know which specific pieces of information should be considered in order to determine the unmet retirement need, the affordability of an investment, and the adequacy of the 401k plan. The superior quality of the information trainees selected, as compared to experts and novices, can be seen in their mean hierarchy scores for the first trial. Trainees selected more information, on average, then members of the other two groups, but more importantly, the information they considered was higher-level information drawn from near the top of the task hierarchies (see Figure 5). Furthermore, over the course of the remaining five trials the trainees actively pruned less relevant variables from their search routines, focusing on only the most informative (high-level) pieces of information. Strong confirmation for this interpretation can be found by comparing the pattern of shading shown in the IUDPs for the three groups (Figures 6, 7, and 8). An inspection of these figures indicates that trainees considered the most informative (high-level) variables when generating a solution for the final problem, and at the same time, they covered a great deal of the landscape contained in the Need, Account, and Affordability hierarchies. In contrast, novices and experts tended to ignore information contained in the Need hierarchy, instead focusing on the mid—to lower level variables contained in the Account and Affordability hierarchies.

In discussing the relationship between information quality and decision quality it is also interesting to note that on average, experts' solutions were many times better than those of novices, however both groups used roughly the same quality of information. (That is, the mean hierarchy scores for these two groups did not differ appreciably.) There was also converging evidence from a think-aloud protocol study of members of these same two groups that suggests they used the task information in different ways (Hershey et al., 1998). Relative to experts, novices appeared to be more *data driven* in their use of the informational cues, using parameters in computations in order to arrive at "exact" solutions to components of the problem. However, given that they often ignored important pieces of information (or even whole branches) within an information hierarchy, novices' "exact solutions" often contained appreciable error. Experts, in contrast, appeared to use task parameters in a more *conceptually driven* fashion. They would often view a piece of information and make an intuitive judgment regarding the adequacy of the parameter based on their prior problem solving experiences (e.g., "After seeing what she earns, I can't imagine that she'll have enough money to invest in a retirement plan"). In part, it was this tendency by experts to make intuitively-based evaluations of individual task parameters that allowed them to process the information quickly and efficiently. More importantly, however, it appears that it was their prior knowledge of the range of acceptable parameters for a variety of variables that allowed them to form a general impression of whether an investment was indicated or not, and on that basis they were able to specify a reasonably accurate, intuitively-based investment amount. There is converging evidence from another study currently in progress in our laboratory that suggests that experts process task information in a more conceptually oriented fashion, and novices employ a more data driven processing strategy (Hershey et al., 1998).

There are many similarities and a few differences between the performance of experts and novices on the first trial of this study and the performance of experts and novices in the Hershey et al. (1990) study. Across both studies experts used fewer total steps and made fewer recursions to solve the first problem. However, the present findings are different in that this sample of experts requested more secondary variables than novices and used an equal number of unique variables to solve the first problem. These differences may be the result of differences in the problems used, or differences in the participants who formed the expert sample.

Our earlier work used an Individual Retirement Account investment decision rather then the 401k decision. The conceptual differences between the two problems are minor, although our task analysis of the 401k problem yielded about 20 percent more variables than for the IRA problem. Perhaps the larger problem space is responsible for the observed differences in the selection of secondary and unique variables. A second possibility is that the experts in the two studies had different amounts of work experience, and therefore, different mental models of the financial problems used in those studies.[4]

Differences between experts' and novices' information use on problems two through six form a different pattern than that seen on the first problem. By the second trial novices, not experts, were making fewer recursions, using fewer total steps, and taking less time to solve the problems. It was also experts, not novices, that requested more primary variables and used more unique variables. We think that these changes are explained by novices simplifying their solutions as they gain some familiarity with the problem domain, whereas experts continued to solve the problems at a consistent level of complexity across trials. These trends are consistent with the idea that experts had clearer mental models of the 401k problem than novices at the outset of the session, and novices developed rudimentary processing strategies as they solved the problems.

Following the educational intervention, trainees, like experts, possessed a clear mental model of the task, but, like novices, they lacked production experience at solving retirement investment decisions. This unique combination of task knowledge without the benefit of first-hand problem solving experience led to some interesting patterns of information use for the first problem, and intriguing patterns of strategy change as they solved subsequent problems. The small number of secondary variables trainees selected on the first trial indicates that their mental model of the problem was sufficiently detailed to specify the key variables to be considered. Moreover, like experts they were found to process task information in a non-recursive and goal-directed fashion, presumably because they had a clear set of objectives to attend to when solving the problem (i.e., assess need, assess affordability, assess account characteristics). Although the trainees considered many more pieces of information than the other two groups on the first trial, they adaptively pruned their information search subroutine to a manageable subset of the most critical pieces of information by the third problem. The substantial drop in hierarchy scores (which indicated use of higher level, aggregate data) over trials (described earlier and seen in Figure 5) indicates that they eliminated less critical variables from consideration. In many respects, by the sixth trial the trainees appeared to be model problem solvers. They were generating high quality solutions quickly and in a non-recursive fashion, using only key pieces of information that they were able to specify at the outset of the trial.

