AUTOMATED VERBAL SUMMARY

FOR DECISION ANALYSIS

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF

ENGINEERING-ECONOMIC SYSTEMS

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN ENGINEERING-ECONOMIC SYSTEMS

Eric Richard Johnson

May 1995

UMI Number: 9535608

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

John Weyant (Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Michael Fehling

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Ross Shachter

Approved for the University Committee on Graduate Studies:

Abstract

While decision theory can provide a justifiable recommendation for action in any situation, this recommendation is satisfying only if the decision maker is confident that she has captured all the important aspects of her situation in a formal statement of the decision problem. Decision analysts can often foster satisfactory problem formulation by presenting analyses and insights from a preliminary model of the situation, thus guiding domain experts to identify places where the model can be improved most efficiently. However, eliciting such preliminary models efficiently and choosing how to present their results is, as yet, part of the unformalized craft of decision analysis. This dissertation formalizes a portion of the decision analysis process and presents a tool that supports it. The Deft decision formulation tool presented here, helps the analyst elicit a decision model from experts parsimoniously, and gives a verbal summary of its behavior in a way that supports critical revision of it.

I begin by reviewing the problem solving literature, to motivate my focus on the decision analysis approach. Next, I illustrate the use of Deft as a group decision support tool in a hypothetical decision scenario regarding the United States' synthetic fuels commercialization policy. I then discuss the architecture of Deft, which allows computer models to be used conveniently without sacrificing decision theoretic rigor. Finally, I motivate and describe both the logic behind Deft's elicitation of values for sensitivity analysis, and the novel way in which these are reported in Deft's verbal summary facility. I defend the former on grounds of parsimony, and the latter by reference to the literature of judgement and decision making, which shows the congruence of Deft's verbal summary to natural human ways of thinking about decision problems.

Acknowledgments

I gratefully acknowledge the support and guidance given me by my adviser, John Weyant. Even as my work has evolved away from his area, energy modeling, he has been a constant source of respect, encouragement, good humor, and good advice. I also thank Ross Shachter and Michael Fehling for serving on my committee. Ross encouraged me to try a variety of ideas early in the process, and one of the early experiments developed into the summarizer described in chapter 7. Michael always encouraged me to apply the fundamental notion of doing the best thing to decision analytic activities, as well as the rest of life. He and Ward Edwards provided me a vision of what it is to do research properly. I value that vision, and I have attempted to honor it here, even as I fell short of it. The kernel of the approach to sensitivity analysis in chapter 6 was communicated to me by Sam Holtzman. It's a shame that our schedules and interests did not allow us to collaborate on the further development of that idea. While each of these contributions has been important to me, the broadest and deepest influence on this dissertation is Ron Howard's approach to decision analysis. His identification of the essence of decision theory has shaped the concepts here, and his continual search for a better way to implement it developed the practices that this work explores. I hope I have taken a step toward research into Ron's ideas that Michael and Ward can respect.

I'd like to thank Susan Clement, Jackie Piozet and Helen Kohn for their help through the years. Their friendship and assistance ensured that I was able to treat this as a purely academic kayak ride, without foundering on administrative shoals. Another continuing source of support and friendship for me has been the group of people I studied with in our first year: Elisabeth Browne, Bill Poland and Andrea Kress. Finally, I would like to remember the delightful smile and the pleasant thoughtful manner of Veronica Chu, which have brightened my days here.

Chapter 1. Introduction

This chapter motivates and sets out the objective pursued in this dissertation. It then gives an overview of how this objective is pursued.

1.1 Complex, uncertain decisions

Situations that a Decision Maker (DM) regards as complex and uncertain are ubiquitous. The motivation of this work is my experience in the energy business, and while the examples given here are from this area, such situations can be found in many other domains.

No DM has perfect knowledge, hence each has some uncertainty. Aspects of the world that are typically uncertain enough and important enough to merit explicit probabilistic consideration in energy business decisions include: the overall demand for power, the price of oil, thermal resource availability and streamflow, and actions of new market players like independent power producers.

There are two main sources of complexity in decisions. The first is the difficulty of identifying important aspects of the situation at hand. Under these circumstances, the natural response is to create a detailed, and hence complex, simulation model that captures one's best understanding of all the possibly relevant factors. The second source of complexity is the fact that many solution approaches convert other features of a problem into complexity. For instance, a common way to solve dynamic problems is to capture a sufficiently large portion of the history of the system into system 'state' (as discussed in Luenberger (1979), chapter 4). This gives such systems the Markov property (future behavior is influenced only by the current state) and allows the use of straightforward solution techniques, but in the process it turns dynamism into complexity. Probabilistic conditioning and discretization allow the specification of joint probability distributions for decision-making under uncertainty, but they do so by generating a large number of cases to be considered.

