STRUCTURING KNOWLEDGE ACQUISITION THROUGH GENERIC TASKS: A CASE STUDY IN HINDSIGHT

by

Rajan Srikanth
Leonard N. Stern School of Business
Information Systems Department
New York University
90 Trinity Place
New York, NY 10006

July 1989

Center for Research on Information Systems
Information Systems Department
Leonard N. Stern School of Business
New York University

Working Paper Series

CRIS #213
STERN #89-82

Note: An earlier version of this paper was presented at the ICIS Doctorial Consortium, Pittsburgh, December 1987.

Abstract

Knowledge Acquisition is widely recognized as the single major bottleneck in the commercialization of Expert Systems technology. The typically ad-hoc choice of techniques for eliciting and representing expert knowledge, makes Expert Systems development expensive and prone to failure. Arguments have been made in the Knowledge Acquisition literature for performing an epistemological or "knowledge-level" analysis to "structure" the knowledge elicitation process. The need of the hour is for an empirical evaluation of these claims. In this paper, we present the results of a study that evaluates an approach to Structured Knowledge Acquisition, that is based on analyzing expert behavior using generic problem-solving tasks. Data from a large Expert Systems project currently nearing completion, has been used for the study.

1. Introduction

Knowledge Acquisition (**KA**) refers to the process by which *relevant knowledge* in the domain of interest, is elicited from a *knowledge source*. "Relevant" knowledge may take the form of: (1) **propositions** about objects of interest ("facts" or "assumptions"), (2) **relationships** between objects ("rules") and (3) **procedures** for using knowledge to solve problems. Human experts, texts, manuals, and cases etc., are examples of knowledge sources.

The development of techniques for *effective* knowledge acquisition (KA) is receiving increasing attention from researchers and practitioners alike. KA is often considered the biggest bottleneck in the commercial exploitation of Knowledge-Based or "Expert" Systems (ES) technology [Barr & Feigenbaum 82]. Of late, several conferences and special issues of prestigious journals have been dedicated to research on this topic (see References). KA research also has strong implications for Information Requirements Determination (IRD) in traditional systems development, and for the design of Decision-aiding or Decision Support Systems (DSS).

1.1. The Problem

Criticism has often been levelled at available Knowledge Acquisition techniques that they are adhoc and do not provide guidance on the following *key aspects* of KA:

1. what knowledge is *relevant*,
2. how it may be *elicited*, and
3. in what form it may be encoded (or *represented*).

In most instances, KA is typically driven "backwards" from the choice of *Knowledge Representation* (**KR**) scheme. For example, Knowledge Engineers are known to ask experts to articulate the "rules" they are using, when infact, part or all of their expertise may not be effectively expressible in that form. KR schemes are constrained by the "language" (rule-based, frame-based, etc.) in which they are implemented. The primitives of these languages do not provide the higher-level constructs necessary to effectively model different kinds of problem-solving. As a result, KA tends to be haphazard and ineffective.

Though several methodologies have been proposed for KA, they have not been subjected to empirical evaluation because traditional experimental strategies are inapplicable. The increasing interest in the commercial exploitation of these technologies, on the other hand, makes empirical evaluation an urgent and pressing necessity.

1.2. Epistemological Analysis

Newell (1982) argues that *epistemological* or "knowledge-level" analysis (of problem-solving behavior) is a necessary base for effective KA. An epistemological analysis makes use of knowledge-level **primitives** for characterizing different types of knowledge, their organization, and use [Clancey 83, Clancey 85]. It is considered a valuable first step in KA, since it provides a "road map" for the knowledge acquisition effort [Chandrasekharan 86, Hayward et al. 86]. Researchers have argued that "structuring" KA in this manner, would serve a purpose similar to Structured Analysis in traditional systems design [Weilinga & Breuker 84]. In particular, it would result in: (1) improved modeling/representation of knowledge, and (2) leveraged (expert) system design.

Several frameworks have been proposed for epistemological analysis to aid knowledge acquisition. They are based on different assumptions about what constitutes the "appropriate" set of *primitives* for knowledge-level analysis. Types of *mental structures* [Olson & Rueter 87, Rousse & Morris 86], *knowledge types* [Gammack & Young 84, Wright & Ayton 87], and *task types* [Chandrasekharan 86, Weilinga & Breuker 84, Stefik et al. 82], are examples of classes of primitives that have been suggested.

1.3. This paper

The objective of this research is to explore, through an empirical study, the benefits of using an epistemological analysis to structure KA. In particular, we are interested in the feasibility of performing an epistemological analysis of relatively complex expertise, and its effects on the "content" of elicited knowledge.

The approach (to epistemological analysis) that we have chosen to investigate, makes use of "generic" types of problem-solving tasks to model expert knowledge [Chandrasekharan 86]. *Classification, evaluation of hypotheses, plan configuration, design,* etc., are examples of **generic tasks**. Strong evidence in Cognitive Psychology literature for the use of "task types" to study problem-solving behavior, and the additional leverage for KA provided by *generic tasks,* are among our reasons for choosing this framework.

The methodology we use in this study, is best described as a qualitative analysis of a case study, in hindsight. The development of An Expert System for Options Pricing (AESOP) - a project we are currently undertaking for the American Stock Exchange - was the case chosen for investigation. Transcripts of knowledge acquisition sessions from the project, were *retrospectively* analyzed using the generic task framework. Prescriptions for KA generated from this analysis, are then compared with the transcripts, to identify benefits that may derive from structuring KA using generic tasks.

The rest of this paper is organized as follows: section 2 presents a brief review of the state-of-the-art in Knowledge Acquisition. In section 3 we outline the "generic tasks framework," and present arguments for using this framework to structure KA. A description of the study follows in section 4, and a discussion of the findings is offered in section 5. Finally, section 6 summarizes the limitations of our methodology, and our conclusions about the findings of the study.

2. Knowledge Acquisition

Expert Systems (ES) are computer programs that exhibit expert-level competence in a particular domain. They do so by providing representations for encoding expert knowledge, and reasoning mechanisms for using this knowledge to solve problems. The process of gathering such knowledge from the expert, encoding, and making use of it to build an ES, is referred to as **Knowledge Engineering**.

Knowledge Acquisition (KA) is the part of Knowledge Engineering that involves "identifying the problem, selecting an expert, eliciting the expertise, codifying it in some representation and, further refining the *knowledge base*" [Clancey 84].

Observation, informal or semi-structured interviews, repertory-grid techniques, and protocol analysis [Ericsson & Simon 84] are the primary techniques used for KA [Waterman 86, Boose 85, Kuipers & Kassirer 83, Grover 83]. A typical KA project consists of a series of "interviews" with an expert, aimed at *eliciting* and refining problem-solving knowledge.

2.1. The Knowledge Acquisition Bottleneck

The process of knowledge acquisition has often been cited as the single major bottleneck in the development of Expert Systems [Barr & Feigenbaum 82, HayesRoth et al. 83]. The source of this problem arises from one or more of the following:

1. Expert knowledge is typically unavailable for introspective recall ("compiled knowledge" [Newell & Simon 72, Anderson 81]),
2. Knowledge that is available for recall, sometimes cannot be effectively articulated, either because it is "spatial" rather than "verbal" [Wickens84, Larkin & Simon 85], or because the representation "language" does not provide the necessary *primitives*.
3. The expert may either be unaware of his real "expertise" [Collins 85], or cognitive biases may affect the process of knowledge articulation [Cleaves 86].
4. Last, but perhaps most important, the Knowledge Engineer may not know what to "look" for, and how best to elicit knowledge.

State-of-the-art KA techniques take the view that human limitations as "eliciters" and encoders of knowledge, lie at the root of all KA difficulties. Computerized systems have been built that attempt to automate knowledge elicitation, and eliminate the mediating role of the Knowledge Engineer. Examples of such systems include TEIRESIAS [Davis & Lenat 82], EMYCIN [vanMelle 79], implementations of the ID3 algorithm [Quinlan 79], SEEK [Politakis & Weiss 84], ETS [Boose 85, Boose & Bradshaw 87], MOLE [Eshelman et al. 87], etc.

