MIS Quarterly

RESEARCH ARTICLE

THE DIFFERENTIAL USE AND EFFECT OF KNOWLEDGE-BASED SYSTEM EXPLANATIONS IN NOVICE AND EXPERT JUDGMENT DECISIONS[1]

By: Vicky Arnold
Dixon School of Accounting
College of Business Administration
University of Central Florida
Orlando, FL 32816
U.S.A.
Vicky.Arnold@bus.ucf.edu

Nicole Clark
School of Computing
University of Tasmania
Hobart, Tasmania 7001
AUSTRALIA
nclark@utas.edu.au

Philip A. Collier
Department of Accounting and Business Information Systems
University of Melbourne
Melbourne, Victoria 3010
AUSTRALIA
pcollier@unimelb.edu.au

Stewart A. Leech
Department of Accounting and Business Information Systems
University of Melbourne
Melbourne, Victoria 3010
AUSTRALIA
saleech@unimelb.edu.au

Steve G. Sutton
Dixon School of Accounting
College of Business Administration
University of Central Florida
Orlando, FL 32816
U.S.A.
Steve.Sutton@bus.ucf.edu

Abstract

Explanation facilities are considered essential in facilitating user interaction with knowledge-based systems (KBS). Research on explanation provision and the impact on KBS users has shown that the domain expertise affects the type of explanations selected by the user and the basis for seeking such explanations. The prior literature has been limited, however, by the use of simulated KBS that generally provide only feedback explanations (i.e., ex post to the recommendation of the KBS being presented to the user). The purpose of this study is to examine the way users with varying levels of expertise use alternative types of KBS explanations and the impact of that use on decision making. A total of 64 partner/manager-level and 82 senior/staff-level insolvency professionals participated in an experiment involving the use of a fully functioning KBS to complete a complex judgment task. In addition to feedback explanations, the KBS also provided feedforward explanations (i.e., general explanations during user input about the relationships between information cues in the KBS) and included definition type explanations (i.e., declarative-level knowledge). The results show that users were more likely to adhere to recommendations of the KBS when an explanation facility was available. Choice patterns in using explanations indicated that novices used feedforward explanations more than experts did, while experts were more likely than novices to use feedback explanations. Novices also used more

[1]Ron Weber was the accepting senior editor for this paper. Fran Ackerman was the associate editor. Shirley D. Gregor, Dan O'Leary, and E. H. J. Vaassen served as reviewers.

declarative knowledge and initial problem solving type explanations, while experts used more procedural knowledge explanations. Finally, use of feedback explanations led to greater adherence to the KBS recommendation by experts—a condition that was even more prevalent as the use of feedback explanations increased. The results have several implications for the design and use of KBS in a professional decision-making environment.

Keywords: Explanations, explanation use, knowledge-based systems, expert systems, intelligent systems

Introduction

Knowledge-based systems (KBS) and other forms of intelligent systems are frequently used to capture key knowledge and expertise from individuals within an organization. They make organizational knowledge usable by other decision makers (Gregor and Benbasat 1999; Hayes-Roth 1997). While intelligent systems have not lived up to expectations (Duchessi and O'Keefe 1992), organizations continue to believe KBS can enhance organizational effectiveness (Gill 1995; Gregor and Benbasat 1999; Wanninger 1998). Accordingly, researchers have focused on methods for improving the usability of KBS through enhanced designs, such as integrating explanation facilities (Arnold et al. 2004a; Dhaliwal and Benbasat 1996; Gregor and Benbasat 1999).

Explanations are generally considered essential to facilitate user interaction with a KBS (Swartout 1987). They can facilitate use of, and reliance on, a system by a decision maker if the system can build user trust and perhaps even "argue" its recommendation. Long-term trust of a system is developed through quality output, but short-term trust comes from the system's ability to explain the rationale underlying its recommendations (Swartout 1983; Ye and Johnson 1995). As a result, explanations have become a core component of most KBS designed for use by professional decision makers (Hayes-Roth and Jacobstein 1994; Mao and Benbasat 2000).

Substantive research has taken place on developing alternative explanation types (for an overview, see Dhaliwal and Benbasat 1996), understanding different users' preferences for alternative explanation types (Hsu and Steinbart 1997; Ye and Johnson 1995), and understanding how various users navigate and integrate KBS explanations (Mao and Benbasat 2000). Nonetheless, Dhaliwal and Benbasat (1996) note that little is known about the behavioral effects of the use of KBS explanations due to the paucity of research examining KBS use as a decision aid for user judgments in working situations. Researchers within the domain have widely noted the specific need for research examining the effect of KBS on decision-maker performance (Dhaliwal and Benbasat 1996; Gregor and Benbasat 1999; Hsu and Steinbart 1997; Mackay and Elam 1992; Ye and Johnson 1995).

This study examines the ways in which users with varying levels of expertise use alternative types of KBS explanations and the impact of such use on users' judgments. Prior research has studied the use of various subsets of explanation types by KBS users of varying knowledge levels. This study examines how professionals with high levels of task experience use a KBS that has a full range of explanation types (i.e., definition, rule-trace, justification, and strategic) and delivery modes (i.e., feedforward and feedback) to provide various levels of knowledge assistance (i.e., declarative, initial problem solving strategy-based, and procedural level).

Understanding how professional decision makers with varying levels of expertise use KBS explanations when a full range of options are available and how this usage impacts their judgments is important for organizations considering design and deployment of KBS. Prior research has demonstrated that the type of explanations selected by users differs based on available options. Further, Gregor (2001) notes that, other than her study, research on KBS explanation use has not employed fully functional systems that would likely engage the user and affect usage behavior, and that no prior studies have examined explanation use within a cooperative problem-solving environment (as is normally found with KBS use by professionals). Additionally, prior KBS studies have not used professional decision makers to make high-level, complex judgments, limiting the understanding of how KBS are used and the related impact on decision making processes for systems deployed in corporate environments.

The research reported in this paper contributes to the growth of empirical research on KBS explanations in three ways. First, it is grounded in the basic findings that novices and experts perform differently in explanation use and inference for judgment formation. To assure expertise is reasonably captured, 64 partner and manager (expert) and 82 staff and senior (novice) insolvency practitioners from major professional services firms participated in KBS training sessions that concluded with the completion of an extensive case analysis. Second, two fully functional KBS for insolvency decision making were used (one with and one without explanations) during the training and experimental sessions to examine the effect of the explanation facility on user judgments. One group used the original version of the software (no explanation facility) and formed the baseline used for comparison. The second group used the expanded version containing a fully functional explanation facility. This allowed participants to become engaged with the KBS through both the data input and recommendation processes of the aid. For participants with the explanation facility-enabled version, access to explanations was available during data input as well as during recommendation evaluation. By comparing the information provided by the group using the KBS without explanations with the group using the KBS with explanations, the impact on decision making can be narrowed down to just the impact of the explanation facility. Third, all user input and interactions with the KBS were traced via a replay process tracing facility (Arnold et al. 2001). Information was gathered on participants' judgments and explanation type, mode, and knowledge components; this was done without interfering in the decision-making processes through an experimental intervention such as protocol analysis. The goal of the design is to focus on the actual *theory-in-use* by experienced professionals rather than focusing on what such professionals may report as their *espoused theory* because the two often are incongruent (Argyris and Schön 1974).

The remainder of this paper consists of the following sections. The first integrates the research with extant theories and provides a conceptual basis for the hypotheses. The second section presents the experimental design, and the third section presents the results. The fourth and final sections document the implications of the results for management practices and future research and the limitations of the study.

Theory and Hypotheses Development

To understand the impact of explanations on decision making when using a KBS, several dimensions must be considered. First, the issues related to design and use of explanation facilities are important to judgment processes. Second, the various factors that impact overall usage of KBS explanations will also impact a user's propensity toward use of and adherence to recommendations provided by a KBS. Third, the user's expertise may impact the types and number of explanations accessed, as well as adherence to recommendations. Figure 1 provides a schematic of the theorized impacts of expertise, explanation mode, explanation knowledge component, and short-term learning on decision-making processes. While other effects beyond the scope of this study may occur, the effects of particular interest in the current study are the bolded items in Figure 1.

Fundamental to understanding the impact of KBS explanations on decision-making processes is the recognition that the provision of different explanation types (definition, rule-trace, justification, and strategic) in alternative modes (feedforward versus feedback) fulfill the need for different levels of knowledge support. Dhaliwal and Benbasat (1996) provide the original theoretical basis for integrating feedforward/feedback modes with different explanation types. Their model prescribes six forms of explanation combinations based on feedforward versus feedback mode across the latter three types of explanations (rule-trace, justification, and strategic).[2] Arnold et al. (2004a) expand on these six forms of explanation to include the other two combinations: feedforward/definition and feedback/definition.[3] They also describe how to access explanations through a flexible hypertext system which, in particular, enables access to feedback explanations for all the subgoals as well as the final recommendations in the system.[4] Arnold et al. (2004a) present a full instantiation of the eight explanation forms embedded within a working KBS. Figure 2 provides a specific description of each explanation type in feedforward versus feedback mode and aggregated by component knowledge levels (declarative, initial problem solving, and procedural).

The model presented in Figure 1 aligns closely with the theory provided by Gregor and Benbasat (1999), particularly with the following propositions:

- Explanations will be used when the user lacks the knowledge needed to contribute to problem solving. The knowledge could be definitional knowledge or knowledge of a problem solving procedure (i.e., use of feedforward as a means of declarative and initial problem solving level knowledge acquisition).

- Novices will use explanations for learning (short- and long-term) more than experts (i.e., use of feedforward as a means of declarative and initial problem solving level knowledge acquisition).

- Experts are more likely than novices to use explanations for resolving anomalies (disagreement) and for verification (i.e., feedback as a means of assessing procedural level knowledge).

The key is to examine these propositions within the context of a decision-making scenario where the behavior of the decision maker can be observed while working in a cooperative problem solving mode with the system.

