

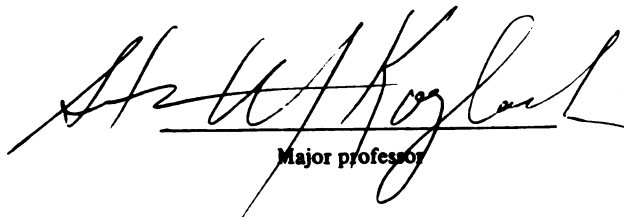


This is to certify that the
thesis entitled

Expertise and intra-task variation
in decision-making
presented by

Dennis John Devine

has been accepted towards fulfillment
of the requirements for
Master of Arts degree in Psychology



Major professor

Date June 23, 1993

**LIBRARY
Michigan State
University**

PLACE IN RETURN BOX to remove this checkout from your record.
TO AVOID FINES return on or before date due.

DATE DUE	DATE DUE	DATE DUE
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____

MSU Is An Affirmative Action/Equal Opportunity Institution

c:\circ\data\due pm3-p.1

ABSTRACT

EXPERTISE AND INTRA-TASK VARIATION IN DECISION-MAKING EXPERTISE AND INTRA-TASK VARIATION IN DECISION-MAKING

By

Dennis John Devine

Past research on the acquisition of expertise has taken the task for granted and ignored task characteristics that both promote and limit performance differences between experts and novices. In this study, three constructs were shown to have relevance to the information seeking aspect of decision-making: A THESIS

using the ACT* theoretical framework -- domain knowledge, alternative labels and task structure. A 2 (Domain Knowledge) x 2 (Alternative Labels) x 2 (Decision Strategy) factorial design was conducted in a lab setting using a 2 (Expertise) x 2 (Task Structure) x 2 (Alternative Labels) x 2 (Domain Knowledge) factorial design. The results showed that experts and novices made decisions about a game of basketball. Submitted to Michigan State University in partial fulfillment of the requirements for the degree of Master of Arts

methodology with experts and novices making decisions about a game of basketball. MASTER OF ARTS ed that, as predicted, experts were more sensitive to the stereotypicality of the task and the presence of alternative labels than were novices. The discussion focuses on the need to consider within-task factors in future research on expertise. Department of Psychology 1993

ABSTRACT

EXPERTISE AND INTRA-TASK VARIATION IN DECISION-MAKING

By

Dennis John Devine

Past research on the acquisition of expertise has taken the task for granted and ignored task characteristics that both promote and limit observed performance differences between experts and novices. In this study, three constructs were shown to have relevance to the information seeking aspect of decision-making using the ACT* theoretical framework -- domain knowledge, alternative labels and task structure. A 2 (Domain Knowledge) x 2 (Alternative Labels) x 2 (Decision Structure) experiment was then conducted in a lab setting using a computerized information-board methodology with experts and novices making decisions about a game of basketball. Results indicated that, as predicted, experts were more sensitive to the stereotypicality of the task and the presence of alternative labels than were novices. The discussion focuses on the need to consider within-task factors in future research on expertise.

ACKNOWLEDGMENTS

Viewing this finished copy as my awards ceremony, I would like to thank all of those "unnamed" individuals who provided material or moral support during the long and tedious process of completing this work. However, special thanks are offered to Steve Kozlowski for overseeing and advising this project, and for imparting so much about the research "process" along the way. I would also like to thank Neal Schmitt and Kevin Ford for their incisive questions and thoughtful feedback while serving on my committee. This paper was improved in so many ways from their comments that it staggers me to think of how far it has come (and how poor it was when it started out). Thanks also to Stephen Gilliland for his help in creating the software used in this study and answering my numerous pesky methodological questions. Finally, I would like to thank my wife, Julie Devine, for her help in proofreading this manuscript and for celebrating its step-by-step completion with me along the way.

Independent Variables

Stimulus Materials

Dependent Variables

Procedure

RESULTS	70
Overview	70
Assessment of Expertise	75
Manipulation of	77
Method Analyses	80
LIST OF TABLES	vi
LIST OF FIGURES	vii
INTRODUCTION	1
Overview of production systems	5
Theoretical Framework	7
Acquiring expertise in ACT*	11
Expert-Novice Literature Review	15
Knowledge Organization	16
Domains with hierachial goal structures	20
Process-tracing	24
Summary of research on expertise	29
ACT* and the Task Environment	30
Summary -- Category Labels	36
Task Structure	37
Labels, Task Structure and Expertise	45
Model and Hypotheses	49
METHOD	58
Participants	58
Independent Variables	59
Stimulus Materials	62
Dependent Variables	64
Procedure	66

RESULTS	70
Overview	70
Assessment of Expertise	75
Manipulation Checks	77
Method Analyses	80
Main Analyses	85
Hypothesis 1: Decision Accuracy	86
Hypothesis 2: Cue Latency	97
Hypothesis 3: Total Search Depth	100
Hypothesis 4: Contextual Search Depth	103
Hypothesis 5: Choice Search Depth	111
Hypothesis 6: Overall Search Variability	118
Hypothesis 7 and 8: Search Pattern	127
DISCUSSION	132
Summary	132
Expertise and Intra-Task Variation	135
Study Issues	139
Task Variable Interactions	144
Conclusions	145
APPENDICES	
APPENDIX A Basketball Knowledge questionnaire . . .	147
APPENDIX B Cue values for search matrices	153
APPENDIX C Rationale for correct alternative choice	155
APPENDIX D Importance Ratings	157
APPENDIX E Post-experimental questionnaire	158
APPENDIX F Basketball Experience questionnaire . .	160
LIST OF REFERENCES	162

LIST OF TABLES

Table 1	Study Hypotheses	56
Table 2	Means and standard deviations of study variables	72
Table 3	Variable intercorrelations	73
Table 4	T-tests, manipulation check for Alternative Labels	78
Table 5	T-tests, manipulation check for Decision Structure	81
Table 6	Results of Method-factor analyses	84
Table 7	Results, Overall CATMOD analysis	87
Table 8	Results, Overall MANOVA	88
Table 9	Decision choices across all conditions	94
Table 10	Univariate ANOVA, Cue Latency	98
Table 11	Univariate ANOVA, Total Search Depth	102
Table 12	Univariate ANOVA, Contextual Search Depth	104
Table 13	Univariate ANOVA, Choice Search Depth	113
Table 14	Univariate ANOVA, Overall Search Variability	119
Table 15	Univariate ANOVA, Choice Matrix Search Variability	125
Table 16	Univariate ANOVA, Search Pattern index	129
Table 17	Hypotheses-Study Result linkages	133
Table B-1	Cue values for Search Matrix A	153
Table B-2	Cue values for Search Matrix B	154

LIST OF FIGURES

Figure 1	An ACT*-based model of Information Acquisition	50
Figure 2	Domain Knowledge x Decision Structure interaction for Decision Accuracy	90
Figure 3	Order x Decision Structure interaction for Decision Accuracy	92
Figure 4	Order x Decision Structure interaction for Cue Latency	99
Figure 5	Domain Knowledge x Alternative Labels interaction for Contextual Search Depth	106
Figure 6	Domain Knowledge x Decision Structure interaction for Contextual Search Depth	107
Figure 7	Order x Decision Structure interaction for Contextual Search Depth	108
Figure 8	Domain Knowledge x Decision Structure x Order interaction for Choice Search Depth with "ill-structured task first" ordering	116
Figure 9	Domain Knowledge x Decision Structure x Order interaction for Choice Search Depth with "well-structured task first" ordering	117
Figure 10	Domain Knowledge x Decision Structure x Order interaction for Overall Search Variability with "well-structured task first" ordering	121
Figure 11	Domain Knowledge x Decision Structure x Order interaction for Overall Search Variability with "ill-structured task first" ordering	122
Figure 12	Domain Knowledge x Alternative Labels interaction for Choice matrix Search Variability	126
Figure 13	Order x Decision Structure interaction for Search Pattern index	130

solving a certain (relatively narrow) class of problems with the aid of education and repeated experience.

The notion of expertise is also of great interest to organisations. One of the main goals of training is to

INTRODUCTION

take novice individuals in an organisation and turn them into experts. Discussing the concept of "expertise" is a bit like trying to describe yourself to a good friend or your spouse -- what can you say that everyday experience hasn't said already? Most of us are familiar with the concept of the "expert" and probably feel like we have a pretty good intuitive understanding of what an expert is. We hear the term used every day, we bump into "experts" everywhere we turn and we even seek them out from time to time.

When it comes right down to it, however, "experts" play an important part in our lives on a day-to-day basis.

Government policies are set and actions are determined by a myriad of "expert" consultants and advisors. We spend over 300 billion dollars a year as a nation on health care, and much of this amount goes toward paying for the services of medical "experts." When something goes wrong with our cars, we seek out the expertise of an auto mechanic who can fix it and prevent our busy lives from grinding to a halt. When we get embroiled in the legal system, we get a lawyer. When we want to hide as much money as possible from the government, we turn to a certified public accountant. In essence, most "professions" in our society exist so that individuals can make a living by becoming proficient at solving a certain

solving a certain (relatively narrow) class of problems with the aid of education and repeated experience.

The notion of expertise is also of great interest to organizations. Implicitly, the function of training is to take novice individuals in an organization and turn them into experts at their jobs. However, experts also play a large role in the indirect functioning of organizations, such as when "subject-matter experts" are called upon to describe their positions during job analysis or when human resource experts generate/identify the criteria by which people will be selected into the organization and against which their performance will be appraised.

Ideally, the potential exists for every position in an organization to be held by an "expert" -- someone who performs the job quickly and accurately, can make intelligent, reasoned decisions and can solve problems in a unique and creative fashion. Most of the important decisions/actions within the context of an organization are determined by someone who, implicitly or explicitly, is expected to be an expert at what she/he does. Given the one thing that everyone seems to agree with about experts -- that they perform better than novices or "average" performers -- an understanding of the process of becoming an expert seems very relevant to improving the functioning of any organization.

A greater understanding of how experts are created will not necessarily result from more research with

"expertise" in the title. A great deal of research has already examined the decisions of experts and much is already known (Chi et al., 1988; Ericsson & Smith, 1992). For instance, it has been fairly well-established that, in most performance domains, "experts" can be identified and the **outcomes** of their decisions/problem-solving efforts can be shown to differ from those of the "novice." In addition, these studies have consistently shown that, in most domains, experts make better decisions and/or judgments than novices (Larkin, McDermott, & Simon, 1980; Chi et al., 1988; Ericsson & Smith, 1992). However, research has largely failed to address how expert and novice decision **processes** are different.

This study will examine how information is acquired by both experts and novices in the process of making decisions. There is already an extensive literature base documenting the manner in which individuals gather and use information when making decisions already exist (e.g., Svenson, 1979; Payne, 1982; Ford et al., 1989). Also, many studies have focused on the differences in how domain experts and domain novices solve problems and make decisions (e.g., Chase & Simon, 1973; Chi et al., 1981; Hershey et al., 1990).

However, few studies on expert-novice differences have examined these differences in conjunction with variation **within** a task (Ford et al., 1989). For the most part, the "task" has been taken for granted in past research and little consideration seems to have been given to identifying

the contextual variables (such as the labelling of alternatives or the stereotypicality of the problem situation) that may play a large role in whether or not the expert performs like an expert. It seems fairly clear that experts will perform better than novices in some situations and perhaps even most. However, little research has addressed the task-based influences that allow, further and ultimately limit this distinction in performance.

In addition, past research on expertise and expert-novice differences has not been based on a theoretical framework of human cognition and performance. Due to the difficulty of obtaining experts in most domains and the arduous nature of process-tracing data collection methods, most studies have not attempted to make predictions about expected differences, only document them. Even then, most studies have used small sample sizes. No study in the literature has utilized the advantages of an information board methodology (i.e., standardization, precise and explicit measures of search variables) in conjunction with a large enough sample size to test the hypotheses generated by this exploratory research.

A variety of common findings have emerged in this literature, but there is a pronounced need to begin integrating findings, deriving predictions from a conceptual framework and testing findings with adequate statistical power. The purpose of this study, then, is to do what the previous discussion has identified as lacking in the

literature: identify a theoretical framework to integrate past findings and generate predictions for the current study and then, using that identified theoretical framework, examine potential task characteristics that may affect the performance of experts and novices. In the end, by identifying aspects of the task which experts are sensitive to but novices are not, we can arrive at a greater understanding of the basis of expert performance.

The remainder of the introduction is organized in the following fashion. First, a description is offered of a model of human cognition, ACT* (Adaptive Control of Thought), developed by John Anderson. This description will review the basic tenets of ACT* theory and will be used to identify three constructs relevant to the study of human cognition and performance, domain knowledge (expertise), alternative labelling and task structure. Following this, each of the corresponding literatures will be reviewed. The last section of this introduction will present a model that attempts to integrate ACT* theory and the findings in the traditional expert-novice literature for the purpose of deriving hypotheses for the current study.

Overview of production systems

Any theory that attempts to explain the wide variety of expert-novice differences present in the literature must be broad enough to account for a number of phenomena across many research domains yet be specific enough to suggest hypotheses for future study. Production-systems models of

human cognition appear to be able to do this. Such models have a number of strengths. For example, production systems, while they look quite complex and horrendous on paper, are homogeneous in format, simple in structure and utilize independent productions that can be combined in a flexible and efficient fashion. Production-systems models retain the stimulus-response flavor of behaviorism but provide for goal-driven, "top-down" effects as well, as goals are hierarchically ordered, satisfied and changed. This adaptive characteristic makes production system models the "principle theoretical medium in which to cast complex theories of human intelligence" (Klahr, Langley, & Neches, 1987, p. ix).

All production systems consist of certain generic components: a short-term working memory, a long-term production memory and a "recognize-act" cycle that allows the two to interact. These elements are basic to any production system model of human cognition.

Working memory (similar to STM, short-term memory) is a limited-capacity processing site where data ("elements") are stored in symbolic code and where operations can be performed to alter old information and create new information. The **production memory** (also known as LTM, long-term memory) is an unlimited-capacity storage site for condition-action rules known as **productions**. Without these rules, the system could not function. Productions are propositions composed of conditions, which describe

configurations of data elements that might appear in working memory, and actions, which specify alterations to the elements in working memory or some form of overt behavior. The **recognize-act cycle** has three stages: matching, conflict resolution, and action. In the matching stage, the contents of working memory are compared to the conditions of stored productions in production memory. In the action stage, if a pattern of elements in working memory satisfies all of the conditions for a given production, then some operation specified in the action portion of the production is performed on the contents of working memory. If the conditions for more than one production are matched with the contents of working memory, some selection rule is used to choose the production which will occur in the conflict resolution stage.

The next section of the paper will address a specific production systems theory, ACT*, to be used as a conceptual framework for this study.

Theoretical Framework

ACT is a production system framework of human memory and cognition. After laying out the fundamentals of human cognition over 15 years ago with ACTE theory (Anderson, 1976), the ACT theoretical framework has been used as the basis for generating several other ACT theories which have revised and expanded the original ACTE theory. ACT is centered around the notion that higher-level cognition in humans is a unitary system (Anderson, 1983).

At the heart of the ACT theoretical framework is the assumption of two qualitatively different types of knowledge: declarative knowledge and procedural knowledge. Declarative knowledge consists of factual information organized into an associative semantic network of nodes and links in long-term memory, while procedural knowledge consists of system-based "rules" for information processing in the form of "If-Then" **productions**. ACT* assumes three memories where these two types of knowledge are stored: declarative memory, working memory and production memory. Declarative memory is a long-term, unlimited capacity storage entity for factual information organized into a propositional network of nodes and links. Each node in declarative memory corresponds to a **cognitive unit** (or "chunk") which can have up to five elements. Information at a given node can be represented in one of three ways, each way preserving a certain type of relation among its elements. **Temporal strings** (or lists) preserve ordinal relations among elements, **spatial images** preserve configurational information and **abstract propositions** preserve semantic relations. Cognitive units, in turn, can be elements of other cognitive units as well, allowing for, in Anderson's terms, a "tangled hierarchy" of knowledge representation that remains flexible.

Working memory is a special, limited-capacity declarative repository where information can be altered and created. Working memory is more of a function than a

particular place. It consists of knowledge structures in declarative memory which are active, as well as knowledge structures which are created through the application of productions. The production system can only modify and create data in its small working memory. The other two memories are just for storage.

Information in WM is matched against the conditional clauses of existing productions to see if any apply. If a production applies, it specifies an action or actions to be performed on WM data or on the environment. Knowledge structures can become activated in WM through the sensory encoding of environmental stimuli, the activation of information in declarative memory or through construction "on-the-spot" by productions in WM itself. Information stays in WM only so long as it maintains a certain level of activation which decays over time if not replenished from source nodes. Sources of activation include objects of perception in the environment and knowledge structures created in WM. Activation levels are assumed to be a continuous property of each node in declarative memory and activation is assumed to spread over links in the declarative network to associated concepts. Newly created knowledge structures in WM have a given probability of being stored in long-term declarative memory.

Production memory is a long-term storage site for the rules (productions) used by the cognitive system that enable initially selected for testing on the basis of a preliminary

it to function. Except for base-level perceptual coding by the senses, nothing gets done without a production. ACT* assumes that all knowledge starts out as declarative knowledge and is only later translated into procedural knowledge that allows it to be "used." According to ACT, novices solve problems through the use of interpretative procedures that do not require the use of domain-specific knowledge (i.e., existing productions). These interpretative procedures fashion productions out of declarative knowledge contained in the instructions and the problem statement through the use of domain-independent productions organized into strategies such as analogy, hill-climbing, means-ends analysis, working backwards, etc. Declarative knowledge is used interpretatively in that it serves as data for these domain-independent procedures (Anderson, 1982; Anderson, 1987). Anderson (1982) provides an example of how geometry problems can be solved with a set of 21 productions incorporating no procedural knowledge of geometry and using only declarative information given by the problem. For a more elaborate discussion of the process by which declarative knowledge is transformed into procedural knowledge, see Anderson (1982) or Anderson (1987).

Productions in ACT* theory utilize both constants (e.g., house, car, Dad, Mary) and local variables (LVstring, LVnumber, LVcolumn, LVrow, etc.) and are indexed for testing on the basis of these two components. Productions are initially selected for testing on the basis of a preliminary

match between their constants and the contents of working memory. Productions that pass this first stage are then tested further to see what extent their antecedent elements have been satisfied. When the conditional portions of more than one production are satisfied by the contents of working memory, the principles of refractoriness, strength and specificity are used to determine which production will be chosen to occur. In effect, the system selects stronger productions before weaker productions, specific productions before more general productions, and avoids having the same production selected endlessly ("looping").

The control of information processing and subsequent behavior is built into the production system itself through the use of hierarchial goals in the conditions of the productions. The production system may only have one goal active at a time, but goals are "stacked" so that when one is accomplished, activation is "popped" to the next higher goal. Productions are organized into sub-routines that have the same sub-goal in their conditional clauses. Thus, according to Anderson (1982), "the hierarchial control of behavior derives from the structure of problem-solving" (p. 372). Tuning occurs after the creation of domain-specific

Acquiring expertise in ACT*

There are two general processes in ACT* that occur as individuals acquire more decision-making/problem-solving experience in a given domain: knowledge compilation and tuning. **Knowledge compilation** occurs through the mechanisms

of composition and proceduralization. **Tuning** takes place after knowledge compilation through the mechanisms of generalization, discrimination and strengthening. When problems are solved through the use of general interpretative procedures, productions are created from a trace left by the problem-solving process. These newly-created productions are stored in a production memory separate from declarative memory. As productions are activated and used, processes act on these productions to make them more efficient and task-specific. **Composition** refers to the process by which adjacent productions in a problem-solving sequence are collapsed into a single production that has the effect of the multiple-production sequence. **Proceduralization** is a complementary process that builds new versions of old productions but with fewer conditions, allowing selection of the same production with less matching information in WM. Composition combines the conditional clauses of several productions into the conditional clauses of one production and then specifies multiple operations to be performed in the action clause. Proceduralization simply eliminates existing conditional clauses. Tuning occurs after the creation of domain-specific productions and alters them to make them more useful and efficient. Generalization procedures are used to make existing productions applicable in a wider variety of situations. Generalization occurs when productions containing local variables are created from specific

productions involving constants. Discrimination occurs when new productions are made by adding conditional clauses to existing productions so as to limit their applicability. Strengthening occurs as productions are successfully used. Each time a production is invoked and feedback indicates that it has been successful, its strength increases incrementally, making the production more likely to be invoked again in the future. Stronger productions are more likely to be selected in the conflict-resolution stage if multiple productions are applicable and their conditions are tested more quickly than the conditions of weaker productions.

Given that ACT presents an architecture for the acquisition of cognitive skill, this implies that the acquisition of cognitive skill amounts to an increase in one or both types of knowledge specified by ACT: declarative knowledge and procedural knowledge.

Therefore, when one is first introduced to a particular task (or domain), problems are approached in a domain-independent manner using declarative knowledge given in the problem statement and/or instructions in an interpretative fashion. In doing this, declarative knowledge is "plugged into" a set of general, interpretative productions which establish goals, identify methods, try these methods and evaluate outcomes in an algorithmic manner. Interpretative productions involve numerous local variables which must be categorized, voluminous conditional clauses which must be

matched to elements in WM and, in the end, choose actions which move the system forward in small, conservative steps. With each step, an increasingly heavy burden is placed on the system to remember what has been "discovered" (i.e., retain it in WM).

Therefore, at any given time within a particular domain, the efficiency of cognitive processes (e.g., problem-solving, decision-making) is related to the amount of domain-related declarative knowledge held by an individual as well as the number and sophistication of domain-related productions available to utilize this declarative knowledge. Thus, according to ACT, **domain expertise** can be viewed as a function of the quantity and quality of these two types of knowledge and, by definition, the extent of domain expertise will influence the nature and efficiency of domain-related cognitive processes.

At this point, the basic content of ACT* theory has been laid out in summary fashion. Those interested in a more detailed discussion of ACT* theory are invited to consult the original source, as expanded and modified in iterative fashion (Anderson, 1976; Anderson, 1982; Anderson, 1987).

In sum, ACT* theory is a model of human cognition, problem-solving and learning that specifies the processes by which individuals acquire cognitive skill. As such, ACT* is also a theory of how individuals acquire expertise in a given domain and ACT* processes can be used to suggest

expert-novice differences. The major processes involved in acquiring expertise are: creation of productions through interpretation of declarative knowledge and then -- after productions have been formed -- composition, procedural knowledge, proceduralization, generalization, discrimination and strengthening. While these processes have been briefly described, their implications for expert-novice differences have not been explored. The next section of the paper will examine these implications.

Expert-Novice Literature Review

The extensive expert-novice literature which exists has traditionally viewed expertise as a function of the extent of one's declarative knowledge within a domain (Spilich et al., 1979; Chiesi et al., 1980; Voss et al., 1980; Means & Voss, 1985; Hershey et al., 1990). ACT* recognizes the declarative knowledge base component of expertise but adds another component -- procedural knowledge. According to ACT, traditional conceptualizations of expertise are deficient in that they fail to address how experts use the declarative knowledge differently.

The study of expertise arose within the research and domains of cognitive science and artificial intelligence in conjunction with attempts to understand and model accurate human problem solving performance. The behavioral decision literature, stemming from these two disciplines, has largely focused on modelling and evaluating the outcomes of expert judgment and decision making. This section will review the

major exploratory studies on expert-novice differences and the implications that these findings have for hypotheses about experts and novices during information acquisition. The distinction between declarative and procedural knowledge was beginning to arise in the mid-1970s but, by and large, the expert-novice literature has failed to address this early distinction.

Knowledge Organization

In a classic study, Chi et al. (1981) asked experts and novices in the domain of physics to sort physics problems into categories. In a finding that has been replicated many times now in other domains (e.g., Hinsley et al., 1977; Schoenfeld & Herrmann, 1982), little overlap was found in the category labels used by the two groups. Novices were found to sort physics problems on the basis of "surface" structures such as literal objects mentioned in the problem or evident spatial relations while experts sorted their problems according to major principles of physics that could be used to solve the problems. It was also found that the basic approach to problem-solving for experts (i.e., physics principles) was cued by the description of the states and conditions of the physical situation.

Chi et al. (1981) found that, when cued for solution procedures with various types of physics problems, verbal protocol analysis revealed that expert responses could easily be put into the form of production routines. These production routines, organized into clusters of "If-Then"

rules used to guide working memory operations, utilized explicit solution methods such as algebraic equations. Novice production routines contained gaps where conditions were not tied to actions and often did not contain equations that could be used to solve the problem. When novice productions did contain formulas, the formulas chosen early in the verbal protocols contained the dependent variable and missing independent variables, while efforts were directed at finding the values for missing independent variables. This was seen as an indication that novices were utilizing a "working backwards" approach to problem solving.

From these results, Chi et al. (1981) hypothesized that knowledge is indexed in memory as a function of how a problem is categorized and suggested that expert-novice differences may be related to poorly formed, qualitatively different or nonexistent categories in the novice. They theorized that there are two components that interact in determining how information about a problem is represented. Initially, problem representation is a function of categorization. Then, upon categorizing the problem, a "problem schema" is invoked based on category membership and knowledge associated with the category is used to fill in the representation. Thus, problem representation is a function of how a problem is categorized and the thoroughness and/or existence of "problem schemata" with relevance to the type of problem at hand.

Novices were also found to organize information according to

After learning a set of computer programming concepts to criterion, McKeithen, Reitman, Rueter, and Hirtle (1981) found that experts displayed more subjective organization than novices in terms of the serial ordering consistency of key programming concepts in cued and free recall. Beginner students in computer programming appeared to recall ALGOL W concepts using a variety of common-language associations and mnemonic techniques. Expert recall was based upon functional similarity in that concepts were recalled contiguously when they occurred in the same program. A multidimensional scaling routine also revealed that the expert "tree" structures derived from recall data were more similar to each other than novice and intermediate were to other "tree" structures in their respective classes.

Adelson (1981) assessed subjective organization of recall to determine if expert computer programmers were more consistent than novice programmers in the way that computer programming concepts were organized in a multi-trial, free recall task. Using a measure developed by Sternberg and Tulving (1977) in which subjective organization is determined by the number of concepts recalled in pairs (i.e., temporally adjacent) on successive trials, it was found that the experts' pair frequency score was 84 percent of the maximum while the novices' pair frequency score was only 26 percent of its maximum. Using a multidimensional scaling routine and a hierarchical clustering technique, novices were also found to organize information according to

syntactic categories while experts organized concepts around membership in sub-programs that could be formed from the total set of concepts. Expert clustering on the basis of procedural similarity was seen to reflect a schematic knowledge organization in long-term memory.

