1988

AN EXPERIMENTAL STUDY IN OVERCOMING HYPOTHESIS-CONFIRMING SEARCH STRATEGIES IN COMPUTERIZED INFORMATION RETRIEVAL SYSTEMS

Sheila M. Jacobs
Oakland University

Robert T. Keim
Arizona State University

AN EXPERIMENTAL STUDY IN OVERCOMING HYPOTHESIS-CONFIRMING SEARCH STRATEGIES IN COMPUTERIZED INFORMATION RETRIEVAL SYSTEMS

Sheila M. Jacobs
School of Business Administration
Oakland University

Robert T. Keim
Department of Decision and Information Systems
College of Business
Arizona State University

ABSTRACT

A recent trend in information retrieval systems technology is the development of on-line information retrieval systems. One objective of these systems has been to attempt to enhance decision effectiveness by allowing users to preferentially seek information, thereby facilitating the reduction or elimination of information overload. Since information systems users may preferentially seek information to confirm their initial beliefs, decision making effectiveness may be dependent on the accuracy of the decision maker's initial hypothesis of causality.

The basic research question addressed in this paper is: Will the use of a knowledge-based DSS (KBDSS), designed to search for and present both confirming and disconfirming evidence, result in enhanced decision effectiveness?

To assess the effect of information retrieval system type on decision effectiveness, a laboratory experiment was conducted in which participants were required to make an initial attribution of causality for a problem, to query either a conventional on-line information retrieval system or a KBDSS for additional information, and then to make a final attribution of causality. The conclusions reached from this experiment provide constructive guidance for information systems designers in overcoming the concept known as confirmation bias, that tendency to seek information that confirms the user's first impression.

1. INTRODUCTION

Due to the evolution of computer technology, information stored by the computer is much more accessible to the individual decision maker, and searching for information is an important element of computer-aided decision making. One objective of information retrieval systems (such as decision support systems and on-line inquiry systems) has been to attempt to enhance decision effectiveness by allowing users to preferentially seek information, thereby facilitating the reduction or elimination of information overload. There is reason to believe, however, that on-line information retrieval systems do not necessarily lead to more effective decision making. When users are seeking information at a point in time subsequent to forming initial beliefs, the issue of information search strategy is relevant.

The decision support systems (DSS) approach has helped to improve the information retrieval capabilities of individual decision makers. This approach advocates (a) a data capture and extraction approach that could combine a variety of data sources (e.g., external as well as internal data), (b) the ability to quickly and easily add and delete data sources, (c) the portrayal of data structures in user terms so the user could understand what data was available, (d) the ability to handle the user's personalized data (e.g., unofficial and personal data), and (e) a full range of data management functions for all of this data (Sprague 1980).

Recent research in information search strategy suggests that individuals may preferentially seek information to confirm their initial beliefs. If this is the case, decision making effectiveness may be dependent on the decision maker's initial beliefs, that is, the initial hypothesis or attribution of causality. Thus, if an individual preferentially seeks confirming information, his or her revised attribution of causality (after the information search) will tend to be similar to the initial attribution of causality, regardless of the accuracy of the initial attribution.

An accumulating body of research indicates that both specialists and non-specialists are often poor decision makers (Slovic, Fischhoff, and Lichenstein 1977; Slovic and Lichtenstein 1971). One reason for this poor deci-

sion making is that individuals have difficulty making use of disconfirming information (Einhorn and Hogarth 1978). Another reason is that decision makers' initial attributions of causality are not necessarily accurate.

Reneau, Wong-On-Wing, and Pattison (1984) examined the relationship between on-line information retrieval systems and information search strategy. The results of their study suggest that on-line information retrieval systems may foster the use of hypothesis-confirming information search strategies, thus impairing decision effectiveness when the initial hypothesis of causality is not accurate. The primary implication of this study is that the current trend for reducing information overload (on-line information retrieval systems) should be re-examined, since this trend may impair decision effectiveness just as information overload may impair decision effectiveness.

It seems that effective computer-based decision support requires an information retrieval system capable of retrieving a subset of all available information, in order to reduce information overload, and supporting an information search strategy that considers all relevant information rather than merely hypothesis-confirming information. Expert system technology has the potential to provide this type of computer-based decision support. An information retrieval system with an expert component (i.e., a knowledge-based decision support system), should be able to support a preferred information search strategy in addition to reducing information overload.