Situations that are typically viewed as uncertain and complex in the energy business include planning short- and medium-term operations, setting rates, designing and pricing new energy products or delivery arrangements (e.g., interruptibility, time-of-use rates, priority service, or load management), entering into a purchase or sale contract, and construction of generation or transmission capacity.

1.2 Objective

This thesis seeks a way for uncertain decision-makers to come to know the best action at small cost when facing complex situations. I begin by discussing the terms in this statement of objective.

The preconditions for this work include the fact that the DM has imperfect knowledge about the world. In other words, she is not certain what will occur under all conditions. She may have beliefs about these outcomes, and some may be better justified than others. This dissertation pursues "adequately justified" belief of what the best action is. I use the word 'knowledge' to refer to adequately justified belief. This definition is consistent with some of the relevant epistemological literature, and it sidesteps the requirement that ostensible knowledge be "true", which would apparently preclude the use of the notion under conditions of imperfect information. I discuss what sort of justification is adequate later.

When I refer to insight, I mean knowledge with broad policy-making ramifications. Although other definitions have been offered, this definition is fairly consistent with most usage in the literature. Insight is a variant of knowledge that would be well suited to our interests in this work.

My work seeks knowledge of what the best action is. I define the best action as the action among those under consideration that is decision-theoretically optimal. I discuss this notion more carefully later, but its essence is that each alternative may be judged equal to a lottery of a very good outcome versus a very bad outcome, and we wish to choose the alternative that is equivalent to the best probability of the very good outcome. I take this characterization as an operational definition of what it is to choose rationally.

I view complexity as being, fundamentally, an attribute of a proposed solution method, not of the problem. The reader is requested to interpret my statement of my objective in this way. In computer science, complexity is taken to address problems, but it is measured in terms of the number of steps in a solution process (see, e.g., Hopcroft and Ullman 1979, Wagner and Wechsung 1986). For this work, a complex situation is one where the only apparent solution methodology has many steps, cases or computations.

My work aims at a balance between the adequacy of the justification of belief that a course of action is best and the cost of the analysis required to justify the belief.

I will speak of "the" decision maker, and in the context of a decision to be taken by an organization, I will understand the term to mean the leader of the organization, or the person to whom the function in question is delegated. The mere fact that technical work related to an analysis is performed by a different person than decision making that uses the analysis does not preclude generation of knowledge in the latter's mind, if the delegation is accompanied by commensurate levels of communication and trust. Whether organizations are or could be constructed in such a way as to ensure that this requirement is met is another matter; but I take it for granted that there exist some situations where goodwill and common interests make a consensual procedure along the lines discussed in chapter 4 possible.

The illustration in chapter 4 refers to a decision facilitator, who has some understanding of decision analysis. This dissertation will recommend that the decision maker employ a decision analytic approach, and hence, that she or her subordinates acquire a certain level of expertise into the subject. This dissertation aims to assist in some difficult aspects of decision analysis, but relies upon the DM to have mastered important parts of the literature on issues where the latter gives explicit recommendations, or to have hired a facilitator with at least this level of knowledge.

1.3 Merits of the objective

The objective of a text may simply be stated, leaving the reader free to set the book down if s/he does not find value in the objective. However, it may be helpful for me to make a few remarks here about why the reader might find the achievement of this objective attractive.

I focus on uncertainty in the mind of the DM, not any uncertainty that ostensibly inheres in things in the world, because the former is the uncertainty that makes decision-making difficult. The fact that a DM may choose to treat portions of her knowledge as certain in a practical decision situation does not undermine this work. However, if the DM is content to view the entire situation as one whose outcomes are known with certainty, my work may not be particularly valuable to her.

My definition of knowledge as adequately justified belief invites inquiry into what form of justification is adequate. This is a helpful framing for DMs trying to make good decisions without spending a lot of time on them. The central lesson inherent in the notion of value of information (Howard 1966b) is that information has value only insofar as it affects actions. The story of an elegant thoroughgoing analysis that is ignored by the client is so common as to be trite in decision analytic circles. If the DM does not believe an analysis to be adequately justified, she will not act upon it, and hence will receive no benefit from it. Knowing an action's optimality makes her more likely to actually do it. Thus it helps analysis contribute to rational action, not mere formally correct thought.