Cooke and Schvaneveldt (1988) also used key words in computer programming to show differences between expert and novice organization. Subjects were given the task of rating the relatedness of 16 programming concepts, one pair at a time. Distance measures for each word pair were derived and scaled according to the PATHFINDER algorithm.

Resulting network solutions for each skill level as a group varied systematically as a function of computer programming knowledge. Relatedness ratings for concept pairs were correlated higher within groups than between groups. Across four skill levels, relatedness ratings for individual concept pairs and derived group network solutions were correlated higher with adjacent skill levels than levels separated by other levels with expert-naive ratings and group networks correlated the least. These two findings indicate a similar conceptual organization within each skill level with novices exhibiting the least intra-group between-groups agreement and novices exhibiting the least intra-group agreement. Finally, subjects could be categorized into skill levels on the basis of their own network structures, providing further evidence of consensual methods of organization at each stage in skill development.

The major results of these studies using sorting and serial order techniques support the notion that experts organize knowledge in terms of **how it is used** while novices organize knowledge on the basis of **syntactical similarity**. Also, it appears that experts are more consistent than novices in the way that they organize domain concepts, and the finding that knowledge organization for any given skill level is more similar to nearby skill levels than those further removed suggests that there may be a general, domain-specific developmental sequence that individuals go through on their way to becoming experts.

Domains with hierarchical goal structures

Several studies have attempted to explicate ("idealize") the hierarchical goal structure of a given domain and then show that experts have a greater understanding of how domain-related actions relate to and change the variables in this goal structure. Much of this work has been done by James Voss, George Spilich, Harry Chiesi and their colleagues. Spilich, Vesonder, Chiesi and Voss (1979) developed an idealized goal structure for the game of baseball that was to be used again in two more studies. Goals in baseball were arranged hierarchically (e.g., winning the game, scoring runs, advancing runners, etc.) and their respective attainment was represented at each level by a pattern of variables (e.g., having won or lost, the score, the number of batters on base, the pitching count, etc.). A **game state** was

then defined as the configuration of goal variables at each level at any given time, while **game actions** were seen as events occurring in the game that changed the value of a goal variable at one or more levels.

Spilich et al. (1979) hypothesized that the high knowledge individual has a more extensive knowledge of game actions as well as a greater understanding of how game actions are related to changes in a game's goal structure variables. Therefore, given a certain game action, high knowledge individuals know more about how such actions may produce a change in the game state.

Using a fictitious account of a half-inning of baseball, Spilich et al. (1979) found that high-knowledge subjects recalled more goal-related text propositions as a whole and, in particular, recalled significantly more propositions about relevant **enabling** aspects of the game setting, specific **game actions** and **auxiliary actions** pertaining to how game actions occurred. Experts also recalled information in the appropriate order more often than novices and recalled game actions in more integrated sequences. It was concluded that high-knowledge individuals have a greater knowledge of how specific game actions are related to the goal structure of the game and are better able to process sequences of game actions in terms of monitoring changes in the value of goal structure variables.

Chiesi, Spilich and Voss (1979) again used the idealized domain approach to baseball to illustrate further

behavioral differences between high- and low-knowledge individuals. In this study, high knowledge individuals were found to be better at recognizing both old and new text propositions, particularly when changes in the text became more important to the outcome of the game. It was also found that high knowledge individuals needed less information to make recognition judgments, were superior at recalling event sequences and anticipated a greater percentage of high-level goal state outcomes.

Voss, Vesonder and Spililch (1980) used baseball once again to look at differences between high- and low-knowledge individuals in terms of how they generate and recall domain information. Subjects were asked to generate an account of a fictitious half-inning of baseball and then recall those accounts two weeks later. It was found that high-knowledge individuals generated richer accounts in that game actions relating to lower-level goal states variables were mentioned more often. On the other hand, in terms of recall, low-knowledge individuals often displayed problems integrating sequences of actions that they themselves had generated. High-knowledge individuals recalled more of their passages correctly and were more likely to do so in the proper order. It was concluded that low knowledge individuals were deficient in the establishment of sub-goals and did not integrate game state change sequences as well as high knowledge individuals.

be discussed later but, for now, we turn to another group of studies that have looked at the

Using an idealized goal structure in the fictional domain of "Star Wars," Means & Voss (1985) found the same sorts of expert-novice differences. After the authors generated a hierarchical structure of goals relevant to actions in the movies "Star Wars" and "The Empire Strikes Back," participants were divided into high- and low-knowledge groups within two different age levels on the basis of a knowledge test about the two movies. Each participant was prompted to offer reasons explaining a number of low-level "basic" actions in the movies. After each correct response, participants were prompted to offer a reason for an action corresponding to a higher-order goal in an attempt to assess their ability to relate actions in the movies to the idealized goal-structure. Experts were able to identify more basic actions, more fully explicate the sub-goal structure and identify more high-level goals of the characters. It was concluded in these four studies that experts were better able to relate the basic actions in a given domain to logical, hierarchically-arranged goals within that domain. While no precise mechanism was offered to account for this conclusion, it was suggested that experts "map" to game actions and changes in goal state variables onto existing knowledge structures which are more complete and correct than the novices' knowledge structures. This "mapping" process will be discussed later but, for now, we turn to another group of studies that have looked at the

acquisition of information in decision making/problem solving. The problem is one which could be solved by Process-tracing and actions in response to the primary

cause. As the current study utilizes a process-tracing approach, studies involving experts and novices that have used this approach with experts and novices are particularly interesting and relevant. The following is a selective review of that literature.

One of the earliest process-tracing studies is one by Simon and Simon (1978). Using verbal protocols in the domain of physics, they found that experts took a quarter of the time that novices did to solve the problems given, and the experts made fewer errors. Experts and novices were also found to use different strategies in solving problems.

Another study by Voss, Greene, Post & Penner (1983) using verbal protocols with experts and novices in political science is illustrative of expert-novice differences in problem representation. In their study, experts and novices were presented with the problem of low crop productivity in the Soviet Union and asked to offer a method to improve it. The problem as stated contained little information that could be used to generate a solution. Experts were found to exhibit two different strategies to develop a more structured representation of the problem -- decomposition and conversion. Experts first decomposed the problem by using stored knowledge to make inferences, add constraints, and assign responsibility for the low productivity to a

small set of variables. At this point, the experts had converted the problem into one which could be solved by specifying remedial actions in response to the primary causes of the problem. Johnson (1980) found expert-novice differences in information usage with a verbal protocol procedure in a medical setting. The task involved rating the desirability of accepting applicants for residencies and internships by experienced physicians and novice undergraduates. Experts completed the ratings task in a little less than half the time that it took novices and experts looked at less than half as much information as novices did. Experts also examined **different** information, spending more time on the application form while novices concentrated on the transcript and letters of recommendation to a greater extent. Finally, experts examined information more actively, moving around in the folder and returning to previously searched information more often.

Johnson (1980) concluded that experts appeared to concentrate on a small subset of perceived diagnostic information and treat the remainder of the information as relatively uninformative. Also, nonlinear cue usage was an important part of experts' predictive validity (15%) when individual subject regression equations were constructed. Johnson, Duran, Hassebrock, Moller, Prietula, Feltovich and Swanson (1981) found further evidence for the interactive use of cues by experts in a study on the

diagnosis of congenital heart defects. Using an actual case study depicting a particular form of pulmonary valve defect in the heart, four "diagnostic" cues in the profile of a patient's case information were dichotomously manipulated as either "strong" or "weak." Sixteen variations of the four cues representing all possible combinations were given to four experts, four trainees and four students and an expert computer diagnostic algorithm known as DIAGNOSER. Experts made more correct diagnoses than the other participants and DIAGNOSER, apparently through the use of two diagnostic cues in an interactive fashion when the cues were in their "weak" form. While the combination of "weak" cues was apparently not enough to convince the novices to make the correct diagnosis, according to Johnson et al., (1981), "The experts, on the other hand, were sufficiently confident on the basis of the remaining data presented in the case and were therefore unaffected by the lack of strong evidence" (p. 270).

Hobus et al. (1987) conducted a study to compare how experienced and inexperienced doctors use contextual information. Experienced doctors and medical students were given short case-histories presented on slides. Each case history contained a picture of the patient, a previous disease history and the usual presenting complaint. The information implicitly provided by the picture and disease history was predicted to be more meaningful to experts and thus processed more elaborately. Specifically, contextual

information was assumed to provide knowledge about the existence of "enabling conditions" in the disease process. As predicted, experts generated significantly more correct hypotheses and were also able to recall more of the contextual information. It was concluded that "experts are better able to utilize information implicitly available in an information-restricted environment than novices are. . . ." (Hobus et al., 1987; p. 475)

In another study, Johnson & Sathi (1988) provided experts (experienced research analysts) and novices (MBA students) with the task of predicting year-end closing prices for 40 securities with available information on 22 dimensions. Half of the securities were accompanied by news item summaries of stories about the company that had appeared in the Wall Street Journal over the course of the year. Regression equations were then calculated for each individual to determine what cues were used to make while decisions. Johnson & Sathi (1988) found that experts took less time per security and looked at fewer attribute dimensions than novices. Experts were also marginally more accurate than novices as a whole. Most importantly, the presence of the news items substantially increased expert accuracy but did not affect either the novices' accuracy or the accuracy of the linear regression model. Upon dividing experts' predictive validity into linear and nonlinear components, it was found that nonlinear cue usage was only significant in

the news item condition and that, in this condition, over 50% of the experts' predictive validity was due to the nonlinear use of this information. On the other hand, novices did not use nonlinear information to any appreciable extent. Johnson & Sathi (1988) concluded that experts concentrated their search on the identification and interpretation of rare events (cue values). Dawson, Zeitz and Wright (1989) looked at expert-novice differences in social cognition. Clinical novices and experts "observed" behavioral examples of three prototypical child "targets": an aggressive target, an inverse target (whose behaviors can best be described as counter-intuitive in each situation) and a random target. Experts were more accurate at labelling the aggressive target, were superior at making predictions for the inverse target and organized their free recall of target behaviors around classes of antecedent events (e.g., adult praise, censure, etc.) while novices did not organize their behavior according to any particular principle. A final study by Hershey et al. (1990) actually compared experts and novices directly using a modified, asymmetrical information board methodology. They theorized that expertise consists of a construct very similar to productions and problem schemata: the "script." Through experience, experts develop problem-solving "scripts" that employ a set of rule-based mental operations to identify and utilize relevant problem parameters in reaching a solution.

"Procedural efficiency" for novices was expected to be low. Expertise was operationalized on the basis of domain-based knowledge for financial problem-solving and created a found hierarchy of variables in three domains (need, suitability and affordability) relating to the decision of whether a married couple should open an IRA. Then, using verbal protocols and problem-solving process maps, they found that experts solved problems in less than half the time of novices using fewer overall steps, utilized information at higher levels in the information hierarchies and looked at fewer pieces of unique information. It was concluded that expert problem-solving routines appeared more goal-oriented and demonstrated superior representation of the problem.

The studies on expert-novice differences using process-tracing methodologies fill in and support the findings in other domains. Experts have tended to be faster, to make fewer mistakes, use less information and use that information more configurally. While more will be said about this later, such findings are generally consistent with ACT*. The next section will review and synthesize the major findings on experts and novices up to this point.

Summary of research on expertise

A number of studies have been reviewed that examined the different outcomes associated with expert and novice performance and have found that, as would be expected, experts tend to be more accurate than novices (Simon and

Simon, 1978; Johnson, 1980; Johnson et al., 1981; Hobus et al., 1987).

Using process-tracing methods, experts have been found to gather and use information differently than novices, as would be suggested by a production systems account of human problem-solving/decision-making. Experts have been found to be faster than novices (Simon & Simon, 1978; Johnson, 1980; Johnson, 1988; Hershey et al., 1990), to use fewer cues than novices (Johnson, 1980; Johnson, 1988; Hershey et al., 1990) and use higher-order information (Hershey et al., 1990) in making their decisions. Finally, experts have been shown to use information in a nonlinear fashion to a greater degree than novices (Johnson et al, 1981; Johnson, 1980; Johnson, 1988). These findings are generally consistent with ACT*, and will be used to support the prediction of various findings in the current study.

ACT* and the Task Environment

The assumptions of the cognitive architecture hold implications for the selection of two other constructs in problem solving/decision making research. According to ACT*, the key aspects of cognition are: 1) activating declarative knowledge structures in working memory, 2) selecting one or more productions to test for applicability and 3) matching working memory knowledge structures to the conditional clauses of domain-related productions. Any phenomenon affecting these activation, selection and

matching processes will influence the nature and efficiency of domain-related cognitive processes. Two such phenomena are the presence of descriptive category labels in the problem-solving/decision-making context and the degree of "structure" (stereotypicality) in the task statement. ACT suggests that category labels might affect cognitive processes by facilitating the activation of declarative knowledge structures in working memory. Task structure, on the other hand, is implicated in the selection of productions for testing and the matching of working memory information to the conditions of productions.

With regard to the **activation** of knowledge structures, ACT specifies three ways in which information can achieve the necessary level of activation to be accessible to working memory: 1) through perception of some object or concept in the external environment resulting in the activation of the declarative knowledge structure corresponding to that object or concept, 2) through creation in working memory itself by manipulating or combining active knowledge structures and, finally, 3) through spreading activation among connected knowledge structures in the existing declarative knowledge base. ACT* postulates a "spreading activation" mechanism and, therefore, implicates the use of category labels in information acquisition,

Category labels should affect the problem-solving/decision-making process by influencing the spread of activation to associated knowledge structures in declarative

memory. Category labels can be thought of as nodes in an associative semantic network. These "category label" nodes are strongly related to other nodes representing the characteristics of the category. Presumably, when category label nodes become active in working memory through the perception during information search, a great deal of activation is generated and spreads out to other "feature" nodes along the strongest links. This will proceed to rapidly bring related features (i.e., concepts or psychology characteristics) to the activation levels necessary to enter consciousness. The information that becomes active in WM can be used as if it were gained from search of the external environment.

Spreading activation is a function of node strength and link strength. Node strength in turn is a function of how frequently a node has been activated in the past and determines the overall quantity of activation which will spread from a given node. Link strength is a function of past amounts of activation which have spread from one node to another along a given link. Link strength determines how the activation at a given node will divide up and spread out along the various links which are connected to the node in question.

For experts, domain-related elements (features, etc) should be strongly linked, with the result that activation is "channeled" down these domain-relevant links and not down other links from a given node, say links relating the two

nodes on the basis of semantic similarity. Thus, for the expert, activating the node for a given category label should bring into working memory various domain-related features and attributes of this node through spreading activation over strong, domain-relevant nodes. For the novice, activating the node for a given category label should bring into consciousness other concepts that bear a semantic, common-language similarity to the label.

For the most part, the traditional cognitive psychology literature has tended to view category labels as ACT* does - as nodes in an associative network of semantic concepts (although see Hintzman, 1986; Kahneman & Miller, 1986 for interesting alternatives). Therefore, there are a number of studies in the literature which are relevant to an understanding of how category labels are used to solve problems and make decisions. In a study by Fiske et al. (1987), subjects were asked to rate the likability of target individuals in the presence/absence of stereotyped occupational labels (e.g., professor, artist) in conjunction with the consistency/inconsistency of available attribute information. Correlations between independent ratings of the labels themselves and the likability of persons with labels plus attribute information indicated that more category-based processing went on in the label-consistent conditions as predicted. In a second study using the same general procedure but with the addition of verbal protocols,

it was found that more attribute traits were mentioned by subjects in the conditions where attribute-based processing was predicted to occur.

Two information board studies have directly looked at the effects of labels on subsequent search behavior as well. Hattrup and Ford (1991) asked subjects to rate the attractiveness of target profiles in terms of their desirability as co-workers. Target profiles were occupational stereotypes, with and without their respective labels, and attributes were either consistent or inconsistent with the occupational stereotype. Results indicated that subjects did indeed search for less individuating information in the presence of occupational labels and took less time to make their ratings. Also of note, these two findings were not affected by the consistency of the attribute information with the elicited category.

The second information board study (Gilliland, Wood, & Schmitt, in press) asked persons experienced in real-estate and economic development to rate the desirability of locating a business in various states. As predicted, less attribute information was accessed when states were labelled as opposed to unlabelled. Furthermore, in this repeated measures design, fewer states were examined and information acquisition across states was more variable when participants received the labelled decision task first but

not when they received this task second after receiving the unlabelled decision task previously.

Several other studies have examined the effect of category labels in the context of differences in domain knowledge (i.e., expertise). These studies suggest that category labels allow the expert various advantages in performing domain-related tasks.

In an interesting study that held constant the amount of knowledge about a domain while examining expert-novice recall, Fiske, Kinder and Larter (1983) found differences in the manner that high and low knowledge subjects in political science organized their recall about the country of Mauritius. Subjects read descriptions of the country in which it was labelled as either communist, democratic or neither along with attributes which were both consistent and inconsistent with this label. It was found that individuals with high knowledge in political science recalled more attributes overall as well as a greater number of attributes that were inconsistent with the stated ideology of the country. On the other hand, low knowledge individuals recalled consistent attributes for the most part. Overall, both groups tended to order the recall of attributes in a similar fashion. Thus, it appears that labels in this case allowed the expert to organize the information provided and construct a superior representation of that knowledge in working memory (i.e., available to be recalled).

Labels should allow the expert to engage in category-based

In a study using examining clinical diagnostic categories, Murphy and Wright (1984) asked expert clinicians, counselors with some experience and novices to list as many diagnostic features as possible for children in three psychopathological categories. It was found that experts had larger categories in terms of number of attributes but also had lower category distinctiveness in that attributes that were added to each category become progressively overlapping. Murphy and Wright suggested that people first focus on relatively distinctive features of concepts in order to separate them from other concepts. Later, with increasing knowledge and attention to real-world covariation, this "discrete" categorization relaxes, presumably reflecting a growing network of domain-related associations in the declarative knowledge structures of the expert.

Summary -- Category Labels

These studies suggest that the presence of categorizing labels, when associated with meaningful stereotypic schemata, affect resulting decision processes. The extent of domain knowledge possessed by an individual should interact with the presence of labels in that high knowledge individuals have rich, well-developed schemata associated with their category labels while low knowledge individuals may have incomplete schemata or none at all (Chi et al., 1981). The presence of category labels and procedural schemata should allow the expert to engage in category-based

processing and reduce information acquisition while low knowledge individuals should adopt a more extensive attribute-based processing approach because of the need to acquire relatively more attribute information.

Thus, "category-based processes" are traditionally thought to involve the activation and use of default features associated with various categories stored in declarative memory. ACT* suggests that, through the mechanism of spreading activation, features that are strongly associated with various categories are activated in declarative memory and become accessible to working memory without the need to acquire such information from the environment. However, category-based processing cannot be used by individuals who do not possess the links between feature information and a given category label (i.e., novices).

Task Structure

ACT* also suggests a number of influences on the selection of a production in the conflict resolution stage of cognition as a result of task structure. According to a recent paper (Anderson, 1992), ACT* specifies five factors that determine whether a production will be chosen from a set to be executed in a given situation: the strength of the individual production, the strength of other, competing productions, the degree to which the conditional elements of the production have been satisfied, and the activation of the knowledge structures representing these conditional

elements. Aspects of the task that affect one or more of these factors are relevant to the study of how people solve problems and make decisions.

Task structure should impact the problem-solving/decision-making process by making certain productions more likely to be selected for testing and, also, by increasing the speed at which conditional elements in such productions are matched against the contents of working memory. When a stereotyped problem situation is encountered, a strong, domain-specific production is indexed for selection by the constants in its conditional clauses. Selection of applicable productions is a probabilistic function of how often a production has been used successfully in the past. Therefore, frequent exposure to common domain-specific situations or patterns results in the creation of "strong" productions which become compiled and proceduralized to a greater extent each time they are invoked.

Stereotypical problem situations result in the creation of productions which benefit the problem-solver at the matching stage of cognition as well. Highly compiled and proceduralized productions are capable of simplifying complex, contingent processes to the point of a single production containing one or two conditional clauses. Information in working memory can be quickly matched to these conditions, as matching time in ACT is a function of knowledge structure activation level in working memory and

the strength of a given production. Therefore, the conditional elements to be tested for a strong, frequently-used production are few in number and tested quickly relative to the conditions of other, less frequently used productions.

Reitman (1965) defined task structure along a continuum representing the number of "open" constraints left unsatisfied by the problem statement. **Well-structured** problems and **ill-structured** are opposite ends of a continuum defined in terms of the number of "constraints" left unresolved by a problem statement. According to Reitman (1965),

"To the extent that a problem situation evokes a high level of agreement over a specified community of problem solvers regarding the referents of the attributes in which it is given, the operations that are permitted, and the consequences of those operations, it may be termed unambiguous or well-defined with respect to that community. On the other hand, to the extent that a problem evokes a highly variable set of responses concerning referents of attributes, permissible operations, and their consequences, it may be considered ill-defined or ambiguous with respect to that community" (p. 151).

The traditional definition of a well-structured problem involves the perception of a common, domain-specific pattern of antecedent stimuli and a consistent relation or action "filling" the open constraint of "RESPONSE." In other words, the defining element of "structure" in a given problem situation is the degree to which that problem is stereotypical. To the extent that a problem situation is familiar and has been experienced in a stereotyped,

routinized fashion, it can be considered well-structured. Highly structured problems involve situations in which open constraints have been filled in a consistent, identical manner over time. Conversely, ill-structured problem situations lack a consistent pattern of constraint resolution.

The "structure" of a given task is partly a function of its domain. Some domains (e.g. mathematics, physical sciences) are considered well-structured because there is a great deal of agreement on the relations between domain elements, acceptable parameter values for these elements, permissible operations on these parameter values and agreement on the consequences of various operations. In domains such as political science, economics and sports, typical problems often have multiple, conflicting goals, causal information about the relationships between variables is ambiguous or conflicting, and one person's solution is another person's mistake. These types of domains are considered ill-structured domains in that there is little consensus about the nature of the problem (i.e., what it is, how to solve it or when it is solved). In other words, the degree of problem structure is a function of how well-specified the goal is, how much information is given to start the problem and how much agreement there is concerning what the goal is or when it has been achieved.

There are a number of studies in the literature on expert-novice differences that examine the effect of how

domain stimuli are "structured" in a task environment. The typical study has found that experts recall a great deal more information than novices when such information is structured according to domain-specific conventions that give it contextual **meaning**. When information is "unstructured" (e.g., randomly configured), expert recall performance decreases to the level of the novice.

In a classic study, Chase and Simon (1973) found differences between experts and novices performance in chess. Chase and Simon used two different (and now standard) tasks to infer that experts and novices did indeed encode and store knowledge in different fashions. The "perception" task asked subjects to reconstruct chess positions using as many glances at the original position as necessary. The second task, a memory task, allowed subjects to reconstruct as much of a chess position as possible after five seconds of view.

What Chase and Simon (1973) found was something that would be replicated many times again in other studies: when chess pieces were arranged in configurations taken from real-life games, experts could reconstruct the chess positions with relative ease -- much faster and more accurately than novices. However, when chess pieces were randomly arranged, expert performance declined to the level of the novice.

Chase and Simon (1973) concluded that the superior expert recall in structured situations was the result of a

greater recognition of chess patterns rather than any basic processing advantage in terms of working memory. This performance advantage for real-life game positions was attributed to a "chunking" mechanism that allowed them to cluster the pieces on the board into various sub-patterns and store these clusters as a **unit**. Presumably, "chunking" (determined by glances in the perception task and pauses in the memory task) was possible in the real-life game positions because sub-patterns on the board matched "chunks" that had previously been experienced and stored by the expert. Experts were found to have larger chunks than novices and experts were found to chunk pieces together in terms of abstract patterns of relations (e.g., attack, defense) rather than proximity, color and piece type such as the novices. Experts were also found to recall more chunks than novices in one of the tasks in spite of the finding that experts suffered from the same constraints imposed by a limited working memory (WM) capacity, suggesting that "chunks" can also be chunked hierarchically.

Chase and Simon interpreted these findings as evidence that expert knowledge shows superior organization compared to novices. This enables experts to match chess configurations with existing stored configurations acquired through experience, chunk "chunks" hierarchically and, in the end, retain more information. Novices, without the benefit of stored patterns, cannot chunk pieces into clusters and are to store individual chess pieces rather

than unitized clusters. Thus, they suffer from the rigorous constraints of working memory. Chi (1978) was able to replicate these same findings with child experts and adult novices in the domain of chess. In this case, child experts were clearly better than adult novices when chess positions were structured but not when unstructured.

In a study by Reitman (1976) involving the game GO, the classic findings of Chase and Simon (1973) were partially replicated. Using the same tasks as Chase & Simon, experts were able to reconstruct game positions in the perception task quicker than novices and were able to recall more pieces in the memory task than were novices **when board configurations were taken from real games** (i.e., "structured").

Engle and Bukstel (1978) were able to replicate the classic findings of Chase and Simon (1973) in the domain of bridge. In their study, subjects were asked to reconstruct bridge hands in the now-standard perception and memory tasks and were asked to play 10 bridge hands. As before, Engle and Bukstel (1978) found that when information was organized ("structured") according to suit in the perception and memory tasks, the two expert subjects performed much better than the novice. This was not the case on the memory task when hands were arranged in an unstructured fashion. In the perception task, the two experts did better in both structured and unstructured conditions but took longer in the latter condition -- perhaps reflecting an on-line

re-configuring of information according to the way schematic information was organized in memory (i.e., by suit).

Charness (1979) also used the domain of bridge to look at the recall of experts and novices and found similar effects of structure (i.e., arrangement of the cards). Tasks in this study included rapid bidding, card recall (immediate and after a short study time), and planning out the play of hands using verbal protocol procedures. As in previous studies, performance was highly and positively correlated with expertise when hands were arranged by suit (i.e., were well-structured). Charness (1979) concluded that skill in bridge consists of having a large store of recognizable bridge hands associated with appropriate actions. It was hypothesized that better players encode the cards in a manner that triggers plausible lines of play. With the requisite vocabulary of stored card patterns, expertise becomes a matter of classifying hands correctly so that the proper strategy is invoked.

Egan and Schwartz (1979) used skilled electrical technicians to show the effects of expertise on recall in yet another domain. After brief exposure to circuit diagrams, expert technicians were able to recall more information than novices when circuits were functional (**meaningful**) but not when they were **random**. IRTs and transitional error probabilities indicated that experts were "chunking" their recall by function, were faster on their

between-chunk transitions and, also, were found to have larger initial chunks during recall as compared to novices.

Summary -- expertise and task structure

There is a good deal of research that suggests that expert performance benefits from having information organized in a meaningful way. "Meaningfulness" is a function of past exposure to common domain-specific patterns (chess, circuitry) and methods of organizing information (bridge). In each domain, recall advantages exist for experts when meaningful patterns are perceived and these same advantages all but disappear when conventional patterns are removed (Chase & Simon, 1973; Reitman, 1976; Engle & Bukstel, 1978; Charness, 1979; Egan & Schwartz, 1979).