Judgment Effects

From a design perspective, Ye and Johnson (1995) make the case for using Toulmin's (1958) model of argumentation as a foundation for developing rule-trace, justification, and strategic explanations. They study use of these explanations in a feedback mode through a simulated KBS and find evidence to support their case. Gregor and Benbasat also support use of Toulmin's model of argumentation in developing explanations for use in KBS. The basis for the model is to provide a clear argumentation strategy: when "listeners" hear a cognitively appealing rationale, they will be more convinced that the position argued is correct.

[2]These three types have been categorized as *process type* explanations in some prior literature.

[3]Definition-type explanations have also been referred to as *answer help* or *terminological* in some prior literature.

[4]Feedback explanations are optionally accessible through hypertext links prior to a recommendation by INSOLVE (the KBS used in this research). As such, feedback explanations do not always report how recommendations were made; rather they describe subprocesses by which recommendations are formulated. This generalizes the nature of feedback explanation compared with much prior research where they can only be accessed after a recommendation is reached.

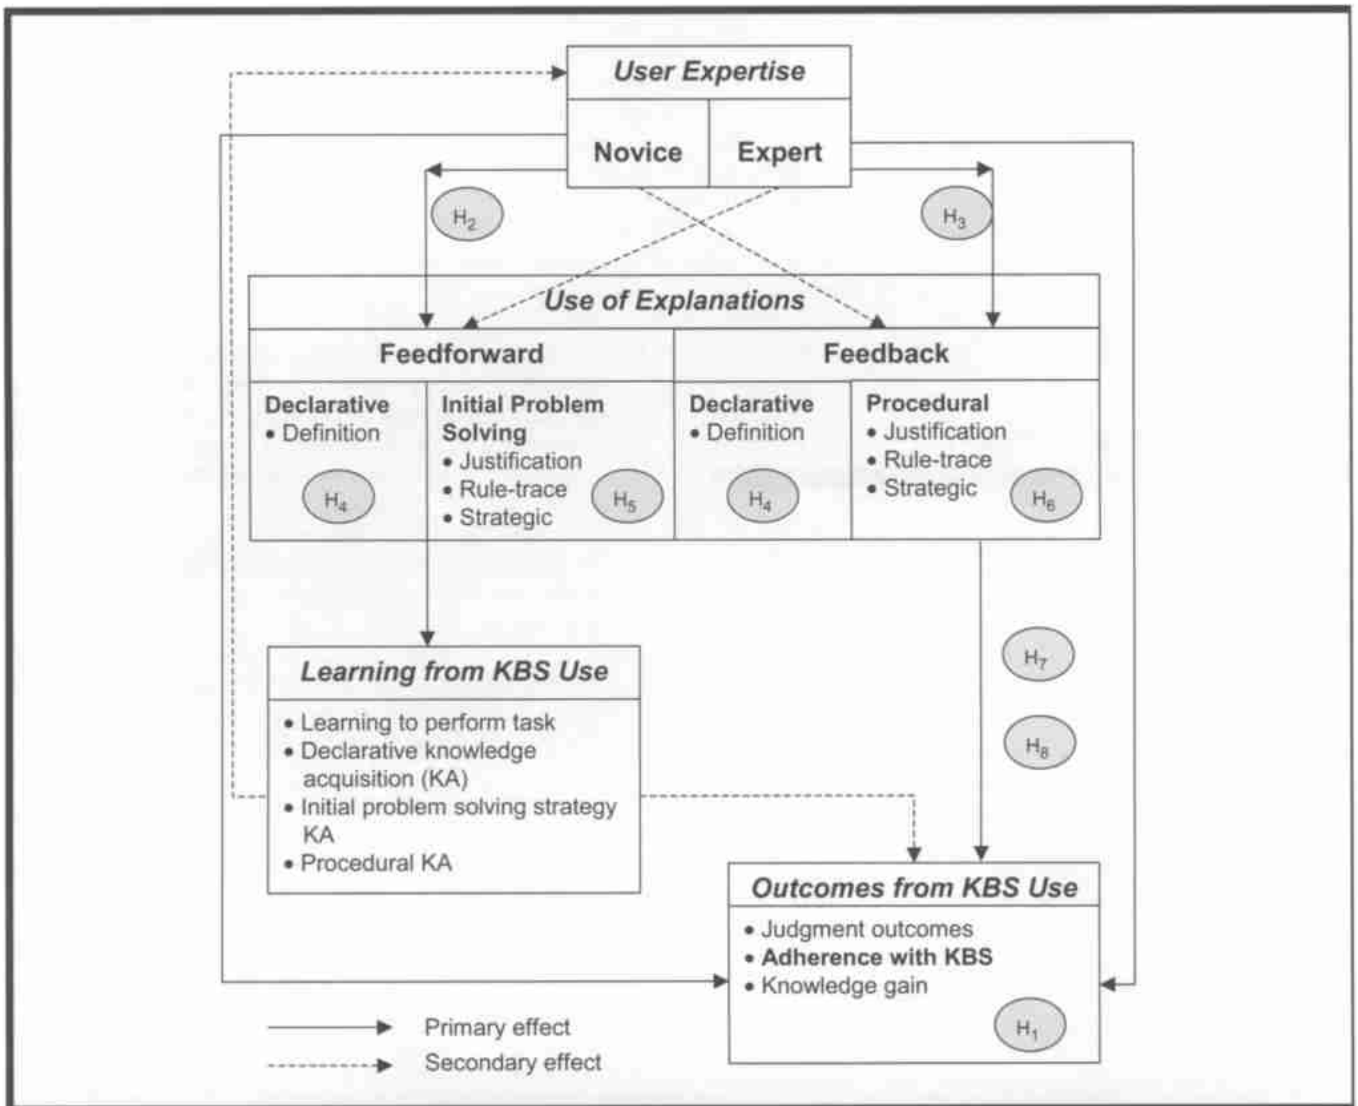

Figure 1. Research Model

While Toulmin's model of argumentation is intended to convince users of the correctness of the KBS's recommendation, novices are more likely to be convinced than experts. This expected effect on novices is consistent with the theory of technology dominance (Arnold and Sutton 1998), which posits that novices who have a limited capability to make a judgment within a given domain will be highly likely to rely on the recommendations of a KBS with or without explanations. The theory implies that the KBS, rather than the user, dominates the decision-making process. The impact on expert users is less clear. The theory posits that a KBS will influence an expert's decision, but an expert may not adhere to the recommendations of the aid. Rather, the theory advocates use of cooperative problem-solving systems that allow the user and KBS to interact and share control of the decision-making process. Evidence of such novice and expert behavior has been found in recent studies (Arnold et al. 2004b; Masselli et al. 2002; Noga and Arnold 2002). The presence of explanations should enhance the impact of a KBS on experts' and novices' decisions. Moreover, if structured around Toulmin's argumentation model, the KBS will provide stronger argumentation in support of recommendations. Thus, successful implementation of a Toulmin-based explanation facility should also pull both novice and expert decision makers' judgments toward the KBS's recommendation. This leads to the first hypothesis.

H1: KBS users will be more likely to adhere to the recommendation of a KBS when explanations are provided.

Explanation Type	Feedforward Explanations	Feedback Explanations
Definition	For each information cue or question, a feedforward definition explanation provides an elaboration of the terminology used. This may draw on dictionaries, standards, expert insights, and examples relevant to the intent of the information cue or question.	When reviewing a conclusion and its rationale, a feedback definition provides an elaboration of the terminology used. Definition explanations take the same form regardless of presentation in feedforward or feedback mode.
Rule-trace	These explanations are straightforward, natural-language representations of knowledge fragments (rules) in terms of the relationships between component information cues that make up the rules. If a user asks *why* information is needed, the KBS provides an explanation for all the rules that use the information cue in formulating a conclusion or recommendation.	When reviewing a conclusion or recommendation, a feedback rule-trace provides a straightforward, natural-language representation of the rule in terms of the procedure that is used to combine the component information cues to establish a value for the conclusions or recommendation. Any available values for the component information cues and conclusions are also disclosed.
Justification	These explanations amplify rule-trace explanations by supplying causal or other relevant information and examples that form the basis for knowledge fragments. Such additional information may have been acquired during the knowledge-acquisition process or from a subsequent information gathering and synthesis process.	A feedback justification provides a rich description of the rationale for the procedure that is used to combine the component information cues to establish a value for the conclusion or recommendation. Any available information cues and conclusions are also disclosed. As such, feedback justifications should provide a well-reasoned and compellingly argued case for the reasoning performed by the KBS.
Strategic	These explanations describe the problem solving method used by the KBS. With feedforward strategic explanations, a user should be able to surmise how a KBS solves problems and derives its recommendations.	A feedback strategic explanation enables a KBS to explain the problem-solving method that has been used in any processing already completed to derive the conclusions or recommendations provided by the system in a specific case.

Declarative Level Knowledge | **Initial Problem Solving Strategies** | **Procedural Level Knowledge**

Figure 2. Explanation Combinations

Feedforward Versus Feedback Explanation Usage

Several factors that appear to impact explanation preferences and use are also likely to affect decision making outcomes. One common attribute of interest in prior studies has been expertise (Gregor and Benbasat 1999). While Ye and Johnson did not directly examine expertise, they did examine their data for differences between experts and novices in *post hoc* tests. While the overall group of participants in that study preferred explanations that provide the rationale for a KBS's recommendation, the small subset of expert participants sought explanations describing how the decision aid used the underlying data to make a recommendation. Subsequently, Mao and Benbasat (2000) focused on expert versus novice differences in using and processing explanations. They anticipated that the substantial differences in cognitive structures and processes used by experts versus novices would carry over to use of a KBS. Their results confirmed their expectation that novice users focused heavily on explanations that aided in understanding the reasoning process. On the other hand, experts who lacked confidence in their own concurring judgment tempered such use with a similar interest in verifying the system's recommendation. A major theoretical contribution of their work was that domain expertise could influence explanation use.