Under the idea of adaptive rationality, an optimal balance of the costs and benefits of analysis is achieved. This notion is attractive, insofar as it releases the tacit assumption in perfect rationality that analysis is costless, while still following the idea of "doing the best thing". Although motivated by this idea, this thesis does not optimize the tradeoff between adequacy of justification and cost; it merely tries to do fairly well. Adaptive rationality is less important when the resources at stake are substantially larger than the costs of analysis, because a simple meta-analytic heuristic will perform adequately: do as much analysis as we know how to do, because it is cheap compared to the decision.

1.4 Methodology

Like decision analysis (Howard 1966a), which plays an important role in this thesis, the general approach of this work is prescriptive; I prescribe a procedure and argue that it is likely to achieve its intended end. A prescription may be contrasted with a norm. Ethical directives are normative, insofar as they specify what an agent should do. Decision theory is, as noted in Holtzman (1985), conditionally normative. It indicates what a DM should do *if* she accepts a given formalization of her preferences and the decision situation. Ethics and decision theory are indifferent to empirical support, as they are essentially analyses of the terms ethical and rational. By contrast, this dissertation is a prescription for how to think about decisions with an argument based on descriptive research that shows that it can be followed, and that it will give the results I claim.

The methodology of this dissertation is two-pronged, drawing strength from both existing practice and empirical research. In the spirit of artificial intelligence research, and computerization research in general, the work gives careful consideration to existing practice, distills its essence, and automates this. The premise of this prong is that existing practice is responsive to the underlying demands of the situation, even if its rationale is not explicit. The second prong is to consider the empirical research bearing on how people think about decisions, and to frame an argument that my prescriptions will work in the terms of that research.

I believe it is important to empirically verify the specific claims that both my decision formulation system, and the decision analytic approach of which it is a part, achieve their stated ends. One goal of this dissertation research is to enable and spur the required empirical research into the effectiveness of the DA Cycle, with and without automated verbal summary.

The remainder of this dissertation is organized as follows. Chapter 2 reviews the literature related to thinking and problem solving. Chapter 3 appraises this literature and motivates my focus on verbal summary in the context of decision analytic modeling. Chapter 4 gives an illustrative example of the use of computer software and ideas developed in this dissertation to formulate a decision problem. Chapter 5 describes the computational components required to support this process, and shows how they are arranged into an architecture. Chapter 6 develops and describes a parsimonious approach to eliciting decision model values needed for sensitivity analysis. Finally, Chapter 7 draws on the psychological literature to develop a way to summarize such a model that will help the DM believe that the formulation of the problem is adequate, and it shows how such summaries can be generated automatically. This result, together with existing tools and approaches for solving well-formulated decision problems can give the decision maker knowledge of what the best action is.

Chapter 2. Literature Review

This chapter reviews the literature of problem-solving approaches and research into thought and understanding.

2.1 Problem-solving

I first review approaches to decision-making and problem-solving (PS) that have been discussed in the literature. The phrase PS is taken to be somewhat more general than decision making, and, although my ultimate focus will be on decision-making, I review approaches under the broader rubric. I treat PS approaches in four sub-sections according to their most salient activity: maximize, model, communicate, and iterate.

2.1.1 Maximize

The problem-solving approaches here are oriented around maximization, each specifying what is to be maximized. I discuss, in turn, decision theory, fuzzy logic, and the Analytic Hierarchy Process.

Decision Theory

Decision theory is a model of rational decision making in which we identify the possible outcomes of each alternative, assign probabilities to each outcome, and assign a preference measure (utility function) to each outcome. We then choose the alternative with the maximal expected utility.

Bernoulli (1730) sets out the basic approach to choice under risk using a utility function. Conditional probabilities are defined in Bayes (1763). Subjective probability theory is set out and defended in Laplace (1812). Von Neumann and Morgenstern (1944) set out an axiomatization of rational individual behavior based on subjective probabilities and utilities. Savage (1954) defends subjective probability and discusses its relationship to utility theory. Luce and Raiffa (1957) address individual decision-making under uncertainty axiomatically, comparing different proposed axiomatizations and settling on a system of 10 axioms. Pratt, Raiffa and Schlaifer (1964) derive four axioms for decision making under uncertainty from principles of consistent behavior and scaling of judgements. Matheson and Howard (1968) and Raiffa (1968) both set out roughly the same axiomatization that guides much of DA practice today: All uncertain prospects are to be thought of probabilistically; enough distinctions are to be made that a preference ordering on them can be established; given three unequally preferred outcomes, an uncertain chance of the best vs. the worst equivalent to the intermediate outcome may be defined; the decision maker is to be indifferent to substitution of one lottery for another she deems equivalent; and, when faced with two lotteries having the same two outcomes, the decision-maker is to choose the one with the greater probability of the better outcome. Howard (1977) discusses a possible sixth axiom that some decision-makers may accept: insensitivity of risk tolerance to wealth (sometimes called the "delta property"). Recently, axiomatic bases have been created for other approaches to decision making: (Saaty 1986 and Smets 1988). This led to identification of criteria that could be used to choose among axiomatic systems. Holtzman (1985) suggests that a formal decision-making procedure should be clearly applicable to small problems and that it should scale up well. Howard (1992) defines a decision composite as a set of decision axioms together with the theorems about decision-making that follow from it. He then establishes desiderata that a decision composite should satisfy, and finds that only standard decision theory (Howard 1966a) satisfies all the desiderata.