To use an analogy from the communications world, removing the Knowledge Engineer might help eliminate "noise" in (knowledge) transmission, but it cannot ensure that the appropriate material is "transmitted," and correctly "interpreted." We believe that the need of the hour, is for a method to assist KA by identifying what knowlede is "relevant," and providing prescriptions for how it may be elicited and encoded. In the next subsection, we present an overview of research that addresses this concern.

2.2. Epistemological Analysis for Knowledge Acquisition

Epistemological analysis or knowledge-level analysis, refers to the process of using a set of "knowledge primitives" to characterize expertise in terms of: (1) its *structural* properties - "form" and "organization," and (2) patterns of acquisition and *usage* (to solve problems). A knowledge-level analysis, performed prior to KA provides a "road-map" for structuring knowledge elicitation, and a framework for choosing appropriate elicitation techniques. This section reviews research that seeks to provide the "appropriate" set of **primitives** for such an analysis.

One group of researchers recommends the use of "types of knowledge" as primitives for analyzing expertise, and structuring KA. Gammack and Young (1984) suggest classifying knowledge along lines of "function" into *concepts* and *relations*, *procedures*, *facts*, *heuristics*, and *classificatory knowledge*. Wright and Ayton (1987), argue for a more generalized functional categorization in terms of *declaritive* and *procedural* knowledge. Olson and Rueter (1987), on the other hand, suggest that primitives describing the "mental organization" of knowledge, are best suited for aiding KA. *Lists*, *tables*, *hierarchies*, and *physical spaces* are examples of such "mental structures." Approaches based on knowledge type are of limited usability for KA. In order to apply them, a greater degree of intimacy with the expertise under investigation, is needed prior to knowledge acquisition, than can be reasonably assumed.

A second group of researchers argue that a framework for describing and eliciting knowledge should be based on observable characteristics of problem-solving behavior. Such an approach avoids some of the practical difficulties encountered with the use of knowledge types. Rousse and Morris (1986) recommend choosing between inferential and verbalization techniques for knowledge elicitation, based on the "awareness" (*implicit* or *explicit*) and "control over problem-solving" observed in the expert's behavior. The proposals of Stefik et al. (1982), and Hayward et al. (1986) are more specific. "Problem types," such as *interpretation*, *diagnosis*, *monitoring*, *prediction*, *planning* and *design*, are suggested as primitives for describing problem-solving behavior. Stefik et al. argue that identification of "problem type" provides prescriptions for knowledge acquisition, enabling the Knowledge Engineering team to "focus on issues that relate to critical steps in reasoning." Studies indicate however, that problem types may not be the "appropriate" level of *primitives* for KA. Each problem type is not necessarily characterized by a unique problem-solving behavior; for example, [Pople 82] shows how diagnosis may be thought of either as a process of *designing / synthesizing* an explanation, or as a task of *classification*. Ambiguity of this kind, weakens the argument for these approaches since unequivocal prescriptions for KA cannot be derived.

Bylander and Chandrasekharan (1987) develop a framework for knowledge acquisition, based on the "theory of generic tasks" [Chandrasekharan 86]. In brief, *generic tasks* are task-level primitives (rather than problem-level primitives) that may be used to analyze human problem-solving behavior. A "task" is a unit of problem-solving activity that is characterizable by the type of *input*, and *output*. If expertise is broken down into its "constituent tasks," the unique properties of each task type may be used to derive specific guidelines for KA. In the next section, we review this approach and advance arguments for selecting it as the method of choice for performing knowledge-level analysis.

3. Theory of generic tasks

A **generic task** is a task-level unit of problem-solving activity that is characterizable in terms of:

1. the types of information *input* and *output*,
2. the nature of knowledge needed for the task - in particular its *form* and *organization*, and
3. the family of *control regimes* or problem-solving strategies, observable in solving such tasks.

Chandrasekharan (1986) identifies six *generic tasks* (the bracketed terms are our simplified explanations): *hierarchical classification* (categorizing), *hypothesis matching or evaluation* (determining "fit"), *knowledge-directed information passing* (invoking schemas), *abductive assembly of hypotheses* (finding "best" explanation), *object synthesis by plan selection and refinement* (planning/design), and *state abstraction* (qualitative simulation).

A brief description of these *generic task* types is offered below, in terms of the input, output, type of knowledge, organization, and problem-solving strategy that characterize each of them.

- **Hierarchical classification** involves determining the relevant hypothesis, or category, given a description of the state (symptoms, etc.). Categories are arranged hierarchically with the child nodes representing more precisely defined subcategories. Problem-solving proceeds by an "establish-refine" strategy until all known facts have been applied. Classifying a patient's disease(s) from the symptoms manifested, is an example of hierarchical classification.
- **Hypothesis matching/evaluation** takes the form of determining how well a given category/hypothesis "fits" data describing a given state. This "degree of fit" or certainty is often computed from the extent to which to which intermediate hypotheses fit with raw data. It is typically symbolic, though numeric values may sometimes be used. The following are examples of hypothesis matching or evaluation: determining the "likelihood" of jaundice given yellow eyes and bilirubin in the blood; or determining the "appropriateness" of using a special-purpose machine given the need for lower costs, and the awareness that there will be high production volume.
- **Knowledge-directed information passing** refers to tasks where knowledge needed to make an inference is deduced from other existing knowledge. Knowledge concerning each "data item" -.defaults, alternative procedures for determining values, etc. - is stored in a "frame." When a value is sought for the data item, the frame is invoked and the information stored in it is utilized. As an example, imagine the following: it is desired to determine if there is an impending increase in the price of a stock option. The "frame" associated with the option's "price" contains the knowledge that it goes up with stock price. To apply this knowledge, the "trend" of the stock price (whether it is going up or down) is inferred from the volume of trading in the stock, and the "rule" *stock price increases with demand.*
- **Abductive assembly of hypotheses** is the determination of the minimal set of hypotheses that "best explains" a given a set of symptoms (or state description). Knowledge of relationships (causal and other), and interdependencies among hypotheses is made use of in this task. Problem-solving alternates between: (1) *assembly* - where the best hypothesis that explains the remaining symptoms is added to the "composite hypothesis" and, (2) *criticism* - where superfluous and incompatible hypotheses are removed from the composite hypothesis. Consider the following example: given that an automobile engine does not start, one may hypothesize that there is no fuel. But on considering that the horn does not work, a second hypothesis "electrical system failure," is added as another possible explanation. Critical examination then "removes" the first hypothesis since electrical failure explains both symptoms.

- **Object synthesis by plan selection and refinement** concerns the design of an "object" (physical or abstract, like a device or plan) to satisfy given specifications. Knowledge about the object is represented as a *component hierarchy*; each child node represents a sub-component of the parent and has an associated set of properties (specifications, procedures, default values for parameters, constraints, etc.). Problem-solving proceeds recursively by either decomposing the object into its components, or "recognizing" and using known solutions. Failure necessitates "replanning" at higher levels. The development of a manufacturing plan to produce a device, involves specifying the component parts, plans to produce them, etc., and is an example of *object synthesis.*
- **State abstraction** involves qualitatively "simulating" the effects of a given change, on the state of the system. "Compiled knowledge" in the form of if-then-else rules is used to abstractly represent the functioning of the system and/or its subsystems. Problem-solving proceeds "bottom-up," with the effects of changes being propogated up through the levels of abstraction. Predicting the effects of o-ring failure on the space-shuttle's booster rockets, is an example of this type of task.

Figures 1 through *5* provide more detailed descriptions, and illustrations of these generic tasks.

3.1. Arguments for using the generic task framework for KA

Several arguments may be made for using *generic tasks* as primitives for *epistemological* or knowledge-level analysis. Firstly, they have strong empirical basis as units for the study of human problem-solving behavior. Research in Cognitive Science has shown that there exist different "problem schemata" (or types of tasks) in the problem-solving behavior of experts [Hinsleyetal78, Reif79, Novak & Araya 80, Chi et al. 82].