If the use of an on-line information retrieval system does foster the use of a hypothesis-confirming information search strategy, thus impairing decision effectiveness when the initial hypothesis of causality is not accurate, then systems designers can affect decision quality via the manipulation of information search strategy. The purpose of this research is to study the effect of information retrieval system type on information search strategies and attributions of causality. An understanding of the effect of information retrieval system type will help enable systems analysts to develop guidelines for the design of computer-based decision aids (potentially with expert system components) that support preferred information search strategies.

2. BACKGROUND AND DECISION TASK

During the second half of the twentieth century, there has been considerable interest in hypothesis-testing as it relates to scientific reasoning. In *The Logic of Scientific Discovery*, Popper (1959) stated that the researcher should not attempt to confirm a hypothesis; rather, the researcher should attempt to maximize the likelihood of hypothesis disconfirmation. Platt (1964) extended Popper's prescriptions for scientific inquiry to include the "strong inference" strategy. Platt agreed with Popper's falsification position but felt that falsification should be attempted in the context of multiple alternative hypotheses. The topic of hypothesis-testing has been of interest in the psychological literature and in philosophy of science research. For example, the research findings of Wason (1960, 1968a, 1968b) have focused much attention on this issue. Of particular interest to psychologists have been the issues of hypothesis confirmation and hypothesis disconfirmation.

2.1 Confirmation Bias

Popper's falsification position and Platt's strong inference strategy both raise questions about the actual behavior of people investigating hypotheses. There is both anecdotal and experimental evidence to suggest that individuals appear to be biased toward seeking confirmatory evidence when testing hypotheses (Mynatt, Doherty, and Tweney 1977). Einhorn and Hogarth (1978) concluded that both highly trained individuals and novices are often poor decision makers. They felt that one reason for the fallibility of human judgment is that people have difficulty making use of disconfirming information. Individuals tend to neither search for nor use disconfirming evidence, and this is one of the main factors related to the difficulty people have in learning from experience.

Snyder (1981, p. 278) defined a confirmatory hypothesis-testing strategy as "the preferential soliciting of behavioral evidence whose presence would tend to confirm the hypothesis under scrutiny." He stated that people seem to believe that the presence of confirming evidence is particularly informative and relevant for testing a hypothesis and that a hypothesis survives according to its ability to accumulate confirming evidence.

The consequences of adopting a confirmatory hypothesis-testing strategy must be considered. Snyder and Swann (1978b) conducted an experiment to investigate one possible consequence: that a confirmatory hypothesis-testing strategy would actually generate the behaviors that would erroneously confirm the hypothesis. Their results indicated that, as a consequence of the hypothesis-testers' confirmatory strategies for gathering evidence, the targets did indeed behave in ways that would erroneously confirm the hypothesis.

Experimental evidence has been gathered to study people's hypothesis-testing strategies and their ability to use disconfirming evidence for making inferences. Some of the research used rule-discovery tasks in the experiments, while other research used truth-value tasks for investigations. A series of investigations by Snyder and his colleagues (Snyder 1981; Snyder and Campbell 1980; Snyder and Gangestad 1982; Snyder and Swann 1978a, 1978b), plus investigations by Trope and Bassok (1983), Harrison, West, and Reneau (1984), and Reneau, Wong-On-Wing, and Pattison (1984), involved testing hypotheses about people.

Taken together, these empirical investigations suggest that individuals regard confirming evidence as more relevant than disconfirming evidence when testing hypotheses; therefore, they show bias in accumulating confirming evidence rather than disconfirming evidence. The hypotheses may then be accepted or rejected based on information in which confirming evidence is over-represented and disconfirming evidence is under-represented.

2.2 Attribution Theory

By definition, a hypothesis-confirming information search strategy requires an initial hypothesis of causation (i.e., a hypothesis must exist before there can be an attempt to confirm it through the acquisition of additional information). Therefore, in order to investigate hypothesis-confirming information search strategies, it is necessary to have an initial hypothesis of causation. An explanation of how individuals generate initial hypotheses of causation is provided by attribution theory. This theory is an important social psychological approach that deals with the factors involved in people's attempts to understand observed events. These events include actions taken by the people themselves as well as actions taken by other people.

Kelley (1967) developed an ANOVA model to explain how people make causal attributions. His fundamental idea is the principle of covariation between causes and effects. He considers the important classes of possible causes to be persons, entities (things or environmental stimuli), and times (occasions or situations).

In order to validate their attributions--that is, to verify whether they have correctly linked cause and effect--individuals use three criteria:

1. *distinctiveness*--responding differently to different stimuli
2. *consensus*--how widespread the behavior or effect is among different people
3. *consistency*--how consistent the behavior or effect is over time

The pattern of high distinctiveness, high consensus, and high consistency is particularly likely to lead to a stimulus (external) attribution. The pattern of low distinctiveness, low consensus, and high consistency is likely to lead to a person (internal) attribution.