Fuzzy Logic

Many authors (including Schmucker 1984, Boettner 1985, Gheorghe et alia 1985, Maeda and Murakami 1988, Gheorghe and Stoica 1987, Buckley 1987, Lebailly, Martin-Clouaire and Prade 1987, Godo et alia 1989, and Sakawa and Yano 1989) discuss "fuzzy systems", and attempt to find some way to optimize a fuzzy view of such systems, as if the underlying state of events were indeterminate. A common argument in defense of this approach is to point out that statements like "John is tall." cannot necessarily be judged true or false, even by persons who are acquainted with the person in question, and who understand the sentence. However, what is at issue here is not any indeterminacy in the state of the world; rather, it is the incomplete determination of the meaning of the words (e.g., tall) used in ordinary language. Careful authors such as Zadeh (1968) and Oden (1979) note that imprecision in the applicability of terms to events is different from incomplete knowledge about which events occur. Furthermore, Lindley (1982) shows the inadmissibility of the fuzzy calculus for uncertainty in situations where resources are at stake, Blanning (1985) shows its disutility for management of rule invocation in any production system that is specified without causally redundant terms in rule antecedents, and Elkan (1993) shows that as a formal system, a standard version of fuzzy logic collapses mathematically to two-valued logic, while empirically, fuzzy logic is not adequate for reasoning about uncertain evidence in large-scale expert systems. As a result, Holtzman (1985) and Howard (1988) insist on definitions of important distinctions that are clear enough to be applied to any given circumstance without use of judgement, even if this takes a nontrivial effort. Klir (1985) and Gaines (1987) emphasize the importance of choosing distinctions well for problem solving, but they are willing to employ unclear distinctions. As yet, it appears that no one has analyzed the costs and benefits of clarity to help us decide under what circumstances achievement of clarity is warranted.

Analytic Hierarchy Process

Saaty (1980) identifies terms whose meaning to people is ostensibly stable, and uses these to elicit the ratio of one factor's impact on a criterion to that of another. He proposes to create multi-level hierarchies in which the relative impact of items at each level upon the items in the level above is calculated. Then the matrices of one level's relative impact on the next are multiplied to ascertain the relative impact of the items at the lowest level on the highest level. For a decision problem, he suggests that the alternatives be placed in the lowest level and something analogous to an objective function be a singleton element at the highest level, and that the alternative with the largest impact on the objective be chosen. Saaty (1982), Gass (1986), Hämäläinen et alia (1986), Watkins et alia (1992) and many others apply AHP to priority-setting, planning, conflict resolution, cost/benefit decision-making, and group decisions. Salo and Hämäläinen (1993) argue that the pairwise comparisons in the AHP should be understood as ratio statements about preference differences between pairs of consequences, because under this interpretation the AHP constitutes a quantitatively meaningful variant of multiattribute value measurement. Pöyhönen, Hämäläinen and Salo (1994) find that the perceived meaning of Saaty's verbal expressions varies among subjects and also depends on the whole set of elements that are being compared; and they propose scales that perform better.

2.1.2 Model

I treat modeling support in five parts: representation of one's ideas about the situation at hand to support model construction, representation of existing computer models to support model reuse, metareasoning to support choice of solver and allocation of resources to competing ends, interactive computer-aided specification of model structure, and model development tools.

Represent modeler's understanding

A very substantial portion of the literature, and practice, is devoted to representing one's ideas about the situation at hand in a computer model, or in some other formalism that assists in the creation of a computer mode. I discuss these here, beginning with comments on the important distinctions, then treating network-based representations (notably decision networks and structured modeling), and textual ones.