In ACT, high task structure can be thought of as analogous to a stimulus situation that activates a frequently-used production containing only a few simple conditional elements that need to be satisfied. The "domain-specific pattern of stimuli" from the traditional definition corresponds to the satisfaction of conditional clauses of a production in ACT*.

Labels, Task Structure and Expertise

In summary, the goal of the cognitive process is to activate information in working memory that corresponds to the conditional clauses of one or more productions in production memory and then choose one production to apply. In other words, to solve problems and make decisions, the system needs a certain amount of declarative knowledge as

well as productions that can utilize this knowledge. **Domain expertise** is now defined as the quantity and quality of two types of knowledge in the system: domain-related declarative knowledge and domain-specific productions applicable to this declarative knowledge. Domain expertise should have an effect on the cognitive process by influencing the amount of information that can enter working memory by the process of spreading activation and, also, by influencing the number of productions that can be applied to information in working memory.

Category labels in ACT are nodes which serve as a source of activation for strongly-associated networks of nodes representing domain-relevant information. According to ACT*, category labels should have an impact on the problem-solving/decision-making process by providing a tightly clustered network of declarative knowledge that can be quickly activated by the presence of a label and brought into working memory through the process of spreading activation. **Task structure** is the degree to which a task situation or problem statement elicits a solution production which has its conditions satisfied without the need for further information search. High task structure should result in the immediate selection of a compiled, highly proceduralized production involving few conditions (possibly only a goal condition) that can be tested rapidly, leaving the problem-solver/decision-maker with a "solution" almost instantaneously.

Up to this point, ACT has been used to suggest the importance of two constructs on the problem solving/decision making process, category labels and task structure. While these constructs have been discussed in relative isolation, it is important to note that the effects of labelling and task structure are strongly tied to the extent of domain expertise that an individual possesses. In other words, the mechanisms underlying the effects of labelling and task structure are such that novices would not be expected to benefit from the presence of labels in a problem/decision context or high task structure. On the other hand, ACT suggests that experts should benefit from both the presence of labels and highly-structured task situations.

Consider first the case of labels. Labels presumably have their effect by augmenting the spread of activation to related concepts and features, thus bringing additional information into working memory without the need to acquire it from the environment. For spreading activation to occur in ACT, declarative knowledge structures must already exist and be linked to other nodes in declarative memory. If an individual has no existing declarative knowledge base (i.e., is a novice), then all knowledge that is used in solving a problem or making a decision must come from perception in the environment or creation in working memory.

This line of reasoning is consistent with Fiske and Neuberg's (1987) continuum of processing theory, which suggests that labels allow experts to use category-based

processes (i.e., use knowledge associated with category labels). Category labels, when combined with information that ~~is~~ available, should allow the expert to rule out certain alternatives or attribute dimensions from further consideration through the substitution of "default" features drawn from memory that are associated with the various labels (Johnson, 1980). This category-based "labelling" advantage should save time and reduce information retrieval. In the case of labels, the expert uses the labels in the problem context to activate related, problem-relevant information in working memory through the aid of strong associations in declarative memory.

The presence of labels in the problem context then implies two benefits to experts that would not seem applicable to novices: less information needs to be acquired from the environment, and the conditional clauses of various productions should be tested faster than those of novices due to the higher level of activation for a given feature.

Similarly, task structure is also a domain-dependent phenomenon. Through past frequent association, experts have associated certain common patterns (problem classes) with certain "successful" actions in the form of productions. Initially, problem-solving routines involve many productions, local variables and hordes of conditional clauses. However, over time, constants replace local variables in the conditions, the routine becomes compiled

and many conditional clauses drop out of the resulting production with proceduralization. With frequent usage, a long string of productions may be proceduralized and compiled into one production with perhaps a goal condition as its only antecedent. When this goal statement, which corresponds to a particular class of problems, becomes activated, the production is quickly invoked on the basis of its constants and the few conditions (if more than one) are rapidly tested.

ACT has thus been used to show the relevance of three constructs on the problem-solving/decision-making process: domain expertise, category labels and task structure. An attempt has been made to demonstrate that the effect of category labels and task structure on information acquisition is largely dependent on past experience. As such, the effects of these two constructs should be examined jointly with domain expertise.

Model and Hypotheses

Figure 1 depicts a conceptual model of the discussion to this point. According to this view, expertise is a result of five ACT* mechanisms that produce the procedural knowledge necessary to acquire skill in a given domain. These mechanisms are composition, proceduralization, generalization, discrimination and strengthening. Over time, and with numerous and varied experiences, these mechanisms will produce some degree of expertise in an individual. The degree of expertise attained then affects

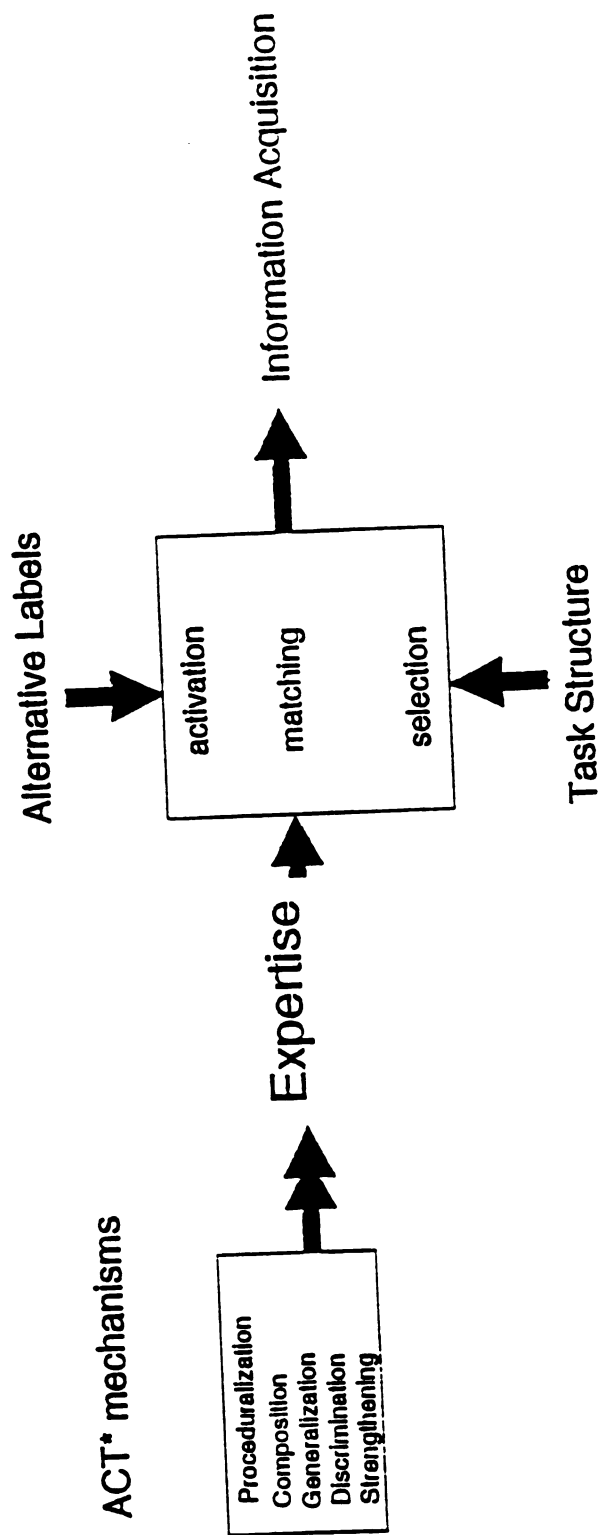


Figure 1. An ACT*-based model of Information Acquisition

the information acquisition process for a given task. Other things being equal, ACT* suggests that experts will generate a set of productions tailored by discrimination, strengthened by frequent usage, compiled into problem-solving routines and condensed by proceduralization. Thus, for the expert, problem-solving routines are triggered rapidly and with little information and are thus able to avoid the constraints of limited working memory capacity.

Relative to novices then, experts should be more accurate, examine cues for a shorter period of time, access less information, and search more variably across alternatives. Past literature has tended to support these findings. Every study reviewed in this paper has found experts to be more accurate than novices (see above) -- as one would expect by definition -- but experts have also been found to take less time in problem-solving (Simon & Simon, 1978; E. Johnson, 1980, 1988), access less information (E. Johnson, 1980, 1988; P. Johnson et al., 1981; Hershey et al., 1990) and use information in a more variable fashion (E. Johnson, 1980, 1988; P. Johnson et al., 1981).

However, the past literature has tended to focus on both process and outcome differences between experts and novices and take the task for granted. Indeed, Ford et al. (1989) found that the process-tracing literature as a whole has tended not to examine variations within the task except for task complexity, which has been operationalized by giving decision-makers a varying number of attributes and/or

alternatives to choose from. Given the robust effect of task complexity (Payne, 1982; Ford et al., 1989), it would seem fruitful to examine other intra-task differences in the context of differences in decision-maker knowledge.

ACT* provides the means for suggesting why this is the case. As shown in Figure 1, ACT* suggests that the intra-task variables will affect experts and novices differently. As discussed above, ACT* suggests that a well-structured task and the presence of labelled alternatives will have information value for the expert but not for the novice. Task structure should provide information to the expert concerning **which** attributes are needed in the choice alternative while alternative labelling provides information about **where** these attributes are most likely to be found (i.e., which alternatives are most likely to possess them).

Presumably, with any problem or decision in a given domain, initial efforts to arrive at a solution involve many productions, hordes of antecedent conditions to satisfy and a great deal of intermediate knowledge to remember and pass on to the next production just to yield a search for the needed attribute(s). Over time, composition, proceduralization, generalization and discrimination will result in relatively few productions, as suggested by the initial part of the model, ACT* --> expertise.

The traditional definition of a well-structured problem is one in which task stimuli are organized according to common, domain-specified conventions and open constraints

are few in number and consistently "filled." Thus, well-structured versions of a task should correspond to (and invoke) strong, compiled, well-tuned productions that are selected and tested easily and quickly. Ill-structured versions of the same task may have to rely on the older, less-efficient production routines which must compete with others to be selected, require more information, are tested more slowly and, in the end, do not allow the expert to behave like an expert.

The implication here is that when expertise and task structure are both allowed to vary, we may see the expert behave radically different on essentially the same task. In effect, Figure 1 predicts that the expert will benefit from the well-structured problem much as would be expected from the more global, "main effect" hypotheses stated above that make no mention of task structure, but even more so. Accuracy should be high on a well-structured version of the task, as the relevant productions have been compiled and tuned through extensive feedback. In addition, information access should be focused on just the important attributes and no more. Since the needed attributes are known in advance and are needed regardless of other environment variables, no other attributes need to be searched. In sum, with well-structured tasks, experts should be highly accurate and should acquire less information, especially with respect to "contextual" information which is

unnecessary due to the strong associations between highly-structured situations and certain specific attributes.

For alternative labels, the benefit to experts is similar to the benefit provided by a well-structured situation but through a different mechanism. It seems reasonable to postulate that, in addition to a domain-specific *solution* production, task performance is also a function of a production or set of generic productions that *search* the environment for information. ACT* suggests that the effect of alternative labelling may result from an interaction between the extent of declarative knowledge possessed by an expert and how the search production process is organized. In essence, the expert can use her/his greater store of conceptual declarative knowledge to order and focus search on alternatives that are most likely to possess the necessary attributes because category labels provided by the task environment can be used to probabilistically associate necessary attributes with choice alternatives on the basis of stored conceptual knowledge in declarative memory. Then, this inferred information can be used to guide information acquisition. Thus, we might expect experts to acquire less information when alternatives are labelled with reference to a category because fewer alternatives should be searched and search, which may occur for many attributes in an effort to label the alternative, may be terminated arbitrarily when the alternative can be categorized. This also implies that expert search with

labelled alternatives will be more variable across alternatives because not all of the available alternatives will be searched.

Table 1 lists the summarized hypotheses for this study. Note that hypotheses are arranged by dependent variable and, within each number, become less sweeping and more qualified by ACT*.

Table 1.

Study Hypotheses.

1. Experts will be more accurate than novices across all situations.
 - 1a. Experts' accuracy will improve in the well-structured task relative to the ill-structured task while the accuracy of novices will remain the same.
2. Experts will spend less time looking at each cue than novices.
3. Experts will access fewer cues in making their decisions than will novices.
4. Experts will access more cues in the contextual sub-matrix than will novices.
 - 4a. When alternative labels are present in the decision situation, experts will access fewer cues in the contextual matrix relative to when alternative labels are not present. Novices will be unaffected by alternative labels.
 - 4b. When the decision situation is well-structured experts will access fewer cues in the contextual matrix relative to when the decision situation is ill-structured. Novices will be unaffected by decision structure.
5. Experts will access fewer cues in the solution sub-matrix than will novices.
 - 5a. When alternative labels are present in the decision situation, experts will access fewer cues relative to when alternative labels are not present. Novices will be unaffected by alternative labels.
 - 5b. When the decision situation is well-structured, experts will access fewer cues in the solution sub-matrix than will novices. Novices will be unaffected by decision structure.

Table 1 (cont'd)

6. Experts will be more variable in information acquisition than will novices.

6a. When alternative labels are present in the decision situation, experts will be more variable in information acquisition relative to when alternative labels are not present. Novice variability will be unaffected by alternative labels.

6b. When the decision situation is ill-structured, experts will be more variable in the choice sub-matrix than will novices. Novices will be unaffected by decision structure.

7. Experts will access information in a more interdimensional fashion than novices.

8. All participants will search more interdimensionally when alternative labels are present than when they are not present.

METHOD

Participants

Two hundred seventy-eight participants were recruited from the introductory psychology pool at Michigan State University to take a test of basketball knowledge. After scoring test results, 123 students were called back (62 students with high domain knowledge and 61 students with low domain knowledge) to participate in the 90 minute self-paced, computer-mediated portion of the experiment. Students received course credit for participating in each part of the study and no monetary incentives were offered.

Design

This study involves a fully-crossed three-way factorial design: **2 (Domain knowledge: high, low) X 2 (Alternative labels: present or absent), 2 (Decision structure: ill-structured, well-structured)**. Each subject participated individually and made two decisions related to game situations in basketball. In each decision situation, subjects were asked to take the role of head coach for an NBA basketball team at the end of a regular-season game. Each decision situation presented the participant with a game situation (e.g., the game is tied with 22 seconds left on the clock, etc) and four "starting" players. The

decision task for each participant was to choose a fifth player (from among four potential alternatives) to be on the floor when play resumes.

In order to make each decision, participants needed to acquire some information from the information matrix. Four of the alternatives in the information matrix (i.e., the existing four players) are provided as a context in which to aid selection of the fifth player. These four alternatives and their attributes represented the **contextual sub-matrix**. Subjects could examine information provided for the alternatives in the contextual sub-matrix but were not allowed to choose any of these alternatives to complete the team. The decision task was completed by choosing a player from among the four allowable alternatives (i.e., the **choice sub-matrix**). Alternatives in the contextual and solution sub-matrices were described along the exact same attribute dimensions. The only functional significance of differentiating the two matrices involved the stipulation that the final choice must come from the choice sub-matrix. A computer was used to record the manner in which participants acquired information for each decision.

Independent Variables

Participants were categorized into high domain knowledge or low domain knowledge on the basis of their performance on a multiple choice questionnaire on basketball knowledge (See Appendix A for this questionnaire). **Domain knowledge** was measured as a between-subjects factor using

this 40-item test. In keeping with the implications of ACT*, an attempt was made to construct the knowledge test so that it would tap both declarative and procedural knowledge.

The basketball knowledge test was constructed as follows: A set of 15 items assessing declarative knowledge in the domain of basketball formed the core of the declarative knowledge component in this study. These items were used in a previous pilot study and KR-20 reliability was estimated at $r = .86$. In addition several items were added to this scale in an attempt to make it more content valid. Following this, items were written concerning role differentiation among basketball positions, strategy and tactics in an effort to tap procedural knowledge of basketball.

The entire test (both declarative and procedural scales) was then given to two members of the Men's Basketball coaching staff at Michigan State University in order to assess the degree of convergence concerning "right" answers for the procedural knowledge items. Two members of the coaching staff responded to this request. Agreement between both coaches and the a priori answer key was good (see Results for further discussion).

The **Alternative labels** factor was manipulated as a between-subjects factor by including a position label for the various alternatives (i.e., point guard, off guard, small forward, forward, power forward, center) or simply assigning each alternative a letter of the alphabet (i.e.,

Player A, Player B, etc.,). Half of the subjects made their decisions with all alternatives labelled by position, half made their decisions without. These labels are used ubiquitously at all levels of basketball to describe player roles and thus should be probabilistically associated with various levels of certain attributes (i.e., performance statistics) in the minds' of individuals with a fair degree of basketball knowledge.

Decision structure was manipulated as a within-subjects factor by having each subject make two decisions. One decision involved a common endgame basketball situation where upcoming behavior of the opposing team could be predicted with some degree of certainty (well-structured) and the other decision involved a more general, late-game situation where there is little agreement on how future play will transpire or which players (and their corresponding attributes) best serve the team (ill-structured).

The two decision situations were chosen to reflect opposite ends of the task structure continuum. The ill-structured condition involves a situation in which open constraints (i.e., upcoming actions by the other team and needed attributes on one's own team) are questionable. The well-structured condition involves a situation in which the open constraints are routinely filled in a predictable manner -- the other team fouls immediately and the team that is ahead needs to have good free-throw shooters on the court. The ill-structured task cannot logically be resolved

without recourse to the contextual matrix to determine which attributes are lacking on the present team members. The well-structured situation provides enough information to identify one attribute which will certainly be needed. Thus, the well-structured decision should provide enough information to trigger a production like the following:

IF: One's team is ahead late in a basketball game

THEN: Put in good foul-shooters

Stimulus Materials

The proposed design involves the use of two computerized information boards in a controlled setting. Each computer information board consisted of eight alternatives (players) X 15 attributes (performance-related information) for a total of 120 pieces of information. Each computerized information board included the following dimensions for every alternative:

- 1) Season field goal percentage
- 2) Season free-throw percentage
- 3) Season three-point field goal percentage
- 4) Field goals made-attempted (game)
- 5) Free-throws made-attempted (game)
- 6) Three-point field goals made-attempted (game)
- 7) Turnovers (game)
- 8) Offensive Rebounds (game)
- 9) Defensive Rebounds (game)

- 10) Steals (game)
- 11) Blocked Shots (game)
- 12) Assists (game)
- 13) Points scored (game)
- 14) Years in the NBA
- 15) Height

Two highly-similar information boards were created for the two decision tasks. Both information boards were constructed by generating cue values for each of the above 15 attribute dimensions for the eight player alternatives after consulting Hollander (1991). The manner in which cues were distributed across the various attribute dimensions was intended to reflect real-life skewed tendencies. Thus, in both information boards, some cue dimensions are positively skewed, some normally distributed and some negatively skewed. Cues in a respective attribute dimension were distributed in like manner for each of the two decision tasks (e.g., season field goal percentage is positively skewed in both decision tasks). Each search matrix was used for one decision task. (See Appendix B for the two search matrices used in the study).

The first decision task represents a common basketball endgame situation for which a consensus exists concerning what should happen and what will happen. This decision task featured a situation in which the subject's team a) had the ball, b) was ahead by four points and c) 22 seconds

remained in the game at the end of a regular season NBA game. Also, the 24 second shot clock was off, so the subject's team could conceivably hold onto the ball and run out the clock without shooting if allowed.

The second decision task involved a similar, but less stereotypic, situation. Participants were asked to choose a fifth player for an endgame situation in which the a) the other team had the ball, b) the subject's team was ahead by one point and c) 2 minutes were left to play. The shot clock was noted as being set at 24 seconds.

To avoid confounding decision structure with a particular search matrix, half of the subjects received one search matrix for their high-structure decision and half the subjects received the other.

Dependent Variables

A number of dependent variables will be examined in this study. These variables are **decision accuracy, cue latency, search depth, search variability** and **search pattern**.

The decision accuracy criterion was operationalized by comparing the player alternative chosen by subjects with the "correct" choice. The "correct" choice was constructed in each solution sub-matrix so that it would be obvious to anyone with moderate basketball knowledge and access to the entire search matrix. Cues in the contextual sub-matrix indicated that the participant's team was shooting well but was sorely in need of a player who could rebound. Given the

general situation, participants needed to access information in the contextual sub-matrix to determine what the problem is before the correct choice can be made. For further information about the decision accuracy criterion, see Appendix C for a rational justification of how the correct choice was arrived at in each of the decision tasks.

Cue latency was operationalized by recording the overall amount of time subjects spend looking at information cues making a single decision. Search depth was operationalized as the number of times a subject accesses an item of information in the either of the two sub-matrices during a given decision task. An information retrieval was counted each time a cue was acquired, even if redundant.

Search variability was operationalized according to Payne's (1976) conceptualization. Using this method, search variability is defined as the standard deviation of the number of items accessed by the subject for each alternative. The resulting value, ranging from zero upwards, reflects the degree to which a subject gathers information in a compensatory fashion. A value of zero indicates that a subject accessed the same number of items for each alternative and this is seen as a strong indication of linear processing. A high value on search variability reflects unequal amounts of information accessed for the different alternatives and has been viewed as an indication of noncompensatory processing of information.

Search pattern refers to the relative amount of interdimensional v. intradimensional information gathering. This measure was operationalized by counting the number of interdimensional transitions made by each subject, subtracting the number of intradimensional transitions made by the subject and dividing this number by the total number of transitions made by the subject. An interdimensional transition occurs when a subject accesses a second item of information pertaining to the same alternative as the previous item of information. An intradimensional transition occurs when a subject accesses a second item of information belonging to the same attribute dimension as the last item of information. The resulting value has a range from -1.00 to 1.00 and reflects the relative tendency for a subject to gather information by searching along alternatives (in this case, by persons) or by attributes (in this case, performance information). A value of 1.0 reflects entirely interdimensional search while a value of -1.0 reflects information acquisition that is entirely intradimensional.

Procedure

A total of 278 undergraduate students responded to a short questionnaire designed to assess their knowledge of the game of basketball, as well as perceptions of importance for various basketball attributes (e.g., offensive rebounds) at various basketball positions (e.g., point guard). Appendix D contains the instrument used measure the

importance perceptions. Arrangements were made to bring back 62 of the highest-scoring students on the test and 61 students scoring in the range slightly above chance on the test (indicating that they had some familiarity with basketball).

The 123 students that were called back participated in the decision-making portion of this study. Each student was randomly assigned to either the "Alternative labels" condition or the "No Alternative labels" condition of the Alternative Labels factor. Following this, each subject was seated in front of an IBM-compatible XT-8088 class personal computer in a small room from which they completed the remainder of the study.

After subjects were seated in isolation, the experimenter booted up a floppy disk containing a decision task program appropriate to the participant's alternative labelling condition (i.e., with position titles or without). Each disk contained, in addition to the two decision tasks, step-by-step instructions for how to access information from the search matrix and how to make a decision for each task (i.e., choose an alternative). Participants were allowed to access as much information as they desired in each matrix in order to decide which player to choose for the upcoming game situation and were also provided with blank sheets of paper with which to take notes during the experiment.

Upon seating, the experimenter gave a brief one-minute overview of what was about to take place. All instructions

needed to complete the decision task were contained within the computer program, so this verbal summary was intended to orient participants and focus their attention. The experimenter also gave each participant a slip of paper at this time which contained a summary of the game situation to be encountered in the first decision task as well as a few other simple reminders. After this, each participant was instructed to hit any key on the keyboard to begin the study. All subsequent information/instructions were given via the computer.

Participants first made a practice decision using a small search matrix in order to familiarize themselves with the process of accessing cues and choosing an alternative. The practice task involved choosing the best basketball player from among four alternatives. The search matrix gave information about each attribute on four dimensions. None of these dimensions was used in the actual decision tasks to avoid biasing the search process on the basis of learning in the practice matrix. When the practice decision task was completed, each participant was instructed by the computer to notify the experimenter that she/he was ready to begin the actual experiment. The task in each of the decision tasks was to choose a player from the solution sub-matrix to complete the team. All search data was stored on floppy disks which were then collected by the experimenter.

Each participant was allowed to work through the two decision tasks at his/her own pace. The order in which each

participant received her/his decision tasks was determined randomly except for the last few participants who were assigned ordering to allow for even numbers of both orderings. In each condition, half of the subjects got the well-structured (ahead, 17 seconds left) task first and half received the ill-structured (behind, 2 minutes left) task first. The experimenter remained in an adjacent room in order to answer any questions and trouble-shoot any computer-related problems. When the first task was finished, the participant informed the experimenter who re-entered the participant's room to load the second task and gave the participant a new summary sheet. Upon finishing the second decision task, subjects were given a short questionnaire asking them to describe the way they went about acquiring information and served also as a check on how well the participant had understood the instructions. (See Appendix E for this post-experimental questionnaire). Finally, a debriefing sheet explaining the general goals of the research was given to each participant.

2

a

t

a

i

d

g

u

T

a

a

r

v

d

n

a

T

e

m

f

RESULTS

Overview

The Results section is divided into three main parts: an examination of the properties of the knowledge test used to assess domain-related expertise in basketball, an examination of the manipulation checks used to assess the influence of the two method factors, Order (order in which decision tasks were received) and Matrix (cue values for the given decision task) and a reporting of the main analyses used to test the hypotheses specified previously in Table 1. The main analyses are broken down into two method-based approaches corresponding to the measurement scale of the dependent variables (categorical versus ratio) and their respective statistical tests (categorical linear modelling versus multivariate analysis of variance).

The approach used to analyze the hypothesis about decision accuracy involves the construction of a linear model from repeated-measures categorical data and was accomplished through the use of the SAS CATMOD procedure. This procedure generates a parameter estimate for each effect specified in the design model on the basis of membership in a population defined by the between-subjects factors in the design (SAS Institute, 1985). Estimated

parameters are used to create a linear model of the response probabilities observed in each cell of the design. The appropriate test of significance for a given effect is the Chi-square value associated with its parameter estimate. These Chi-square values can be interpreted in a fashion analogous to that of the F test when using ANOVA.

For hypotheses involving interval-level (or better) dependent measures, a repeated-measures Multivariate Analysis of Variance was used (MANOVA). A multivariate procedure was used in this study as a result of the covariation among the dependent variables, average $\underline{r} = .225$ ($\underline{n} = 91$). Table 2 contains means and standard deviations for all variables in the study and Table 3 reports variable intercorrelations.