In addition to explaining how individuals generate initial hypotheses of causality, attribution theory also illustrates and explains the sources of bias, error, or imperfection that distort attributions of causality. The attribution theory literature documents sources of systematic bias that lead individuals to misinterpret events. For example, the "fundamental attribution error" indicates that an observer will tend to make an internal attribution relative to the observed behavior of an actor. In other words, the observer will underestimate the impact of external (situational/environmental) factors when determining a cause for the actor's behavior. The "actor-observer bias" deals with how actors and observers differ in their susceptibility to the fundamental attribution error. The actor tends to attribute his or her own behavior to external (situational/environmental) factors when the outcome of the behavior is negative; the observer tends to attribute the actor's behavior to internal factors, such as stable dispositions, attitudes, and personality traits. In other words, the type of attribution made (internal or external) may depend on the perspective (actor or observer) of the person making the attribution.

2.3 The Decision Task

This study involves the decision task depicted in the model shown in Figure 1. The participant is presented with a description of an event occurrence. This description contains distinctiveness, consistency, and consensus criteria, as defined by attribution theory. Based on these criteria, the participant is asked to make an initial attribution of causality. This initial attribution, following attribution theory, will be an internal attribution (the person described in the scenario is responsible for the event occurrence), an external attribution (the person is not responsible), or an attribution somewhere between these two extreme points.

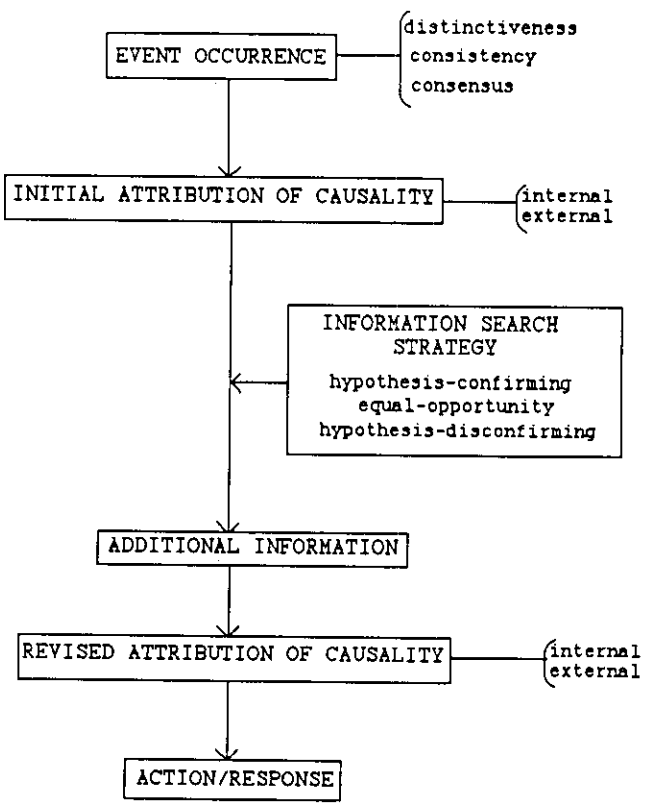

Figure 1. Model of the Decision Task

The initial attribution of causality is actually a hypothesis; thus, the experimental task involves testing a hypothesis about a person. In order to perform this hypothesis-testing task, the participant employs one of the information search strategies listed in the model in order to obtain additional information related to the event occurrence and to the person about whom the hypothesis has been formed. A computerized information retrieval system contains the additional information. Finally, the participant makes a revised attribution of causality for the event occurrence. This study deals with the decision task prior to the point of action/response.

3. HYPOTHESES

The basic research hypothesis of this study is that a knowledge-based decision support system (KBDSS), designed to search for and present both confirming and disconfirming evidence in a semi-structured decision task, will result in enhanced decision effectiveness. The following specific hypotheses represent ways in which the basic research hypothesis can be tested:

H1: A KBDSS, designed to search for and present both confirming and disconfirming evidence in a semi-structured decision task, will result in a greater change to a subject's initial hypothesis of causality than a conventional information retrieval system.

H2: A KBDSS, designed to search for and present both confirming and disconfirming evidence in a semi-structured decision task, will result in a change to a subject's initial hypothesis of causality from a more extreme attribution toward a more neutral attribution.