The necessity of using the appropriate distinctions in an analysis has been noted in many parts of the literature. Decision analysis explicitly emphasizes the importance of helping the DM identify the right variables for analysis when formulating a decision problem, e.g., in Howard et alia (1975), Matheson and Howard (1968) and Holtzman (1985). Schrattenholzer (1985) opines that finding a small but comprehensive set of variables is important for modeling enterprises such as developing long-term global energy use scenarios. Klir (1985) sets out an epistemological hierarchy describing the way people make sense of systems, in order to design a General Systems Problem Solver. The lowest level is unanalyzed sensation; next, distinctions that identify data are defined, systems that generate these data are formulated, and analogies among these systems are identified. Gaines (1987) interprets this hierarchy in the light of personal construct psychology (Kelly 1955),[1] identifying data definitions with personal constructs, and hypothesizing that people construct such hierarchies to minimize the flow of uncertainty or surprise from one level to the next. In this account, the problem solver (explicitly or implicitly) formulates each level of the hierarchy so that the behaviors in the adjacent levels are described in as simple a way as possible. Thus, the problem solver would try to find distinctions that characterize experience simply, and that allow simple models to be built of their interactions. The statistical literature discusses the problem of missing variables (e.g., Huff 1954 p. 130, Heise 1975 p. 44), underscoring the importance of identifying the appropriate distinctions. Weyant (1990) identifies conditions that account for the successful use of quantitative models in formulating energy policy in the US Congress. These include inclusion in the analytical framework of all impacts thought to be important by the decision makers and compatibility of the options and criteria considered in the analysis with those used in the policy debate. Finally, a crucial aspect of the use of neural networks is the identification of input features to use in training the network - much more time is spent on the attempt to identify an appropriate set of distinctions than on training a network once the right set is found Smyth (1993). All of these articles point to the importance of formulating and discussing decision problems using the most appropriate distinctions.

Decision networks, introduced in Howard and Matheson (1976), are directed graphs containing nodes of three basic kinds: decision, event, and utility. A decision node contains a set of alternatives. An event node contains the decision-maker's conditional probability distributions over the outcomes of an event. There must be exactly one utility node; it specifies the decision maker's utility function. Node types are distinguished by shape: decision nodes are rectangular,

[1][Kelly 1955] sets out personal construct psychology and describes repertory grids, which is a methodology for eliciting "personal constructs" by comparison and contrasts among triples of elements.

event nodes are oval, and the utility node is octagonal. Arrows into utility or event nodes specify conditioning -- the contents of these nodes are specified for each possible value of the conditioning (predecessor) nodes. Event nodes are called deterministic if their distributions are all degenerate (i.e., the value of the corresponding variable is completely determined in each condition); otherwise they are called chance nodes. Deterministic event or utility nodes are distinguished by doubled borders. Informational conditioning arrows into a decision node indicate that the outcome of predecessor nodes will be known when the decision is made, thus allowing the alternative that maximizes the expected value of the utility function to be specified conditional on these outcomes.

Owen (1978) describes the use of decision networks to structure a decision problem backward-recursively from the utility node, and Holtzman (1985) defines Intelligent Decision Systems, which computerize this approach. An algorithm that identifies the optimal set of choices for all decisions in decision networks with discrete variables is given in Shachter (1986). Pracht and Courtney (1986) describe a visual user interface for structuring, representing and maintaining mental model knowledge. The semantics of Pracht's arcs is causal instead of probabilistic, but the support given the modeler appears to be similar to that provided by Holtzman. Bradshaw et alia (1989) and Bradshaw and Boose (1990) take a step toward integration of repertory grids (which represent, among other things, the degree to which situations instantiate the modeler's personal constructs) with decision networks by transforming grids whose elements are observations into a grid with frequency distributions of outcomes of interest. McGovern et alia's (1991) Intelligent Decision System for strategic decision making acquires problem knowledge, including decision network model structure, directly from the decision maker's description of the problem in simplified natural language. Recently there has been much work at dynamic creation of decision network models from a database: Srinivas, Russell and Agogino (1989), Wellman (1990), Heckerman and Horvitz (1990), and Goldman and Breese (1992).

In Structured Modeling (SM) (Geoffrion 1985), groups of related variables are nodes in a directed graph, and arcs among them represent definitional reference. Contents of nodes are specified in relational data tables, allowing direct interface to a database management system. Models may be integrated by graph union. Geoffrion (1989) shows that SM's textual representation of models, called a modular outline or schema, also supports model integration. This article introduces SML (Structured Modeling Language), which is used to represent schemas. Geoffrion (1992) describes SML fully. Geoffrion (1991) and Neustadter, Geoffrion et alia (1992) describe modeling environments based on SML.