Multiple dimensions can explain the differences between novice and expert users. At its basic level is the mode of delivery: feedforward versus feedback. Feedforward refers to declarative explanations and feedback refers to procedural explanations that are typically provided after a recommendation (Dhaliwal and Benbasat 1996). A feedforward explanation is described as declarative because it provides an explanation about how inputs to a KBS are used in terms of relevant information cues and their relationships. Conversely, a feedback explanation is tailored to describe processing in terms of

> how a decision was made, not made or remains unknown, with additional disclosures of the case-specific values of any dependent information cues that led to the current state of the decision. In particular, if one line of reasoning led to a decision being made, then only that line of reasoning will be reported in a feedback explanation; unsuccessful lines of reasoning will remain undisclosed (Arnold et al. 2004a, p. 10).

Dhaliwal and Benbasat theorize that while novices will be more likely than experts to access explanations, they will focus primarily on feedforward explanations that provide the requisite declarative knowledge and initial problem-solving strategies necessary to understand the decision-making task modeled by the KBS. Feedforward explanations assist users to understand why the KBS is requesting certain information and how that information will be used by the system to make recommendations. Hence, feedforward explanations assist users with limited knowledge in a domain to understand better a KBS's queries and to respond to these queries with higher-quality inputs. Conversely, as noted earlier, experts are expected to use more feedback explanations as they attempt to understand how or why a KBS produced a given recommendation. This perspective is consistent with Gregor and Benbasat's propositions predicting that novices will pursue feedforward explanations, and experts will use explanations more for resolving anomalies and for verification (i.e., feedback). Similarly, Argyris and Schön (1974) suggest that experienced professionals become reluctant to accept a theory of action that differs from their own and will question the inconsistent theory with which they are presented. Given a recommendation from a KBS that is inconsistent with the decision that the experienced user would make, the user should pursue feedback explanations to reconcile why the contrary recommendation was made. This leads to the second and third hypotheses.

H2: When using a KBS with explanation facilities, novices will choose more feedforward explanations than experts.

H3: When using a KBS with explanation facilities, experts will choose more feedback explanations than novices.

While Dhaliwal and Benbasat emphasize the importance of feedforward and feedback explanations, the empirical literature has focused on feedback. For instance, Mao and Benbasat (2000) note that feedforward explanations are not generally available in existing KBS; therefore, their experimental study examined only output (i.e., feedback) explanations. Similarly, Ye and Johnson examined whether explanations would convince users to accept the system's recommendation and used a simulation of a working KBS. The KBS generated a recommendation and then made one or more of three feedback explanation types available to users. Because users did not input data into the system, feedforward explanations were deemed unnecessary.

Only two published studies have addressed feedforward explanation provision empirically. Mao and Benbasat (2001), in an extension of their earlier study, provide evidence for the impact of domain knowledge on the use of explanations, finding limited support for a differential effect of expertise on the use of feedforward explanations. Gregor (2001) provides *definition* explanations in a feedforward mode and finds that individuals using a personal financial planning KBS access such explanations when available during the input process.

Alternative Knowledge Level Explanation Usage

The knowledge component supported by alternative types of explanations and the user's expertise are perceived to interact and influence decision makers' use of explanations. Indeed, one potential benefit from KBS explanations is the transfer of knowledge from KBS to user—not just as a teaching tool, but to facilitate use of the KBS (Dhaliwal and Benbasat 1996). Such knowledge in KBS research generally consists of declarative and procedural components (Anderson 1993). Lamberti and Wallace (1990) find that low-skilled users are more satisfied with declarative explanations, while Dhaliwal and Benbasat suggest that more-experienced users are more interested in how and why decisions are made (i.e., more procedural-level knowledge). Mackay and Elam (1992) note the need for researchers of KBS to be more attentive to learning, particularly from a working perspective (i.e., a KBS used to aid in decision making).

The foundation for the expectation that novices prefer feedforward and experts prefer feedback explanations is based on Anderson's (1993) ACT-R theory of learning. Anderson theorizes that expertise is formed by combining declarative- and procedural-level knowledge and develops in a three-stage process. The first two stages relate to declarative learning. The first stage is the acquisition of *declarative knowledge* (knowledge of facts represented through definitions); and the second stage is the acquisition of general rules for "how to" knowledge represented as *initial problem solving strategies*. The third stage leads to procedural learning and entails the development of *procedural knowledge* through cyclically applying and then refining initial problem solving strategies through experience and observation (i.e., development of usable how-to knowledge). These three stages of learning are similar to the life cycle of learning put forth by Kolb (1984), who suggests a learner first seeks to understand basic facts and objective rules, before pursuing interpretive skills on more subjective information and finally achieving a consistency in decision making through refined interpretive skills.

In applying ACT-R theory, various forms of explanations facilitate the transfer of each knowledge component to support each of the three stages in the learning process, including learning for working processes (i.e., Mackay and Elam's view of supporting learning during actual KBS use for successful task completion). Overlaying ACT-R's learning process on the eight forms of explanation provision infers that as novices attempt to understand a decision making task modeled by a KBS, they would first turn to definition explanations to acquire declarative knowledge. As a declarative knowledge foundation is established, novices will additionally seek non-definitional feedforward explanations to facilitate an understanding of the procedural processes (i.e., initial problem-solving strategies). As decision makers develop more expertise, they will turn to non-definitional feedback explanations that explain how a KBS derived a conclusion to resolve any perceived anomalies between the KBS and decision makers' own procedural knowledge base. From an ACT-R perspective, this focus on procedural knowledge-oriented explanations occurs as more-expert decision makers attempt to reconcile recommendations from the KBS with their own preferred solution. The reconciliation between the two solutions should lead users to examine the underlying problem resolution strategy of the KBS and compare it with their own applied problem-solving strategies to resolve conflicts and/or refine their own problem solving strategies. Consistent with Dhaliwal and Benbasat, novice users with little domain knowledge are expected to pursue explanations that provide declarative knowledge (i.e., definition explanation type). After acquiring declarative knowledge, novice users are also expected to show a preference for explanations that provide initial problem-solving strategies (i.e., feedforward rule-trace, justification, and strategy explanations). Experts are expected to prefer procedural knowledge (i.e., feedback rule-trace, justification, and strategy explanations). All three conditions are consistent with the theory provided by Gregor and Benbasat. This leads to the fourth, fifth, and sixth hypotheses.

H4: When using a KBS with explanation facilities, novices will choose more declarative knowledge explanations than experts.

H5: When using a KBS with explanation facilities, novices will choose more initial problem-solving strategy-based explanations than experts.

H6: When using a KBS with explanation facilities, experts will choose more procedural knowledge explanations than novices.

As noted earlier, the theory of technology dominance (Arnold and Sutton 1998) posits that novices will tend to defer to a KBS due to the relative expertise of the KBS, while experts are less likely to rely on it. The theory suggests that a cooperative problem-solving approach between an expert user and KBS should lead to better decision making (for empirical support of the theory, see Arnold et al. 2004b; Masselli et al. 2002; Noga and Arnold 2002). The theory is premised on the belief that a well-constructed KBS can provide peer-level advice for the user and help influence the user's decision-making processes in a positive manner.

The theory underlying explanation design suggests that much of the influence of a KBS on experts' decision-making processes is based on the clarity of argumentation and strength of the explanation. Additionally, the more explanations that a user views, the more convincing the argument should be. Recall that expert decision makers theoretically prefer feedback explanations (Dhaliwal and Benbasat 1996) and choose to use available explanations when they perceive an anomaly in the KBS recommendation. This is also the stage in decision making where Arnold and Sutton (1998) theorize that a KBS may enhance an expert's judgments. Thus, for explanations to be effective with experts, the argumentation must be effective (i.e., based on use of Toulmin's model for explanation development), they must be motivated to view the explanations, and their judgment should shift toward the KBS recommendation. Accordingly, compared with experts who do not access explanations, experts who access explanations are expected to render decisions more consistent with the KBS recommendation. Moreover, the more explanations experts view, the more convinced they should be by the argumentation. The next hypotheses are formally stated as follows:

H7: Experts who use feedback explanations when using a KBS with explanation facilities will be more likely to adhere to the recommendation.

H8: Experts who use a KBS with explanation facilities will be increasingly likely to adhere to the recommendation as the number of feedback explanations viewed increases.

Research Method

A between-subjects design was implemented using two treatment conditions to test the eight hypotheses. An insolvency case was used in both treatments with a total of 80 insolvency professionals having access to a KBS without an explanation facility in the first treatment and 66 professionals having access to the same KBS with the addition of an explanation facility in the second treatment. The two treatments were used to test the differential effect of the availability of explanations in a KBS forming the experimental manipulation (no explanations versus full explanation capability) for H1.

A between-subjects design was also used for H2 through H8, but only using data from the participants in the second treatment (i.e., full explanation capability available). H2 through H6 examine the differential use of alternative explanation categories (feedforward, feedback, declarative, initial problem solving, and procedural) based on expertise level (novice versus expert). H7 and H8 focus on the expert subset of the sample to examine the differential impact of feedback explanation usage (use versus non-use of feedback explanations) on experts' judgments.

Participants

All participants were experienced insolvency professionals from large professional services firms located in either Sydney or Melbourne. The professionals ranged in experience from junior-level staff (i.e., staff and seniors) to senior-level staff (i.e., partners, directors, and managers). The experiment was held at participating firms' offices as training sessions. The purpose of the training session was to introduce the participants to INSOLVE and its potential usability as a high-level KBS to support complex decision-making processes. Professionals' participation was based on availability during the time period in which the training session was being held at a given firm's offices. A partner in each firm sent advanced notice of the session to potential participants requesting their participation.

In each session, all participants were provided with hands-on training with the same version of INSOLVE (i.e., no explanations versus full explanation capability). Because groups rather than individuals were assigned to treatments, it was not possible to balance out the number of participants randomly assigned to each experimental cell (e.g., a two-by-two design with expert versus novice and no explanations versus full explanation capability). Firms that provided a different number of participants than planned exacerbated the imbalance between novices and experts. In addition, one office notified the experimenters the night before the scheduled training session that none of the participants would be available for the session because they had just been called out on a major insolvency case. Nonetheless, 80 insolvency practitioners (43 novices and 37 experts) participated in the "no explanations available KBS" treatment, and 66 insolvency practitioners (39 novices and 27 experts) participated in the "full explanation capable KBS" treatment. Table 1 shows demographic information on the participants.