H1 pertains to the effect of the KBDSS on the subject's attribution of causality (i.e., did the attribution change) and H2 pertains to the direction of the change in the subject's attribution of causality. The two information retrieval systems used in this experiment contain a balance of relevant internal and relevant external information. An examination of all relevant information should lead the subject away from an extreme attribution (e.g., internal, as manipulated by one of the independent variables defined below) and toward a more neutral attribution (e.g., both internal and external causes for the problem). Therefore, a significant adjustment to the initial causal attribution, in the expected direction, is considered to be enhanced decision effectiveness. The implication if H1 and H2 are supported by the results of the experiment is that a more effective decision was made because more factors were considered by the decision maker.

4. VARIABLES

4.1 Independent Variables

4.1.1 Information Retrieval System

The basic information system is a detailed log of a large systems development project containing information on project personnel, hardware, end users, company standards, economic conditions, and organizational climate. In addition to factual information, the project log contains subjective/evaluative information, such as the project manager's assessments of the effectiveness of project personnel and assessments of the effect of economic conditions on the project team.

The project log was developed in two formats: a menu-driven information retrieval system which presents the user with a checklist of available topics so the user can decide what and how much information to retrieve, and a knowledge-based decision support system which itself decides what and how much information to retrieve after querying the user about the situation/problem to be investigated. The menu-driven information retrieval system is representative of on-line inquiry systems and data retrieval capabilities of decision support systems. The KBDSS is representative of the knowledge-based systems discussed by Blanning (1983), Henderson (1987), Courtney, Paradice, and Mohammed (1987), and Goul and Tonge (1987).

The knowledge-based decision support system asks the user a series of multiple-choice questions relevant to the subject at hand. The system continues to ask questions until it has enough information to determine which items in the project log should be viewed by the user. The system selects all relevant information for the user, using an equal-opportunity search strategy (i.e., it selects both confirmatory and disconfirmatory information). This confirmatory and disconfirmatory information is in fact a balance of items supporting internal causal attributions and items supporting external causal attributions. These classifications (internal and external) are based on the ratings of the information items by information systems faculty members and graduate students and by students similar to the subjects. The KBDSS was developed using the EXSYS Expert System Development Package.

The information in the project log is balanced in terms of items supporting internal causal attributions and items supporting external causal attributions. In addition, all information in the project log is confirmatory so that the treatment effect is as strong as possible. Therefore, if a subject uses a hypothesis-confirming information search strategy, the initial hypothesis of causality should be supported. If a subject uses a disconfirmatory or equal-opportunity strategy (the latter being the strategy of the KBDSS), the initial hypothesis should be changed, at least to some extent.

4.1.2. Scenario for Initial Hypothesis of Causality

Both the fundamental attribution error and the actor-observer bias suggest that individuals show bias when making attributions (e.g., those in the role of observer are more inclined to make internal attributions). The basic research hypothesis of this study is that the use of the KBDSS will result in a change to the initial causal attribution. Attribution theory suggests that the nature of the initial attribution of causality (internal or external) may be a confounding variable due to the attribution biases. Therefore, initial hypothesis, or attribution, of causality is a second independent variable in this experiment and the data is analyzed for an interaction effect between the information retrieval system type and the initial hypothesis of causality.

The participants each received one brief work history dealing with a particular employee on a systems development project and a problem associated with this employee. Three versions of the work history were developed using Kelley's (1967) ANOVA model of how people make causal attributions. One version suggested an internal initial hypothesis of causality for the subjects in the two internal hypothesis treatments. The second version suggested an external initial hypothesis. The third version of the work history did not suggest any initial hypotheses of causality (i.e., it was neutral).

4.2 Dependent Variable

4.2.1 Change in Causal Attribution from Pretest to Posttest

In order to measure the change in causal attribution from pretest to posttest, two dependent measures were used.

The first measure is a seven-point Likert-type scale. Using this scale, the subject indicated to what extent he or she felt the subordinate was responsible for causing the described problem. The scale was presented to the subject twice--first as a pretest measure and seemed as a posttest measure. The change in causal attribution from pretest to posttest was the dependent variable for analysis.

The second dependent measure was adapted from previous studies (Harrison, West, and Reneau 1984; Reneau, Wong-On-Wing, and Pattison 1984). It is a measure of internality (an internality scale) computed as:

ABILITY + EFFORT - TASK DIFFICULTY - LUCK

This measurement of internality was based on the Weiner (1974) suggestion that four possible attributions can be made for success and failure: ability and effort (internal attributions) and task difficulty and luck (external attributions). The subject was twice asked to allocate a total of 100 points to these four factors--once as a pretest measure and once as a posttest measure. The change in internality from pretest to posttest was the dependent variable for analysis.