Multivariate Analysis of Variance procedures (MANOVA) accounts for the intercorrelation among dependent variables when testing hypotheses. For dependent variables which are highly correlated, multivariate ANOVA procedures yield F tests for each effect calculated over the ~~set~~ of dependent variables. Only effects which are significant in the Multivariate analysis or explicitly hypothesized should be examined in the subsequent univariate tests (Cole & Grizzle, 1966).

Before the hypotheses can be meaningfully tested, it is first necessary to evaluate the psychometric properties of the knowledge test used to assess domain knowledge in basketball. We now turn to a discussion of this process.

Table 2

Means and standard deviations of study variables

<u>Variable</u>	<u>Mean</u>	<u>SD</u>
1) Contextual search, well-struct. ¹	11.89 cues	15.54
2) Contextual search, ill-struct. ²	15.49 cues	18.89
3) Choice search, well-struct.	23.05 cues	19.13
4) Choice search, ill-struct.	27.45 cues	19.82
5) Cue latency, well-struct.	3.59 secs	2.00
6) Cue latency, ill-struct.	3.71 secs	2.06
7) Search pattern, well-struct.	0.10	.72
8) Search pattern, ill-struct.	0.14	.71
9) Search variability, well-struct.	2.72 cs/alt	2.07
10) Search variability, ill-struct.	3.24 cs/alt	2.23
11) Choice search var., well-struct.	1.57 cs/alt	1.51
12) Choice search var., ill-struct.	1.99 cs/alt	1.97
13) Decision accuracy, well-struct.	0.45	.50
14) Decision accuracy, ill-struct.	0.43	.50

¹"Well-struct." refers to measurement of the variable in the well-structured decision task

²"Ill-struct." refers to measurement of the variable in the ill-structured decision task

Table 3

Variable intercorrelations

<u>Var</u>	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	<u>6</u>	<u>7</u>	<u>8</u>	<u>9</u>
1	100	64	59	41	04	-09	14	14	09
2		100	37	39	-02	08	07	11	10
3			100	69	04	-05	27	21	73
4				100	-12	-07	16	24	55
5					100	63	02	-02	-09
6						100	-02	05	-08
7							100	54	31
8								100	27
9									100
10									

Note: correlations above .19 are significant at $p = .05$
 correlations above .23 are significant at $p = .01$

Table 3 (cont'd)

Variable intercorrelations

<u>Var</u>	<u>10</u>	<u>11</u>	<u>12</u>	<u>13</u>	<u>14</u>
1	05	32	24	-12	11
2	03	33	21	15	17
3	45	61	32	-06	14
4	67	46	48	01	26
5	-17	-23	-20	-07	00
6	-10	-16	15	13	16
7	16	27	12	-13	09
8	30	22	20	07	10
9	65	70	37	00	09
10	100	39	68	-03	16
11		100	49	08	04
12			100	01	08
13				100	-06
14					100

Note: correlations above .19 are significant at $p = .05$
 correlations above .23 are significant at $p = .01$

1.

2.

3.

4.

5.

6.

7.

8.

9.

10.

11.

12.

13.

14.

15.

16.

17.

18.

19.

20.

21.

22.

23.

24.

25.

26.

Assessment of Expertise

As noted above, the basketball knowledge test was designed to tap both declarative and procedural knowledge in the domain of basketball. Items intended to tap declarative knowledge concerned with knowledge of the rules used in basketball. As such, "objectively" correct answers exist and there is no need to assess reliability. This is not true for procedural knowledge, which had to do with "how" basketball is played and won, so an effort was made to assess agreement between the scoring key and two subject-matter experts.

Concerning the procedural items, disagreement between the scoring key and the two members of the coaching staff occurred on only six of the 23 items in the procedural scale. Of these items, three items garnered different responses from the three different judging sources and were dropped. For the other three items, the "correct" response was agreed upon by two of the three sources so these items were kept with the "agreed-upon" response judged the correct answer. Thus, there was fairly good agreement concerning the "correct" responses to the procedural knowledge items.

Viewing the test as a measure of basketball knowledge, it appears to have good psychometric properties. The KR-20 estimate of internal consistency was 0.92 and the standard error of measurement was correspondingly low as a result, $SEM = 2.70$. When the worst four items were deleted and the remaining 40 entered as a block into a regression equation

designed to predict an additive composite of a six-item Basketball Experience measure (see Appendix F), the resulting multiple correlation was very high, $R = .84$, with $R\text{-square} = .71$ and Adjusted $R\text{-square} = .66$. These values indicate that a great deal of the variance in one's reported lifetime involvement with basketball can be accounted for by performance on the basketball knowledge test.

In summary, it appears that the 40-item measure of basketball knowledge "taps" a single construct (basketball knowledge) and measures it well. In addition, the overall test score is highly related to a composite measure of individual experience throughout life with the game of basketball, providing some preliminary evidence of convergent validity. Finally, the low standard error of measurement provides evidence to support to the conclusion that individuals categorized as "experts" did indeed have true knowledge scores significantly higher than those individuals categorized as "novices." When a 90% confidence interval was created around each individual's score by multiplying the standard error of measurement by 1.65, even the highest-scoring novices (20) and lowest-scoring experts (29) did not have overlapping confidence intervals (24.45 = upper limit of novice distribution, 24.55 = lower end of expert distribution). Therefore, we can be relatively confident that experts did indeed know more than novices.

Manipulation Checks

Before discussing the results of the tests of the hypotheses, it is necessary to examine the impact of the manipulations used in the study and assess the influence of two method variables included in the design, the order in which decision situations were received and (Order) and the particular cue values in each of the decision situations (Matrix).

The Alternative Labels manipulation was expected to yield strong associations between basketball positions and certain performance attributes (statistics) in the minds of experts. Novice associations were expected to be weaker and less differentiated across position. This assumption was tested by asking both experts and novices to rate the degree to which the following 10 attributes were important (1= Very unimportant to 5 = Very important) to both the point guard and center positions in basketball: points scored, field goal percentage, free throw percentage, three-point field goal percentage, steals, assists, turnovers, blocked shots, offensive rebounds, and defensive rebounds.

Comparisons (t -tests) were then conducted for the perceived importance of each of the 10 statistics as they related to both point guard and center. Table 4 presents the results of this analysis. For the point guard position, mean ratings of importance for experts were significantly different than novices on seven of the 10 statistics at $p = .05$ and for each of the seven statistics, the experts rated

Table 4

T-tests, manipulation check for Alternative Labels**Point Guard**

<u>Variable</u>	<u>Nov.</u>	<u>Ex.</u>	<u>df</u>	<u>t</u>	<u>p</u>
Points scored	3.49	3.14	178	1.85	.066
Field Goal %	3.45	3.66	168	-1.14	.260
3-pt FG %	3.27	3.43	174	-0.93	.355
Steals	3.44	4.30	177	-5.28	.000
Assists	3.65	4.87	173	-9.24	.000
Off. Rebounds	2.52	1.44	172	6.68	.000
Def. Rebounds	2.98	1.58	172	8.19	.000
Blocked Shots	2.86	1.31	175	8.28	.000
Turnovers	3.86	4.83	171	-7.08	.000
Free-Throw %	3.54	4.51	171	-6.69	.000

Table 4 (cont'd).

T-tests, manipulation check for Alternative Labels**Center**

<u>Variable</u>	<u>Nov.</u>	<u>Ex.mean</u>	<u>df</u>	<u>t</u>	<u>p</u>
Points scored	3.91	4.10	179	-1.17	.243
Field Goal %	3.71	4.21	171	-3.24	.002
3-pt FG %	3.10	1.30	176	11.45	.000
Steals	3.10	1.81	176	7.85	.000
Assists	3.31	1.70	178	10.24	.000
Off. Rebounds	4.09	4.86	173	-5.80	.000
Def. Rebounds	3.83	4.97	176	-8.02	.000
Blocked Shots	3.81	4.69	175	-5.75	.000
Turnovers	3.39	3.58	171	-1.13	.262
Free-Throw %	3.76	4.14	174	-2.51	.014

the importance in a more extreme fashion, either very high or very low. Results were nearly identical for the center position, with mean ratings of importance for experts significantly different than novices for eight of the 10 performance attributes and experts once again more extreme in their perception of importance. The results of this manipulation check support the conclusion that, as presumed, experts have stronger associations between basketball position labels and performance dimensions than do novices.

The results of the manipulation check for Decision Structure similarly support the idea that information acquisition would be focused on fewer attributes for all subjects in the well-structured condition. Across all participants, the number of attributes listed as being weighed in the decision process was significantly greater in the ill-structured task, $t(113) = 7.47$, $p < .000$, ($M = 3.68$ cues in the ill-structured, 4.95 cues in the well-structured). Thus, it appears that most subjects did indeed perceive the need to consider more information in making their decisions for the ill-structured task. Table 5 displays the results of this manipulation check.

Method Factor Analyses

Having verified that the study manipulations were indeed having the effect desired, it was then necessary to examine the influence of two method factors inherent in the design. Order refers to the order in which individuals received their two decision tasks -- i.e., receiving the

Table 5

T-tests, manipulation check for Decision Structure

Means					
<u>t-test</u>	<u>Novice</u>	<u>Expert</u>	<u>df</u>	<u>t</u>	<u>p</u>
Well-struct Attributes	3.55	3.81	112	0.89	.374
Ill-struct Attributes	4.59	5.29	112	-1.69	.094
<u>t-test</u>	<u>Ill-</u>	<u>Well-</u>	<u>df</u>	<u>t</u>	<u>p</u>
Novices	3.55	4.59	055	-5.21	.000
Experts	3.81	5.29	057	-5.50	.000

well-structured situation first followed by the ill-structured situation, or vice versa. Matrix refers to the particular information cues received for each particular decision task. Both factors were between-subjects.

The Order factor was composed of two levels -- those individuals who received the well-structured problem first followed by the ill-structured problem and those individuals who received the ill-structured task first followed by the well-structured task. The Matrix factor was also composed of two levels -- those individuals who received matrix B for their well-structured task and matrix A for their ill-structured task (Diskset 1) and those individuals who received matrix A for the well-structured task and matrix B for their ill-structured task (Diskset 2).

The order in which participants received the two decisions was determined randomly, with the result that roughly half of the participants received the well-structured task first and half received it second. Levels of the Matrix variable were not randomly assigned due to the complexity of the design and the anticipated non-effect of the factor. However, both method variables were examined for their effects prior to conducting the main analyses.

The null hypothesis used in the preliminary method analyses was that the various effects for both Order and Matrix would be zero. If this is the case, it is safe to ignore these factors in further analyses and treat variance associated with the two factors as additional error

variance. If either factor is found to have one or more significant effects in the method analyses, it is necessary to include that factor in subsequent analyses.

For the dependent variable Decision Accuracy, a SAS repeated-measures linear model estimation procedure for categorical data was conducted involving the two method variables Order and Matrix, along with the repeated measure that was confounded with the two, Decision Structure. For the ratio-level dependent measures (contextual search depth, solution search depth, search pattern, overall search variability, solution matrix search variability and item latency), a Multivariate Analysis of Variance (MANOVA) was conducted involving the two method factors and the confounded factor of interest, Decision Structure.

The results of both manipulation checks were consistent. In both analyses, Order was involved in a significant interaction with Decision Structure and thus was included in subsequent analyses whereas there were no significant effects found for Matrix. Table 6 displays the results of these two analyses. In the categorical analysis for Decision Accuracy, the chi-square test of the Order x Decision Structure parameter was found to be statistically significant in the linear model ($p < .02$). No other parameters in the method-factor saturated model was found to be significant. In the multivariate analysis of variance, the Order x Decision Structure effect was again the only

Table 6

Results of Method-factor analyses

<u>Effect</u>	<u>χ^2 Value</u>	<u>p value</u>	<u>F test</u>	<u>p value</u>
Order (O)	0.91	0.34	0.33	0.92
Diskset (S)	0.58	0.45	0.34	0.84
Decision	0.27	0.60	3.13	0.01
Structure** (DS)				
O x S	1.90	0.17	1.02	.41
O x DS**	5.84	0.02	4.17	.00
S x DS	0.56	0.46	1.81	.10
O x S x DS	2.90	0.09	0.62	.716

Note: significance at $p < .01$ indicated by ****

method effect to achieve significance, $F(6, 114) = 4.17$, $p < .01$).

The conclusion from these analyses is that the dependent measures in this study were not affected by the particular cue values given to the participants (Matrix) but were affected by the order in which they received their decision situations (Order). On the basis of the Order x Decision Structure factor in both the CATMOD and MANOVA method analyses, Order was included in the main CATMOD and MANOVA analyses for the tests of the specified hypotheses. As there was no significant effects involving Matrix, it was dropped from subsequent analyses.

Main Analyses

The overall CATMOD and MANOVA analyses were then conducted with four factors and their interaction terms: Domain Knowledge, Alternative Labels, Decision Structure and Order. For the main CATMOD analysis, Decision accuracy was the dependent variable. For the main MANOVA, six dependent variables were entered into the model for each of the two decision tasks: contextual matrix search depth, solution matrix search depth, cue latency, overall search variability, solution matrix search variability and search pattern. In all, there were 12 dependent variables for the MANOVA.

In the main CATMOD analysis, two effects were found to be significant at $p < .05$, Domain Knowledge and the Domain Knowledge x Decision Structure. In the overall MANOVA, the

three main effects (Domain Knowledge, Alternative Labels and Decision Structure) were found to be significant at $p < .01$, in addition to two higher-order interactions which were significant as well, the Order x Decision Structure interaction effect ($p < .01$) and the Domain Knowledge x Decision Structure x Order three-way interaction ($p = 0.01$). Table 7 contains the results of the overall CATMOD analysis while Table 8 contains the results for the overall multivariate ANOVA.

Following this, six univariate repeated-measures ANOVA analyses were conducted, one for each of the six pairs of dependent variables specified in the overall multivariate MANOVA (e.g., contextual search depth for the ill-structured task v. the well-structured task). Except for the specific hypotheses predicting Domain Knowledge x Alternative Labels interactions, only the five proceeding effects (three main plus two interaction) discussed above were entered into the model for the univariate tests.

Hypothesis 1: Decision Accuracy

Hypothesis 1 predicted that experts would be more accurate than novices in making their player choices, and that experts would benefit more from the presence of a well-structured decision. This hypothesis was tested using the SAS CATMOD procedure previously described. Four factors were specified in the model: Domain Knowledge, Alternative Labels, Decision Structure and Order. Table 7 shows the chi-square values and p values for the various effects in

Table 7

Results, Overall CATMOD analysis

<u>Effect</u>	<u>df</u>	<u>Chi-square</u>	<u>p value</u>
Intercept	1	379.37	0.0000**
Domain Know. (DK)	1	11.93	0.0006**
Alt. Labels (AL)	1	0.36	0.5459
Dec. Struc. (DS)	1	0.29	0.5926
Order (O)	1	0.60	0.4394
DK x AL	1	0.29	0.5887
DK x DS	1	21.83	0.0000**
DK x O	1	0.70	0.4030
AL x DS	1	1.54	0.2141
AL x O	1	2.57	0.1090
DS x O	1	5.49	0.0191*
DK x AL x DS	1	3.11	0.0776
DK x AL x O	1	0.00	0.9946
DK x DS x O	1	0.48	0.4881
AL x DS x O	1	0.00	0.9889
DK x AL x DS x O	1	2.68	0.1018

DK = Domain Knowledge
 AL = Alternative Labels
 DS = Decision Structure
 O = Order

Note: significance at $p < .05$ indicated by **

Note: significance at $p < .01$ indicated by ***

Table 8

Results, Overall MANOVA

<u>Effect</u>	<u>df</u>	<u>F test</u>	<u>p value</u>
Constant	6	167.90	0.000
Domain Know. (DK)	6	3.18	0.006
Alt. Labels (AL)	6	4.79	0.000
Dec. Struc. (DS)	6	3.33	0.005
Order (O)	6	0.46	0.838
DK x AL	6	1.62	0.150
DK x DS	6	1.38	0.229
DK x O	6	0.32	0.927
AL x DS	6	1.20	0.312
AL x O	6	1.88	0.091
DS x O	6	4.29	0.001
K x L x S	6	0.69	0.660
K x L x O	6	0.70	0.650
K x S x O	6	2.97	0.010
L x S x O	6	0.56	0.769
K x L x S x O	6	1.62	0.147

DK = Domain Knowledge
AL = Alternative Labels
DS = Decision Structure
O = Order

Note: significance at $p < .05$ indicated by ***
Note: significance at $p < .01$ indicated by ****
Note: Df (error term) = 114 for all tests

the saturated model. The main effect of Domain Knowledge was highly significant, $X^2 (1) = 11.93$, $p = .0006$, as was the Domain Knowledge x Decision Structure interaction, $X^2 (1) = 21.83$, $p < .0000$. In addition to these two predicted effects, an Order x Decision Structure interaction was found to be significant as well, $X^2 (1) = 5.49$, $p = .0191$.

An analysis of the marginal response probabilities showed that, overall, experts made the correct choice in 53.2% of the decision situations, while novices were correct in 33.6% of their decisions. The nature of the finding that experts were more accurate than novices over all decisions, however, is qualified by the Domain Knowledge x Decision Structure.

Figure 2 shows that experts were relatively more accurate in the well-structured decision while novices were relatively more accurate in the ill-structured decision. As predicted, experts were much more likely to make the correct choice in the well-structured task, $p = 67.7\%$, as opposed to the ill-structured task, $p = 45.9\%$. On the other hand, novices showed just the opposite pattern, being relatively more accurate in the ill-structured task, $p = 38.7\%$ and less accurate on the well-structured task, $p = 21.3\%$. In the ill-structured decision, experts and novices showed similar probabilities of choosing the right alternative (46% v. 39%, respectively) but experts were more than three times more likely to get the well-structured problem correct than were novices (68% v. 21%), once again consistent with ACT*.

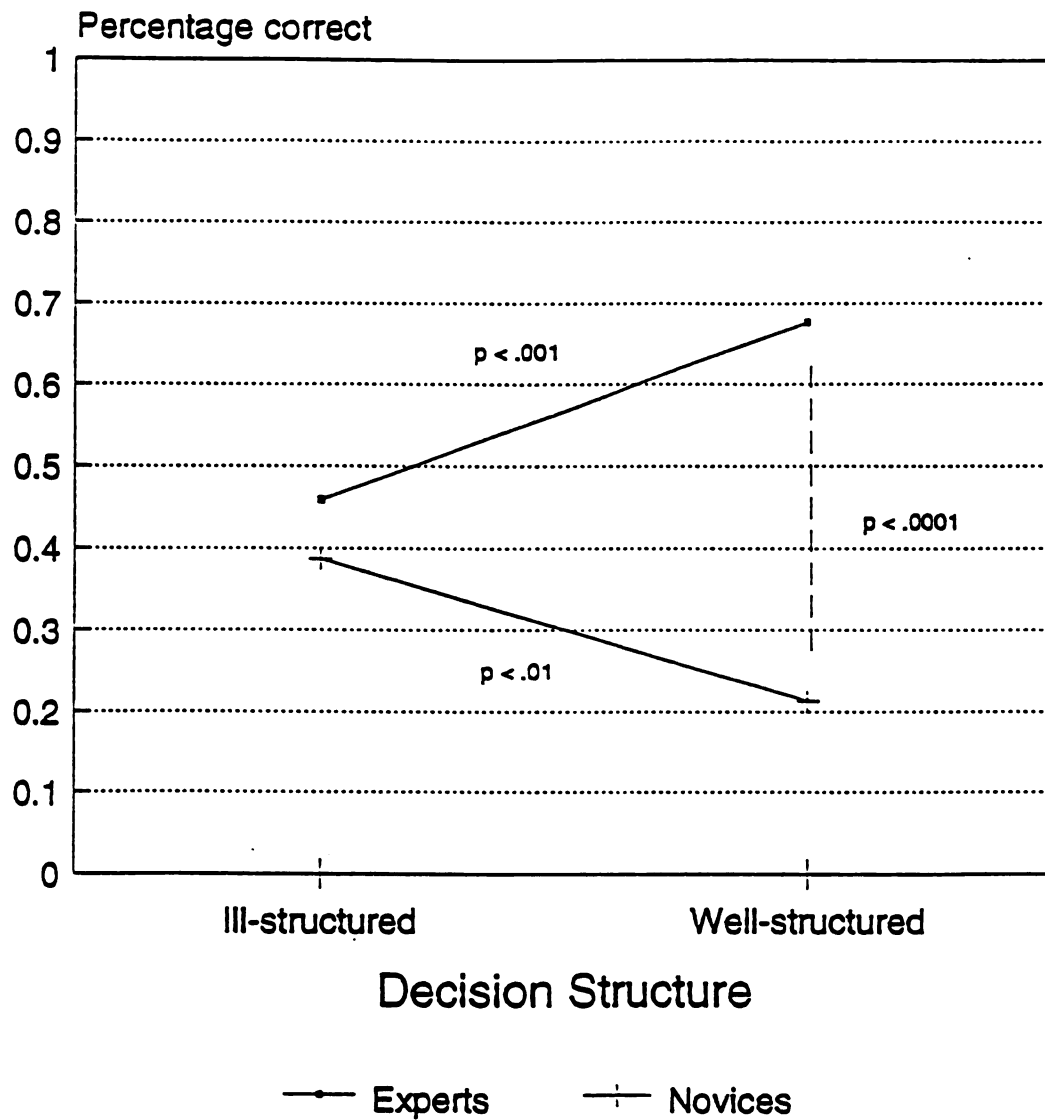


Figure 2. Domain Knowledge x Decision Structure interaction for Decision Accuracy

Decomposition of this interaction revealed that experts were significantly more accurate than novices in the well-structured condition, $\chi^2 (1) = 42.71$, $p < .0001$ (68% versus 21%) and significantly more accurate than their own performance in the ill-structured decision, $\chi^2 (1) = 12.10$, $p = .0005$ (68% versus 46%). Novices, surprisingly, were significantly **less** accurate in their well-structured decisions than they were in their ill-structured decisions, $\chi^2 (1) = 8.24$, $p = .0041$, (21% versus 39%). With respect to Hypotheses 1 and 1a, strong support was found for the prediction that experts would be more accurate than novices overall and particularly so in well-structured decisions.

Figure 3 shows marginal response probability graphed as a function of which decision task was undertaken first and decision structure. This figure shows that **all** individuals tended to be more accurate in making their second decision than in their first, regardless of the structure of the particular decision task. Individuals who received the ill-structured task first got it right 37.7% of the time as a group and then went on to improve to 54.1% in the well-structured task. Individuals who received the well-structured task first made the correct choice 35.5% of the time and then improved to 46.8% in the following ill-structured task.

Across participants in the ill-structured decision, the difference in accuracy for those who got it first (38%) and those who got it second (47%) was not significant,

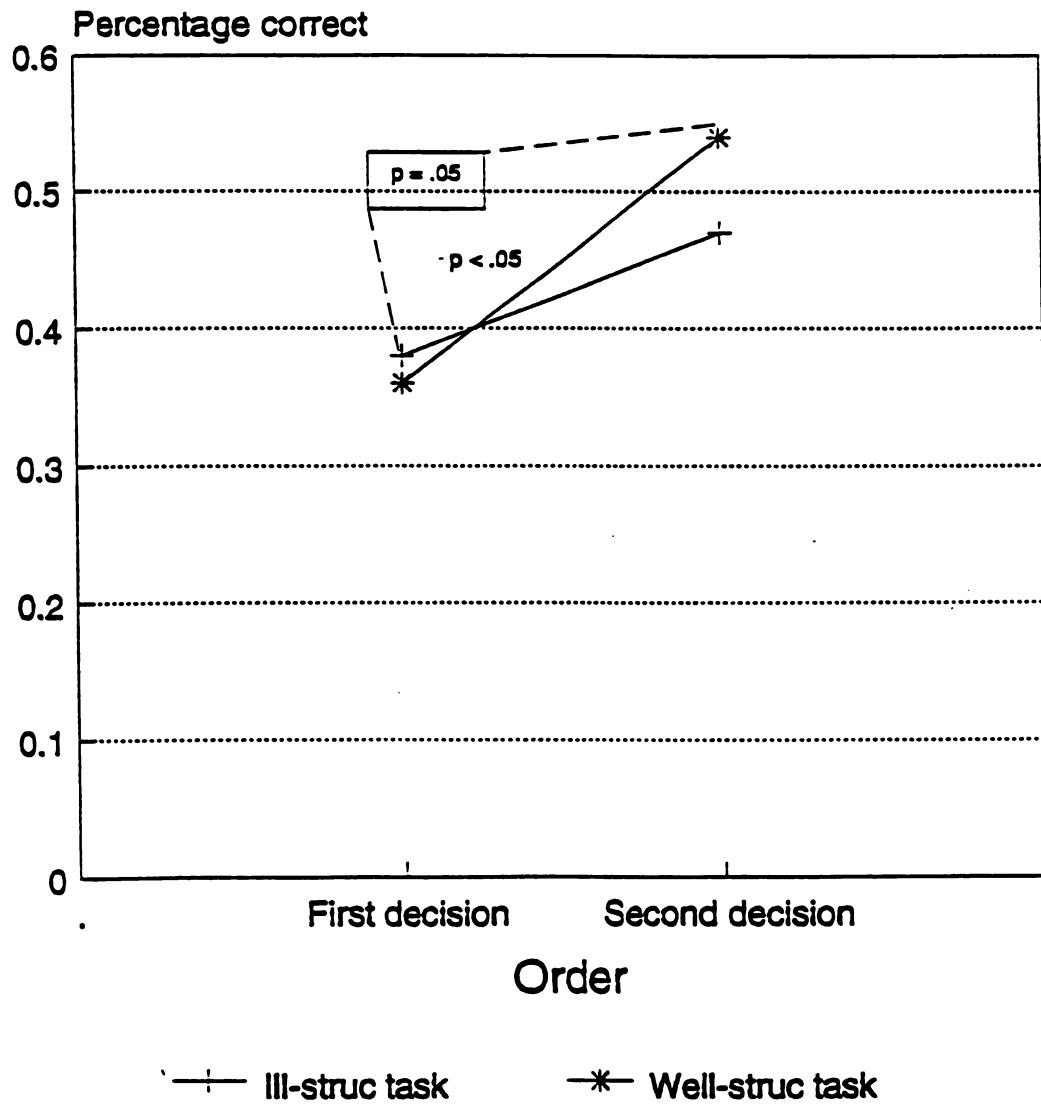


Figure 3. Decision Structure x Order interaction for Decision Accuracy

$\chi^2 (1) = 1.25$, n.s. In the well-structured decision task, the difference in accuracy was significant, $\chi^2 (1) = 5.23$, $p = .0223$ (36% versus 54%, for receiving the task first and second, respectively). Within participants, the difference in accuracy when moving from the well-structured decision first to the ill-structured decision (36% versus 47%) was not significant, $\chi^2 (1) = 1.43$, n.s., while for those getting the ill-structured first and then the well-structured (38% versus 54%) the difference in probability of a correct choice was significant, $\chi^2 (1) = 3.79$, $p = .05$.