Experimental Task

Insolvency was selected for the experimental task of interest based on three criteria: (1) a large number of insolvency professionals make the same types of decisions; (2) the decisions require a high level of expertise and the risk associated with poor decision choice can be high; and (3) the task is complex and depends on subjective judgment processes. These attributes are desirable because they facilitate delineation of expertise between various participants and provide a domain in which a KBS can serve a useful purpose for both novice and expert decision makers.

Insolvency practice is prevalent in virtually all countries across the world that evolved under British law and has existed for over a century. It often involves a Chartered Accountant (CA) taking over management of a company that is in financial distress and evaluating the various alternatives for the company's future. In many cases, this may require immediate liquidation. In other cases, the CA may decide that the best return for creditors and other stakeholders can be obtained by keeping the company in business. This latter case is referred to as "trading-on" the company. It may continue until either (1) the business can be sold as a going concern or (2) management of the company can be returned to its directors.

The insolvency case was based on one developed by King (1988) for use with experienced accounting professionals making a going-concern decision. The factors of interest to an auditor in a going-concern decision are similar to those considered by insolvency professionals. The case was substantially modified to fit more closely to the insolvency environment and to balance as closely as possible the positive and negative information. This balance of evidence should allow the decision maker to be able to justify either a decision to liquidate or trade-on. The case was pretested with an insolvency partner, one insolvency senior, and a faculty colleague. The case was again revised and pretested a second time by an insolvency partner and manager, this time using the KBS that was later used in the experimental sessions. Supportive feedback from the insolvency practitioners on the realism and appropriateness of the case resulted in this version of the case being applied in all experimental conditions.

The case materials comprised three sets of information used in a sequential decision-making environment. The initial set of information consisted of a complete set of financial statements along with baseline data about a company in financial distress. After reviewing this initial information, participants provided a likelihood estimate that they would elect to continue trading-on the company. Participants then progressed sequentially to the second and third sets of data. To increase the realism of the case, participants were informed that two different staff members had been given the responsibility to seek the additional evidence—one collected data supporting liquidation and one collected data supporting a trade-on (continue operating the company). Participants provided a revised likelihood estimate of continuing to trade-on after reviewing each additional set of information. The information was split into three sets, in part to reduce the cognitive load that could be overwhelming given the complexity of an insolvency decision. Further, an insolvency decision is a recurring decision that is made constantly throughout an insolvency engagement with a reversal of plans possible at any time as additional information becomes available.

Experimental KBS: INSOLVE

The KBS selected for use in the experiment was INSOLVE. INSOLVE was originally designed to replicate the decision-making processes of expert insolvency practitioners (Collier et al. 1999; Leech et al. 1998, 1999). It was the product of a 4-year project funded by the Australian Research Council and the Institute of Chartered Accountants in Australia (IT Chapter). In all, 23 insolvency experts were used to develop the knowledge acquisition phase of INSOLVE. The resulting cognitive model was validated and refined based on the feedback of six of the experts and was subsequently programmed into a working prototype. Formal validation of INSOLVE was completed with 17 experts (6 of whom were not involved in the knowledge acquisition phase). INSOLVE received high scores for validation. (For a complete discussion of the development and validation of INSOLVE, see Collier et al. 1999; Leech et al. 1998.)

The model is designed to be a collaborative KBS for use in cooperative problem solving where the KBS periodically exchanges control with the user. The user is required to enter both qualitative and quantitative estimates at various stages for the KBS to work effectively and may also adjust the weightings for different factors in the knowledge base. Figure 3 provides two examples of input screens that arise in the course of entering data on the user's initial impressions of the client's financial situation. Once the user enters the command turning control over to the KBS, INSOLVE prompts the user for additional information that would be useful before formulating a recommendation and providing a brief report (generally about two screens in length) on the selected recommendation. Figure 4 displays a consolidation of a sample two-screen report generated by INSOLVE. At any point, the user can add or change the information input into the system and request a new

Table 1. Participant Demographic Information

Variable	Full Explanation Capability Available Group	No Explanation Available Group	Combined
Number of Participants:			
Staff/Seniors	39	43	82
Managers/Partners/Directors	27	37	64
Total	66	80	146
Age in Years			
Staff/Seniors	26.18	25.67	25.91
Managers/Partners/Directors	35.11	35.58	35.38
Insolvency Experience in Years			
Staff/Seniors	3.48	3.67	3.58
Managers/Partners/Directors	12.12	10.26	11.04
Primary Decision-Making Responsibility			
Staff/Seniors	8.87	2.16	5.35
Managers/Partners/Directors	66.58	68.19	67.51

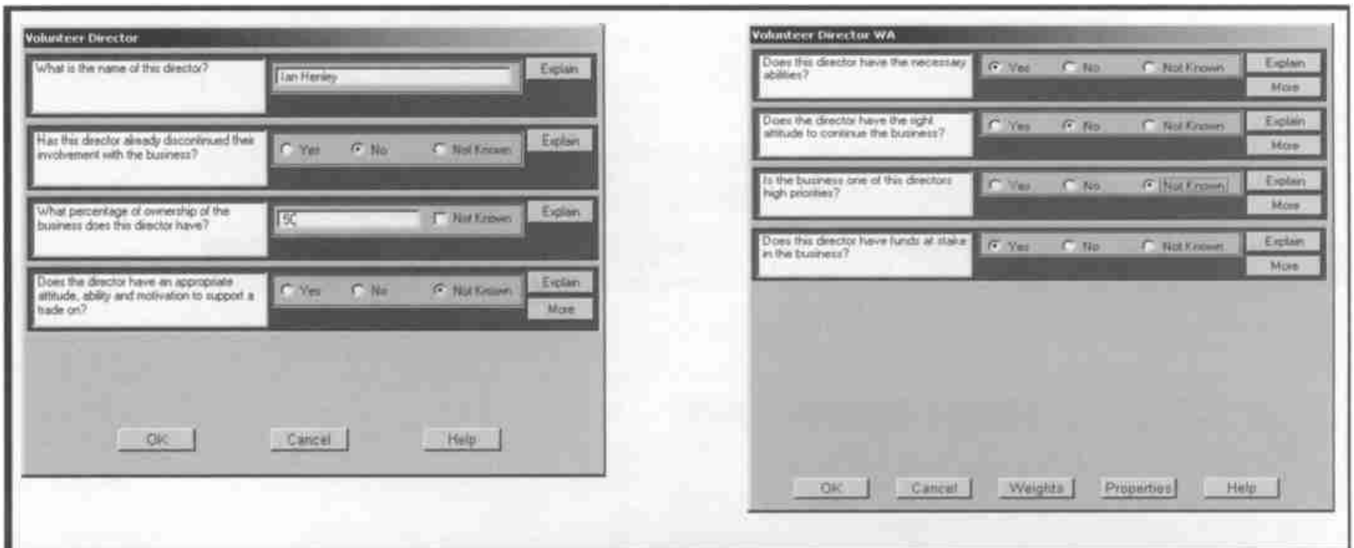

Figure 3. Sample Input Screens from INSOLVE KBS

recommendation. In summary, INSOLVE is a high-level, reliable KBS and is representative of an advanced-level KBS desirable for practice.

The Australian Research Council provided additional funding to develop a fully functional explanation facility within INSOLVE. The explanation facility was built using the large reservoir of residual elaborations from the knowledge acquisition phase of INSOLVE. This rich residual data includes elaborations about relevant factors, knowledge fragments, and process fragments that are ideal for creation of definition and justification explanations (see Figures 5 and 6 for examples). Rule-trace explanations simply represent an explicit rendering of the rules and other procedures in the underlying knowledge base (see Figure 7 for an example). Strategic explanations are based on the underlying task structure of INSOLVE: each recommendation in INSOLVE has an associated

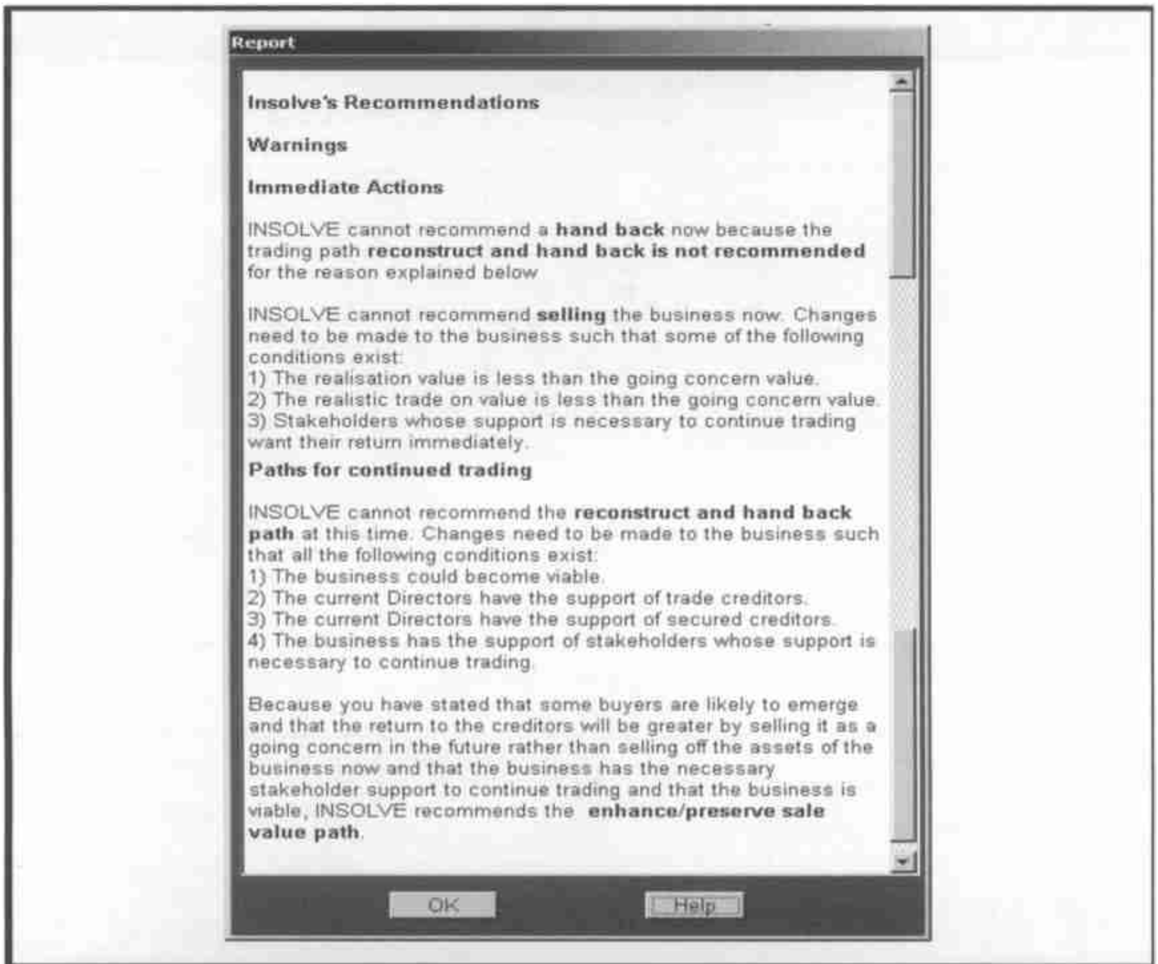

Figure 4. Sample Report Generated by INSOLVE KBS (Consolidation of Two Screens)