Second, *generic tasks* have "descriptive power"; they can be used as "building blocks" to describe even complex reasoning processes. Different kinds of expertise may be modelled as "compound" *generic tasks.*

Thirdly, they provide leverage for KA. By identifying the "type" of problem being studied, the Knowledge Engineer obtains guidelines on what sort of problem-solving behavior to expect, and the nature of knowledge underlying it. To quote Chi et al., "the knowledge useful for a particular problem is indexed when a given (physics) problem is catagorized as a specific type." The theory of generic tasks, applies this insight to <u>all</u> problem-solving behavior by explicitly associating a certain form, organization of knowledge, and problem-solving strategy with each *generic task.*

4. The study

In this paper, we have thus far advanced theoretical arguments in favor of epistemological analysis for structuring knowledge acquisition, and in particular, for the use of the generic task approach. We attempted to *empirically* verify the benefits of this approach through an exploratory study. The sections that follow, report the results of this study.

HIERARCHICAL CLASSIFICATION

Input: Description of state (symptoms)

Output: Relevant categories or hypotheses

Typically

- knowledge is expressed as hypotheses
- is organized as a classification hierarchy, child nodes represent subhypotheses or subcategories of the parent
- each node contains knowledge that establishes its relevance (hyp. matching)
- problem-solving proceeds by an "establish-refine" strategy until all observations are explained

an example:

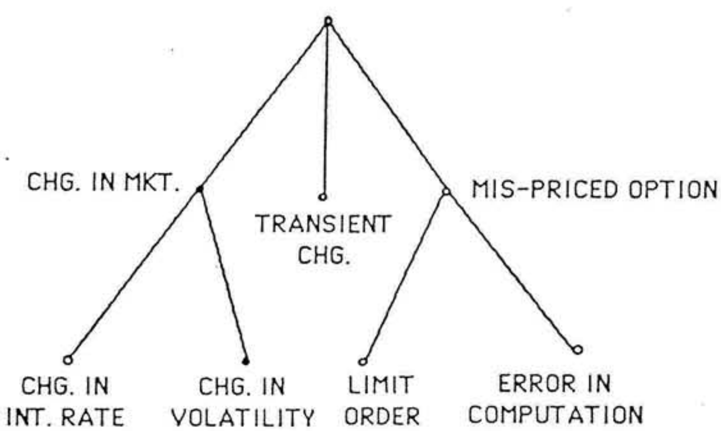

GIVEN:

* an increase in order flow for a near term at-the-money option
* that is sustained even when the price is increased marginally
* the absence of any limit order, and
* no change in interest rates

to determine that the relevant hypothesis is that there has been a change in volatility

Figure 1: Generic tasks - Hierarchical classif

HYPOTHESIS MATCHING OR EVALUATION

Input: A set of observations and a hypothesis

Output: Certainty of hypothesis, or
degree of fit / appropriateness

Typically

- knowledge is in the form of symbolic abstractions of degree of fit
- is organized as a hierarchy of such abstractions
- overall degree of fit is computed from intermediate degrees of fit or raw data

example:

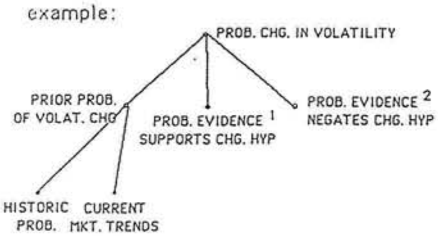

GIVEN:
* a set of observations and the hypothesis that there might have been a change in volatility of the stock, and
* a "high" prior probability of volatility chg.
* that evidence[1] "strongly" supports the hyp.
* that evidence[2] "weakly" negates the hyp.

to conclude that the hypothesis is supported strongly

Figure 2: Hyp. matching

KNOWLEDGE-DIRECTED INFORMATION PASSING

Input: a datum or fact, to be determined

Output: a value/state for the datum

Typically

- knowledge takes the form of default values, procedures for inferring values, etc.
- is organized as a frame in a frame hierarchy
- value for datum is computed by accessing frame for default values or procedures to infer value

example:

Given:
* the piece of knowledge "if over-exposed then narrow bid-ask spread"
* to determine over-exposure by invoking the frame for the datum, and
* finding the following rules to infer a value
  "if sell orders vastly outnumber buys
  then you are over-exposed"
  "if you are inventory shows you are long on stocks and the stock price is decr
  then you are over-exposed"

to infer that you are over-exposed because the stock price just took a big dip, and so to narrow the bid-ask spread

Figure 3: Knowledge-directed information passing

ABDUCTIVE ASSEMBLY OF HYPOTHESES

Input: a set of symptoms

Output: a minimal set of hypotheses that
covers the data

Typically

- knowledge takes the form of hypotheses and relationships between them (causal, special-case, mutual exclusion, etc)
- problem solving is a two-stage process
  - hypotheses generation, followed by
  - assembly - a subset of the generated hypothesis is identified that satisfies the maxim of "best coverage"

example:

GIVEN:
* increasing demand for an option series and
* the piece of knowledge that
  "increase in demand for an option indicates that there may have been a change in volatility of the underlying stock"

conclude that
  there has been a change in volatility

also given
* increasing demand for the underlying stock
* and the piece of knowledge
  "increase in demand for the stock results in increasing stock prices"

conclude:
  that the stock prices would increase

In assembling these two hypotheses together conclude
  that the change in demand for options is a result of change in stock price, not a change in volatility - no deviant change

Figure 4: Abductive assembly

OBJECT SYNTHESIS BY PLAN SELECTION AND REFINEMENT

Input: a set of specifications or
desired characteristics

Output: an object/plan that meets the reqmt

Typically

- knowledge takes the form of components, plans that set values for parameters of these, and constraints
- is organized as a component hierarchy
- object design proceeds by generation of a plan at each level, design of sub-plans, and backtracking when constraints are violated

example:

GIVEN
that there is a large demand for a particular series,
to come up with a plan for re-pricing
To re-price, one must
- determine if there has been a change in volatility
- if yes
  - determine new volatility
  - re-price all options using new volatility

  else
  - increment price of affected option only
- if new prices result in overexposure, adjust prices
- represent any active limit or spread orders

Figure 5: Object synthesis

4.1. Study Design

A modified form of the case study method was used. Data collected during "traditional" KA in an expert systems project, served as the starting point for this study. This data was "retrospectively analyzed" using the generic task framework. The specific objectives of the study were: (1) to establish the feasibility of using generic tasks to describe complex problem-solving, and (2) to uncover differences in the "content" of knowledge that (may be) elicited using the two approaches.

4.1.1. The AESOP project

An expert systems research project, sponsored by the American Stock Exchange and currently underway at New York University, was chosen for investigation. In this project, we have undertaken to build, implement and evaluate An Expert System for Options Pricing (AESOP)[1]. The objective of knowledge acquisition in this project, was to elicit and encode the knowledge and reasoning processes, used by an expert "market maker" to generate price quotes for stock options.

4.1.2. Method

Data collected during the knowledge acquisition phase of the AESOP project was used in this study. This data was in the form of transcripts of audio recordings, from 12 knowledge acquisition sessions of 1-2 hours each, conducted over a period of 4 months. A mix of strategies varying from informal, semi-structured interviews, to think-aloud protocols and "forward simulations," were used in the true spirit of "traditional," unstructured approaches to KA.

From an analysis of the transcripts, a general model of the expert's option pricing behavior was developed (*Figure 6*). A description of this behavior in terms of its "constituent tasks," was obtained by analyzing this model using the framework of *generic tasks*. The prescriptions of the generic task framework were then used to draw inferences about the knowledge and problem-solving strategy, underlying expert behavior. For each constituent task, this involved a description of: (1) the "form" of knowledge (facts, hypotheses, relationships, heuristics, or procedures, for example), (2) its "organization" (hierarchical, frame-based, script-like, causal structures, list, etc.), and (3) the "control" or problem-solving strategy (establish-refine, hypothesis generation, testing and criticism, forward and backward chaining, constraint propogation, synthesis, decomposition, etc.).