These results appear to support both Hypotheses 1 and 1a. Consistent with Hypothesis 1, experts across all situations were significantly more accurate than novices (53% versus 34%) and, further, experts were quite a bit more accurate than novices in the well-structured situation (68% versus 21%), as predicted by Hypothesis 1a. Interestingly enough, experts are more accurate in the well-structured task (relative to the ill-structured task) while novices are actually less accurate. The finding that all individuals tended to be more accurate in their second decision task regardless of the structure of that task was independent of the other effects. Thus, support for Hypotheses 1 and 1a appears relatively sound.

Table 9 provides descriptive information concerning the percentage breakdown of choices in the two decisions as a function of Domain Knowledge, Alternative Labels, and Matrix. The Matrix variable is shown in this table because

Table 9

Decision choices across all conditions**Ill-Structured Decision**

<u>Domain Knowledge</u>	<u>Label</u>	<u>Set</u>	Player #			
			<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>
Novice	Yes	1	58.3	08.3	25.0	08.3
Novice	Yes	2	00.0	05.3	31.6	63.2
Novice	No	1	26.3	05.3	47.4	21.0
Novice	No	2	09.1	36.4	18.2	36.4
Expert	Yes	1	14.3	07.1	42.9	35.7
Expert	Yes	2	06.3	37.5	31.3	25.0
Expert	No	1	13.6	13.6	50.0	22.7
Expert	No	2	10.0	10.0	40.0	40.0

Note: "Set" refers to which Diskset was received

Table 9 (cont'd).

Well-Structured Decision						
			Player #			
<u>Domain Knowledge</u>	<u>Label</u>	<u>Set</u>	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>
Novice	Yes	1	16.7	00.0	16.7	66.7
Novice	Yes	2	21.1	21.1	52.6	05.3
Novice	No	1	26.3	15.8	21.1	36.8
Novice	No	2	00.0	18.2	45.5	36.4
Expert	Yes	1	78.6	07.1	07.1	07.1
Expert	Yes	2	00.0	81.3	06.3	12.5
Expert	No	1	40.9	04.5	09.1	45.5
Expert	No	2	00.0	90.0	00.0	10.0

Total Percentages (across all conditions):

	Player #			
	1	2	3	4
Novice	20.0	13.1	33.6	33.6
Expert	21.8	28.2	24.2	25.8

the two tasks had different correct answers as a function of the cue matrix used. For the Matrix level "Diskset 1," Player 3 (Center) is the correct choice for the ill-structured task and Player 1 (Point Guard) is the correct choice for the well-structured task. For Matrix condition "Diskset 2," Player 4 (Power Forward) is the correct choice for the ill-structured task and Player 2 (Point Guard) is the correct choice for the well-structured task.

An examination of Table 9 serves to highlight two interesting points: 1) Novices agreed slightly more often when alternative labels were present and 2) Experts agreed to a greater extent than novices in the well-structured decision but not in the ill-structured situation. In the ill-structured decision, the average percentage for the novices "top choice" across the four cells is 49% and 30% for the runner-up choice. For experts, these values are 43% and 33%, respectively. In the well-structured decision, these values are 51% and 25% for novices, 74% and 18% for experts, reflecting the much greater agreement among experts in the well-structured task. Interestingly enough, in the ill-structured decision task, in three of the four cells for both novices and experts, the alternative chosen most often was the a priori "best" choice. This is in stark contrast to the well-structured decision. Across the four cells, the "number one" choice of novices was never the "right" one, while experts again chose the correct alternative most often in three of the four cells.

When percentage choice for each alternative is averaged over decision tasks, matrices and labelling (see bottom of Table 9), the picture becomes clearer. Novices tended to choose alternatives 3 and 4 more often *regardless of the decision task* while experts spread their choices out to a greater extent over the range of alternatives. More will be said about this finding in the discussion.

Hypothesis 2: Cue latency

Hypothesis 2 predicts that experts will spend less time looking at each information cue than will novices. This hypothesis is an intuitive one, falling out of the traditional literature that has found experts to be faster than novices at solving problem in their domain of expertise. The appropriate test statistic for evaluating this hypothesis is the F test for the main effect of Domain Knowledge in the repeated-measures ANOVA for mean cue latency (see Table 10). The resulting value, $F(1, 115) = .257$, was not significant, indicating that experts did not differ significantly from novices in the amount of time that they took to look at cues across the two decision tasks. Thus, no support was found for Hypothesis 2.

However, an Order x Decision Structure interaction was found, $F(1, 115) = 15.09$, $p < .001$. Figure 4 displays this interaction. As can be seen in the figure, regardless of the structure of the decision, individuals tended to spend more time looking at cues in their first decision than in their second. Examining within individuals, for

Table 10

Univariate ANOVA, Cue Latency

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	9.02	1.30	.257
Alt. Labels	01	3.06	0.44	.508
Dom. Know x Alt. Labels	01	3.04	0.44	.509
Between-subjects Error	115	797.75	/ 6.94	
<hr/>				
Decision Structure	01	0.71	0.52	.472
Dec. Struc x Dom. Know	01	3.36	2.46	.119
Dec. Struc x Order	01	20.59	15.09	.000
Dec. Struc x Dom. Know. x Order	01	0.06	0.04	.833
Within-subjects Error	115	156.93	/ 1.36	

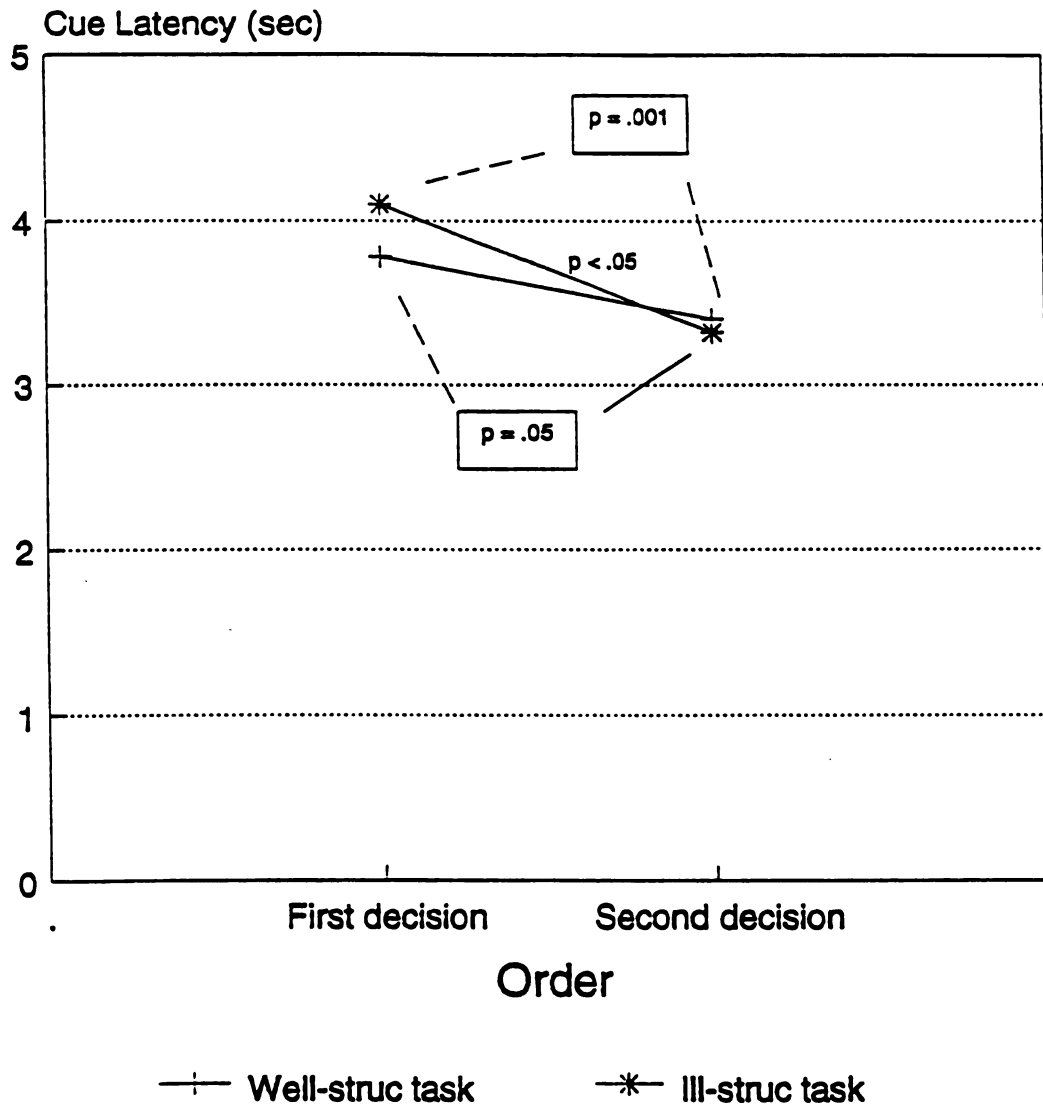


Figure 4. Order x Decision Structure interaction for Cue Latency.

participants who received the well-structured decision first, the mean number of seconds spent looking at each cue was 3.78 seconds in the well-structured task, 3.33 seconds in the following ill-structured decision. This difference was significant, $t(61) = 2.00$, $p = .05$. For individuals who received their decision tasks in the opposite order, 4.10 seconds was the mean latency for cues in the first (ill-structured) task and this value decreased to 3.40 seconds in the second (well-structured) task. This difference was also significant, $t(60) = 3.59$, $p = .001$. Looking across participants within the particular decisions themselves, the mean difference in item latency for individuals in the well-structured decision (3.78 seconds versus 3.40 seconds) was not significant, $F(1, 121) = 1.07$, n.s., while for the ill-structured decision, the mean difference in item latency was significant, $F(1, 121) = 4.40$, $p = .038$.

Thus, while there is no support for the hypothesis that experts would spend less time looking at each cue, the order in which decision tasks were received once again was found to affect the process of information acquisition.

Hypothesis 3: Total Search Depth

Hypotheses 3-5b concern the number of information cues accessed by an individual in the process of decision-making. On the basis of the voluminous expert-novice literature, it was hypothesized that generally experts would acquire fewer informational cues than novices when making decisions within

their domain of expertise. However, the production-systems framework highlights the importance of information provided to the expert by both the decision context (structure) and the labelling of alternatives. This view suggests these aspects of the decision environment will have informational value to experts but not to novices. Experts, then, will acquire fewer cues than novices when labels are present for decision alternatives and when a decision situation is well-structured, respectively. When these two phenomena are not present, experts are predicted to behave much like novices.

Hypothesis 3 is the general prediction that experts will search fewer cues than novices across the entire matrix of cue values in each decision. A repeated-measures univariate ANOVA was performed on the sum total of all cues accessed in each decision, collapsing over the contextual and choice sub-matrices. See Table 11 for the results of this analysis. The main effect for Domain Knowledge in this analysis was significant, $F(1,115) = 7.54$, ($p = .007$), but analysis of the cell means indicated that experts accessed more information than novices, not less as predicted. Experts, on average, acquired 13 more information cues in each task than did novices (45.13 cues versus 32.02 cues, respectively). Surprisingly then, no support was found for Hypothesis 3. Further analyses were conducted to address the more fine-grained "contextual" versus "choice" distinction in cue values included in each overall matrix

Table 11

Univariate ANOVA, Total Search Depth

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	11048.70	7.54	.007
Alt. Labels	01	12703.87	8.67	.004
Dom. Know x Alt. Labels	01	6205.48	4.23	.042
Between-subjects Error	115	168509.82		
<hr/>				
Decision Structure	01	3228.61	12.24	.001
Dec. Struc x Dom. Know	01	836.37	3.17	.078
Dec. Struc x Order	01	315.18	1.20	.277
Dec. Struc x Dom. Know. x Order	01	784.62	2.98	.087
Within-subjects Error	115	30325.59		

Hypothesis 4: Contextual Search Depth

Hypothesis 4 states that experts will access more cues than novices in the contextual search matrix. The rationale for this hypothesis stems from research findings which suggest that experts concentrate their information-processing efforts on setting up an internal representation of the decision/problem situation and attempting to identify information that will make the problem representation more complete. In addition, Hypothesis 4a predicts that the difference between expert and novice search depth in the contextual matrix will be less when contextual alternatives are labelled, as information provided to the expert by alternative labelling will help to fill in the problem representation. Similarly, Hypothesis 4b suggests that experts will benefit from a well-structured decision situation in which a complete and accessible task representation already exists in memory and so the difference between expert and novice acquisition in this condition should be less as well.

In general, Hypothesis 4 predicting experts would acquire more contextual information than novices was supported by the results of the repeated-measure univariate ANOVA for contextual search depth (see Table 12). The main effect for Domain Knowledge was significant, $F(1,115) = 15.28$, $p < .000$, with experts accessing over twice as many cues as novices in the contextual matrix ($M = 18.67$ cues for experts, 8.63 cues for novices). Indeed, the other two main

Table 12

Univariate ANOVA, Contextual Search Depth

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	6205.44	15.28	.000
Alt. Labels	01	3077.66	7.58	.007
Dom. Know x Alt. Labels	01	1903.14	4.69	.032
Between-subjects Error	115	46713.12 / 406.20		
<hr/>				
Decision Structure	01	785.47	7.33	.008
Dec. Struc x Dom. Know	01	487.66	4.55	.035
Dec. Struc x Order	01	617.06	5.76	.018
Dec. Struc x Dom. Know. x Order	01	113.65	1.06	.305
Within-subjects Error	115	12328.46 / 107.20		

effects were found to be significant as well but were qualified by higher-order interactions. The main effect of Decision Structure was significant, $F(1,115) = 7.33$, $p = .008$, as was the effect of Alternative Labels, $F(1,115) = 7.58$, $p = .007$. Analysis of the condition means indicated that more cues were acquired on average in the ill-structured situation than in the well-structured situation ($M = 15.49$ cues versus 11.89 cues, respectively), and more cues were accessed by individuals who did not receive labels for their alternatives than for those individuals who did ($M = 17.29$ cues without labels versus 10.02 cues with labels).

The appropriate statistic for evaluating Hypotheses 4a and 4b is the F statistic for the Domain Knowledge \times Alternative Labels (4a) and Domain Knowledge \times Decision Structure (4b) interaction effects in the repeated-measures univariate ANOVA for search depth in the contextual matrix. As Table 12 shows, both of these two interactions had a significant effect on the amount of information accessed in the contextual matrix, Domain Knowledge \times Alternative Labels, $F(1,115) = 4.69$, $p = .032$, Domain Knowledge \times Decision Structure, $F(1,115) = 4.55$, $p = .035$, as well as the now-familiar Order \times Decision Structure, $F(1,115) = 5.76$, $p = .018$. Figures 5-7 (respectively) decompose these three interactions.

As Figure 5 shows, novices were basically unaffected by Alternative Labels ($M = 7.98$ cues accessed with labelled alternatives, $M = 9.28$ cues accessed without). This

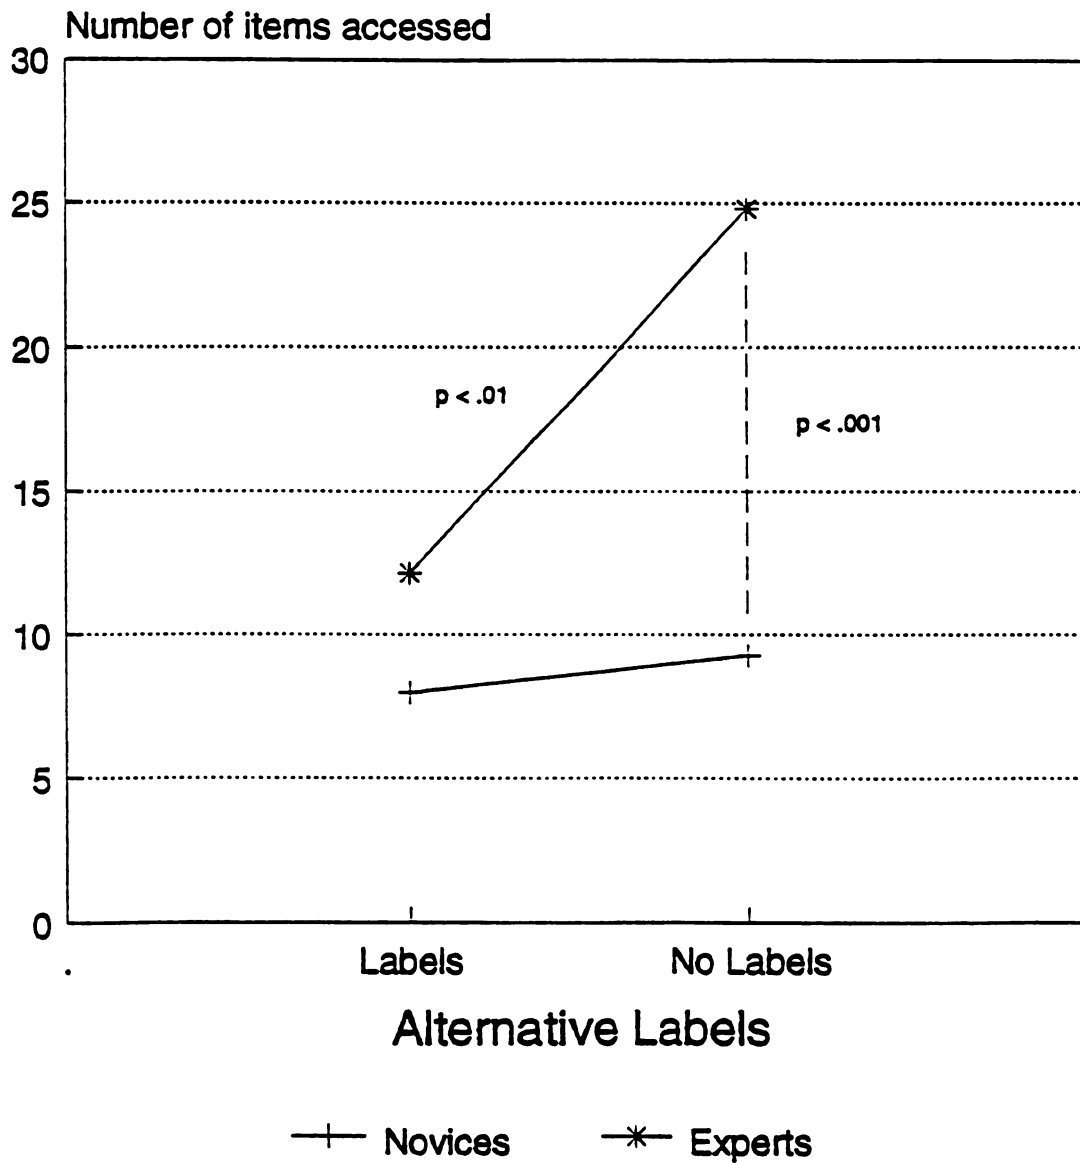


Figure 5. Domain Knowledge x Alternative Labels
interaction for Contextual Search Depth

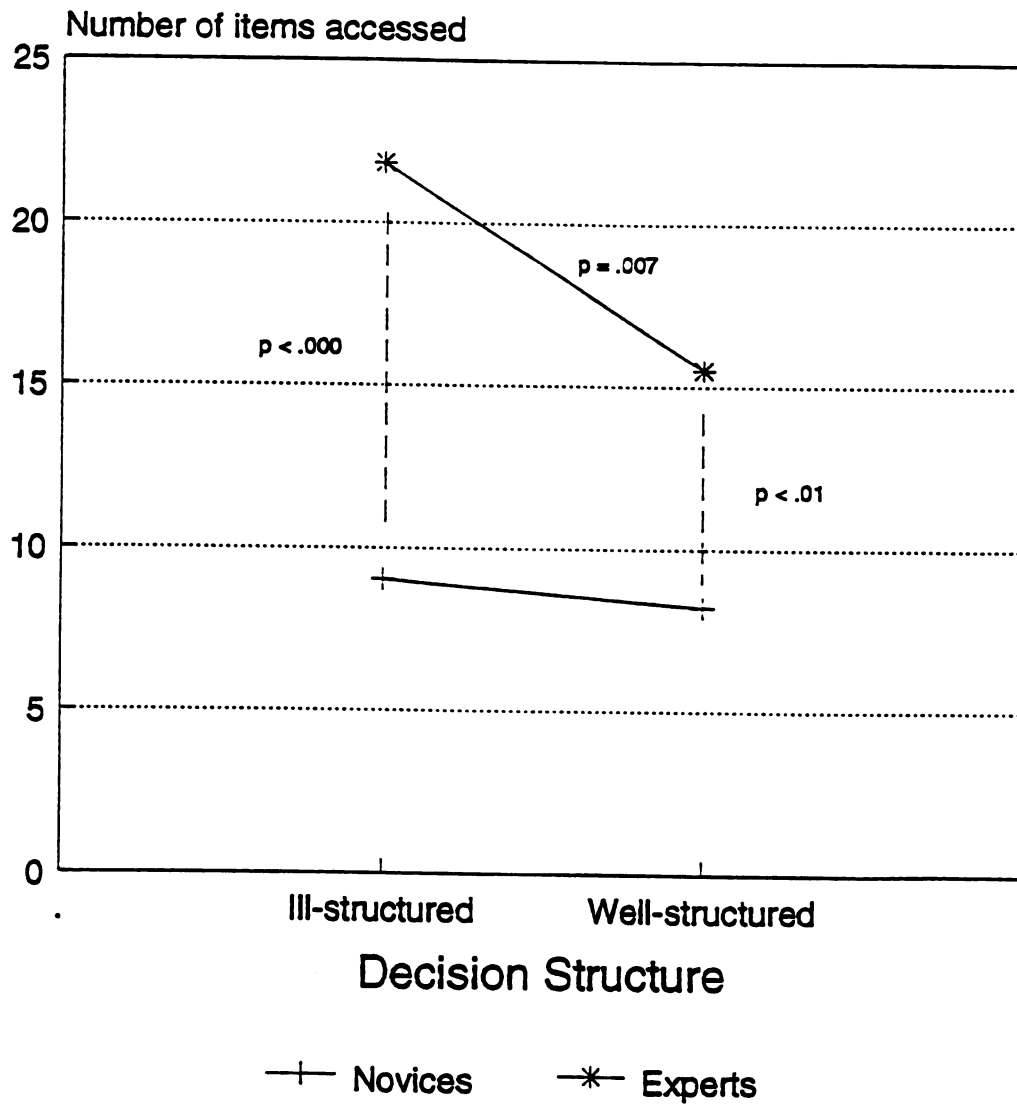


Figure 6. Domain Knowledge x Decision Structure interaction for Contextual Search Depth

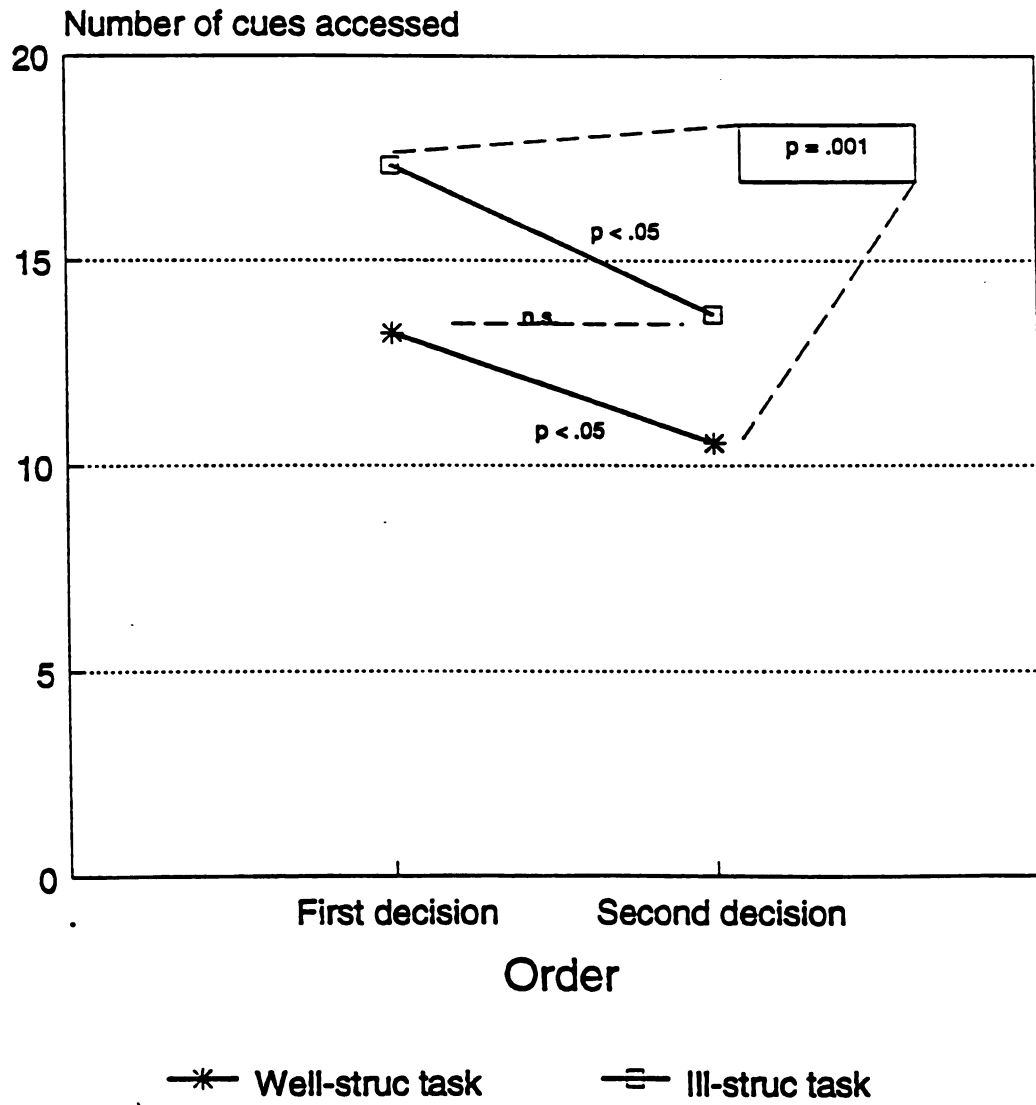


Figure 7. Order x Decision Structure interaction for Contextual Search Depth

difference was not significant, $t(59) = -0.40$, $p = .691$. Experts, on the other hand, were quite sensitive to the presence or absence of alternative labels. Experts who did not receive labelled alternatives searched for over twice as much information in the contextual matrix compared to experts with labelled alternatives ($M = 24.80$ cues without labels versus 12.13 cues with labels). This difference was significant, $t(60) = 3.22$, $p = .002$. In addition, expert search depth in the no label condition was significantly greater than novice search depth in the no label condition [$t(60) = 3.79$, $p < .000$], with experts accessing over twice as much information ($M = 24.8$ cues for experts, 9.28 cues for novices). Finally, expert search depth in the labelled condition was not significantly different from that of novices, $t(59) = -1.36$, $p = .179$ ($M = 12.13$ cues for experts, 7.98 cues for novices).