Paradice and Courtney (1988) construct linear statistical models of direct and indirect relationships from a knowledge base of relationships in a managerial domain. Brown and Lewis (1989) describes the HELM language, which conceptualizes, organizes and specifies the model schema of large-scale linear programming problems. Kirkwood's (1991) ADAM translates an analyst's model specification file into a Pascal program that implements the decision tree analysis.

Represent models for reuse

Many authors have proposed to support choice among solvers by finding a match between problem and solver characteristics.

By analogy with the field of software reuse (Prieto-Diaz and Freeman 1987 and Goguen 1986), many authors simply assume that models ought to be reused, and offer languages Geoffrion (1989), runtime systems Muhanna and Pick (1988), Kottemann and Dolk (1988), Geoffrion (1991), and model retrieval tools Dutta and Basu (1984), Mannino et alia (1990) for that purpose. Some explicitly argue in favor of model reuse. Holtzman (1985, p. xi) notes that "[R]esponsibly analyzing a significant decision usually involves over a hundred hours of intense work ... [A] professional decision analysis is a major effort that is, unfortunately, beyond [the] means [of most individuals]." Reusing previous code can, according to Goguen (1986), make programming easier and more reliable and cost effective. Bankes (1993) notes that models that use engineering theory and data to predict the behavior of systems can profitably be reused.

Will (1975) first set out the notion that models are a resource that should be managed, just as a database manages data. The main purpose of Model Management Systems (MMSs) is to support storage, identification, and retrieval of preexisting modules that are useful to the task at hand. Construction and verification of problem-solving systems from these modules is often discussed in the context of MMSs, but these are more properly the task of modeling systems, not MMSs. Dolk and Konsynski (1984), Blanning (1986) and Lenard (1986) suggest taking Will's "model base" metaphor seriously and employing database technologies to represent models for model management. Dutta and Basu (1984) represents model parameters and the conditions under which a model can be meaningfully run. Kottemann and Dolk (1988) suggest languages for specifying how component models are to be integrated.

Lenard (1987), Bradley and Clemence (1988), Dempster and Ireland (1989) and Mannino, Greenberg and Hong (1990) represent a collection of model as 'objects' that perform various model management functions.

Liang (1986) represents a set of data as a node, a set of functions as an edge, and a basic model as a combination of two nodes and one connecting edge. Moray (1987) represents possible models in a graph with a-kind-of arcs, to help the modeler find a model that is just disaggregated enough to solve the problem at hand.

Howard (1968) suggests categorizing problems according to complexity, uncertainty, and time dependence of variables. Paradice and Courtney (1986), Liu and Tomsovic (1986), Kottemann (1986), Tanniru and Murray (1987), Wang et alia (1988), Banerjee and Basu (1990), and Mili and Szoke (1992) categorize inputs required by solvers. Mili and Cioch (1990) give a framework for documenting the relationship of decision models to problem situations to support model retrieval. Eck et alia (1990) represent the inputs, outputs, variables manipulated, and pre- and post-conditions of solvers. Banerjee and Basu (1990) characterizes the objective, algebraic degree of constraints, and solution space topology of models. Schoppers (1991) identifies five dimensions of problem requirements: response time, program size, processing power, attentiveness, and degradation of results. Watkins et alia (1992) notes the following problem attributes: volume of data, need for tradeoffs, and whether data is numeric.

But there is another side to this issue, exemplified by Schrattenholzer (1985), and Bankes (1993), which gives explicit consideration to the disadvantages of model reuse, noting that reused models tend to grow as logic is added to them to handle new situations, that this increases the amount of baggage in subsequent decisions, making it hard for the DM to identify the important aspects of the model's results.

Metareasoning

A related issue is the allocation of resources to different problem-solving activities, when there is more than one such. This activity is sometimes called metareasoning, as it entails consideration of what reasoning approach to take. Its goal is to achieve a good balance between the effort put into an analysis and its quality. This balance is sometimes called adaptive rationality.