The original transcripts were then "retrospectively" analyzed, to identify occurrences of information pertaining to each constituent task. The knowledge elicited for each task, was compared with the descriptions obtained earlier from the generic task framework, to yield insights concerning differences in content.

[1]The author, and two members of the faculty, are currently involved in the AESOP project.

4.2. A general model of the expert's option pricing behavior

A stock option (or "option," for short) is the right to buy or sell the underlying stock (or *security*) at a specific price, called the "strike price," before a specific date - the third Saturday of the "expiry month." An option *series'* is designated by specifying its expiry month, strike price, and whether it is a "put" or "call" (the right to sell or to buy, respectively). For example, "XYZ March 50 Put" is an option series that designates the right to sell 100 shares of XYZ Corp. for $50, anytime before March 21.

All trading in options on a given stock is routed through a "market maker." His job consists of determining "fair" prices to charge buyers ("ask" price), or to offer sellers ("bid" price) for different option series', and to post them as "quotes" on the board. A "fair" market is one where: (1) all customer orders (for buying and selling options) are represented, (2) prices are so fixed that an equilibrium level of "activity" exists in the market, and (3) exchange regulations concerning pricing, and "spreads" between bid and ask prices, are not violated. The market maker also trades in the options himself, buying and selling to hedge his risks, and making some profits in the process. A formal model known as the Black-Scholes formula, provides an estimate of the "theoretical value" of each option series', based on the price of the underlying stock, prevailing interest rate, and "volatility" (the expected rate of change of the stock price). Using the theoretical value as a starting point, the expert quotes prices from a consideration of market forces, his market "opinion" (whether it is going "up" or "down"), the inventory position in his trading account, and exchange regulations. His pricing policy is to maximize profit while maintaining a fair market.

From observations of our expert, we formulated a model of a market maker's options pricing behavior (*Figure 6*). External events that may indicate a fundamental change in the market, trigger options (re)pricing. A change in the price of the stock, an increased or decreased "order flow" for certain series' of options, and a news flash, are examples of such events. When such events occur, the expert evaluates their "significance," and if necessary, revises his opinion of the market. He also reviews the inventory in his options trading account, and determines the risk exposure of his current position. Oppurtunities for profit, created by market conditions, are then sought out and exploited as far as possible.

If the expert judges that the external event is not indicative of a change in the market, he tries to maintain the current market equilibrium and also hedge his risk exposure. He may adjust the prices or spreads of a few options, or buy/sell options or stock for his own trading account. If he decides that the market is in fact moving, he re-evaluates his assumptions - particularly estimates of volatility - and recomputes the theoretical prices, using the Black-Scholes formula, in an attempt to establish a "new" equilibrium.

In either case, revisions in option prices are transformed into new quotes by generating appropriate bid-ask spreads. Before a quote is posted, it is typically checked to ensure that it does not create a potential for arbitrage, and conforms to exchange regulations.