The entire interaction is thus accounted for by the significant difference of expert search depth in the no label condition compared to experts in the labelled condition and novices in the no label condition. With respect to Hypothesis 4a, experts accessed more cues than novices across all conditions, but distinctly benefitted from the presence of alternative labels which allowed them to reduce their search depth by 50% on average. Thus, there is good support for Hypothesis 4a.

Figure 6 shows an almost identical pattern of results in the Domain Knowledge x Decision Structure interaction for

search depth in the contextual matrix. Once again, experts were sensitive to the presence of a knowledge-dependent source of information while novices were not. Expert search depth in the ill-structured decision ($M = 21.84$ cues) was significantly greater than expert search depth in the well-structured decision ($M = 15.5$ cues), $t(61) = 2.80$, $p = .007$) as well as Novice search depth in ill-structured decision ($M = 9.03$ cues), $F(1,121) = 15.85$, $p < .000$). Also, expert search depth in the well-structured decision ($M = 15.5$ cues) was significantly greater than novice search depth in the well-structured decision ($M = 8.21$ cues), $F(1,121) = 7.10$, $p = .009$.

These analyses indicate that, while experts accessed more information than novices in both types of situations, the discrepancy is greatest in the ill-structured conditions with experts acquiring roughly 33% more cues. This finding is consistent with Hypothesis 4b which predicted that experts would be much more sensitive to the benefits of a well-structured decision situation than will novices.

Figure 7 displays the Decision Structure interaction involving the Order method factor. As apparent in the figure, contextual search depth for individuals who received the well-structured task first was not affected by the particular structure of the decision in either the well-structured ($M = 13.21$ cues) or ill-structured tasks (13.68 cues). However, individuals who received the ill-structured task first ($M = 17.32$ cues) and the well-structured task

second ($M = 10.54$ cues) were affected by the structure of the task. This interaction of Order \times Decision Structure was significant, $F(1,115) = 5.76$, $p = .018$.

T-tests of the cell means for this interaction indicated that, for individuals receiving the ill-structured problem first, contextual search depth in their ill-structured task ($M = 17.33$ cues) was significantly greater than contextual search depth in their well-structured task ($M = 10.55$ cues), $t(60) = 3.46$, $p = .001$, and across participants, was significantly greater than the ill-structured contextual search depth for individuals who received the task second ($M = 13.68$ cues), $F(1, 118) = 5.255$, $p = .024$. A t-test of the means for individuals receiving the well-structured decision first indicated that contextual search in the well-structured decision ($M = 13.21$ cues) was not significantly different than their contextual search in the ill-structured decision ($M = 13.68$ cues), $t(61) = -0.26$, n.s, but was significantly different than the well-structured contextual search depth of individuals who received this task second ($M = 10.54$ cues), $F(1,118) = 6.77$, $p = .01$.

Hypothesis 5: Choice Search Depth

Hypothesis 5 was based on the assumption that experts would acquire more information in the contextual sub-matrix and would use that information to identify a narrow set of attributes needed in the choice alternative and a narrowed set of alternatives to examine, thus reducing their

information acquisition in the choice sub-matrix in comparison to novices. Table 13 contains the results of the univariate ANOVA for choice matrix search depth. Hypothesis 5 was tested by examining the significance of the main effect of Domain Knowledge in the choice matrix search depth univariate ANOVA, which was not significant, $F(1,115) = 1.65$, $p = .202$, indicating that experts did not differ significantly from novices in the number of cues they accessed among choice alternatives. At this point, there is no support for Hypothesis 5 and the notion that information acquisition for experts in the solution matrix will be more constrained than that of novices.

Not surprisingly, the two main effects for the other task factors were significant in the ANOVA for choice matrix search depth. For the Alternative Labels factor, individuals in the labelled conditions searched for significantly less information ($M = 21.96$ cues) than individuals in the no label conditions ($M = 28.49$ cues), $F(1,115) = 4.28$, $p = .041$. The main effect for Decision Structure, $F(1,115) = 9.96$, $p = .002$, resulted from the fact that, across all conditions, individuals searched for significantly less information in the well-structured task ($M = 23.05$ cues) than in the ill-structured task ($M = 27.45$ cues). However, this effect is heavily qualified by a higher-order interaction.

Hypothesis 5a is similar to Hypothesis 4a in predicting that expert choice matrix search depth will be reduced in

Table 13

Univariate ANOVA, Choice Search Depth

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	982.92	1.65	.202
Alt. Labels	01	2551.21	4.28	.041
Dom. Know x Alt. Labels	01	807.96	1.35	.247
Between-subjects Error	115	68620.82 / 596.70		
<hr/>				
Decision Structure	01	1142.98	9.96	.002
Dec. Struc x Dom. Know	01	140.43	1.22	.271
Dec. Struc x Order	01	4.30	0.04	.847
Dec. Struc x Dom. Know. x Order	01	500.15	4.36	.039
Within-subjects Error	115	13196.14 / 114.75		

the presence of labelled alternatives. The appropriate statistic for evaluating this hypothesis is the F test for the Domain Knowledge x Alternative Labels test in the repeated-measures univariate ANOVA. Contrary to prediction, expert search in the choice matrix was not significantly affected by the presence of labels, $F(1, 1115) = 1.35$, $p = .247$. Thus, there was no support found for Hypothesis 5a.

Hypothesis 5b is similar to Hypothesis 4b in suggesting that a well-structured decision will benefit experts in that the need to acquire cues is reduced by the structuring of the task. Hypothesis 5b specifies that experts will access fewer cues relative to novices when the decision problem is well-structured but not when it is ill-structured. The appropriate test of this hypothesis is the significance of the F test for the Domain Knowledge x Decision Structure effect in the repeated-measures univariate ANOVA for choice matrix search depth. This effect was not significant, $F(1, 115) = 1.22$, $p = .271$, but the Domain Knowledge x Decision Structure effect was included in a significant three-way interaction involving Domain Knowledge, Decision Structure and Order, $F(1, 115) = 4.36$, $p = .039$. (See Figures 8 and 9).

As suggested by the figures, the locus of the three-way interaction is the significance of the Domain Knowledge x Decision Structure interaction for individuals who received the ill-structured task first, $F(1, 59) = 7.13$, $p = .010$, and the non-significance of the Domain Knowledge x Decision

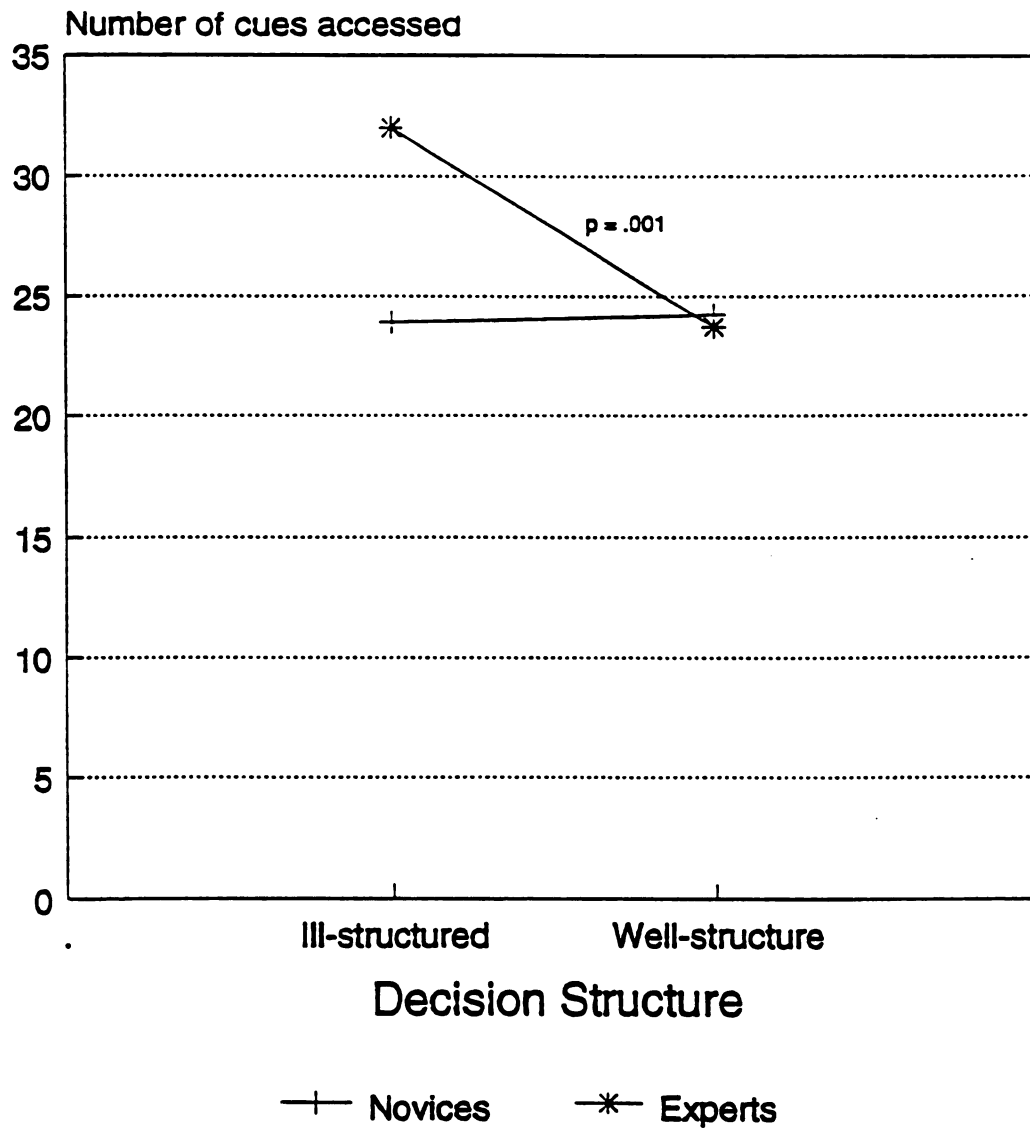


Figure 8. Domain Knowledge x Decision Structure interaction for Choice Search Depth with "ill-structured task first" ordering

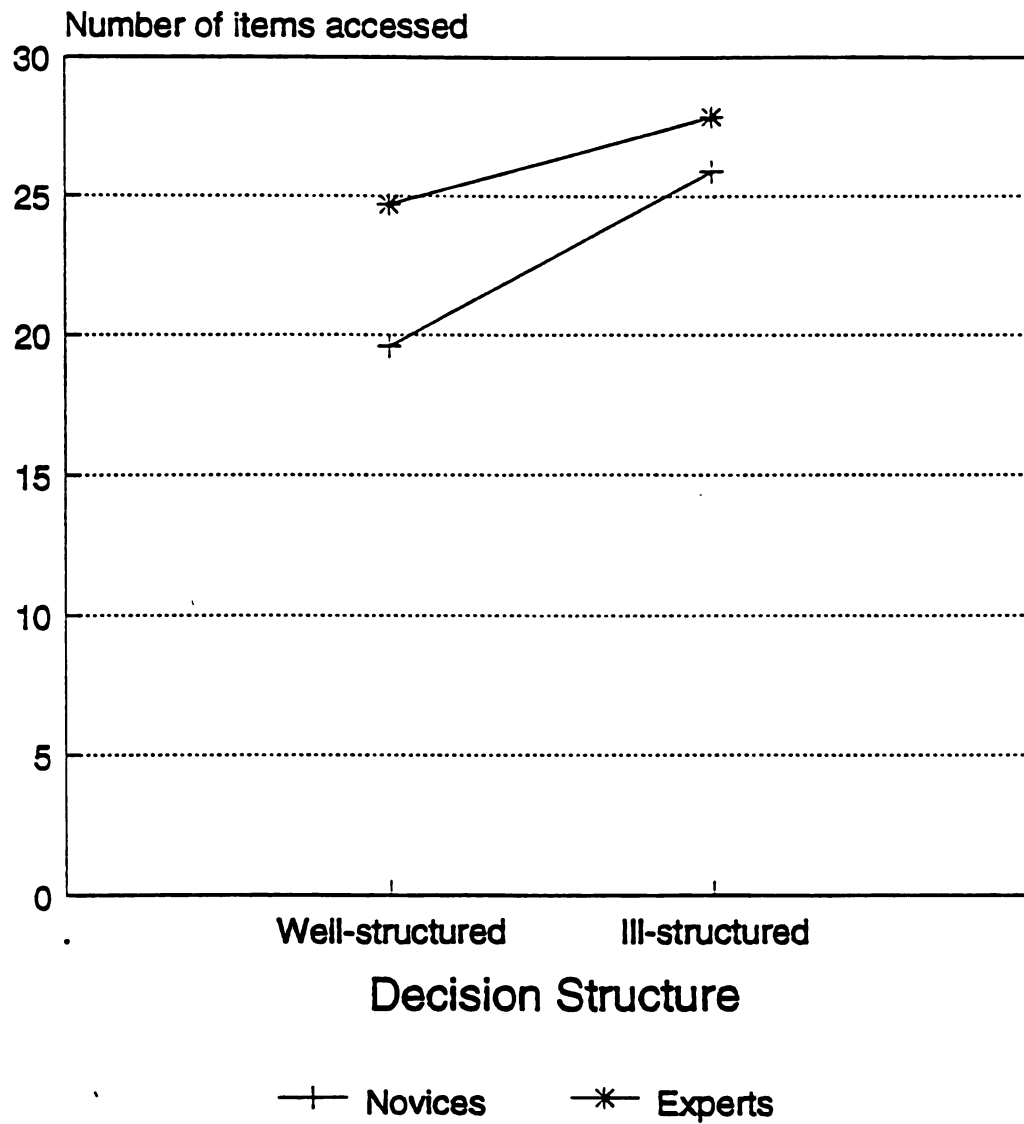


Figure 9. Domain Knowledge x Decision Structure interaction for Choice Search Depth with "well-structured task first" ordering

Structure interaction for individuals who received it second, $F(1,60) = .51$, $p = .478$. Thus, the predicted Domain Knowledge x Decision Structure interaction does occur in situations in which individuals receive the ill-structured task first, but not for individuals who received the well-structured task first.

Decomposition of the significant two-way interaction for the "ill-structured task first" ordering indicated that expert search depth in the ill-structured decision task ($M = 32.00$ cues) was significantly greater than expert search depth in the well-structured task ($M = 23.71$ cues), $t(30) = 3.55$, $p = .001$. All other differences between cell means were non-significant. Thus, the three-way interaction occurred as a result of the sensitivity of experts to the benefits of the well-structured decision task when they received the ill-structured task first but not when they received this task second.

In sum, there appears to be some support for Hypotheses 5b. In the analysis of search depth in the choice sub-matrix, the Domain Knowledge x Decision Structure interaction was not significant, but the observed three-way interaction between Domain Knowledge, Decision Structure and Order indicates that the Domain Knowledge x Decision Structure interaction does exist in some situations. These results provide support for the notion that decision structure influences the amount of information concerning

choice alternatives accessed by experts to a greater degree than it does novices.

Hypothesis 6: Overall Search Variability

Hypothesis 6 specifies that experts will be more variable than novices in their information acquisition across all alternatives. The appropriate test for this hypothesis is the F statistic for the repeated-measures univariate ANOVA for the standard deviation of information cues accessed across all alternatives (see Table 14 for the results of this ANOVA). The result of this test was non-significant, $F(1,115) = 0.34$, $p = .559$, indicating that there is no support for the hypothesis that experts, across all conditions, would show greater variability in the number of information cues accessed across alternatives.

Hypothesis 6a qualifies Hypothesis 6 on the basis of the logic set forth in the ACT*-based model. According to the model presented in the introduction, alternative labels provide the expert with information concerning **which** alternatives to search to find the desired attribute levels, implying that not all alternatives will be searched equally. Novices would not be expected to be able to narrow search for desired attributes on the basis of alternative labels, or should fail if they try.

Hypothesis 6b predicts that experts should be able to narrow the set of important attributes needed in the choice alternative from the streamlined production invoked by the well-structured decision. Since there is little overall

Table 14

Univariate ANOVA, Overall Search Variability

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	25690.15	0.34	.559
Alt. Labels	01	103570.18	1.38	.242
Dom. Know x Alt. Labels	01	46099.63	0.61	.435
Between-subjects Error	115	8621486.82 / 74969.45		
<hr/>				
Decision Structure	01	157261.49	10.67	.001
Dec. Struc x Dom. Know	01	5591.51	0.38	.539
Dec. Struc x Order	01	25755.15	1.75	.189
Dec. Struc x Dom. Know. x Order	01	117796.80	7.99	.006
Within-subjects Error	115	1694789.01 / 14737.30		

information to be acquired in this situation, experts would then presumably search every possible choice alternative for the important attribute(s). In the ill-structured decision, evaluation of individuals was predicted to be more holistic and, as more information needs to be acquired, experts are predicted to gather information about a potential alternative until enough has been gathered to identify the role such an alternative is playing and evaluate how well the role is being filled. This should in turn lead to more variable search across alternatives, as search for alternatives is halted at arbitrary points when enough is known about them.

Thus Hypotheses 6a and 6b, derived from a production systems framework, suggest that experts will show greater variability in information access not in all situations but only in those situations where the decision environment interacts with their knowledge base to provide extra information that is not available to the novice. Hypothesis 6a was not supported, however, as the Domain Knowledge x Alternative Labels interaction effect was not significant, $F(1,115) = 0.61$, $p = .435$. Hypothesis 6b did receive partial support, however, in that the Domain Knowledge x Decision Structure x Order interaction was again significant, $F(1,115) = 7.99$, $p = .006$. Figures 10 and 11 display this three-way interaction.

Once again, the locus of the three-way interaction stems from the predicted Domain Knowledge x Decision

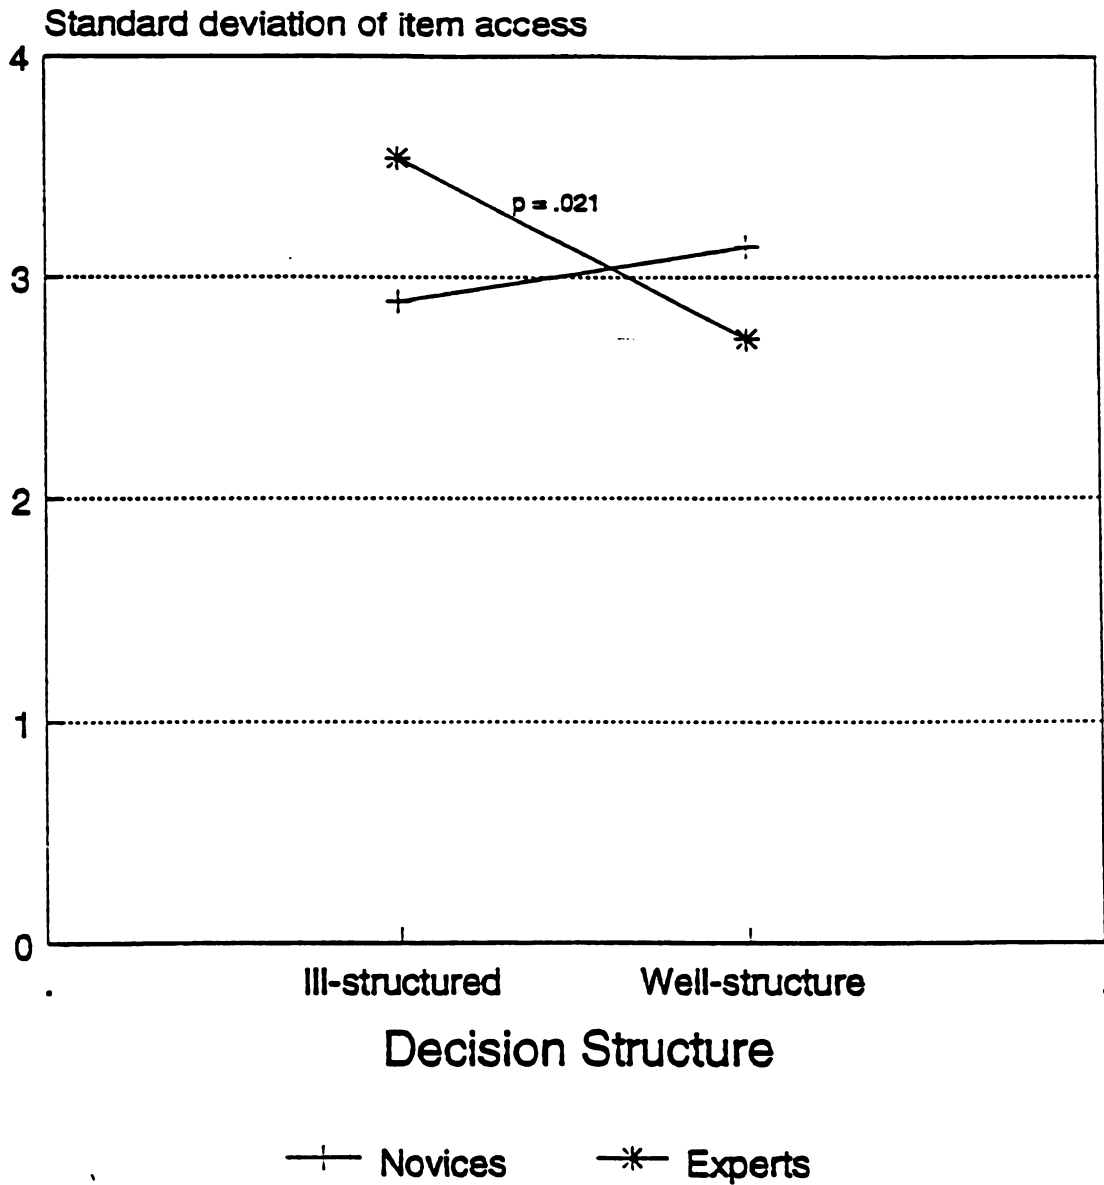


Figure 10. Domain Knowledge x Decision Structure interaction for Overall Search Variability, with "ill-structured task first" ordering

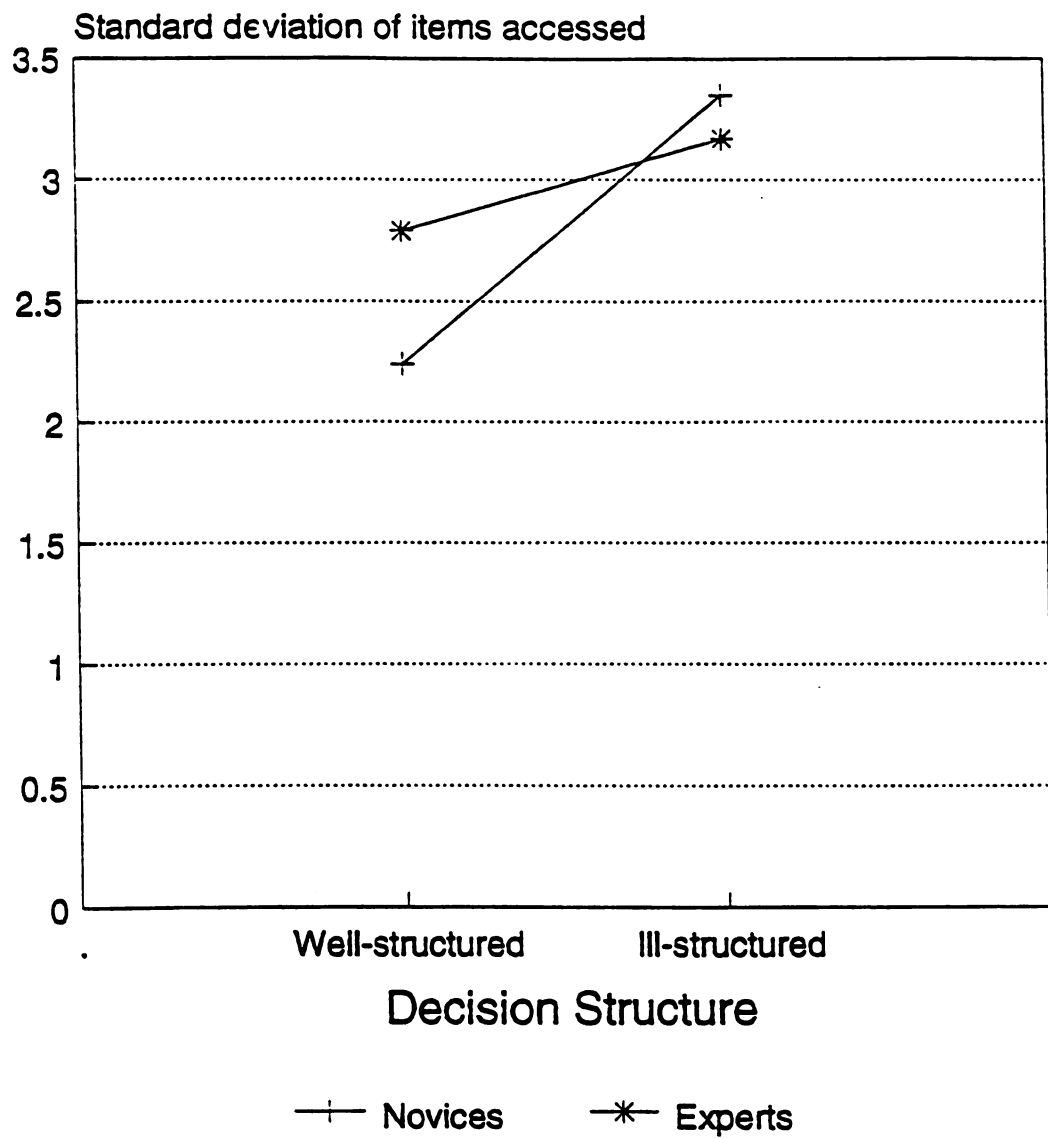


Figure 11. Domain Knowledge x Decision Structure interaction for Overall Search Variability, with "well-structured task first" ordering

Structure interaction occurring when the ill-structured decision was received first but not when the well-structured decision was received first. The Domain Knowledge x Decision Structure effect approaches significance with the well-structured first ordering, $F(1,60) = 2.45$, $p = .123$, but actually achieves significance for the ill-structured first individuals, $F(1,60) = 6.42$, $p = .014$. As predicted by Hypothesis 6b, in the "ill-structured first" conditions experts are considerably more variable in information acquisition across alternatives in the ill-structured condition ($SD = 3.54$ cues) than in the well-structured condition ($SD = 2.72$ cues), $t(30) = -2.45$, $p = .021$. While there were no other significant differences, novices revealed just the opposite tendency, displaying more variability in the well-structured task and less in the ill-structured task.