5. RESEARCH METHOD

5.1 Subjects

One hundred ninety-five graduate and undergraduate business students participated in this study. Business students were chosen as participants in this research because there is evidence to support the use of students as surrogates for business people in experimental tasks (Ashton and Kramer 1980).

The subjects were enrolled in senior level and graduate information systems and accounting information systems classes at a large urban university. The subjects had studied the systems development process and thus were as homogeneous as possible on the understanding of the task scenario: work on a systems development project. Forty-three percent of the subjects had information systems work experience and 55 percent had experience in supervising subordinates. Fifty-six percent of the participants were male and 58 percent were at least 23 years old.

5.2 Experimental Task

The participant was asked to assume the role of the manager of a large systems development project and was told that there had been some problems on this project. The participant was given a written description of one of these problems, a problem involving a senior programmer whose programming group had failed to meet a deadline for a program critical to the first phase of the system under development. As a result of the missed deadline, a significant number of overtime hours had to be authorized, having a serious impact on the manager's (participant's) budget. The participant's specific task was to determine to what extent the subordinate was responsible for causing the described problem; that is, the participant was asked to make an attribution of causality.

5.3 Experimental Procedures

Each participant took part in one of ten experimental sessions, each lasting about 30 minutes. The sessions were held in a microcomputer lab, where each participant was assigned to an individual workstation. The subjects using the KBDSS were separated from the subjects using the conventional information retrieval system. Subjects were randomly assigned to one of six treatment groups, based on two types of information retrieval system and three scenarios for initial hypothesis of causality.

The experimental procedure was as follows:

1. Subjects listened to a standard introductory script read by the administrator of the experiment.

2. Subjects read written instructions and one of the three versions of the problem description.

3. Following the instructions appearing on the terminal screens, subjects answered questions about the cause of the described problem. These questions were the dependent measures, used initially as the pretest measures.

4. Subjects read instructions on the use of the information retrieval system and then accessed this system in order to investigate the problem with the subordinate on the systems development project.

5. The participants using the KBDSS were asked by the system to answer questions pertaining to the subject of the investigation. The KBDSS then displayed the relevant information items from the project log. The subjects using the conventional information retrieval system were dealing with a menu-based system that permitted them to decide which information and how much information to view. Due to the size of the database and due to the fact that it contained irrelevant information, very few subjects viewed all of the screens (the system automatically recorded each screen accessed). This is consistent with a decision problem where information load makes comprehensive search barely conceivable, if not impossible. Such decision problems are the focus of KBDSS research.

6. After the investigation was completed, the subjects answered the same questions asked before the information system was queried. These questions were now the posttest measures.

7. The subjects completed demographic data questionnaires.

8. Debriefing sessions were held after all ten experimental sessions were completed.

6. RESULTS

The results of this study are consistent with other studies referenced above that show evidence of confirmation bias in hypothesis-testing. The data analysis, for the 70 subjects for whom the manipulation of initial attribution of causality was successful, indicates that these subjects basically retained their initial hypotheses (attributions) after obtaining additional information, although in some cases (e.g., when the KBDSS was used) these attributions became less extreme. Table 1 illustrates this point.

Table 1.

Cell Means for System x Attribution Interaction (n = 70)

	Likert-type Scale		Internality Scale	
	Conventional	KDBSS	Conventional	KDBSS
Internal	2.7500	3.4444	47.5000	40.5556
External	4.2857	4.5000	-3.7143	-22.5000
Neutral	4.0000	2.8000	35.7143	32.0000

For example, on the Likert-type scale, the internal attribution scenario group subjects all have pretest scores between 1 and 3 inclusive (1 is an extreme internal attribution and 4 is a neutral attribution). The mean posttest Likert-type score for the internal attribution group using the conventional information retrieval system is 2.7500 and the mean score for the group using the KBDSS is 3.4444 (see Table 1). Both posttest scores are below the neutral score of 4. The external attribution scenario group subjects all have pretest Likert-type scores between 5 and 7; their posttest mean scores are 4.2857 (conventional) and 4.5000 (KBDSS). These scores are both above the neutral point (although less extreme than the scores for the internal attribution group subjects). Similar results are found on the internality scale. Positive numbers on this scale indicate internal attributions and negative numbers indicate external attributions. The posttest cell means for both internal attribution scenario groups are positive (47.5000 and 40.5556) and the posttest cell means for both external attribution scenario groups are negative (-3.7143 and -22.5000) (see Table 1).