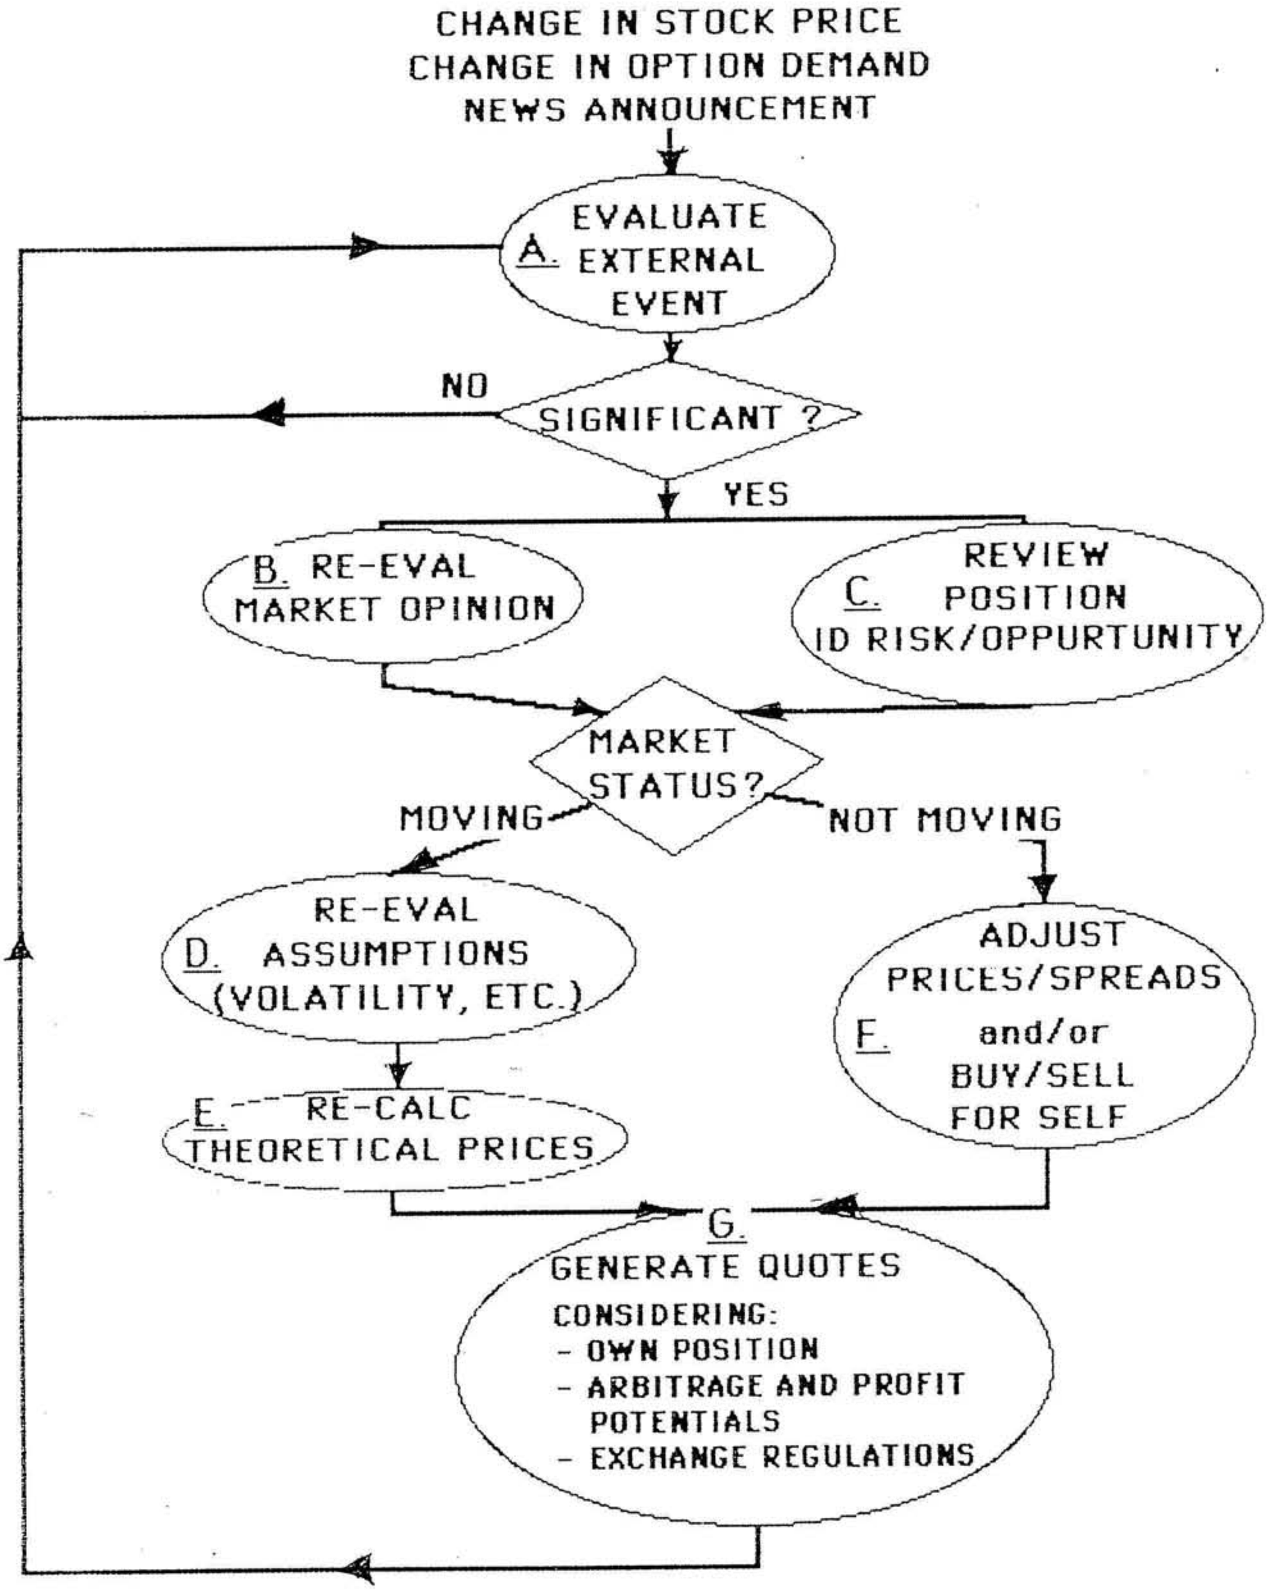

Fig.6 Options pricing behavior of expert

4.3. Description in terms of generic tasks

We used the model shown in *Figure 6* as a general description of the expert's option pricing, and analyzed it in terms of generic problem solving tasks. *Figure 7* summarizes the results of this analysis.

When an external event occurs, the expert first categorizes it (a classificatory task), and then evaluates the hypothesis that it is significant, under the circumstances. If the event is judged significant, he generates hypotheses about what might have caused the change, and "abductively assembles" them to find the "best" explanation. Based on this explanation, he visualizes the future state of the market, and forms an opinion - a state abstraction task. He then determines how much "exposure" he is willing to undertake, by invoking the relevant methods from a schema, and taking into account expectations of future market states. He also classifies his current inventory position in terms of "how exposed" it is, and "synthesizes" a plan to buy/sell stock or options, to make his inventory position more satisfactory.

Having made his position secure, the expert turns his attention to the market at large. If in his opinion, the market as a whole is moving, he must find a new set of prices that will help establish a market equilibrium. To do this he synthesizes a plan to re-evaluate the market volatility. Once the volatility has been determined, he recomputes the theoretical values of all the options from the Black-Scholes formula. If on the other hand, the expert judges that the external event does not indicate a fundamental movement in the market, he attempts to maintain the current market equilibrium, and minimize his risk exposure. He selects from a set of alternative plans, and decides to either increase the spread, change bid and asked prices, buy/sell stock, or do a combination of these. In order to choose among the alternatives, he simulates the "effects" of each (a state-abstraction task), and determines his preferences.

A revision in option prices implies that new quotes must be generated. The expert determines the appropriate "spreads" by invoking schemas and considering customer orders and exchange regulations. If potential for aribtrage is detected in the quotes, he reprices to eliminate arbitrage against him, and tries to exploit arbitrage that is in his favor. The quotes are then posted to the "board" and the options pricing cycle queisces, awaiting the next "event."

4.4. Prescriptions derived from generic task framework

The generic task framework postulates that each task "type," is characterized by a certain form and organization of knowledge, and problem-solving or "control" strategy. If a certain task "type" can be identified in the problem-solving behavior of an expert, "prescriptions" for eliciting that part of the expertise, may be derived from these postulates.

In the last section, we performed an epistemological or "knowledge-level" analysis of our expert's options pricing behavior, using generic tasks. *Figure 8* summarizes the prescriptions for KA that were derived for each *constituent task* identified in this analysis. The underlying implication is that these prescriptions leverage knowledge acquisition, by providing guidance to the Knowledge Engineer about <u>what</u> knowledge he or she must "look for."

	TYPE OF GEN.TASK
A. Evaluate external event	
- categorize external event	heir. classification
- determine if it is significant	hypothesis eval.
B. Re-evaluate market opinion	
- generate explanation for event	generate hyp. & abductively assemble
- visualize implications for future	state-abstraction
C. Review position	
- determine risk preference	schema
- obtain all relevant info.	invocation.
- classify current position	classification
- if unsatisfactory, synthesize a plan to remedy it	plan synthesis
- look for profit pot. in current market	
- formulate strategy, buy/sell	
D. Re-evaluate assumptions	
- if market is seen as moving, synthesize a plan to reevaluate volatility, effects of dividend income, etc.	plan synthesis selection & execution
E. Recalculate theoretical prices	
- run Black-Scholes model	
F. Adjust prices/spreads ...	
if market is not seen as moving,	
- synthesize a plan to limit exposure and hedge risks in face of external event	plan synthesis, state-abstractior plan selection &
G. Generate quotes	
- generate spreads	schema invocation
- look for violation of exchange regulations - min/max spreads, limit-orders, spread-orders, pricing in eigth's/sixteenth's	
- look for potential arbitrage synthesize & execute a plan to eliminate such oppurtunities	plan synthesis, selection & execution

Fig. 7 Description in terms of generic tasks

A.1. Categorize external event

Type of task: hierarchical classification

Form of knowledge: Classes of external events, and characteristics that distinguish each event.

Organization of knowledge: Event classes arranged in a hierarchy from the general to the specific.

Control strategy: Proceeds by establishing a general category, refining it to a more specific classification, and so on.

A.2. Determine if external event is significant

Type of task: hypothesis evaluation

Form of knowledge: Rules (of an if-then-else form), and facts or hypotheses with degrees of belief associated with them.

Organization of knowledge: A collection of rules, facts and hypotheses.

Control strategy: "Chain backward" for information needed to evaluate (intermediate) hypotheses, and "forward" using available information, to generate hypotheses.

B.1. Generate explanation for external event

Type of task: abductive assembly of hypotheses

Form of knowledge: If-then-else rules (with associated certainty factors), representing past experience or rules of thumb.

Control strategy: Alternates between generating hypotheses, and consolidating by retaining only those that explain more of the observations, or are more plausible.

B.2. Visualize future market conditions

Type of task: state abstraction/simulation

Form of knowledge: Heuristic rules, scripts, or schemas.

Organization of knowledge: Collections of rules, or deeper "causal models".

Control strategy: Propogate assumed changes, and their consequences till no further inferences may be made.

C.1. Determine acceptable "exposure" (risk) level

Type of task: knowledge-directed info retrieval/schema invocation

Form of knowledge: Schemas containing alternative methods for assessing acceptable risk level, under different circumstances.

Organization of knowledge: A hierarchy of related schemas.

Control strategy: Schemas are triggered by current market opinion, current preferences, etc. Once invoked, a schema's methods are used to establish the acceptable risk level.

C.2. Classify current inventory position.

Type of task: hierarchical classification

Form of knowledge: Hypotheses about 'state of exposure' of current position.

Organization of knowledge: Hierarchical from the general "position type" hypotheses, to more specific ones.

Control strategy: Proceeds top-down by successively refining the classification of the current position.

C.3. Synthesize plan to remedy unsatisfactory inventory position.

Type of task: plan synthesis

Form of knowledge: Descriptors of current and goal state, known procedures and strategies for accomplishing particular goals, and heuristics or rules of thumb.

Control strategy: Use "known" procedure or strategy to accomplish desired goal, or decompose into subgoals and synthesizing plans for them.

D.1. Re-evaluate volatility.

Type of task: plan synthesis

Form of knowledge: Consists of alternative pre-defined, sequences of actions (since the problem is well defined and recurrent).