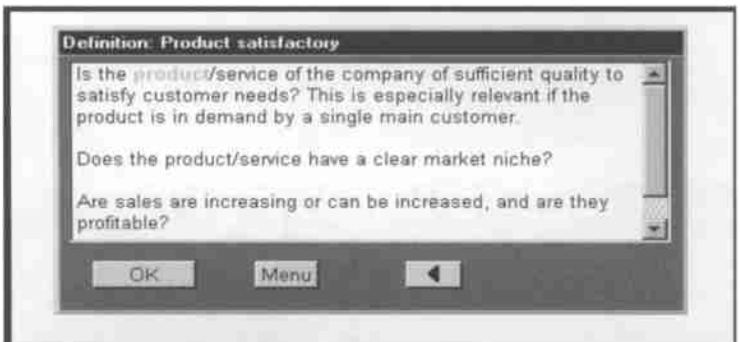

Figure 5. Definition Explanation in INSOLVE

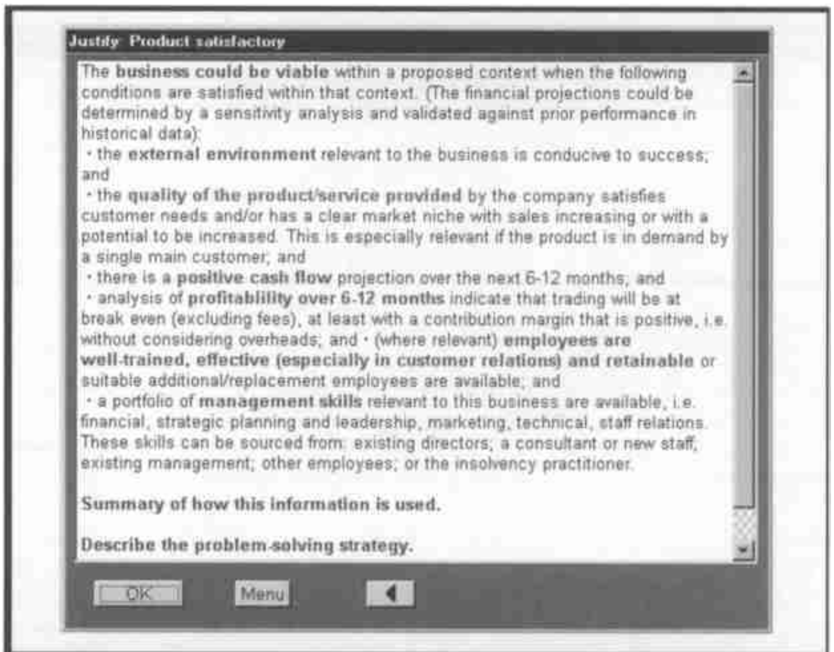

Figure 6. Sample Justification Explanation in INSOLVE

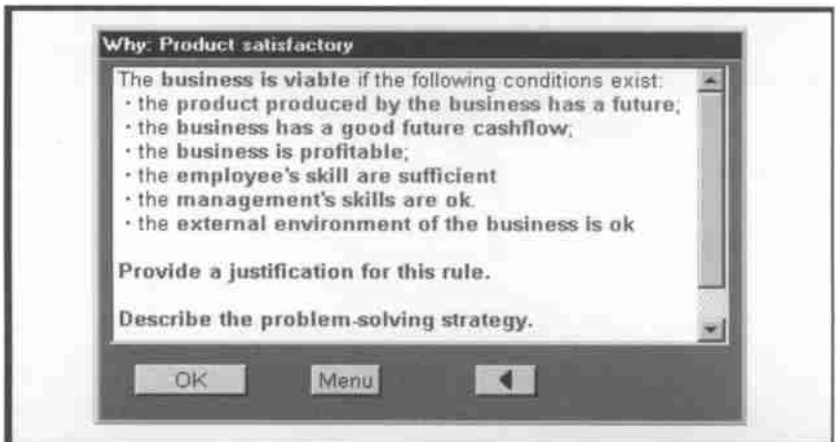

Figure 7. Sample Rule-Trace Explanation in INSOLVE

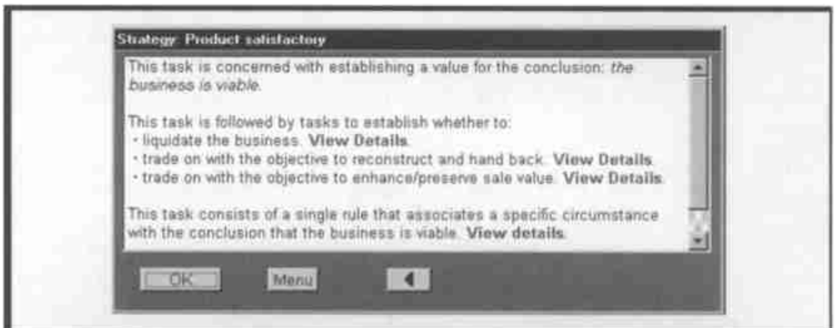

Figure 8. Sample Strategic Explanations in INSOLVE

Table 2. Number of Explanations per Explanation Type in INSOLVE

Explanation Type	Number of Explanations
Definition	144
Rule-Trace	325
Justification	322
Strategy	306

task, which has the purpose of establishing a value for an intermediate conclusion or final recommendation (see Figure 8 for examples). Finally, *Toulmin warrants* form the basis for justification and are based on "general examples" and "being told" (Chandrasekaran et al. 1989). Table 2 shows the number of explanations by type and mode that are embedded within INSOLVE.

Each explanation type (definition, justification, rule-trace, and strategic) is available in both a feedforward and feedback mode. Definition explanations take the same form regardless of presentation in feedforward or feedback mode. For the other three types of explanations, the text for the feedforward explanations describes INSOLVE's problem-solving knowledge in declarative terms (i.e., in terms of relevant information cues and their relationships). Feedback explanations are expressed in procedural terms and describe how a recommendation is made (or why it is not made or remains unknown). Consistent with Mao and Benbasat's (2001) findings that explanations are more likely to be accessed and more effective when contextualized within the problem solving process, all explanations are made available in INSOLVE via either hypertext or through a *button* located on-screen in the immediate proximity of the item to be explained. (For a complete discussion of the development of the explanation facility within INSOLVE, see Arnold et al. 2004a.) While the number of access points depends on the user's strategy in using INSOLVE, an efficient user might have approximately 60 to 65 different points of access displayed while completing the case. Each access point provides linkage to each of the four types of explanations either in feedback or feedforward mode (i.e., a total of 245 to 260 explanations). Given the nested nature of the explanation facility links, a user could, in the most extreme situation, access all 1,097 explanations contained within INSOLVE.

Of course, the previously observed *production paradox*, where users are more focused on working than on learning, would predict that the actual number of explanations viewed will be low (Carroll and McKendree 1987).

Experimental Procedure

The experimental procedure was embedded within a series of training sessions that introduced the participants to INSOLVE and to the potential usability of a high-level KBS for supporting complex decision making. The sessions began with a 20-minute introduction to INSOLVE consisting of an explanation of the development process, an overview of the system validation results, and instruction on basic use of the KBS. The participants then completed a tutorial using INSOLVE to generate a series of recommendations related to a company in financial distress. For those participants using the KBS with full explanation facilities, the tutorial also exposed the participants to the types of explanations that appear in both a feedforward and feedback mode. The purpose of the training from an experimental perspective was to ensure user familiarity with the KBS and to allow the user to go through an assessment of the usability and value of the KBS prior to starting the actual experimental case. This familiarity is considered a necessary precursor to the willingness of experts to use a KBS in the decision making process (Arnold and Sutton 1998). It is also important to note that during the training session users begin to formulate opinions as to the reliability and value of the system that, in turn, will influence their trust and acceptance of the system.

After completing the tutorial and asking questions of the roving instructors, participants completed the experimental case using INSOLVE. While the focus of the scheduled training session from the participants' perspective was to receive an introduction to intelligent decision aids and how they could be used in the insolvency decision domain, the focus during case completion was on task performance and effective decision making. For the post-training case completion, the participants were instructed to focus on decision performance while using INSOLVE to the degree it was helpful. A log of each user's activity was maintained through a "replay process tracing" facility that captured exact data entries, options selected (including explanation access), and time between key events. (For a more complete description, see Arnold et al. 2001.)