How do we interpret results in which college sophomores with six hours of financial training outperform degreed accountants with four years of work experience? While we recognize the need to be cautious, we do not want to ignore a potentially important outcome. We think one important part of the answer lies in the power of the conceptual task analysis that served as the foundation of the financial training. It would be unreasonable, by any stretch of the imagination, to argue that the trainees in the present study were financial planning experts. Although no tests of generalization to other financial problems were conducted, we doubt that trainees would have shown much improved aptitude for other types of financial planning problems.

However, there are a number of complex life planning tasks—such as retirement planning, saving for a child's education, or making the decision to purchase or rent housing—where it would be preferable for the individual to be able to conduct analyses on his or her own behalf. This is particularly true given the fact that many (if not most) individuals in contemporary society possess insufficient resources to enlist the aid of an expert financial planner. Unfortunately, the same individuals who are least likely to have a coherent working knowledge of the dynamic complexities of long-term financial investing are at the greatest risk of making poor financial planning decisions. It is precisely these individuals who could benefit the most from brief educational interventions—such as the one used in the present study—in order to provide them with rudimentary task-specific mental models of personal finance and investing.

A second part of the answer to the question of how trainees outperformed experts lies, we believe, in trainees heavy use of high level, aggregated variables from all three conceptual dimensions contained in the problem space (retirement need, investment affordability, and account characteristics) whereas the experts used less high level information and mostly ignored the retirement need issue (see Figures 7 and 8). The experts inattention to whether or not the person described in a scenario already had sufficient retirement resources would have been a source of error in two of the six scenarios. Also, far more experts than trainees carried out their own computations and aggregations of lower level data to arrive at higher level variables to solve the problem—another source of potential error since neither sophisticated financial calculators nor computer support was available. In contrast, the trainees requested and used most of the highest level information available to solve the problems, performing relatively few computations and aggregations on their own. Perhaps these observed computational differences can be attributed to the clear conceptual structure of the task trainees developed in the pre-test session, experts' rigidity in calculating values of their own based on years of number crunching experiences, or a combination of the two. Regardless of the reason, however, it is clear that the accuracy of trainees' solutions benefited greatly from the high-level variables they used to solve the problems.

Perhaps the most significant contribution of the present study to the literature on problem solving is the identification of a training technique that can be used to enhance cognitive performance in complex domains. Unfortunately, the field of cognitive psychology has been relatively silent regarding novel approaches to training in real-world, ill-structured problem domains. The majority of training techniques have focused on teaching individuals to employ "weak" or "domain-general" problem solving strategies (Anderson, 1987) such as means-ends analysis, working backward, and analogical reasoning. Whereas such strategies certainly have their areas of application, it would be unlikely that one could apply them (either individually or collectively) to derive an optimal solution when faced with an information-rich, cognitively complex task. For tasks such as these, we would propose that there is no substitute for a coherent and well-integrated mental representation of the relationships between elements in the problem space. That being the case, we encourage future studies to examine the practical limits of conceptual model training in complex domains. For example, work from our laboratory has revealed that both young and old subjects acquire a great deal of information by attending a conceptual model training. However, in terms of the quality of post-training problem solving performance, it appears that older subjects may not benefit as much as younger individuals (Hershey et al., 1998; Walsh & Hershey, 1999). We believe that understanding the specific conditions under which an educational intervention will (or will not) be effective is a timely and important applied goal.

NOTES

Accepted for publication: 24 November 1999. This work was supported by a grant awarded to David A. Walsh from the National Institute on Aging (#RO1–AG08106). The authors wish to acknowledge the helpful comments of Ward Edwards, Steve Read, and John L. Horn on an earlier draft of this manuscript.
Address correspondence to: Douglas A. Hershey, Department of Psychology, Oklahoma State University, 215 N. Murray Hall, Stillwater, Oklahoma 74078; E-mail: hershey@okstate.edu.

1. Other papers that have resulted from this study include an analysis of participants' metacognitive performance (Hershey & Wilson, 1997), and an analysis of age-differences in information search (Walsh & Hershey, 1993; experiment 2).

2. A separate analysis was conducted to determine whether subjects' performance improved over trials. This analysis failed to indicate a practice effect—not an unexpected finding given that subjects did not receive feedback on the quality of their solutions.

3. IUDPs were developed for the sixth trial because this was the only trial (by design) in which all subjects encountered the same problem. The group information similarity analysis (below) is also computed using data based on the sixth trial.

4. Unfortunately we do not have detailed descriptions of the daily work activities of either group, nor self-reports of their familiarity with the problems, although we see the strong need to collect this information in the future.

REFERENCES

Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review, 94*, 192–210.

Bacon, F. (1597). *Meditationes sacræ*. De Hæresibus.

Chase, W. G. & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing* (pp. 215–281). New York: Academic Press.

Chi, M. T. H., Feltovich, P. J. & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121–152.

Chi, M. T. H., Glaser, R. & Farr, M. (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum.