Control strategy: The plan that is most appropriate - given environmental factors - and least "expensive", is used.

F.1. Limit exposure and hedge risks.

Type of task: plan synthesis and state abstraction

Form of knowledge: State descriptors, heuristics, known strategies or procedures, and preferences.

Control strategy: Strategies are selected for application after simulating their effects on future exposure, and selecting one based on preferences.

G.1. Generate spreads.

Type of task: knowledge-directed info retrieval/schema invocation

Form of knowledge: Schemas containing methods for determining/modifying spreads under different conditions.

Organization of knowledge: Organized as an exception heirarchy of schemas.

Control strategy: A basic spread is adjusted based on current inventory position in the series, and expected future states.

G.2. Look for potential arbitrage and eliminate.

Type of task: plan synthesis

Form of knowledge: Assumptions about market conditions, known procedures/strategies for detecting potential arbitrage, and eliminating different arbitrage conditions.

Control strategy: Recognize existence of relevant arbitrage condition and select/configure strategies based on preference for risk exposure.

Figure 8: Prescriptions for knowledge acquisition derived from the generic task framework

Figure 9: A sampling from the retrospective analysis of the knowledge acquisition transcripts

Incidents pertaining to G. *Generate Quotes*

"We would just spread our bid in there where we would want to buy and sell around that theoretical value so that we, based on our evaluation of these options, were buying a little cheaper than what they were worth, or selling them a little bit more expensive than what they are worth. If we get to pick the right volatility, that would correspond to the real marketplace, out there we should be getting buyers and sellers who should be meeting a lot of people's idea of what they want to do in this.."

"..., if he (the customer) put in an order; the market when he came in was 1/2 or 3/4, he came in and wanted to buy at 5/8-a dollar, the market would then become 5/8-3/4 ... I do have an obligation to represent him in the market place, and if there is a seller that comes in, to make sure he buys it.."

".. $2.45 is what Black-Scholes has set .. That is the theoretical value. So understanding what the minimum increments are, what the option trades are, we would quote them up and down from $2.45."

"I may (also) want to be more flexible than what the theoretical values are, so I may want my bid in June 65's to be atleast $2 higher than May 65's. The reason for that could be many but it could be because of inventory that I have already in place and I want to move inventory out based on that ..."

"I trade options under $4, from zero to $1, (the spread) would be 1/8 point wide; from a dollar to $2, there would be 3/16th wide; from $2 to $4, I would keep them 1/4 point wide.

What do you mean by this 'spread?'

This is the spread from bid to offer. Above that ($4), I would spread them 3/8th's and go up to 1/2.

.. Is your decision based on this 'volatility?'

Based on liquidity, mostly liquidity. My ability to attract order flow. I am not so altruistic that I keep a narrow and narrower quote because it is nice to do that, and beneficial to the customer. You try to find a medium where you are going to be able to not harm yourself by quoting something so narrow that you are not going to make any money. Yet, you do not want to quote something so wide as to discourage customer order flow, and you do discourage order flow if he sees an option quoted 3 bid-offered 3 1/2. 'I am not going to buy this option because I am never going to be able to sell it ..'"

"Now the stock is up another 1/4th of a dollar, now I look into my inventory. I look and see if I am 'long', which means I own XYZ, and I am looking to sell to make 1/4 of a dollar. I would restructure my price accordingly. I may go only 8 1/4 - 8 1/2 and offer it at fair value at this point in time."

" .. what you also want to do is not put so much of a shock value into the quote changes because one strike is getting a demand over there, ... the farther out months are much more sensitive in terms of the dollar pricing because they are higher priced options, and any changes in the formula, any variables, will have a greater effect on them - on the farthest out series. That quote was 4 1/2-3/4, I might be very reluctant all of a sudden to change it to 5 1/4 so I might just go 4 3/4-5 and see what that got me. Not suddenly change it .."

" ..you did not want to spread this thing too much, so you moved the bid also up?

Right. I was tying the bid to a 1/4 point less than the ask; the ask being one trading increment higher than where we had been selling them during the afternoon..

If for some reason, the demand held up and you had to raise the asking price again, would you have brought the bid up again?

Yes. If at a point in time, I decided that demand was great enough than my supply, then I would have perhaps have raised the offering price to 2 5/8 and the same way I would have raised the bid price to 2 3/8, keeping that spread 1/4 dollar wide."

"Specialists and market-makers have an obligation to maintain a 'fair and orderly' market. A 1/4 of a dollar wide on a $2 option is still approximately 10% of the bid, which is very sufficient to make a profit on without gouging. By the rules, we could make a market of 2-2 1/2. However, by having a wide mark like that, you discourage business. People dont like to say that if I buy at $2.50, if I decided immediately that I dont like my purchase, (and) I want to sell, I am going to 25% of my money without anything happening! Without any economic changes in the market place ..

Is there a fixed rule that you use?

I use the rule in my own mind, that under $5 - no more than 1/4 point wide, from $5 to $10 - 3/8 dollar wide, $10 to $20 - 1/2 a point, and $40 and above - $1."

"You may find a couple of discrepancies. To protect myself, I spread some options which are very, very inliquid and inactive. I spread them out a little wider than I normally would. Over here we have these puts which are 6 1/2-7, so here's an option where it is 1/2 dollar wide. ... there never is any activity in these options, just because of fluctuations in the stock. It protects me to keep the spread out a little bit wider."

" .. If I quote an option at 6 1/2-7, I dont have to focus my efforts on changing that market until a much greater incremental change in the underlying securities, than I normally would. It allows me to focus on trading the options where there is going to be (a lot of activity) .."

Incidents pertaining to A. *Evaluate external events*

".. there would be more to it than just looking at the underlying price (of the stock). (For example:) the last sale of the stock would be 67 and the stock would be quoted at a 67 bid-offered at 67 1/2; 20,000 shares wanted-500 shares offered. Which would give you a fair indication that the stock was going higher."

"... You are looking for significant movement on that. You are looking for an indication of real price change...

Does that come from your broker on the floor?

A lot of times, it just comes from them. Sometimes it is difficult to get accurate information from the market. You wait to see it happen.."

"As a rule of thumb, option activity goes from most of it in the near term, to the next month, etc., going out to the least. ... Off the top of my head I would say it was about 50-25-15-10 ... You are also going to get the most activity on the strike prices that are close to the price of the stock. The closest. Because that is where the greatest leverage is involved. There are exceptions. One customer can change all that around on one order. But as a rule of thumb, the near term at the money option, is going to have the greatest activity in it, and calls probably trade 70% of the volume."

"Volatility changes don't have to happen that many times in a day. We are just talking about quote changes. We key off that one option when we are talking about changing those variables that we are using in the formula, for caluculating it. That may happen 2 or 3 times a time or it may not change at all. It may not change for days.

Right. But do you have to track that it might happen?

Yes we look for significant order flow that is deviant from the activity of the stock. If we are getting demand in the option because we are getting demand for the stock then that doesn't change. Everything is just going hand-in-hand.. That is not acting deviant. When a stock is not changing and an option is changing much more significantly than the stock in one way or another, then you look at the method of calculating that (options) value and see if you can correct that."

"If XYZ is 67 1/2-3/4 and I have a quote up there that is accurately reflecting where my markets are, and the stock trades at 1/2 or 3/4, I am not going to change anything. I wouldn't change anything until the stock went from 1/2-3/4 to 3/4-8. At a time when it moved up a 1/4 of a dollar and the offer went up a 1/4, and it was trading that way, then I would reprice my options to reflect the increase in the value of the stock by a 1/4 of a dollar."

".. 1/8 of a dollar would be significant enough in a cheaper stock..."

Incidents pertaining to F. *Adjust prices/spreads*

".. if I sold the stock at 68 1/2, I would buy those options at 8 1/2, and what I would effectively be doing right there is I would be creating a credit balance in my account because I am selling $6850 worth of stock and buying $850 worth of options, creating a $6000 credit balance which I collect interest on through the expiration of that option, which is about 2 weeks from now. That is called an interest rate arbitrage. I would make a profit on that. That is a trading strategy. At that point in time, I would clear the book of that 8 1/2 offer."