The case materials were provided in three sealed envelopes clearly labeled for their sequence of completion. The first envelope contained the complete set of financial statements and baseline data for the experimental case. After reviewing the information, participants could optionally enter the information into INSOLVE to acquire a recommendation—which all participants chose to do. After considering the feedback from INSOLVE, participants manually recorded their likelihood estimates of trading-on (using a 0 to 100 percent scale) on an attached response sheet at the end of the case materials and placed the response sheet in the envelope. The same process was repeated for each of the additional two sets of information representing the additional evidence discovery from two colleagues. Demographic information on each participant was collected at the end of the case. The second and third envelopes of additional evidence were randomized in order of presentation between evidence supporting liquidation or evidence supporting trading-on. This randomization of information was used knowing that certain conflicting evidence would cause certain values to change based on the most recently available information. This, in turn, would cause the KBS to render different recommendations (i.e., liquidate versus trade-on) in accordance with the last envelope of evidence received. This controlled for the possibility that insolvency practitioners might be more accepting of a trade-on versus a liquidate decision (or vice versa) when engaging the KBS. This experimental control mechanism minimized the possibility that experimental results for H1 (effect of KBS explanations on user judgment) might be driven by participants' predisposition toward a certain outcome.

Measurement of Dependent Variables

There are two key types of dependent variables: adherence and explanation accesses. Adherence is measured based on the user's final likelihood estimate (0 to 100 percent) of liquidating the company versus the recommendation provided by INSOLVE. The recommendation provided by INSOLVE was either liquidate (0) or trade-on (1).[5]

Accesses of the different types and modes of explanations are determined by a simple count of the number of explanations the user accessed during the session. They are easily determined by examining the replay process tracing files. In segmenting the explanation accesses into separate counts for H2 and H3, *feedforward explanations* include all explanations that report on information cues and their relationships, and *feedback explanations* include all explanations that report on procedures that are used to form conclusions. For H4 through H6, *declarative knowledge explanations* represent the number of definition explanations accessed in either feedforward or feedback mode; *initial problem solving strategy-based explanations* are rule-trace, justification, and strategy explanations accessed in a feedforward mode; and *procedural knowledge explanations* are the same three types in a feedback mode (as represented in Figure 2 via the boxed groupings).

The examination of the use of explanations from both perspectives in the extant literature (i.e., feedback versus feedforward and declarative versus initial problem solving versus procedural) yields a risk of over-reliance due to the same data being used in two alternative sets of analyses. An adjustment to counter this issue is to use a Bonferroni adjustment. Because an OMNIBUS ANOVA testing the overall effect of all explanations does not exist, the appropriate amount of adjustment in the critical value is arguable. This is further complicated because each hypothesis looks at an independent count of accesses, and there is no theorized interaction between the dependent variables of interest in the five hypotheses. As such, the most conservative position is taken by assuming that an OMNIBUS ANOVA has been run (e.g., expertise will affect the overall number of explanations accessed) and then treating the five hypotheses as sub-hypotheses under this illusionary OMNIBUS test. This conservative position leads to a critical value of .01 for each of the five hypotheses given a family critical value of .05 being adopted otherwise in the research.

Results

The data were collected via the series of experimental sessions and provide the basis for testing the eight hypotheses. Data from both the "explanations provided" and "no explanations provided" sessions were used to test the first hypothesis, while only the data from the "explanations provided" sessions were used to test the remaining hypotheses. Table 3 presents mean data on the number of accesses of various explanation types, average time viewed, and average word count per second for both expert and novice participants. Note that the interest in the experimental sessions is on using the explanations to complete a work task. As such, the number of explanations viewed by each participant is low relative to the total available explanations (as would be expected).

Judgment Effects

The purpose of H1 is to test whether insolvency professionals using a KBS with explanations will demonstrate increased adherence to the recommendations of the KBS in comparison to professionals using a KBS without explanations. The comparison with the group using the KBS without explanations isolates the effect of explanations from that of simply using the KBS. This is related to the effect of using argumentation as a basis for explanations and the perceived likelihood that users will place greater faith in a KBS as demonstrated by increased adherence with the recommendation of the KBS. To test for differences from the use of the KBS with explanation facilities provided (H1), ANCOVA was used with the final likelihood estimate of trading-on as the dependent variable and the initial likelihood estimate as a covariate. Availability of explanation facilities and order of information presentation were the independent variables. (A control variable was used to capture the effect of the randomized order of information presentation and the related trade-on versus liquidate recommendation by the KBS, which would cause participants to adjust their decisions in opposite

[5] Recall that the cases were designed so that the information presentation order was randomized. While all participants received precisely the same information, the order of information would alter the recommendation from the KBS between trade-on and liquidation.

Table 3. Experimental Access and Use of Explanations

	Novices						Experts						Combined					
	Number of Accesses		Average Attendance Time (Seconds)		Average Word Count Per Second		Number of Accesses		Average Attendance Time (Seconds)		Average Word Count Per Second		Number of Accesses		Average Attendance Time (Seconds)		Average Word Count Per Second	
	FF	FB	FF	FB	FF	FB	FF	FB	FF	FB	FF	FB	FF	FB	FF	FB	FF	FB
Definition	152	11	13	6	5.0	6.1	59	5	14	6	4.7	5.4	211	16	14	6	4.9	5.9
Rule-Trace	73	33	8	9	6.0	2.8	43	30	6	5	7.0	7.7	116	63	7	7	6.3	4.3
Justification	43	11	11	224	11.3	0.4	20	20	9	22	13.9	5.1	63	31	10	94	12.1	1.1
Strategic	46	14	9	15	7.6	3.7	22	6	8	13	7.1	3.0	68	20	9	15	7.4	3.5
Total	314	69	11	44	6.4	1.0	144	61	10	11	6.6	5.4	458	130	11	29	6.4	1.8

FF: Feedforward
FB: Feedback

directions.) H1 is supported if a significant shift occurs in the difference from the initial likelihood estimate to the final likelihood estimate, reflecting increased likelihood of following the KBS recommendation. The statistical results show that the interaction of explanation facility and order is significant ($F(1,138) = 9.270$, $p = .003$): order is key in that it causes the decision to move in opposite directions so the judgments of participants receiving a trade-on versus a liquidate recommendation should be further apart than for those participants using the KBS without explanations. This test result indicates that providing an explanation facility led to greater adherence by the insolvency professionals to the KBS's recommendation, regardless of whether the KBS recommended trade-on or liquidate. The main effect for order is also significant ($F(1,138) = 145.833$, $p < .001$), and the initial likelihood estimate covariate is significant ($F(1,138) = 14.196$, $p < .001$).

Feedforward Versus Feedback Explanation Usage

H2 and H3 relate to the use of feedforward versus feedback explanations. The extant literature posits that novices will use more feedforward explanations (H2) as they attempt to understand better the problem domain, while experts will use more feedback explanations (H3) as they attempt to understand better how and why a KBS provides a given recommendation. A MANOVA was conducted to consider the overall impact of expertise on the two dependent variables (i.e., feedforward and feedback). The statistical results of the MANOVA show that expertise significantly impacts the use of feedforward explanations ($F(2,63) = 30.221$, $p < .001$) and also significantly impacts the use of feedback explanations ($F(2,63) = 7.466$, $p = .001$). These results support both H2 and H3.

There is a concern in the MANOVA, however, in that the two dependent variables are highly correlated ($r^2 = .540$, $p < .001$). A Roy-Bargman procedure was used to determine the impact of expertise on the two dependent variables when controlling for the correlation between feedforward and feedback explanation use. The procedure was completed by first entering feedback into an ANOVA, confirming significance of expertise, and then using an ANCOVA to test the effect of expertise on feedforward explanations when controlling for use of feedback explanations (i.e., as a covariate). The ANOVA with the number of feedback explanation accesses (dependent variable) and expertise (independent variable) was used to test H3. The results support the hypothesis ($F(2,63) = 265.192$, $p = .001$). Novices accessed on average 1.77 feedback explanations, while experts accessed 2.26. The ANCOVA with the number of feedforward explanation accesses (dependent variable), number of feedback explanations (covariate), and expertise (independent variable) was used to test H2. The results indicate that expertise impacts the number of feedforward explanation accesses ($F(2,62) = 19.984$, $p < .001$). Novices accessed on average 8.05 feedforward explanations, while experts accessed on average 5.33. The number of feedback explanations covariate is also significant ($F1,62) = 28.386$, $p < .001$).

Alternative Knowledge-Level Explanation Usage

H4, H5, and H6 are conceptually based on Anderson's ACT-R theory of learning. The relationships specified in the three hypotheses are based on a translation of the theory to the generation of explanations that facilitate knowledge transfer to the user (Dhaliwal and Benbasat 1996; Gregor and Benbasat 1999). As novices attempt to aggregate knowledge about decision-making processes, they are theorized to learn first through declarative knowledge (H4) and then to pursue an understanding of initial problem solving strategies (H5). Thus, novices are expected to access more explanations related to these two kinds of knowledge. Experts, on the other hand, are more interested in resolving questions about the KBS reasoning processes and will focus more on procedural-level knowledge (H6). A MANOVA was conducted to consider the overall impact of expertise on the three dependent variables.. The MANOVA indicates declarative ($F(2,63) = 33.019$, $p < .001$), initial problem solving strategies ($F(2,63) = 16.9$, $p < .001$), and procedural explanation usage ($F(2,63) = 6.575$, $p = .003$) are all significantly impacted by expertise providing support for H4, H5, and H6.

There is a concern in the MANOVA, however, in that the three dependent variables are highly correlated; the correlation between declarative and initial problem solving is .457 (< .001), declarative and procedural is .277 (.013), and initial problem solving and procedural is .596 (< .001). A Roy-Bargman procedure was again used to determine the impact of expertise on the three dependent variables when controlling for the correlation between the three kinds of knowledge explanation use. The procedure was a three-step process in which variables were entered based on the strengths of their correlations with each other. In step one, procedural explanations (the most highly correlated variable) were tested in an ANOVA to confirm the significance of expertise. In step two, the initial problem-solving strategy explanations (the second highest correlated variable) were tested using an ANCOVA to test for the effect of expertise while controlling for procedural explanations (i.e., using procedural explanations as a covariate). Step three repeated the ANCOVA using the final variable, declarative explanations, to test for the effect of expertise while controlling for both procedural and initial problem solving strategy explanations. The ANOVA with the number of procedural explanation accesses (dependent variable) and expertise (independent variable) was used to test H6. The results support the hypothesis ($F(2,63) = 6.575$, $p = .003$). Novices accessed on average 1.49 procedural-based explanations, while experts accessed on average 2.07.[6] The ANCOVA with the number of initial problem solving strategy explanation accesses (dependent variable), number of procedural explanations (covariate) and expertise (independent variable) was used to test H5. The results indicate that the independent variable expertise is significant ($F(2,62) = 9.529$, $p < .001$) and the covariate of procedural explanations is significant ($F(1,62) = 36.104$, $p < .001$), thereby providing support for H5. Novices accessed on average 4.15 initial problem-solving strategy-based explanations, while experts accessed on average 3.15. The ANCOVA with number of declarative explanation accesses (dependent variable), number of initial problem solving strategy explanation and procedural explanation accesses (covariates), and expertise (independent variable) was used to test H4. The results indicate that the independent variable expertise is significant ($F(2,61) = 12.396$, $p < .001$) and the covariate of initial problem solving strategy is significant ($F(1,61) = 8.579$, $p = .005$), thereby providing support for H4. Novices accessed on average 4.18 declarative knowledge explanations, while experts accessed 2.37.