Ericsson, K. A. & Crutcher, R. J. (1990). The nature of exceptional performance. In P. B. Baltes, D. L. Featherman & R. M. Lerner (Eds.), *Life-span development and behavior* (Vol. 10, pp. 187–217). Hillsdale, NJ: Erlbaum.

Ericsson, K. A. & Smith, J. (1991). Prospects and limits of the empirical study of expertise: An introduction. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise: Prospects and limits* (pp. 1–38). New York: Cambridge University Press.

Figenbaum, E. A. (1989). What hath Simon wrought? In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 165–182). Hillsdale, NJ: Erlbaum.

Gentner, D. & Stevens, A. L. (1983). *Mental models*. Hillsdale, NJ: Erlbaum.

Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306–355.

Glaser, R. (1989). Expertise and learning: How do we think about instructional processes now that we have discovered knowledge structures? In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 269–282). Hillsdale, NJ: Erlbaum.

Glaser, R. & Chi, M. T. H. (1988). Overview. In M. T. H. Chi, R. Glaser & M. J. Farr (Eds.), *The nature of expertise* (pp. xv–xxviii). Hillsdale, NJ: Erlbaum Associates.

Gregg, D. W. (1992). Human wealth-span: The financial dimensions of successful aging. In N. E. Cutler, D. W. Gregg & M. P. Lawton (Eds.), *Aging, money, and life satisfaction* (pp. 169–182). New York: Springer.

Hegarty, M. (1991). Knowledge and process in mechanical problem solving. In R. J. Sternberg & P. A. Frensch (Eds.), *Complex problem solving: Principles and mechanisms* (pp. 253–285). Hillsdale, NJ: Erlbaum.

Hershey, D. A. (1990). The role of knowledge and experience in structuring problem solving performance (Doctoral dissertation, University of Southern California). *Dissertation Abstracts International, 51*, Sec. 7B.

Hershey, D. A. (1995). Influence of age and gender on estimates of long-term financial growth functions. *Aging and Cognition, 2*, 231–250.

Hershey, D. A. & Farrell, A. H. (1999). Age differences on a procedurally-oriented test of practical problem solving. *Journal of Adult Development, 6*, 87–104.

Hershey, D. A., Morath, R. & Walsh, D. A. (1991, November). The use of scripts to solve financial planning problems. In C. Berg (Chair), *Age differences in cognitive strategies*. Symposium conducted at the meeting of the Gerontological Society of America, San Francisco.

Hershey, D. A., Walsh, D. A., Brougham, R., Carter, S. & Farrell, A. (1998). Challenges of training pre-retirees to make sound financial planning decisions. *Educational Gerontology, 24*, 447–470.

Hershey, D. A., Walsh, D. A., Mauldin, L., Frostholm, L. & Stone, S. (1998). *A verbal protocol analysis of complex financial problem solving performance*. Manuscript in preparation.

Hershey, D. A., Walsh, D. A., Read, S. J. & Chulef, A. S. (1990). The effects of expertise on financial problem solving: Evidence for goal-directed, problem-solving scripts. *Organizational Behavior and Human Decision Processes, 46*, 77–101.

Hershey, D. A. & Wilson, J. A. (1997). Age differences in performance on a complex financial decision making task. *Experimental Aging Research, 23*, 257–273.

Hershey, D. A., Wilson, T. L. & Mitchell-Copeland, J. (1996). Conceptions of the psychological research process: Script variation as a function of training and experience. *Current Psychology, 14*, 293–312.

Holyoak, K. J. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise* (pp. 301–335). New York: Cambridge University Press.

Lesgold, A. & Lajoie, S. (1991). Complex problem solving in electronics. In R. J. Sternberg & P. A. Frensch (Eds.), *Complex problem solving: Principles and mechanisms* (pp. 287–316). Hillsdale, NJ: Erlbaum.

Norman, D. A. (1983). Some observations on mental models. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 15–33). Hillsdale, NJ: Erlbaum.

Sternberg, R. J. (1995). Expertise in complex problem solving: A comparison of alternative conceptions. In P. A. Frensch & J. Funke (Eds.), *Complex problem solving: The European perspective* (pp. 295–321). Hillsdale, NJ: Erlbaum.

Voss, J. F., Wolfe, C. R., Lawrence, J. A. & Engle, R. A. (1991). From representation to decision: An analysis of problem solving in international relations. In R. J. Sternberg & P. A. Frensch (Eds.), *Complex problem solving: Principles and mechanisms* (pp. 119–158). Hillsdale, NJ: Erlbaum.

Walsh, D. A. & Hershey, D. A. (1993). Mental models and the maintenance of complex problem-solving skills in old age. In J. Cerella, J. Rybash, W. Hoyer & M. L. Commons (Eds.), *Adult information processing: Limits on loss* (pp. 553–584). New York: Academic Press.

Walsh, D. A. & Hershey, D. A. (1999). *Age and training differences on a complex retirement planning task*. Manuscript in preparation.