"How do you decide which one (option series') to choose to lower your price if you want to lower your prices? Do you choose the near in the money options? What is your rationale for that? Suppose it is a low day and you are not getting any orders?

I am not necessarily going to force order flow in, just because I dont have any order flow. I may want to force order flow to hedge myself against order flow. I did not have the order flow in the 60's and 70's option series that were trading on the security. All I had was one. I might change my option prices to generate order flow for a hedging basis."

"... I was quoting the option 2 1/4 bid-offered at $2.50. The reason for that is that there was a demand for them. There were buyers coming in at 2 7/16ths all afternoon. We sold for our own account 100 options at 2 7/16ths. So to provide continuity for the last sale, we quoted them 2 1/4 bid-offered at 2 1/2 dollars. At that point, regardless there was no other way to hedge that other than to buy the underlying security. We quoted at 2 1/4 bid, we were willing to buy a certain amount of the options back at that point.

How did you arrive at the specific number of 2 1/4?

I attempted to maintain a certain spread on a $2 option. I did not want to widen that spread out by more than 1/4 of a dollar at that point in time..."

" .. what would happen if you started to get order flow for some reason?

Then I would narrow the quote.

Is that the only reason you would narrow the quote? What other reasons would you narrow the quote?

If I had a customer, bid or ask a limit order, that was willing to buy them in that range, I would represent his bid or ask order in the marketplace, which would narrow the quote..."

4.5. Retrospective analysis of knowledge acquisition sessions

In order to compare knowledge elicited in the conventional approach with the prescriptions of the last section, we retrospectively analyzed the audio transcripts of the KA sessions. They were systematically coded to identify occurrences of information pertaining to each constituent task.

An examination of the coded transcripts revealed a surprisingly small and unequal number of "incidents," pertaining to each of the constituent tasks. The number of incidents related to *G* (the Generate Quotes task), far surpassed those related to *A* (Evaluate external event), *D* (Re-evaluate assumptions), and *F* (Adjust prices); *B* (Re-evaluate market opinion), and *C* (Review position) were almost non-existent.

In general, information relating to any one given task was episodic, and found dispersed in discontinuous bits and pieces across the entire duration of the KA sessions. *Figure 9* provides a sampling of the results of this analysis.

5. Discussion of findings

The objective of this study was to empirically examine the benefits of structuring knowledge acquisition using generic tasks. In particular, we were interested in: (1) establishing the feasibility of using generic tasks to describe complex problem-solving, and (2) uncovering the comparitive differences in content of knowledge that (may be) elicited using generic tasks. For this purpose, we used a modified form of the case study method with data collected from an expert systems project. This section offers a discussion of the findings of the study.

A review of *Figure 6 & 7*, and the description provided in Section 4.3, demonstrates the feasibility of describing complex problem-solving behavior in terms of its constituent *generic tasks*. We found generic tasks to be "appropriate" for performing a knowledge-level analysis; it was possible to analyze a model of expert behavior and identify the constituent tasks, without much difficulty. However, in our opinion, the primitives provided by the generic task framework are not "adequate" to describe <u>all</u> problem-solving. Decision-making tasks such as articulation of utility functions, and choosing among alternatives for instance, are not included. Additions to the repertoire of "generic tasks" will have to be made if it is to be used to describe a wide variety of expertise.

A comparison of the prescriptions generated (*Figure 8*), with the results of the retrospective analysis of the AESOP knowledge acquisition sessions, offers some powerful insights. The discussion of these insights will be confined to findings about the "content" of elicited knowledge, since that is our primary concern in this study.

<u>First</u>, our study clearly demonstrates that traditional knowledge acquisition tends to be lopsided in its "coverage" of expert knowledge. In this particular case, the continued focus on the "Generate Quotes" task (perhaps for obvious reasons), resulted in insufficient attention to the other tasks.

Second, even for tasks that were given adequate attention, traditional knowledge acquisition scores low on "clarity"; a clear picture of how different pieces of knowledge tie together is not easily obtained. For example, consider the following portions of the transcript under the heading "Incidents pertaining to G. *Generating Quotes*," in *Figure 9*:

1. "I trade options under $4, from zero to $1, (the spread) would be 1/8 point wide; from a dollar to $2, there would be 3/16th wide; from $2 to $4, I would keep them 1/4 point wide.

What do you mean by this 'spread?'

This is the spread from bid to offer. Above that ($4), I would spread them 3/8ths and go up to 1/2."

2. "I may (also) want to be more flexible than what the theoretical values are, so I may want my bid in June 65's to be atleast $2 higher than May 65's. The reason for that could be many but it could be because of inventory that I have already in place and I want to move inventory out based on that ..."

3. ".. Is your decision based on this 'volatility?'

Based on liquidity, mostly liquidity. My ability to attract order flow. I am not so altruistic that I keep a narrow and narrower quote because it is nice to do that, and beneficial to the customer. You try to find a medium where you are going to be able to not harm yourself by quoting something so narrow that you are not going to make any money. Yet, you do not want to quote something so wide as to discourage customer order flow, and you do discourage order flow if he sees an option quoted 3 bid-offered 3 1/2. 'I am not going to buy this option because I am never going to be able to sell it ..'"

Each of these quotes demonstrates the expert's use of a different method for setting the bid-ask spread, for generating a quote. The transcript does not contain any information that indicates the different circumstances that might arise, and how the expert chooses which of these methods to use. As a result, it is unclear how these different methods relate to each other. The prescriptions from the generic task framework (G.1 *Generate Spreads* in *Figure 8*), on the other hand, indicate that these methods should be organized in a schema. If knowledge acquisition was guided by use of these prescriptions, it would be imperative to explicitly specify conditions under which each method is used.

Third, we also observe less "depth" in knowledge elicited by conventional approaches; detail necessary for effectively utilizing a piece of knowledge is often not obtained. As an illustration, consider the following (from *Figure 9*, "Incidents pertaining to A. *Evaluate external event*):

1. ".. there would be more to it than just looking at the underlying price (of the stock). (For example:) the last sale of the stock would be 67 and the stock would be quoted at a 67 bid-offered at 67 1/2; 20,000 shares wanted-500 shares offered. Which would give you a fair indication that the stock was going higher."

2. "... You are looking for significant movement on that. You are looking for an indication of real price change...

Does that come from your broker on the floor?

A lot of times, it just comes from them. Sometimes it is difficult to get accurate information from the market. You wait to see it happen.."

Quotations 1 & 2 show the expert describing two different types of external events, and how he determines if they are significant. In quotation 1, the expert provides an example to illustrate how he judges if the stock is going higher. However, there is no further elaboration of this concept, and so no generalizable knowledge is obtained about inferring stock movement from stock prices and demand figures. Likewise, in the quotation 2, the expert's declarations about "looking for significant movement," and "an indication of real price change" are not further investigated. If knowledge acquisition had been structured using generic tasks, the prescriptions for A.1 *Categorize external event* (*Figure 8*), would emphasize the importance of precisely defining the hierarchy of event classes, and establishing the means for discriminating between categories.

These observations concerning "coverage," "clarity," and "depth," lead us to conclude that the "content" of elicited knowledge can be significantly improved by using the generic task framework to structure knowledge acquisition.

We conclude this discussion by briefly touching upon another important advantage of performing epistemological analysis: it "focusses" the knowledge elicitation process. Retrospective analysis of the AESOP knowledge acquisition transcripts, revealed earlier that information related to any given task was found dispersed through the length of the transcripts. The unequal coverage, low clarity, and insufficient depth observed in the elicited knowledge, may all be attributed to the use of an approach that could not ensure focussed knowledge elicitation. An epistemological analysis, on the other hand, provides a "road-map" for structuring knowledge acquisition, and could therefore, help keep knowledge elicitation focussed.