H7 and H8 relate to the ability of explanations provided by KBS to convince expert users of the rationality/feasibility of the KBS's recommended problem solution. The KBS's explanations should have such an effect if the decision maker is motivated to use the explanations, the argumentation of the explanations is convincing, and the expert decision maker adjusts his/her judgment.

H7 examines whether experts who view feedback explanations will make decisions that are more consistent with the KBS recommendation. An ANOVA was used to test H7 using a dichotomous variable (access or not access feedback explanations) as the independent variable and the difference in final likelihood estimate as the dependent variable. The results of the test for H7 are significant ($F(2,24) = 4.847$, $p = .017$), and the hypothesis is supported.

H8 examines whether the number of feedback explanations viewed influences the expert decision makers' judgment in the direction of the KBS recommendation. A Pearson's correlation test of the relationship between number of feedback explanations viewed and the difference in final likelihood estimate was used to test H8. The results support H8 ($r^2 = .357$, $p = .036$).

Table 4 provides a summary of the results for the tests of each of the eight hypotheses. All hypotheses are supported. Overall, substantial support exists for the underlying theory.

Discussion and Implications

The research reported in this paper takes a broad look at the use and influence of explanations on expert and novice decisions in a complex judgment task. To our knowledge, this is the first experimental test of the impact of a KBS implementing a fully functional explanation system using both feedforward and feedback modes of explanations in combination with all four types of explanations (definition, justification, rule-trace, and strategy). Prior published research has largely omitted both feedforward explanations and definition explanations.

[6]The amount of information provided in the recommendation report (approximately two screens in length) may have reduced the number of questions (or perceived anomalies) that would lead to feedback explanation accesses.

Table 4. Summary of Results

Hypothesis	Statistical Test	Results*
H1: KBS users will be more likely to adhere to the recommendation of a KBS when explanations are provided.	ANCOVA	Supported ($p = .003$)
H2: When using a KBS with explanation facilities, novices will choose more feedforward explanations than experts.	ANCOVA with Roy-Bargman Step-Down	Supported ($p < .001$)
H3: When using a KBS with explanation facilities, experts will choose more feedback explanations than novices.	ANOVA with Roy- Bargman Step-Down	Supported ($p = .001$)
H4: When using a KBS with explanation facilities, novices will choose more declarative knowledge explanations than experts.	ANCOVA with Roy-Bargman Step-Down	Supported ($p < .001$)
H5: When using a KBS with explanation facilities, novices will choose more initial problem-solving strategy-based explanations than experts.	ANCOVA with Roy-Bargman Step-Down	Supported ($p < .001$)
H6: When using a KBS with explanation facilities, experts will choose more procedural knowledge explanations than novices.	ANOVA with Roy-Bargman Step-Down	Supported ($p = .003$)
H7: Experts that use feedback explanations when using a KBS with explanation facilities will be more likely to adhere to the recommendation.	ANCOVA	Supported ($p = .017$)
H8: Experts that choose more feedback explanations when using a KBS with explanation facilities will be more likely to adhere to the recommendation.	Pearson's Correlation	Supported ($p = .036$)

*A family level significance of .05 is used as a critical value in assessing the acceptance of the hypotheses. As noted earlier in the paper, a conservative measure of the Bonferroni adjustment is used for hypotheses 2 through 6 where the critical value is .01 (i.e., .05/5).

Our results provide evidence that the availability of a fully functional explanation facility influences both novices' and experts' judgments. Perhaps more importantly, significant differences exist in the forms of explanations preferred by users of different knowledge levels. The results indicate that novices will have stronger preferences for feedforward and definitional explanations. Experts, on the other hand, will have a greater interest in procedural-based explanations that are the type that have generally been available in prior experimental studies. This is important to designers and implementers of KBS, given that prior research has suggested that few KBS used in practice provide feedforward and/or definitional explanation facilities (Dhaliwal and Benbasat 1996; Mao and Benbasat 2000; Ye and Johnson 1995).

Designers of KBS for use by novice decision makers should consider the implications of the research for the design of explanation facilities in such systems. For novices, explanations may be key to both short- and long-term learning. This study has focused on short-term learning or, in other words, learning to perform. As such, feedforward explanations (as well as definition explanations in feedback mode) may be critical for novices to use a KBS successfully. On the other hand, the results also indicate that novices tend to accept the KBS's recommendations and move toward adherence. Arnold and Sutton (1998) suggest that such dominance by the KBS may lead to poor decision making. Implementers of KBS for use by novices should be cognizant of this potential negative effect. This study has focused on adherence to the KBS and not on optimal decision performance. Future research should focus on the performance aspects of KBS adoption by novice users.

Beyond the results provided through testing the hypotheses, the raw data should be considered in terms of usage of different forms of explanation. Experts demonstrated greater interest in procedural-based explanations (as theorized and hypothesized). The average number of explanations accessed by experts was 2.37 declarative-based explanations, 3.15 initial problem solving strategy-based explanations, and 2.07 procedural-based explanations. Thus, the mean data indicate that experts had a strong interest in explanations of how the KBS used certain data and made recommendations—a phenomenon that has not been examined in the prior literature. One potential explanation may be that experts are comparing the underlying information in the KBS (e.g., initial problem-solving strategies) with their own knowledge to establish cognitive fit. This has been theorized as a precondition to acceptance of a KBS by experts (Arnold and Sutton 1998). This phenomenon should be examined in future studies. If this phenomenon holds under specific examination, the implications for designers and implementers of future KBS are significant, particularly in terms of designing explanations to facilitate early adoption of KBS by expert users. Care should be taken, however, in comparing the raw means between forms of explanations, as a significant difference existed in the number of opportunities to view various forms of explanations—particularly procedural explanations that are most frequently accessed at the time a recommendation report is requested.

The results of the study related to experts using the KBS, when taken in aggregate, suggest that both feedforward and feedback explanations may be important to acceptance by expert decision makers. The different modes along with the four types of explanations provide an array of views into the information cues, their relationships, and the recommendations of the KBS. The experts' use during our experimental sessions may indicate curiosity as to the inner workings of the KBS. The fact that the presence of a KBS with explanations strengthened the influence of the KBS on experts' decision outcomes suggests that the need for a fully capable explanation facility is fundamental to acceptance of a KBS by expert decision makers—an area that has been problematic across a wide array of KBS implementations in knowledge-based organizations. Additionally, training for experts on the use of the system that incorporates access and use of the explanation facility may enhance their initial perceptions of a new KBS and promote ultimate acceptance. These issues should be explored in future research.

Overall, based on the empirical results of the test of hypotheses, managers wishing to leverage effectively the use of a KBS in a professional decision-making environment should ensure that the KBS is designed with comprehensive explanation facilities. These facilities should help users understand the knowledge contained in the KBS and the procedures used by the KBS to formulate specific conclusions and recommendations. To facilitate use and user acceptance, KBS should be designed to provide definition, justification, rule-trace, and strategic explanations in both feedforward and feedback. The use of contextualization in the provision of explanations also lends additional support to the findings of Mao and Benbasat (2001) that contextualization enhances the likelihood of explanation use.

This study has several limitations. First, the experiment was conducted in a laboratory environment, which facilitates strong experimental control but also necessitates some limitation in the richness of the data available to the KBS users. Usage patterns in a field setting may be different. Second, the patterns of behavior were studied while the KBS was still relatively new to users. While this fits the emphasis that prior researchers have put on explanations being most important for short-term acceptance and use, longer-term use may result in a different pattern of explanation system use. Third, the novices in this study may not perform like novices in other studies because they were fairly experienced (i.e., consisting of staff and senior-level insolvency practitioners from professional services firms). The novice participants have experience close to the level of experts in some other studies. In comparing results between studies, these inconsistencies in participants' backgrounds should be considered in interpreting results. Fourth, the KBS used in the study is a fully functional KBS typical of the type of systems used in practice. Due to legal and maintenance issues, however, the software has never been licensed for use in practice. Nonetheless, given the inquiries that the developers of INSOLVE have received from professional firms interested in its availability, there is reason to believe the system is in a form that allows implementation. Additionally, as discussed in the methods section of this paper, the developers have rigorously validated the system with professionals from large insolvency practices.

Acknowledgments

We gratefully acknowledge financial support from the Australian Research Council, the Institute of Chartered Accountants in Australia (IT Chapter), and the University of Connecticut CIBER. We are also very appreciative of the generous support from the many accounting firms that participated in the research and for helpful comments that we received in workshops at University of Central Florida, University of Connecticut, University of Melbourne, and Rutgers University. We especially appreciate comments from Izak Benbasat, Christine Earley, Clark Hampton, Ron Weber, the associate editor, and three anonymous reviewers on earlier versions of this manuscript.

References

Argyris, C., and Schön, D. A. *Theory in Practice: Increasing Professional Effectiveness*, Jossey-Bass Publishers, San Francisco, 1974.

Anderson, J. R. *Rules of the Mind*, Erlbaum, Hillsdale, NJ, 1993.

Arnold, V., Clark, N., Collier, P. A., Leech, S. A., and Sutton, S. G. "Explanation Provision and Use in an Intelligent Decision Aid," *Intelligent Systems in Accounting Finance & Management: An International Journal* (12:1), 2004a, pp. 4-27.