D'Ambrosio et alia (1987), Fehling and Breese (1988), Ruokangas (1988), Hayes-Roth et alia (1989), Coté and St-Denis (1992), and Waldspurger et alia (1992) discuss resource allocation to and coordination of different PS activities, especially allocation of CPU time in dynamic situations. Fjeldstad and Konsynski (1986), Courtney, Paradice and Ata Mohammed (1987), and Jacob, Moore and Whinston (1989) note that computer models have different strengths from people, and attempt to partition problem-solving responsibilities accordingly. Fehling and Breese (1988), Moore et alia (1989), Tse and Fehling (1989), Schoppers (1991), and Ogasawara and Russell (1993) address metareasoning to choose a solution approach. Horvitz and Breese (1990) and Goldman and Breese (1992) discuss the problem of ideally apportioning resources between solution planning and problem solving, with particular reference to probabilistic inference problems. Rotmans and Vrieze (1990) views input assumptions from an uncertain distribution as a choice among problem-solving resources, thus treating experimental design and sensitivity analysis as metareasoning. Bankes (1993) suggests explicit experimentation and choice among different sizes and levels of resolution for decision models.

Interactive specification of problem structure

Leal and Pearl (1977) describe an interactive computer program that elicits a decision tree from a DM. Lagomasino and Sage (1985), Malakooti (1988) and Chu et alia (1989) prompt the user for preference and probability assessments, and formulate the problem in a way that allows these to be imprecise or incomplete. Krishnan (1989) and Ma et alia (1989) support interactive formulation of LP models.

2.1.3 Communicate

The systems described here have at their root comparison or communication between different points of view regarding the system at hand. The four subsections here address computer systems that interact with the human analyst, computer systems that explain their results, approaches based on comparative analyses, and systems that support human discussion.

Interactive

Mili (1988) 'looks over the shoulder' of the decision-maker, second-guessing her, and criticizing her actions and giving advice when appropriate. Wellman et alia (1989) identifies four potential modeling errors in medical decision trees: impossible strategies, dominated strategies, unaccountable violations of symmetry, and omission of apparently reasonable strategies. Raghavan's (1990) system acts as a devil's advocate: raising pointed questions and generating challenges, arguments and criticism for the decision-maker. That of Raghavan (1991) plays roles such as expert, devil's advocate, critique, playing dumb, and careful listener. Silverman (1991) proposes an automated critic with a knowledge base of possible human errors and criticism strategies could help the user prevent or eliminate these errors.

Symbiotic DSSs work on the problem at hand "alongside" the human analyst, without being explicitly told to do so. Shoval (1986) studies DSSs that perform a search and evaluate their findings without informing the user before the search is completed, and finds them to be effective. McCoy and Boys (1987) predicts the human operator's forthcoming tasks, anticipates upcoming decisions, formulates any necessary decision or execution aids, and analyzes the differences between expected and actual operator actions. Manheim et alia (1990 and 1991) does independent analyses using a second copy of a steel mill scheduling system, based on its observation of the user's use of it; using both the results of the user's work on the scheduling task and its own explorations, it provides advice to the user. Castillo, Dolk and Kridel (1991) provides an "artificially intelligent modeling expert" as well as an "artificially intelligent domain expert" for assisting the user in developing and analyzing process models.

Explanatory

There has been a great deal of work on explanation in logic-based "artificial intelligence" systems. I sample this literature and other approaches to explanation and summary in the literature.

Sinha et alia (1984), Holtzman (1985) and many others' systems give a reasoning trace to justify their actions. Swartout (1985), Neches, Swartout and Moore (1985) and Swartout and Smoliar (1987) require that design processes be recorded in machine readable form in order to justify an expert system's behavior with a causal argument based on its design rationale. Chandrasekaran et alia (1986) explains a planner's behavior by representing plans as devices and referring to a functional representation of human understanding of how devices work. Molokova (1986) defines a language wherein the user's scope and depth of interest in an explanation can be specified, and explains models constructed by an expert system according to such specifications. Bridges and Johannes (1988) use an augmented phrase structured grammar based on analysis of justifications written by people to organize justifications of system-generated plans. The diagnostic expert system for radiologists in Mutalik et alia (1988) chooses one of four strategies (Pursue, Rule-in-rule-out, Conflict, and Not-enough-information) to explain how far the system has resolved competing diagnoses. The explanation component of Wick and Slagle's (1989 explanation) ES shell explains ES actions accurately, objectively and noninteractively after the fashion of a journalist. Bruffaerts et alia (1989) generates a proof tree for how-, why-, why-not- explanations and uses a uniform logic-based object-oriented formalism for knowledge and meta-knowledge to generate conceptual explanations. Jamieson (1989) flexibly links explanations with conclusions generated by a causal reasoning system to explain its concepts and reasoning. Koussev et alia (1989) writes separate logs of successful and unsuccessful rule firings and records some aspects of variable-rule interrelations to answer why and why-not questions about its inferences. Wick and Thompson (1989) select elements from a database of explanation-fragment templates, instantiate their variables according to current data, and use the A* search algorithm to connect them into the shortest complete explanation. The system in Bailey and Duban (1990) employs qualitative causal reasoning about pathophysiological systems to justify its diagnostic and therapeutic advice.