The hypotheses for this research experiment involved adjustments to initial attributions of causality. The subject read a brief work history that suggested an initial causal attribution. The dependent variable measured the change in causal attribution from pretest to posttest.

An examination of pretest scores indicated that some subjects did not make the expected initial causal attribution (i.e., the causal attribution suggested by the work history). Therefore, data analyses were performed twice, first using all subjects with usable responses participating in the experiment (188 out of 195 subjects) and second using only the 70 subjects who made the expected initial causal attributions on both pretest scales.

The results of the data analysis provide minimal support for the first hypothesis, that a KBDSS will result in a greater change to a subject's initial hypothesis of causality than a conventional information retrieval system. There are no significant results on any statistical tests involving the internality scale. The ANOVA's for the Likert-type scale show no significant main or interaction effects for system type. However, the calculations of the cell means do suggest that subjects in the internal and external attri-

bution scenario groups using the KBDSS do change their initial hypotheses of causality more than their counterparts using the conventional information retrieval system.

The second hypothesis, that a KBDSS will result in a change to a subject's initial hypothesis of causality from a more extreme attribution toward a more neutral attribution, is supported by all of the calculations of cell means for the internal and external attribution scenario groups on both the Likert-type and the internality scales. Six of the twelve *t* tests for the differences between pairs of population means also support the hypothesis.

The first hypothesis pertains to the effect of the KBDSS on the subject's attribution of causality (i.e., did the attribution change) and the second hypothesis pertains to the direction of the change in the subject's attribution of causality. Since the two information retrieval systems used in this experiment contain a balance of relevant internal and relevant external information, an examination of all relevant information should have led the subject away from an extreme attribution (e.g., internal, as manipulated by the scenario for initial attribution of causality read by the subject) and toward a more neutral attribution of causality. Therefore, a significant adjustment to the initial causal attribution, in the expected direction, is considered to be enhanced decision effectiveness. The implication if the hypotheses are supported by the results of the experiment is that a more effective decision was made because more factors were considered by the decision maker. Since the first hypothesis did receive minimal support and since the second hypothesis received more substantial support, there is some indication that more factors were considered and thus a more effective decision was made.

These results taken together indicate some support for the conceptual hypothesis that the use of a knowledge-based decision support system, designed to search for and present both confirming and disconfirming evidence in a semi-structured decision task, will result in enhanced decision effectiveness. The cell means calculated in the various data analyses generally provide support for the hypotheses, but the statistical tests are usually not significant at the .05 level.

7. DISCUSSION

This study sought to build on the findings of the Reneau, Wong-On-Wing, and Pattison (1984) study which examined the relationship between on-line information retrieval systems and information search strategy. The results of their study suggest that on-line information retrieval systems may foster the use of hypothesis-confirming information search strategies, thus impairing decision effectiveness when the initial hypothesis of causality is not accurate. The primary implication of the study is that the current trend for reducing information overload (on-line information retrieval systems) should be re-examined, since this trend may impair decision effectiveness just as information overload may impair decision effectiveness.

The intent of the present study was to suggest and test a possible solution to the problem indicated by the earlier study. That is, the present study sought to discover if the use of an information retrieval system with an expert component would be able to support a preferred information search strategy in addition to reducing information overload. In other words, this study sought to discover if the use of a knowledge-based decision support system could enhance decision effectiveness by compensating for the user's (potential) confirmation bias in a hypothesis-testing task.

The results of this study indicate some support for the basic research question, but they indicate stronger support for the contentions in the psychological literature that confirmation bias is very difficult to overcome. For example, Wason (1968a, p. 313), in a review of experiments on rule-discovery tasks, concluded:

> In spite of the small samples used in these experiments, there would appear to be compelling evidence to indicate that even intelligent individuals adhere to their own hypotheses with remarkable tenacity when they can produce confirming evidence for them.

In the present study, even the subjects using the KBDSS received confirming evidence for their initial hypotheses (the KBDSS presented a balance of confirming and disconfirming evidence). Perhaps this confirming evidence helped the subjects "adhere to their own hypotheses."

Alloy and Tabachnik (1984, p. 119) discussed the joint influence of expectation-based processing and data-based processing when using covariation information to make causal attributions. They said if:

> generalized expectations and situational information converge on different causal attributions, the lay attributor must either reinterpret, misremember, or discount contradictory situational information and make an attribution in line with generalized expectations, or set aside strongly held beliefs about causality in favor of situational information instead.... people faced with this dilemma generally reinterpret situational information favoring the attribution suggested by their generalized expectations.