6. Limitations of study and conclusions

6.1. Limitations of the methodology

The methodology we have used for this study is unconventional, and not without its share of problems. We discuss first the limitations of the study design, and next, the drawbacks of the methods used for data collection and analysis.

Our study was designed as a "case study in hindsight." We are aware of atleast two limitations of using such a design: (1) the limitations of case studies - limited generalizability of findings, and (2) the unreliability of using prescriptions born from hindsight, as determinants of future success. Analysis in hindsight involves speculation; findings from such analyses must be used with circumspection. In defense of our study design however, we point out that the case study method is the only means for empirically investigating knowledge acquisition. Since a case, by definition, cannot be replicated, the only way to study two different "treatments" is by a carefully performed "retrospective analysis."

The methods used for data collection and analysis also have certain limitations. Observations made during one "treatment" (conventional knowledge acquisition) was used as the basis for formulating the model of expert problem-solving. This model formed an integral part of the second "treatment"; it served as the basis for the epistemological analysis using generic tasks, from which prescriptions for KA were derived. As a consequence, the two "data sets" were not completely independent. Data analysis involved comparison of this "data" (prescriptions worded in very general terms) with results of the retrospectively coded transcripts of the KA sessions. No objective measure was used for the comparison, since in our opinion, none was warranted or feasible under the circumstances of this study.

6.2. Conclusions

Inspite of its methodological limitations, we believe that this study serves a useful purpose. Foremost, it provides preliminary evidence for the feasibility and benefits, of structuring knowledge acquisition through an epistemological, or "knowledge-level" analysis using generic tasks. In addition, the study demonstrates the use of an unconventional study design, for empirically evaluating knowledge acquisition techniques. Such an evaluation is critical for finding "good" techniques to help ease the Knowledge Acquisition "bottleneck," and facilitates the commercial exploitation of Expert Systems technology.

References

[Anderson 81] John R. Anderson (editor).
Cognitive Skills and Their Acquisition.
Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1981.

[Barr & Feigenbaum 82]
Barr, A. and E.A. Feigenbaum.
The Handbook of Artificial Intelligence.
Pitman Books, London, 1982.

[Boose 85] Boose, J.H.
A knowledge acquisition program for expert systems based on personal construct psychology.
Int. J. of Man-Machine Studies 23:495-525, 1985.

[Boose & Bradshaw 87]
Boose, J.H., and J.M.Bradshaw.
Expertise transfer and complex problems: using AQUINAS as a knowledge-acquisition workbench for knowledge-based systems.
International Journal of Man-Machine Studies 26:3-28, 1987.

[Chandrasekharan 86]
Chandrasekharan, B.
Generic Tasks in Knowledge-based Reasoning: High-Level Building Blocks for Expert System Design.
IEEE Expert (Fall):23-30, 1986.

[Chi et al. 82] Chi, M.T.H., P.J. Feltovich and R. Glaser.
Categorization and Representation of Physics Problems by Experts and Novices.
Cognitive Science 5:121-152, 1982.

[Clancey 83] Clancey, W.J.
The epistemology of a rule-based expert system: A framework for explanation.
Artificial Intelligence 20(3):215-251, 1983.

[Clancey 84] Clancey, W.J.
Knowledge Acquisition for Classification Expert Systems.
In *Proceedings ACM'84 Annual Conference: The Fifth Generation Challenge*, pages 11-14. 1984.

[Clancey 85] Clancey, W.J.
Heuristic classification.
Artificial Intelligence 27:289-350, 1985.

[Cleaves 86] Cleaves, D.A.
Cognitive Biases and Corrective Techniques: Proposals for Improving Elicitation Procedures for Knowledge-Based Systems.
In *Knowledge Acquisition for Knowledge-Based Systems Workshop*, pages 9.0-9.11. Banff, Canada, 1986.

[Collins 85] Collins, H.M.
Changing Order: Replication and Induction in Scientific Practice.
Sage Publishers, London, 1985.

[Davis & Lenat 82] Davis, R. and D. Lenat.
Knowledge-Based Systems in Artificial Intelligence.
McGraw-Hill, New York, 1982.

[Ericsson & Simon 84]
Ericsson, K.A. and H.A. Simon.
Protocol Analysis: Verbal Reports as Data.
Bradford Books/MIT Press, Cambridge, MA, 1984.

[Eshelman et al. 87]
Eshelman, L. et al.
MOLE: a tenacious knowledge-acquisition tool.
International Journal of Man-Machine Studies 26:41-54, 1987.

[Gammack & Young 84]
Gammack, J.G. and R.M. Young.
Psychological Techniques for Eliciting Expert Knowledge.
Research and Development in Expert Systems.
Cambridge University Press, London, 1984, pages 105-112.

[Grover 83] Grover, M.D.
A Pragmatic Knowledge Acquisition Methodology.
In *Proceeding of the Eighth International Joint Conference on Artificial Intelligence*, pages 436-438. 1983.

[HayesRoth et al. 83]
Hayes-Roth, F., D.A. Waterman and D.B. Lenat.
Building Expert Systems.
Addison Wesley, 1983.

[Hayward et al. 86]Hayward, S.A., B.J. Wielinga and J.A. Breuker.
Structured Analysis of Knowledge.
In *Knowledge Acquisition for Knowledge-Based Systems Workshop*, pages 18.0-18.6. Banff, Canada, 1986.

[Kuipers & Kassirer 83]
Kuipers, B., and J. Kassirer.
How to discover a knowledge representation for causal reasoning by studying an expert physician.
In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 49-56. 1983.

[Larkin & Simon 85]
Larkin, J.H., and H.A. Simon.
Why a Picture is Worth Ten Thousand Words.
Technical Report, Carnegie-Mellon University, Pittsburg, PA, August, 1985.

[Newell & Simon 72]
Newell, A. and H.A. Simon.
Human Problem Solving.
, , 1972.
.

[Novak & Araya 80]
Novak, G.S., Jr., and A.A. Araya.
Research on expert problem solving in physics.
Technical Report Tech.Rep.NL-40, University of Texas, Dept. of Computer Sciences, Austin, TX, June, 1980.

[Olson & Rueter 87]
Olson, J.R., and H.H. Rueter.
Extracting expertise from experts: Methods for Knowledge Acquisition .
Journal of Expert Systems (Summer):, 1987.

[Politakis & Weiss 84]
Politakis, P. and S.M. Weiss.
Using Empirical Analysis to Refine Expert System Knowledge Bases.
Artificial Intelligence 22(1):23-48, 1984.

[Pople 82] Pople, H.W.
Heuristic Methods for Imposing Structure on Ill-Structured Problems.
Artificial Intelligence in Medicine.
Westview Press, 1982, pages 119-190.

[Quinlan 79] Quinlan, J.R.
Discovering rules by induction from a large collection of examples.
Expert Systems in the Micro-electronic Age.
Edinburgh University Press, England, 1979, pages 168-201.

[Rousse & Morris 86]
Rousse, W.B., and N.M. Morris.
On looking into the black box: Prospects and limits in the search for Mental Models.
Psychological Bulletin 100(3):349-363, 1986.

[Stefik et al. 82] Stefik, M. et al.
The Organization of Expert Systems, A Tutorial.
Artificial Intelligence 18:135-173, 1982.

[vanMelle 79] van Melle, W.
A Domain-Independent Production Rule System for Consultation Programs.
In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 923-925. 1979.

[Waterman 86] Waterman, D.A.
A Guide to Expert Systems.
Addison-Wesley Publishing Company, Reading, MA, 1986.

[Weilinga & Breuker 84]
Weilinga, B.J and J.A. Breuker.
Interpretation of Verbal Data for Knowledge Acquisition.
In *Proceedings ECAI-84*, pages 41-50. 1984.

[Wright & Ayton 87]
Wright, G. and P. Ayton.
Eliciting and Modelling Expert Knowledge.
Decision Support Systems 3:13-26, 1987.