Hypothesis 6: Choice Matrix Search Variability

An attempt was made to analyze search variability within the two component matrices as well. Due to the common occurrence of no search in the contextual sub-matrix for over 50% of the individuals in the study, results from such an analysis are questionable. However, no such problem existed for search variability with respect to the four potential solution alternatives, as every participant searched at least some subset of information in the choice sub-matrix. Returning to Hypothesis 6a and 6b, the effects

of interest were the Domain Knowledge x Alternative Labels and the Domain Knowledge x Decision Structure interactions.

As Table 15 shows, two effects were found to be significant in the repeated-measures univariate ANOVA, the main effect of Structure and the Domain Knowledge x Alternative Labels interaction effect. Not surprisingly, individuals in the ill-structured decision task tended to be more variable in accessing information across alternatives than when on the well-structured task, $M = 2.43$ cues in the ill-structured task, $M = 2.06$ cues for the well-structured task.

Hypothesis 6a received partial support from the significant Domain Knowledge x Alternative Labels interaction, $F(1, 115) = 5.66$, $p = .019$. Figure 12 displays this two-way interaction. The surprising element in this interaction is the behavior of novices. While experts are relatively unaffected by the absence of alternative labels, novices react to the absence of alternative labels by becoming considerably more compensatory in their information acquisition. The difference in standard deviation among novices in the no label condition ($SD = 1.42$ cues) is significantly different from the standard deviation of information acquisition for experts in the no label condition ($SD = 2.53$ cues), $t(60) = -2.58$, $p = .011$ as well as novices in the labelled condition ($SD = 2.63$ cues), $t(59) = 3.19$, $p = .002$. No other differences were significant.

Table 15

Univariate ANOVA, Choice Matrix Search Variability

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	113282.90	2.09	.151
Alt. Labels	01	170789.43	3.15	.079
Dom. Know x Alt. Labels	01	307096.53	5.66	.019
Between-subjects Error	115	6237205.72 / 54236.57		
<hr/>				
Decision Structure	01	83013.68	3.96	.049
Dec. Struc x Dom. Know	01	8925.86	0.43	.515
Dec. Struc x Order	01	15864.21	0.76	.323
Dec. Struc x Dom. Know. x Order	01	60691.11	2.90	.092
Within-subjects Error	115	2410425.11 / 20960.22		

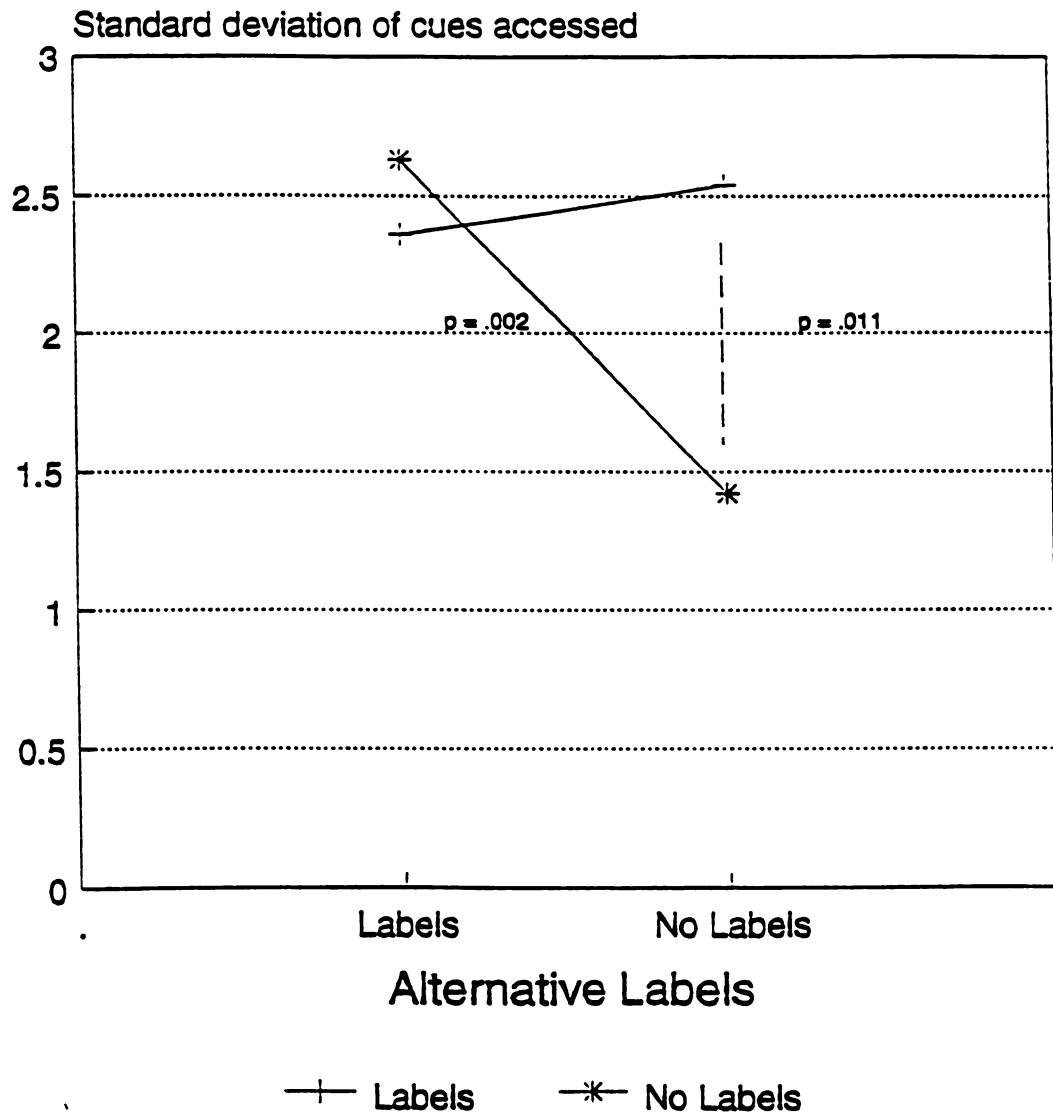


Figure 12. Domain Knowledge x Alternative Labels
interaction for Choice matrix Search Variability

In acquiring information for choice alternatives, experts appear not to be influenced by the presence of labelled alternatives, retaining high search variability. However, surprisingly, it appears that **novices** become a great deal more linear (i.e., compensatory) in their information acquisition for choice alternatives when deprived of labels, suggesting that novices use the labels as well.

Overall then, there is no support for the more global hypothesis (6) that experts will be more variable than novices in acquiring information cues across alternatives, but they do exhibit greater search variability than novices in the ill-structured decision, as hypothesized (6b), when the decisions are received in a certain order. Concerning the effect of alternative labels (6a), novices appear to alter their behavior with the presence of labels rather than experts, with the former group searching in a decidedly more linear fashion without labels.

Hypotheses 7 and 8: Search Pattern

The final dependent variable examined in this study was search pattern. Hypothesis 7 specified that experts would search in a more intradimensional fashion than would novices as a result of a more attribute-oriented search process. The appropriate test statistic for this hypothesis is the F statistic for the main effect of Domain Knowledge in the repeated-measures univariate ANOVA. (See Table 16 for the results of this analysis). This value did not achieve

significance, $F(1, 115) = .73$, $p = .396$. Thus, there is no support for the notion that experts will search more intradimensionally than novices across all decision situations.

However, the robust Order x Decision Structure interaction was again significant, $F(1, 115) = 4.16$, $p = .044$ and the decomposition of this effect indicated that, once again, search behavior was quite different in the ill-structured task simply as a function of whether it was received first or not. Figure 13 displays this interaction. In the well-structured task, it hardly matters whether individuals received the decision first or second -- they searched in almost identical patterns regardless ($M = 0.092$ for "well-structured first" ordering, $M = 0.099$ for "well-structured last" ordering, $F(1, 121) = .003$, $p = .958$). The huge difference in search behavior occurs in the ill-structured task -- when individuals received this task first, they searched in a moderately interdimensional fashion ($M = 0.267$) as opposed to when they received this task second, when there appears to be no predominance of one search pattern over the other ($M = 0.007$). This difference was significant, $F(1, 121) = 4.23$, $p = .042$, as was the difference between the ill- and well-structured tasks for individuals receiving the "ill-structured first" ordering, $t(60) = 2.34$, $p = .022$. No other differences were significant.

Table 16

Univariate ANOVA, Search Pattern index

<u>Effect</u>	<u>df</u>	<u>SS/MS</u>	<u>F</u>	<u>p</u>
Dom. Know.	01	0.55	0.73	.396
Alt. Labels	01	4.02	5.26	.024
Dom. Know x Alt. Labels	01	2.13	2.79	.098
Between-subjects Error	115	87.86 / 0.76		
<hr/>				
Decision Structure	01	0.13	0.57	.453
Dec. Struc x Dom. Know	01	0.66	2.91	.091
Dec. Struc x Order	01	0.94	4.16	.044
Dec. Struc x Dom. Know. x Order	01	0.61	2.68	.104
Within-subjects Error	115	26.05 / 0.23		

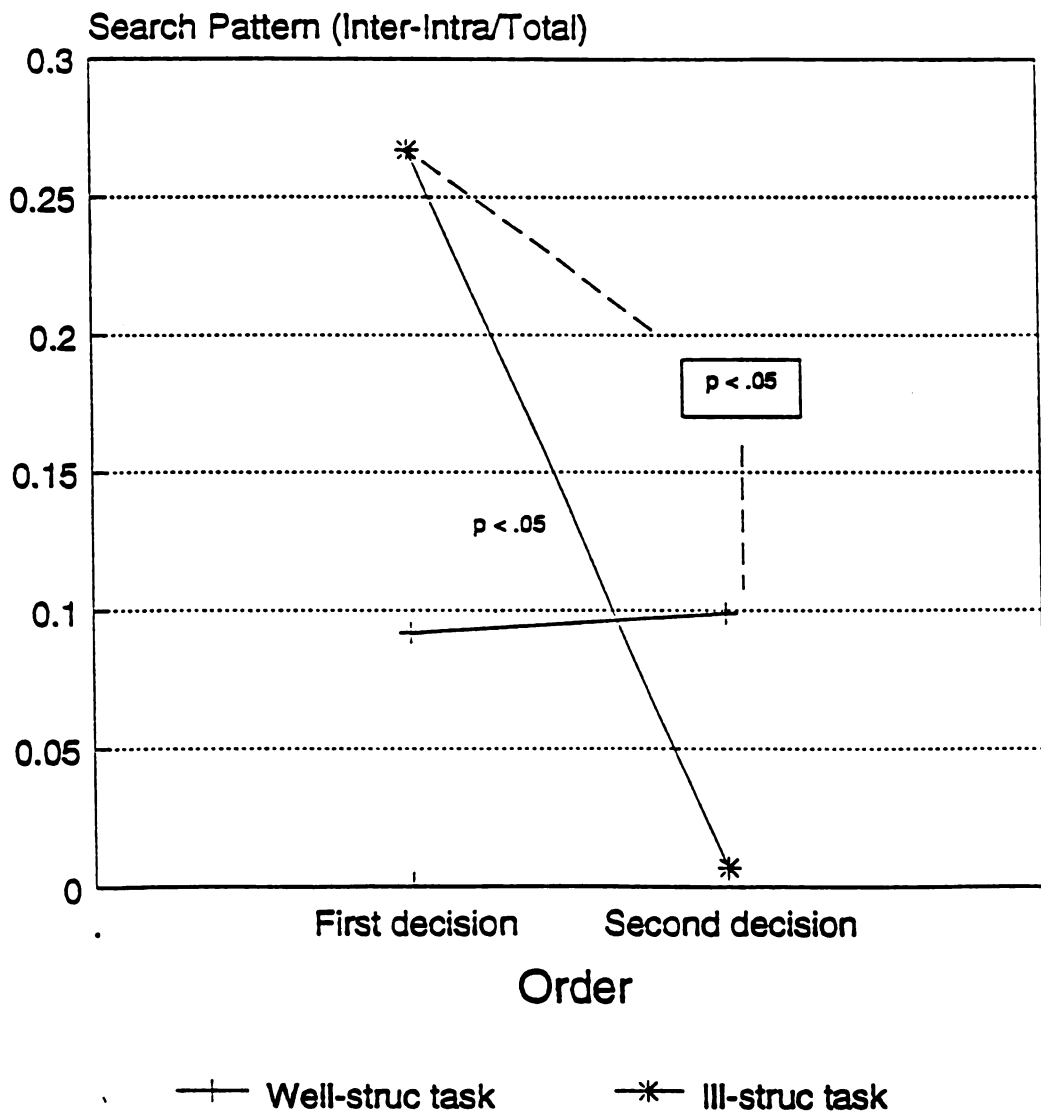


Figure 13. Order x Decision Structure interaction for Search Pattern index

Finally, Hypothesis 8 states that all individuals should search more interdimensionally (i.e., within-player) when alternative labels are present, as the position titles provide a strong suggestion concerning how organize incoming information. The appropriate test of this hypothesis is the F statistic for the main effect of Alternative Labels in the repeated-measures univariate ANOVA for search pattern. Hypothesis 8 received support from this test, F (1,115) = 5.26, p = .024, with all individuals tending to acquire information more interdimensionally (by player) when alternative labels were present than when they were absent ($M = .240$ with Labels, $M = -0.010$ without Labels).

DISCUSSION

The major finding in this study was that, as predicted, the decision processes and outcomes of experts would be sensitive to minor variations within a task while those of novices were not. Previous research on the differences between experts and novices has largely ignored the issue of intra-task variation. This study was intended to begin addressing this issue. In this section, a summary of the study findings is offered followed by a more in-depth exploration of the implications of viewing expertise as a function of both the type of task and the particular configuration of task variables. After this, some comments will be offered concerning several issues arising from the results of the study.

Summary

Table 17 provides a summary of the hypotheses tested in this study. In general, the broader hypotheses predicting unqualified main effects for expertise were rejected while support was found for the alternative set of hypotheses which specified that intra-task variation would affect the behavior of experts.

As expected for decision accuracy, experts were found to be more accurate than novices across all situations and

Table 17

Hypotheses-Study Result linkages**Hypothesis**

<u>##</u>	<u>Predicted Effect</u>	<u>Observed</u>	<u>Support</u>
1	DK	DK	Yes
1a	DK x DS	DK x DS	Yes
2	DK	NONE	No
3	DK	-DK	No
4	DK	DK	Yes
4a	DK x AL	DK x AL	Yes
4b	DK x DS	DK x DS	Yes
5	DK	-DK	No
5a	DK x AL	NONE	No
5b	DK x DS	DK x DS x O	Partial
6	DK	NONE	No
6a	DK x AL	DK x AL-- Choice only	Partial
6b	DK x DS	DK x DS x O	Partial
7	DK	NONE	No
8	AL	AL	Yes

Notes

DK = Domain Knowledge

DK x AL = Domain Knowledge x Alternative Labels

DK x DS = Domain Knowledge x Decision Structure

DK x DS x O = Domain Knowledge x Decision Structure x Order

particularly so in the well-structured situation, in which novices became more inaccurate with respect to their performance in the ill-structured condition. Contrary to prediction, experts did not spend less time looking at cues accessed, and surprisingly, they did not access fewer cues than novices when both sub-matrices were considered together.

However, as predicted, experts did access more cues providing contextual information (i.e., about the existing team) than did novices, and the number of contextual cues that experts accessed was strongly affected by the structure of the decision and the presence/absence of alternative labels as predicted. This was not true of novices. Concerning the number of cues accessed for potential choice alternatives, experts were predicted to access fewer cues than novices in general but they did not. However, the number of cues that experts did access was again affected by the decision structure as predicted when the decisions were received in one of the two orders.

Concerning search variability, experts were predicted to be more variable than novices in general and particularly so in the ill-structured situation and when labels were present. Results indicated that experts were not in fact more variable over both decisions, but were more variable than novices in the ill-structured situation as hypothesized. Surprisingly, the absence of labels did not affect experts' search variability but it did affect

novices, who acquired information quite a bit more linearly when alternatives were not labelled. Finally, experts did not acquire information more intradimensionally than novices across decisions.

Expertise and Intra-Task Variation

The major purpose of this study was to contrast a set of hypotheses concerning the effects of intra-task variation that stem from consistent findings in the expert-novice literature with a second set of "qualified" hypotheses which were specified from ACT*. A second purpose of this study was to examine the acquisition of information that did not directly bear on the choice of alternatives through inclusion of a contextual sub-matrix in the information board methodology. I will first address the implications of the lack of support for the first set of hypotheses and the corresponding support for the ACT*-derived hypotheses, then discuss the utility of including a contextual sub-matrix in future research.

The first set of hypotheses specified generally that, for a given task, experts would behave differently than novices in a systematic fashion across all conditions of that task. With respect to novices, experts were predicted to be more accurate, access less information, be more variable in their search and search more intradimensionally regardless of the degree of task structure present in the situation. In fact, there were no "unqualified" main effects of Domain Knowledge found in this study.

The "unconditional" hypotheses were uniformly rejected in favor of the more "conditional" set of hypotheses predicting that the advantages of expertise would be greater or lesser depending on the task context.

The results of this study support the notion that, to understand the concept of expertise, we can no longer take for granted the "task." Intuitively, the notion of expertise seems linked to some domain in which an "expert" performs all (or most) tasks relatively well. However, the literature on expert-novice differences has focused on tying the notion of expertise to the **task**. Tasks have generally been viewed as a set of behaviors that yield an outcome in some specific form -- e.g., a numerical answer with lots of decimals, or a verbal expression of preference for one apartment another. In essence then, solving a physics problem or choosing an apartment has been viewed as **one** task regardless of "micro" level differences in the variables that, as a configuration, determine how well "structured" the task is and to what extent previous knowledge can be utilized.

However, the "task" is not the lowest unit of analysis useful in understanding expertise. Tasks have attributes which can vary (Beach & Mitchell, 1978). This study examined two of those task characteristics, the presence/absence of alternative labels and the degree to which a task is stereotypical. Other task factors which may affect observed performance include the location of the

needed information (e.g., memory v. environment), feedback history, prior identification of alternatives, as well as the number of relevant alternatives and attributes. ACT* suggests that expert performance develops around productions that are sensitive to the level of these attributes. Thus, expertise may be dependent on the particular **configuration** of task attributes in addition to the task itself.

The present study served to illustrate how changes in the context of the same task can greatly affect how a domain "expert" goes about performing the task. Given the task of simply choosing a player to complete a team in the domain of basketball, small changes in certain variables of the choice situation (i.e., score, time remaining, which team had the ball) greatly affected the process and product of expert choices. Even though the two decisions involved the same number of stimuli, same amount of information and same choice requirement, going from four points ahead with only seconds remaining in the game to four points behind with several minutes left to go resulted in experts acquiring more information about the decision context, more information about the choice alternatives, and acquiring information more variably across the choice alternatives. In the end, they also made less accurate choices.

The distinction between a "task" and specific characteristics of the task suggests that we need to do more than just examine how experts acquire and use information about direct solutions to the task -- i.e., the choice

alternatives in a decision making study. Most studies in the past have simply given participants information that directly concerned potential choices -- i.e., location of an apartment or price of a particular consumer brand. However, information acquisition on choice alternatives may be greatly affected by the context of the choice.

Accordingly, the present study divided the information provided to the decision-makers into two halves on the basis of whether it was directly or indirectly concerned with a solution (choice alternative). The utility of this procedure is demonstrated by considering the differences in the results for the contextual (indirect) and choice (direct) sub-matrices. In the contextual matrix, experts searched for a great deal more information than novices but this was not true in the choice matrix. The implication here is that finding expertise for a particular task may depend on where you look.

The ongoing discussion has been centered around the notions that tasks are inherently heterogenous and that we need to begin to look at various aspects of a given task (degree of structure, context) to better understand expertise. Experts may in fact only behave like the typical conception of an expert when the task is analogous to a "problem" -- when the task is typical and familiar, when there is consensual agreement on the rules for identifying and evaluating necessary attributes, and when this information can be retrieved quickly from memory.

In contrast, experts may perform quite similar to novices when task variables are arranged to handicap expert performance, such as the typical "decision" involving pre-specified alternatives and attributes, the need to acquire all or most information from the environment, and a lack of consensus about what the correct choice is. The results of this study support the notion that expertise cannot be well understood without examining how experts respond to **patterns** of variables in the task environment.

Study Issues

In general, the effects of Decision Structure and Alternative Labels were as predicted -- expert search processes benefitted when situations contained stimuli with alternative labels and/or were well-structured. What was not expected was that some findings would be contingent on the order in which tasks were received.

Three-way interactions involving Domain Knowledge, Decision Structure and Order occurred for search depth in the choice matrix and for overall search variability. Two-way interactions between Decision Structure and Order also occurred in the analyses for decision accuracy, contextual search depth, cue latency and search pattern. Effectively, then, search behavior was affected by the order in which decision tasks were received in every analysis undertaken. The following question then becomes relevant: To what extent does this interaction qualify the conclusion that domain knowledge and task factors interact?

An inspection of the Order interactions represented in Figures 3-13 seem to suggest that this consistent effect may result from learning. As can be seen, there is a similarity in the pattern of Order interactions across analyses. In each case, the interaction occurs because the first task appears to have elicited more effort than the second task. Although all study participants received a practice decision task allowing them to familiarize themselves with the basic information board procedure, the order interactions suggest that participants learned a great deal about how to effectively acquire information in the search matrix in the course of the first decision. The result of this learning appears to be a streamlined information acquisition procedure in the second decision task.

A number of interesting questions can be raised on the basis of the results of this study. One particular issue is why experts accessed so much information? It might be expected (even though not specific stated in this study) that in situations which appear to maximally benefit the expert (i.e., well-structured situations and/or labels), experts should search for less information than novices. This did not occur. In the contextual sub-matrix overall, experts searched for more information as predicted and decreased their search depth for well-structured situations, but they still searched for more information than novices. The same is true for the contextual sub-matrix with regard to the presence of alternative labels. In the choice

sub-matrix, the situation is similar but here we also see the first instance of experts actually searching less than novices, and then the difference is not statistically significant (experts in well-structured, task received second). Thus, the finding that experts acquired more information than novices appears quite robust. Why?

A closer analysis of the Decision Structure manipulation check may provide some insight into the answer to this question. As can be seen from Table 4, experts considered a greater number of attributes important to the selection of a choice alternative than did novices across both decision tasks. More directly, in response to the statement, "I accessed more information than I needed in order to make good decisions," 31 experts either agreed (19) or strongly agreed (12) while only 13 novices agreed and a mere two strongly agreed. Similarly, in response to the statement, "This study was fun," 25 experts agreed with the statement and 29 experts strongly agreed, while 29 novices agreed with the statement and only 9 strongly agreed. Thus, the most promising explanation of why experts acquired so much information seems to be that they enjoyed the task and may have acquired information even after they had made their decision about the correct alternative -- much as one checks the box scores even after the final score of the game is known.

A second issue which arises is why novices focused their selection on Alternatives 3 and 4. Experts, across

all situations, were less inclined to favor any one particular alternative. The most likely explanation seems to be that novices appeared to be searching for the best all-around alternative while experts were looking for something in particular. If the particular attributes which are important in a given situation are unknown, it does seem reasonable to choose an alternative on the basis of attributes which would appear to be important in any situation -- e.g., height, scoring ability, etc. The manner in which choice alternatives are numbered roughly corresponds to a pattern of increasing overall proficiency across all attribute dimensions. In addition, it also appears that how choice players were labelled may have had an effect as well, as novice choices were more concentrated on the alternatives with semantically-positive labels (i.e., "power forward").

The effect of semantic connotation may also help to explain the poor performance of novices on the well-structured task compared to the ill-structured task. Again, a possible clue is provided by the distribution of novice choices over all tasks as shown in Table 9. Novices, in general across both tasks, tended to choose alternatives that had scored more points, were taller and had semantically-positive titles (e.g., "Power Forward"). Coincidentally, this type of player was the correct choice in the ill-structured task but not in the well-structured task, where a different "role" was needed. The better

novice performance on the ill-structured task may well reflect a spurious match between their prejudices and the demands of the ill-structured situation.

Another point which arises is why novice choice search variability decreased to such an extent when alternatives were not labelled. As will be recalled, expert choice search variability was not much affected by the presence of labels. One potential explanation is that although position labels were not strongly linked to basketball performance dimensions for novices, such labels were still used (perhaps incorrectly) to make inferences about where to look for certain information and how to evaluate the information that was acquired. Interestingly, one implication of this notion is that labelling alternatives may impact novice information acquisition process to a greater degree than experts in a decision making/problem solving situation.

Finally, an effort was made in this study to tap not only the traditional element of expertise, declarative knowledge, but also procedural knowledge as well. To this end, items were written and included in the basketball knowledge test that tapped essentially basketball strategy and tactics. These items, while not having "correct" answers in any official game sense, were judged to indeed have responses that were substantially better than others, and in this sense were consensually validated. Future studies might consider using other (and better) methods of capturing procedural knowledge.

Task Variable Interactions

The basic model presented in Figure 1 suggests that the information value of a well-structured decision and the presence of alternative labels is independent of each other. The former provides information about what attributes to search, the latter provides information on where to search (i.e., which alternatives are most likely to provide these desired features). Conceptually then, there is no reason for these task characteristics to interact with Domain Knowledge.

However, as indicated in Table 7, the three-way interaction parameter for Domain Knowledge x Alternative Labels x Decision Structure was marginally significant in the model for Decision accuracy, $X^2 = 3.11$, $p = .078$. The locus of the marginal three-way interaction stems from the fact that experts with labelled alternatives in the well-structured task were substantially more accurate than experts without alternative labels in the same task. Novices, appear to do even worse in the well-structured task when alternative labels are present, possibly for the reasons mentioned above.

Experts in the well-structured task who received alternative labels were very accurate, $p = 80.0\%$, as opposed to Experts in the well-structured task who did not receive alternative labels, $p = 56.3\%$. Apparently, when experts have the benefit of alternative labels and a well-structured decision, they do indeed perform like experts -- four times

more accurate than novices in the corresponding condition (80% versus 19%, respectively).

However, contrary to the proposed model in Figure 1, this findings suggests that the effects of alternative labelling and a well-structured task are somewhat contingent on one another. Experts do not seem to benefit as much from having labelled alternatives when the task is ill-structured. There appear two potential explanations for this phenomenon: 1) the information provided by alternative labels and a well-structured decision is redundant to some degree and/or 2) the presence of both a well-structured situation **and** labelled alternatives aids the representation of information in working memory and thus reduce the amount lost to decay or overload. At this point, it is not clear if either (or both) hypotheses are correct, but future research might fruitfully address this question.