This supports the notion that confirmation bias is very difficult to overcome and helps explain why the *equal-*

opportunity knowledge base used in the present study did not have a major impact on initial attributions of causality.

There is another possible explanation for the relative lack of significance of the type of information retrieval system. Some of the empirical studies involving the truth-value task indicated that when the context of the task is concrete, subjects tend to perform better. Realism apparently induces insight into the logical structure of a task and thus improves performance.

In the present study, some of the results suggest that subjects using the conventional information retrieval system may have searched for and/or attended to disconfirming information. Perhaps the realism of the experimental task, coupled with the fact that the subjects were familiar with the task scenario (a systems development project), affected the results of the experiment. That is, perhaps the subjects using the conventional information retrieval system were less inclined to disregard disconfirming evidence than was expected due to the degree of realism in the experimental task. If this is the case, it is understandable that the impact of the type of information retrieval system would be diminished. A preliminary analysis of the data collected in this study indicated a problem with the experimental manipulation of the subjects' initial attributions of causality. Only 70 of the 188 subjects made the expected initial causal attributions on both the pretest Likert-type scale and the pretest internality scale. The most serious aspect of this problem with the manipulation of initial attribution of causality is associated with the subjects in the external attribution scenario groups. Only 15 of 62 subjects reading the external scenario made external attributions of causality on both pretest scales. This problem was not anticipated for two reasons. First, the scenarios for all three attribution groups exactly followed the formats used in the previous studies. Second, a pilot test was conducted prior to any of the experimental sessions and the expected initial causal attributions were made by the majority of pilot test subjects.

It is possible that the external attribution scenario subjects tended to make internal attributions of causality on the pretest measures due to the attribution biases previously described. The external group subjects may have made the fundamental attribution error; that is, they may have underestimated the impact of situational factors and overestimated the role of dispositional factors when evaluating the problem concerning the subordinate on the systems development project. Since all of the subjects were in the role of *observer* and the subordinate was the *actor* in the problem scenario, the subjects may also have made erroneous attributions of causality due to the actor-observer bias, that is, due to the tendency of the observer to attribute the actor's behavior to internal factors, such as stable dispositions, attitudes, and personality traits.

There is another attribution bias which may have affected the experimental manipulation. McArthur (1972) noted that people consistently underutilize consensus information (relative to consistency and distinctiveness information) when making causal attributions. In the present study, the external attribution subjects may have underutilized the information indicating that the subordinate's peers were also missing deadlines on this systems development project (the consensus information).

Mynatt, Doherty, and Tweney (1977) reported a problem with an experimental manipulation similar to the problem experienced in the present study. They also attempted to manipulate their subjects' initial hypotheses. Forty-five subjects participated in their experiment, but only 20 of them developed the correct initial hypothesis.

The cognitive difficulties associated with disconfirming information have been well documented. Formal training in the use of disconfirming evidence would seem to be a necessary but not sufficient condition for overcoming confirmation bias, since highly trained individuals (e.g., statisticians and scientists) make mistakes similar to those of novices. Future research should be aimed at providing concrete methods, such as computer-based decision aids, to compensate for confirmation bias in hypothesis-testing.

8. REFERENCES

Alloy, L. B., and Tabachnik, N. "Assessment of Covariation by Humans and Animals: The Joint Influence of Prior Expectations and Current Situational Information." *Psychological Review*, Vol. 91, 1984, pp. 112-149.

Ashton, R. H., and Kramer, S. S. "Students as Surrogates in Behavioral Research: Some Evidence." *Journal of Accounting Research*, Vol. 18, 1980, pp. 1-15.

Blanning, R. W. "What is Happening in DSS?" *Interfaces*, Vol. 13, No. 5, 1983, pp. 71-80.

Courtney, J. F., Jr.; Paradice, D. B.; and Mohammed, N. H. A. "A Knowledge-Based DSS for Managerial Problem Diagnosis." *Decision Sciences*, Vol. 18, 1987, pp. 373-399.

Einhorn, H. J., and Hogarth, R. M. "Confidence in Judgment: Persistence of the Illusion of Validity." *Psychological Review*, Vol. 85, 1978, pp. 396-416.

Goul, M., and Tonge, F. "Project IPMA: Applying Decision Support System Sesign Principles to Building Expert-Based Systems." *Decision Sciences*, Vol. 18, 1987, pp. 448-467.