Arnold, V., Collier, P. A., Leech, S. A., and Sutton, S. G. "The Impact of Intelligent Decision Aids on Expert and Novice Decision-Makers' Judgments," *Accounting and Finance* (44:1), 2004b, pp. 1-26.

Arnold V., Clark, N., Collier, P. A., Leech, S. A, and Sutton, S. G. "Unveiling the 'Black Box' of Human-Computer Interaction: The Need for *Replay Process Tracing* Capability," *Advances in Accounting Behavioral Research* (4), 2001, pp. 381-405.

Arnold, V., and Sutton, S. G. "The Theory of Technology Dominance: Understanding the Impact of Intelligent Decision Aids on Decision Makers' Judgments," *Advances in Accounting Behavioral Research (*1), 1998, pp. 175-194.

Carroll, J. M., and McKendree, J. "Interface Advice Issues for Advice-Giving Expert Systems," *Communications of the ACM* (30:1), 1987, pp. 14-31.

Chandrasekaran, B., Tanner, M. C., and Josephson, J. R. "Explaining Control Strategies in Problem Solving," *IEEE Expert* (4), 1989, pp. 9-24.

Collier, P. A., Leech, S. A., and Clark, N. "A Validated Expert System for Decision Making in Corporate Recovery," *International Journal of Intelligent Systems in Accounting, Finance & Management* (8:2), 1999, pp. 75-88.

Dhaliwal, J. S., and Benbasat, I. "The Use and Effects of Knowledge-Based System Explanations: Theoretical Foundations and a Framework for Empirical Evaluation," *Information Systems Research* (17:3), 1996, pp. 342-362.

Duchessi, P., and O'Keefe, R. M. "Contrasting Successful and Unsuccessful Expert Systems," *European Journal of Operations Research* (61), 1992, pp. 122-134.

Gill, G. T. "Early Expert Systems: Where Are They Now?," *MIS Quarterly* (19:1), 1998, pp. 51-81.

Gregor, S. "Explanations from Knowledge-Based Systems and Cooperative Problem Solving: An Empirical Study," *International Journal of Human-Computer Studies* (54), 2001, pp. 81-105.

Gregor, S., and Benbasat, I. "Explanations from Intelligent Systems: Theoretical Foundations and Implications for Practice," *MIS Quarterly* (23:4), 1999, pp. 497-530.

Hayes-Roth, F. "Artificial Intelligence: What Works and What Doesn't," *AI Magazine*, Summer 1997, pp. 99-113.

Hayes-Roth, F., and Jacobstein, N. "The State of Knowledge-Based Systems," *Communications of the ACM* (37:3), 1994, pp. 27-39.

Hogarth, R. M., and Einhorn H. J. "Order Effects in Belief Updating: The Belief-Adjustment Model," *Cognitive Psychology* (24), 1992, pp. 1-55.

Hsu, K., and Steinbart, P. J. "Factors Affecting the Use of Different Types of Explanations Provided by Expert Systems," *Advances in Accounting Information Systems* (5), 1997, pp. 215-241.

King, V. A. *Auditor Decision Making: An Analysis of the Effects of Ordering and Tests of the Predictive Validity of the Contrast/Surprise Model*, unpublished Ph.D. Dissertation, University of Arkansas, 1988.

Kolb, D. A. *Experiential Learning: Experience as the Source of Learning and Development*, Prentice-Hall Inc., Englewood Cliffs, NJ, 1984.

Lamberti, D. M., and Wallace, W. A. "Intelligent Interface Design: An Empirical Assessment of Knowledge Presentation in Expert Systems," *MIS Quarterly* (14:3), 1990, pp. 279-311.

Leech, S. A., Collier, P. A., and Clark, N. "Modeling Expertise: A Case Study Using Multiple Insolvency Experts," *Advances in Accounting Information Systems* (6), 1998, pp. 85-105.

Leech, S. A., Collier, P. A., and Clark, N. "A Generalized Model of Decision-Making Processes for Companies in Financial Distress," *Accounting Forum* (23:2), 1999, pp. 155-174.

Mackay, J. M., and Elam, J. J. "A Comparative Study of How Experts and Novices Use a Decision Aid to Solve Problems in Complex Knowledge Domains," *Information Systems Research* (3:2), 1992, pp. 150-172.

Mao, J., and Benbasat, I. "The Effects of Contextualized Access to Knowledge on Judgement," *International Journal of Human-Computer Studies* (55), 2001, pp. 787-814.

Mao, J., and Benbasat, I. "The Use of Explanations in Knowledge-Based Systems: Cognitive Perspectives and a Process-Tracing Analysis," *Journal of Management Information Systems* (17:2), 2000, pp. 153-179.

Masselli, J., Ricketts, R., Arnold, V., and Sutton, S. G. "The Impact of Embedded Intelligent Agents on Tax Compliance Decisions," *Journal of the American Tax Association*, Fall 2002, pp. 60-78.

Noga, T., and Arnold, V. "Do Tax Decision Support Systems Affect the Accuracy of Tax Compliance Decisions?," *International Journal of Accounting Information Systems* (3:3), 2002, pp. 125-144.

Swartout, W. R. "Explanation," in *Encyclopedia of Artificial Intelligence* (1), S. C. Shapiro and D. Eckroth (eds.), John Wiley & Sons, New York, 1987, pp. 298-300.

Swartout, W. R. "XPLAIN: A System for Creating and Explaining Expert Consulting Programs," *Artificial Intelligence* (21:3), 1983, pp. 285-325.

Toulmin, S. E. *The Uses of Arguments*, Cambridge University Press, Cambridge, UK, 1958.

Wanninger, L. "Profitable Electronic Commerce: Frame Work, Examples, Trends," in *Proceedings of the 11th International Bled Electronic Commerce Conference: Electronic Commerce in the Information Society* (Vol. 2), G. J. Doukidis, J. Gricar, and J. Novak (eds.), Bled, Slovenia, June 8-10 1998, pp. 3-27.

Ye, L. R., and Johnson, P. E. "The Impact of Explanation Facilities on User Acceptance of Expert Systems Advice," *MIS Quarterly* (19:2), 1995, pp. 157-172.

About the Authors

Vicky Arnold is Ernst & Young Professor of Accounting at the University of Central Florida and a principal fellow at the University of Melbourne. She is editor of *Advances in Accounting Behavioral Research* and associate editor of the *International Journal of Accounting Information Systems*. She recently coedited a monograph for the Information Systems section of the American Accounting Association entitled *Researching Accounting as an Information Systems Discipline*. She is currently treasurer of SIG-ASYS of the Association for Information Systems. Her research interests are in judgment and decision making and the impact of information technology on decision making by individuals, organizations, and society. Her research appears in *Accounting and Finance, Advances in Accounting Behavioral Research, Advances in Accounting Information Systems, Auditing: A Journal of Practice and Theory, Behavioral Research in Accounting, Critical Perspectives on Accounting, Database, International Journal of Accounting Information Systems, International Journal of Intelligent Systems in Accounting, Finance and Management, Issues in Accounting Education, Journal of American Taxation Association, Journal of Information Systems*, and *Research on Accounting Ethics*, among others.

Nicole Clark is a lecturer of Computing at the University of Tasmania, Australia. She has a computer science background specializing in software engineering. Her current research interests are in automated intelligent systems, software engineering collaboration with industry, and software engineering education. Her research appears in *Advances in Accounting Information Systems, International Journal of Intelligent Systems in Accounting, Finance and Management, Accounting Forum*, and *Australian Computer Journal*, among others.

Philip A. Collier is an associate professor of Business Information Systems at the University of Melbourne, Australia. He has a computer science background specializing in intelligent decision aids applied to a variety of industrial and business contexts, including chemical plant monitoring and diagnosis, paper manufacturing, and insolvency and corporate recovery. His current research interests are in automated decision support systems, intelligent systems, IT infrastructure, and electronic commerce. His research appears in *Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, Accounting and Finance, Advances in Accounting Information Systems, International Journal of Intelligent Systems in Accounting, Finance and Management, Accounting Forum, Advances in Accounting Behavioral Research,* and *Critical Perspectives on Accounting,* among others.

Stewart A. Leech is a professor of Accounting and Business Information Systems at the University of Melbourne, Australia. He served as chair of both the Information Systems Section (1998-99) and Artificial Intelligence/Emerging Technologies Section (2001-02) of the American Accounting Association. He was cochair for Placement for ICIS 2000 and program chair for the International Research Symposium on Accounting Information Systems (2000), and is currently the Vice-President Asia Pacific of SIG-ASYS of the Association for Information Systems. His current research interests are in the areas of intelligent decision aids, decision making in corporate recovery, and enterprise systems. His research appears in *Decision Support Systems, Abacus, Australian Economic Review, International Journal of Accounting, Accounting and Business Research, Accounting and Finance, British Accounting Review, Journal of Business Finance and Accounting, Australian Accounting Review, Accounting Forum, Journal of Information Systems, Advances in Accounting Information Systems, International Journal of Intelli gent Systems in Accounting, Finance and Management, International Journal of Digital Accounting Research,* and *Advances in Accounting Behavioral Research,* among others.

Steve G. Sutton is KPMG Professor of Accounting at the University of Central Florida and a professorial fellow of Accounting and Business Information Systems at the University of Melbourne. He currently serves as the editor of the *International Journal of Accounting Information Systems* and formerly served as a departmental editor for *DataBase.* His research interests are in the areas of the impact of advanced information technologies on individuals, organizations, and society. His current focus is on the impact of intelligent decision aids on human judgment and decision making, and assurance models for electronic commerce in business-to-business relationships. His research appears in *Accounting and Business Research, Accounting and Finance, Auditing: A Journal of Practice and Theory, Behavioral Research in Accounting, Database, Decision Sciences, International Journal of Intelligent Systems in Accounting, Finance & Management, Journal of the American Taxation Association, Journal of the Association for Information Systems, Journal of Emerging Technologies in Accounting,* and *Journal of Information Systems,* among others. He is also the author or editor for four research monographs.