Conclusions

Past research has viewed expertise as a function of the task but has not devoted much attention to variations in the same task. In addition, the literature on expertise has been exploratory and descriptive and has generally failed to test its many scattered findings with inferential methods.

This study was an attempt to begin addressing these issues. Experts were gathered in sufficient numbers to statistically test hypotheses. Unlike past research, expertise was viewed as a function of declarative knowledge and procedural knowledge. In addition, several factors than

can vary **within** a task were examined along with expertise. Finally, an addition was made to the standard information board methodology that incorporated cues which both directly and indirectly provided information about the correct alternative.

The results of this study support the conclusions that expertise may exist for, and be limited to, a configuration of variables present in the task -- not just the generic "task". This conclusion suggests a view of expertise that is somewhat different from the existing literature, which has tended to ignore variability within the task and has instead focused on performance differences across tasks. Future research would do well to focus on the interaction between the "top-down," schematic, categorical knowledge that an expert possesses and the various "bottom-up" aspects of the task and its environment which both aid and hinder the expert in constructing a representation of the problem and performing like an "expert" should.

APPENDICES

APPENDIX A**Basketball Knowledge questionnaire**Rules and terminology

1. In addition to "full" time-outs in the NBA, another type of time-out exists that lasts for how long?
 - a. 20 seconds
 - b. 30 seconds
 - c. 45 seconds
 - d. 1 minute
 - e. 2 minutes
2. In professional basketball, a "triple-double" refers to which of the following?
 - a. when a player commits three double-fouls in one game
 - b. when a players scores three consecutive two-point baskets
 - c. when a player gets double-digit statistics in the following categories: points, rebounds and assists.
 - d. when a player makes two consecutive three-point baskets
 - e. when a player two free-throws three times in a game
3. How many referees are used in for a professional basketball game?
 - a. one
 - b. two
 - c. three
 - d. four
4. A "pick" is when which of the following occurs?
 - a. a good pass picks open a strong defense
 - b. a defensive player gets between an offensive player and the basket
 - c. an offensive player has the ball stolen from him by a defensive player
 - d. an offensive player shields a defensive player with his body
 - e. a defensive player guards the basket by himself against multiple offensive players on a fast-break
5. How high is the rim of the basket from the floor?
 - a. 10 feet
 - b. 7 feet
 - c. 9 feet
 - d. 12 feet
6. Which of the following types of defenses is illegal in the NBA?
 - a. man-to-man
 - b. zone
 - c. press
 - d. double-teaming

7. How are the top seven college prospects chosen in the NBA draft?

- a. by multiple coin tosses
- b. by computer formula taking into account strength of schedule
- c. in order of the worst team (i.e., worst team chooses first, second-worst team chooses second, etc.)
- d. with a lottery

8. After how many fouls does a player "foul out" of a professional basketball game?

- a. four
- b. five
- c. six
- d. none of the above

9. Which of the following players is commonly referred to as the "sixth man?"

- a. the first player to come off the bench
- b. the coach
- c. the assistant coach
- d. a player who forgets to check in at the scorer's table

10. How many free-throws does a player get when he is fouled in the act of shooting a three-point shot AND makes it?

- a. none
- b. one
- c. two
- d. three

11. What happens when a defensive player steps in the lane too soon right before an offensive player misses a free throw?

- a. nothing
- b. the offensive team gets the ball out of bounds.
- c. the offensive team gets to shoot over
- d. the offensive team automatically gets the point

12. How long does a team have to inbound the ball before a violation occurs and the other team gets the ball?

- a. 4 seconds
- b. 5 seconds
- c. 6 seconds
- d. 8 seconds
- e. 10 seconds

13. How many minutes are there in an NBA quarter?

- a. 8 minutes
- b. 10 minutes
- c. 12 minutes
- d. 15 minutes
- e. 20 minutes

14. How many total games are played by a team in an NBA regular season?
- a. 65 games
 - b. 74 games
 - c. 80 games
 - d. 82 games
 - e. 84 games
15. About how far is the three-point line from the basket in the NBA?
- a. 19 feet
 - b. 20 feet
 - c. 21 feet
 - d. 22 feet
 - e. 23 feet
16. How long does a team have to get the ball over the time line before a violation occurs and the other team gets the ball?
- a. 5 seconds
 - b. 10 seconds
 - c. 15 seconds
 - d. 20 seconds
 - e. 30 seconds
17. The free-throw line is how many feet from the basket?
- a. 10 feet
 - b. 12 feet
 - c. 12.8 feet
 - d. 14 feet
 - e. 15 feet
18. A turnover is committed when a player on offense is "closely guarded" for how long?
- a. 3 seconds
 - b. 5 seconds
 - c. 10 seconds
 - d. 12 seconds
 - e. 15 seconds
19. The term "spot-up" implies what kind of shot?
- a. an open jump shot
 - b. a turn-around shot from the low block
 - c. a closely guarded hook shot
 - d. a drive down the base-line
20. How long is the shot clock in professional basketball?
- a. 24 seconds
 - b. 25 seconds
 - c. 30 seconds
 - d. 45 seconds
 - e. there is no shot

Strategy and Tactics

21. In which of the following situations would playing a zone defense be much better than playing a man-to-man defense?
- a. when the other team is shooting well from the outside
 - b. when the other team is slower and smaller
 - c. when the other team passes very well
 - d. when the other team is taller and quicker

22. Which of the following players is likely to get the most steals in a game?

- a. the point guard
- b. the small forward
- c. the power forward
- d. the center

23. A penetration move by a point guard usually WILL NOT end up in which of the following?

- a. a shooting foul
- b. an open jump shot
- c. a "dish" and a slam
- d. a turn-around jumper from the low post

24. If your team is playing on the road in front of a noisy, hostile crowd and the other team has scored the last 9 points of the game (4 of them coming on steals and dunks), what might you do?

- a. call a time out
- b. use a pressing defense next time
- c. try a fast-break
- d. go for a three-point shot

25. Which of the following is LEAST likely to work against a zone defense?

- a. a three-point shot
- b. a pull-up jumper
- c. a drive to the basket
- d. a base-line turn-around

26. When is it a good time NOT to play a man-to-man defense?

- a. when the other team shoots very well from the outside
- b. when the other team has one excellent scorer
- c. when the other team runs the fast break well
- d. when your team is in foul trouble

27. A good thing to do to when playing at home and your team seems unmotivated is _____?

- a. pull out all five players and put in subs for a few minutes
- b. hold the ball on offense
- c. commit an intentional foul
- d. draw a technical foul

28. The last person you would expect to see bringing the basketball up the court is a _____?

- a. a point guard
- b. an off guard
- c. a small forward
- d. a center

29. In order to use a press defense, which of the following must happen?

- a. the other team must be slow
- b. the other team must have one bad ball-handler
- c. your team must make a basket
- d. the inbounding player must be guarded

30. Which position is least likely to get an offensive rebound?

- a. point guard
- b. off guard
- c. small forward
- d. power forward
- e. center

31. In order to consistently take advantage of a fast-break offense, which of the following things is most basic?

- a. get offensive rebounds
- b. get defensive rebounds
- c. play defense
- d. shoot pull-up jumpers
- e. make long outlet passes

32. Which of the following makes for the best combination of talents among five starters on a basketball team?

- a. a ball-handler, an outside shooter, a low-post man and two rebounders
- b. three shooters and one rebounder and one shot blocker
- c. three scorers and one rebounder and one free-throw specialist
- d. four scorers and one rebounder
- e. two ball-handlers, one shooter and two rebounders

33. What type of risky pass can sometimes offer an open shot against a zone defense?

- a. no-look
- b. behind-the-back
- c. skip
- d. half-court

34. If your team is behind by 12 points with 2 minutes left in the game, which of the following positions should be over-represented on the floor (i.e., more than one)?

- a. point guard
- b. off guard
- c. small forward
- d. power forward
- e. center

35. Which of the following violations is probably least seen in the NBA?
- a. over-the-back
 - b. on-the-arm (shooting)
 - c. blocking
 - d. 3 seconds in the lane
 - e. charging
36. Which of the following statistics are most important to a team trying to preserve a lead?
- a. assists and offensive rebounds
 - b. three point shooting and steals
 - c. defensive rebounds and free-throws
 - d. field-goal shooting and three-point shooting
 - e. assists and free-throws
37. An average free-throw shooter in the NBA will shoot near which of the following percentages over the course of a season?
- a. 68%
 - b. 70%
 - c. 77%
 - d. 85%
 - e. 88%
38. What area would you find a power forward (on offense) playing near most of the time?
- a. the low post
 - b. the top of the key
 - c. the high post
 - d. the wing
39. The player who traditionally is the best outside shooter on a professional basketball team is the _____?
- a. point guard
 - b. off guard
 - c. small forward
 - d. center
40. Which player has his back to the basket most often in a traditional basketball offense?
- a. point guard
 - b. off guard
 - c. small forward
 - d. center

Appendix B

Table B-1

Cue Values for Search Matrix A

Attribute Dimension										
	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	<u>6</u>	<u>7</u>	<u>8</u>	<u>9</u>	<u>10</u>
A	16	49	82	33	6-11	3-3	1-3	4	0	01
B	22	54	79	46	8-12	1-2	5-6	1	0	01
C	23	60	68	17	9-16	5-8	0-0	1	1	03
D	11	52	74	20	4-8	3-4	0-0	1	1	02
E	14	50	61	16	5-10	1-2	1-2	2	0	01
F	05	48	91	30	1-2	3-3	0-1	0	0	00
G	06	48	72	11	3-8	0-0	0-0	0	2	04
H	08	55	75	40	3-7	1-1	1-1	0	1	01

	<u>11</u>	<u>12</u>	<u>13</u>	<u>14</u>	<u>15</u>
A	02	03	09	05	02
B	03	02	02	07	07
C	03	03	01	03	09
D	04	02	02	11	09
E	00	01	00	03	08
F	02	00	01	09	04
G	11	02	00	04	10
H	03	02	01	01	03

Attribute Dimensions

1 = Points scored (game)	9 = Blocked Shots
2 = Season field-goal %	10 = Offensive Rebounds
3 = Season free-throw %	11 = Defensive Rebounds
4 = Season three-point FG %	12 = Turnovers
5 = Field-goals attempted-made	13 = Assists
6 = Free-throws attempted-made	14 = Years in the NBA
7 = Three-point field-goals at./m.	15 = Height (inches over six feet)
8 = Steals	

Appendix B (cont'd)

Table B-2

Cue Values for Search Matrix B

Attribute Dimension										
	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	<u>6</u>	<u>7</u>	<u>8</u>	<u>9</u>	<u>10</u>
I	13	47	84	37	5-8	1-1	2-2	2	0	01
J	13	50	80	48	4-7	3-4	2-3	3	1	02
K	19	57	65	10	8-14	3-5	0-1	0	0	02
L	31	55	77	18	11-21	9-12	0-0	1	1	01
M	02	46	93	29	0-0	2-2	0-0	1	0	00
N	10	49	73	51	3-6	1-2	3-6	0	0	02
O	13	52	77	35	5-11	2-3	1-4	2	0	01
P	09	50	70	15	4-8	1-2	0-0	0	3	05

	<u>11</u>	<u>12</u>	<u>13</u>	<u>14</u>	<u>15</u>
I	02	02	05	04	01
J	02	04	04	04	04
K	03	02	01	07	09
L	05	04	03	06	10
M	00	00	00	06	00
N	04	01	00	03	04
O	01	00	01	11	08
P	09	02	02	07	09

Attribute Dimensions

- | | |
|------------------------------------|------------------------------------|
| 1 = Points scored (game) | 9 = Blocked Shots |
| 2 = Season field-goal % | 10 = Offensive Rebounds |
| 3 = Season free-throw % | 11 = Defensive Rebounds |
| 4 = Season three-point FG % | 12 = Turnovers |
| 5 = Field-goals attempted-made | 13 = Assists |
| 6 = Free-throws attempted-made | 14 = Years in the NBA |
| 7 = Three-point field-goals at./m. | 15 = Height (inches over six feet) |
| 8 = Steals | |

APPENDIX C

Rationale for correct alternative choice

To win in the well-structured situation, the subject's team does not have to do anything -- including shoot the ball. If desired and allowed, the subject's team can simply hold on to the ball and run out the clock. The opposing team needs to get the ball back to have any chance of winning. There are only four events that could get them the ball back: 1) the team with the ball fails to inbound it within five seconds, 2) a steal, 3) a turnover or 4) a foul and a missed free-throw and a rebound.

These four events implicate two attributes which are important in choosing a fifth player on the decision-maker's team: free-throw shooting ability and ball-handling ability. These two skills correspond most closely with the attribute dimensions "Season Free-Throw Percentage" and "Turnovers" (game). **The logical choice would be to select is the alternative who is best on both of these attributes.** To avoid a trade-off dilemma, the correct alternative in both decision matrices has the best values on the team with respect to both of these attributes. In Search Matrix A, the correct alternative in the high-structure decision task is Alternative F. In Search Matrix B, the correct alternative is Alternative M.

In the low structure decision task, the "correct" choice is less obvious but still exists. Given that the subject's team is behind by four and the other team has the ball, the subject's team needs to score several baskets in the next two minutes while preventing the other team from scoring. Two minutes in the NBA -- with time-outs, free-throws and half-court inbounding - can last a long time. There is no immediate need to panic. Necessary attributes include the ability to score, steal, block shots and rebound (with more emphasis on defensive rebounds). The four players already chosen include good scorers and stealers. The other team will likely attempt to use up as much of the 24-second clock as possible before shooting each time they get the ball, but teams that hold the ball usually do not get off good shots so there should be many opportunities to snare defensive rebounds. In addition, every time there is a shot and the other team gets the ball back, they can hold it for another 24 seconds. This puts a premium on choosing a player who can get defensive rebounds.

APPENDIX C (cont'd)

The logical alternative is to select is the player that has the most rebounds (offensive and defensive, or just defensive). The correct alternative in both decision matrices has the most offensive rebounds, defensive rebounds and total rebounds of anyone on the team, with no other alternative remotely comparable. In Search Matrix A, the correct alternative for the low structure decision task is Alternative G. In Search Matrix B, the correct alternative is Alternative P.

APPENDIX D**Importance Ratings**

For the following two basketball positions, please indicate how important the following statistics/attributes are to each position by marking your choice in the appropriate blank on the opscan sheet using the scale below. The two positions you will be considering are "Point Guard" and "Center." Please use the following 5 point scale to indicate how important each statistic is to each position:

- 1 = Not very important to this position
- 2 = Somewhat important to this position
- 3 = Moderately important to this position
- 4 = Fairly important to this position
- 5 = Very important to this position

If you really do not have any idea about the importance of a given statistic/skill to one of the positions, please respond with the following

- 9 = Don't know

Point Guard

- 45. Points per game
- 46. Field Goal %
- 47. Three-point field goal %
- 48. Steals
- 49. Assists
- 50. Offensive rebounds
- 51. Defensive rebounds
- 52. Blocked Shots
- 53. Avoiding turnovers
- 54. Free-throw %

Center

- 55. Points per game
- 56. Field goal %
- 57. Three-point field
- 58. Steals
- 59. Assists
- 60. Offensive rebounds
- 61. Defensive rebounds
- 62. Blocked Shots
- 63. Avoiding turnovers
- 64. Free-throw %

APPENDIX E**Post-experimental questionnaire**

- 1 = Strongly disagree
- 2 = Disagree
- 3 = Neither Agree nor Disagree
- 4 = Agree
- 5 = Strongly Agree

1). The instructions for using the computer were clear and understandable. _____

2). It was easy to use the computer to access player information. _____

3). This study was fun. _____

4). I accessed more information than I needed in order to make good decisions. _____

5). I felt overwhelmed by all of the information I could look at. _____

6). I knew what player attributes were important in the situation where my team was AHEAD before I started looking.

7). I knew what player attributes were important in the situation where my team was BEHIND before I started looking.

8). I had a general "strategy" for accessing player information.

9). In the space below, please circle those player attributes you "weighted" most heavily in making your decisions:

When your team was AHEAD:

Points	Season Field Goal %	Season Free-Throw%
Steals	Blocked Shots	OFF.REB DEF.REB FG MADE-ATT
FT MADE-ATT	3ptFG MADE-ATT	Season 3pt FG% Height
Years/NBA	Turnovers	Assists

APPENDIX E (cont'd)

When your team was BEHIND:

Points	Season	Field Goal %	Season	Free-Throw%
Assists	Steals	Blocked Shots	OFF.REB	DEF.REB
FG MADE-ATT	FT MADE-ATT	3ptFG MADE-ATT		
Season 3pt FG%	Height	Years/NBA	Turnovers	

9). Please comment on your general "strategy" for accessing information when your team was AHEAD:

10). Please comment on your general "strategy" for accessing information when your team was AHEAD:

What was the score of the game when you were ahead?

What was the score of the game when you were behind?

APPENDIX F**Basketball Experience questionnaire**

65. In terms of high school basketball (school team or intramurals), I:

- a. did not play on a basketball team in high school
- b. played high school basketball for one year
- c. played high school basketball for two years
- d. played high school basketball for three years

66. I watch _____ college basketball games on TV per week during college basketball season.

- a. less than one game a week
- b. one or two games per week
- c. three-four games a week
- d. five or more games a week

67. I watch professional basketball on TV _____ times per week during the NBA season.

- a. less than one game a week
- b. one or two games per week
- c. three-four games a week
- d. five or more games a week

68. Currently, I play basketball (informally or intramurals) about _____ times a week.

- a. less than one
- b. once or twice
- c. three-four
- d. more than four

69. Concerning the basic rules of basketball (i.e., rules common to all levels of basketball), I would say that I know them:

- a. not at all
- b. not very well
- c. very well
- d. well enough to be a ref
- c. well enough to understand most of the game

APPENDIX F (cont'd)

70. In terms of knowledge of basketball strategy and tactics, where would you place yourself on the following scale?

- a. I know you need to score more points to win
- b. I know what some of the common plays/defenses are
- c. I know what each of the players should be doing and when someone has taken a good/bad shot
- d. I can often predict what the coach/team will do in an upcoming situation
- e. I could do a decent job coaching a boys or girls team

LIST OF REFERENCES

LIST OF REFERENCES

- Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. Memory & Cognition, 9, 422-433.
- Anderson, J.R. (1976). Language, memory, and thought. Hillsdale, N.J.: Erlbaum.
- Anderson, J.R. (1982). Acquisition of cognitive skill. Psychological Review, 89, 369-406.
- Anderson, J.R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.
- Anderson, J.R. (1987). Skill acquisition: Compilation of weak-method problem solutions. Psychological Review, 94, 192-210.
- Anderson, J.R. (1992). Automaticity in ACT* theory. American Journal of Psychology, 105, 165-180.
- Anderson, J.R., Kline, P.J., & Beasley, C.M. (1978). A theory of the acquisition of cognitive skills. Technical Report No. ONR 77-1. Department of Psychology, Yale University.
- Beach, L.R., & Mitchell, T.R. (1978). A contingency model for selection of decision strategies. Academy of Management Review, 3, 439-449.
- Charness, N. (1979). Components of skill in bridge. Canadian Journal of Psychology, 33, 1-16.
- Chase, W.G., & Simon, H.A. (1973). Perception in chess. Cognitive Psychology, 4, 55-81.
- Chi, M. (1978). Knowledge structures and memory development. In R.S. Siegler (Ed.), Children's Thinking: What develops? Hillsdale, NJ: Erlbaum.
- Chi, M.T.H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

- Chi, M.T.H., Glaser, R., M.J. Farr, (Eds.). (1988). The nature of expertise. Hillsdale, NJ: Erlbaum.
- Chiesi, H.L., Spilich, G.J., & Voss, J.F. (1979). Acquisition of domain-related information in relation to high and low domain knowledge. Journal of Verbal Learning and Verbal Behavior, 18, 257-274.
- Cole, J.W.L., & Grizzle, J.E. (1966). Applications of multivariate analysis of variance to repeated measurements experiments. Biometrics, 22, 810-827.
- Cooke, N.J., & Schvaneveldt, R.W. (1988). Effects of computer programming experience on network representations of abstract programming concepts. International Journal of Man-Machine Studies, 29, 407-427.
- Dawson, V.L., Zeitz, C.M., & Wright, J.C. (1989). Expert-novice differences in person perception: Evidence of experts' sensitivities to the organization of behavior. Social Cognition, 7, 1-30.
- Egan, D.E., & Schwartz, B.J. (1979). Chunking in recall of symbolic drawings. Memory & Cognition, 7, 149-158.
- Engle, R.W., & Bukstel, L. (1978). Memory processes among bridge players of differing expertise. American Journal of Psychology, 91, 673-689.
- Ericsson, K.A., & Smith, J. (1992). Prospects and limits of the empirical study of expertise: An introduction. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise (pps. 1-38). New York: Cambridge University Press.
- Fiske, S.T., Kinder, D.R., & Larter, W.M. (1983). The novice and the expert: Knowledge-based strategies in political cognition. Journal of Experimental Social Psychology, 19, 381-400.
- Fiske, S.T., Neuberg, S.L., Beattie, A.E., & Milberg, S.J. (1987). Category-based and attribute-based reactions to others: Some informational conditions of stereotyping and individuating processes. Journal of Experimental Social Psychology, 23, 399-427.
- Ford, J.K., Schmitt, N., Schechtman, S.L., Hults, B.M., & Doherty, M.L. (1989). Process-tracing methods: Contributions, problems, and neglected research questions. Organizational Behavior and Human Decision Processes, 43, 75-117.

- Gick, M.L. (1986). Problem solving strategies. Educational Psychologist, 21, 99-120.
- Gilliland, S.W., Wood, L., & Schmitt, N. (in press). Cost-benefit determinants of decision process and accuracy. Organizational Behavior and Human Decision Processes. (In press).
- Hattrup, K., & Ford, J.K. (1991). Stereotypes in decision making: Category biases in the acquisition of information. Paper presented at the 6th annual conference of the Society of Industrial-Organizational Psychology (SIOP), St. Louis, MO.
- 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 & Human Decision Processes, 46, 77-101.
- Hinsley, D.A., Hayes, J.R., & Simon, H.A. (1977). From words to equations: Meaning and representation in algebra word problems. In M.A. Just & P.A. Carpenter (Eds.), Cognitive processes in comprehension (pps. 89-108). Hillsdale, NJ: Erlbaum.
- Hintzman, D.L. (1986). "Schema abstraction" in a multiple-trace memory model. Psychological Review, 93, 411-428.
- Hobus, P.P.M., Schmidt, H.G., Boshuizen, H.P.A., & Patel, V.L. (1987). Medical Education, 21, 471-476.
- Hollander, Z. (1991). The 1992 complete handbook of pro basketball. New York: Signet.
- Johnson, E.J. (1980). Expertise in admissions judgment. Unpublished doctoral dissertation, Carnegie-Mellon University.
- Johnson, E.J. (1988). Expertise and decision under uncertainty: Performance and Process. In M.T.H. Chi, R. Glaser & M.J. Farr (Eds.), The nature of expertise, (pp.209-228). Hillsdale, NJ: Erlbaum.
- Johnson, P.E., Duran, A.S., Hassebrock, F., Moller, J., Preitula, M., Feltovich, P.J., & Swanson, D.B. (1981). Expertise and error in diagnostic reasoning. Cognitive Science, 5, 235-283.
- Kahneman, D., & Miller, D.T. (1986). Norm theory: Comparing reality to its alternatives. Psychological Review, 93, 136-153.

- Larkin, J., McDermott, J., Simon, D., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. Science, 208, 1335-1342.
- McKeithen, K.B., Reitman, J.S., Rueter, H.H., & Hirtle, S.C. (1981). Knowledge organization and skill differences in computer programmers. Cognitive Psychology, 13, 307-325.
- Means, M.L., & Voss, J.F. (1985). Star wars: A developmental study of expert and novice knowledge structures. Journal of Memory and Language, 24, 746-757.
- Murphy, G.L., & Wright, J.C. (1984). Changes in conceptual structure with expertise: Differences between real-world experts and novices. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 144-155.
- Neches, R., Langley, P., & Klahr, D. (1987). In R. Neches, P. Langley, & D. Klahr, (Eds.), Production system models of learning and development. Cambridge, MA: The MIT Press.
- Newell, A., & Simon, H.A. (Eds.) (1972). Human problem solving. Englewood Cliffs, NJ: Prentice Hall.
- Novick, L. (1988). Analogical transfer, problem similarity and expertise. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 510-520.
- Payne, J.W. (1982). Contingent decision behavior. Psychological Bulletin, 92, 382-402.
- Payne, J.W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. Organizational Behavior and Human Performance, 16, 366-387.
- Reitman, W. (1965). Cognition and thought. New York: Wiley.
- Reitman, J.S. (1976). Skilled perception in Go: Deducing memory structures from inter-response times. Cognitive Psychology, 9, 336-356.
- Reitman, J.S., & Rueter, H.H. (1980). Organization revealed by recall orders and confirmed by pauses. Cognitive Psychology, 12, 554-581.

- SAS Institute Inc. (1985). SAS User's Guide: Statistics, Version Five edition. Cary, NC: SAS Institute Inc.
- Schank, R., & Abelson, R. (1977). Scripts, plans, goals and understanding. Hillsdale, NJ: Erlbaum.
- Schoenfeld, A.H., & Herrmann, D.J. (1982). Problem perception and knowledge structure in expert and novice mathematical problem solvers. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 484-494.
- Simon, D., & Simon, H.A. (1978). Individual differences in solving physics problems. In R.S. Siegler (Ed.), Children's Thinking: What develops? Hillsdale, NJ: Erlbaum.
- Spilich, G.J., Vesonder, G.T., Chiesi, H.L., & Voss, J.F. (1979). Text processing of domain-related information for individuals with high and low baseball knowledge. Journal of Verbal Learning and Verbal Behavior, 18, 275-290.
- Sternberg, R.J., & Tulving, E. (1977). The measurement of subjective organization in free recall. Psychological Bulletin, 84, 539-556.
- Svenson, O. (1979). Process descriptions of decision-making. Organizational Behavior and Human Performance, 23, 86-112.
- Voss, J.F., Greene, T.R., Post, T.A., & Penner, B.C. (1983). Problem solving skill in the social sciences. In G. Bower (Ed.), The psychology of learning and instruction: Advances in research and theory. New York: Academic Press.
- Voss, J.F., Vesonder, G.T., & Spilich, G.J. (1980). Text generation and recall by high-knowledge and low-knowledge individuals. Journal of Verbal Learning and Verbal Behavior, 19, 651-667.