Harrison, P. D.; West, S. G.; and Reneau, J. H. *Initial Attributions and Information-Seeking by Superiors and Subordinates in Production Variance Investigations*. (Working Paper). Tempe, AZ: Arizona State University, 1984.

Henderson, J. C. "Finding Synergy between Decision Support Systems and Expert Systems Research." *Decision Sciences*, Vol. 18, 1987, pp. 333-349.

Kelley, H. H. "Attribution Theory in Social Psychology." In D. Levine (ed.), *Nebraska Symposium on Motivation*. Lincoln, NE: University of Nebraska Press, 1967.

McArthur, L. A. "The How and What of Why: Some Determinants and Consequences of Causal Attribution." *Journal of Personality and Social Psychology*, Vol. 22, 1972, pp. 171-193.

Mynatt, C. R.; Doherty, M. E.; and Tweney, R. D. "Confirmation bias in a Simulated Research Environment: An Experimental Study of Scientific Inference." In P. N. Johnson-Laird and P. C. Wason (eds.), *Thinking: Readings in Cognitive Science*, Cambridge: Cambridge University Press, 1977, pp. 315-325.

Platt, J. R. "Strong Inference." *Science*, Vol. 146, 1964, pp. 347-353.

Popper, K. R. *The Logic of Scientific Discovery*. New York: Basic Books, 1959.

Reneau, J. H.; Wong-On-Wing, B.; and Pattison, D. D. *Hypothesis Confirming Information Search Strategies and Inquiry-Type Computerized Information Systems*. Working Paper, Arizona State University, Tempe, 1984.

Slovic, P.; Fischhoff, B.; and Lichtenstein, S. "Behavioral Decision Theory." *Annual Review of Psychology*, Vol. 28, 1977, pp. 1-39.

Slovic, R., and Lichtenstein, S. "Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment." *Organizational Behavior and Human Performance*, Vol. 6, 1971, pp. 649-744.

Snyder, M. "Seek, and Ye Shall Find: Testing Hypotheses about other People." In E. T. Higgins, C. P. Herman, and M. P. Zanna (eds.), *Social Cognition: The Ontario Symposium*, Hillsdale, NJ: Lawrence Erlbaum, Vol. 1, 1981, pp. 277-303.

Snyder, M., and Campbell, B. H. "Testing Hypotheses about other People: The Role of the Hypothesis." *Personality and Social Psychology Bulletin*, Vol. 6, 1980, pp. 421-426.

Snyder, M., and Gangestad, S. "Hypothesis Testing Processes." In J. H. Harvey, W. J. Ickes, and R. F. Kidd (eds.), *New Directions in Attribution Research*, Hillsdale, NJ: Lawrence Erlbaum, Vol. 3, 1982.

Snyder, M., and Swann, W. B., Jr. "Behavioral Confirmation in Social Interaction: From Social Perception to Social Reality." *Journal of Experimental Social Psychology*, Vol. 14, 1978a, pp. 148-162.

Snyder, M., and Swann, W. B., Jr. "Hypothesis-Testing Processes in Social Interaction." *Journal of Personality and Social Psychology*, Vol. 36, 1978b, pp. 1202-1212.

Sprague, R. H. "A Framework for the Development of Decision Support Systems." *MIS Quarterly*, December 1980, pp. 1-26.

Sprague, R. H., and Carlson, E. D. *Building Effective Decision Support Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1982.

Trope, Y., and Bassok, M. "Information Gathering Strategies in Hypothesis-Testing." *Journal of Experimental Social Psychology*, Vol. 19, 1983, pp. 560-576.

Wason, P. C. "On the Failure to Eliminate Hypotheses in a Conceptual Task." *Quarterly Journal of Experimental Psychology*, Vol. 12, 1960, pp. 129-140.

Wason, P. C. "On the Failure to Eliminate Hypotheses: A Second Look." In P. N. Johnson-Laird and P. C. Wason (eds.), *Thinking: Readings in Cognitive Science*, Cambridge: Cambridge University Press, 1968a, pp. 307-314.

Wason, P. C. "Reasoning About a Rule." *Quarterly Journal of Experimental Psychology*, Vol. 20, 1968b, pp. 273-281.

Wason, P. C. "Self-Contradictions." In P. N. Johnson-Laird and P. C. Wason (eds.), *Thinking: Readings in Cognitive Science*, Cambridge: Cambridge University Press, 1977, pp. 114-128.

Weiner, B. "Achievement and Motivation as Conceptualized by an Attribution Theorist." In B. Weiner (ed.), *Achievement Motivation and Attribution Theory*. Morristown, NJ: General Learning Press, 1974.