I consider two subclasses of arithmetic explanation: deterministic numerical simulation models, and probabilistic reasoning systems. Kosy and Wise (1984) generates explanations from the quantitative relationships among variables that comprise a financial model. Hasling et alia's (1984) NEOMYCIN makes abstract strategic explanations (those which articulate a general principle) by representing strategic knowledge explicitly and separately from domain knowledge. Cooper (1984) explains the critical qualitative causal and quantitative probabilistic factors that affect the relative likelihood of diagnoses in the domain of hypercalcemic disorders that are generated by the system, volunteered by the user, or both. Helman and Bahuguna (1986) explains the structure and domain of a numerical computer simulation of inventory control for the novice user. Rennels et alia (1987) suggests choosing among four distinct strategies for multiattribute decision making, each of which makes restrictive assumptions about the nature of the domain, as a basis for explanation in medical AI systems. Langlotz et alia (1988) finds any asymmetries in tree structure or inequalities among analogous decision variables that are responsible for a difference in expected utility among branches of a decision tree, selects an explanation technique, applies it to the most significant variables, and converts this to English-language text that justifies its recommendation; Langlotz (1989) applies this system to selected medical decision models in the literature. Jimison (1988) emphasizes only variables that deviate significantly from what is typically observed or which have high value of information or sensitivity in the ID graphics and text summary of complicated decision models and applies this system to a consultation process for patients with angina. Bosch and Weyant (1989) discuss a *human* process for explaining model results. They recommend that a common data set of electric utility cost parameters be established and that all parties testifying before CEC be required to run it through their preferred model, and that the parties be required to document their model's implementation of system constraints to help the presiding judge attribute differences in projections presented by the parties to the use of different models, to different modeling conventions, or to different resource assumptions. Sember and Zukerman (1989) explain inferences drawn by a Bayesian belief network in response to changes in the causal and the evidential support of a given node. Elsaesser (1989) explains Bayesian inference. Suermondt (1992) identifies influential evidence, analyzes conflict among findings, and investigates the ID pathways through which the influential findings affect the probability distributions of the variables of interest to justify diagnoses of emergency conditions during anesthesia. Klein and Shortliffe (1991) uses interpretation-concept query, value-tree pruning and presentation strategies, difference-function-traversal strategies, and model-traversal strategies, all within the Interpretive Value Analysis framework, to explain decision-theoretic choices in the domains of marketing, process control, and medicine.

There has been much less work on computer-generated summary, where summary may be taken to be a short text presenting a work's general sense. Kukich (1985) describes a prototype system that identifies in English significant events found in a flood of electronic data about markets. Novak (1987) presents a method for the automatic generation of simple verbal comments that may help the user understand extensive results from finite element analysis. Jimison (1988) presents a representation for uncertainty that supports computer-generated graphical summary of medical decision models.

Comparative analyses

These approaches employ a dialectical approach, where differing viewpoints are compared and improved in light of each other.

The notion that criticism is a good way to enhance one's understanding has a long history. The idea was clearly stated in Hegel, under the name of dialectic. Under this approach, contrasting ideas, called the thesis and antithesis, are compared, and a new idea that is better than either, called the synthesis, is created as a result. Lakatos (1970) espouses an approach to science that relies heavily on critical comparison of scientific ideas to advance human knowledge. Feyerabend (1974) advocates an anti-systematic approach to science to ensure a broad range of ideas as fodder for this sort of critical endeavor. Longino (1990) argues that, notwithstanding the influence of social and cultural values in the very structuring of knowledge, the objectivity of scientific inquiry can be maintained by understanding it as a process mediated and directed by criticism from differing points of view.

Longino's ideas can be applied to the decision context, where the tension between the subjectivity of probabilities and the desire for widely defensible decisions takes the same form as the debate over "objectivity" in science. Howard (1968) encourages the generation of a computer model and comparison of its results to the modeler's judgements. Mitroff and Betz (1972) advocate identifying differences among experts regarding choices, system states and behavior, and utilities, to ensure a complete problem formulation. Modeling experts (Gass 1977) encourage generation of a computer model and comparison of its results to the modeler's notions of how the world behaves, calling the process verification, or calling it validation if hard data is employed in the process. The model and the judgements are both said to be susceptible to improvement when they are compared: Hogan (1978) claims that this comparison builds understanding; and Yu (1990), that it can improve model formulation. Fischhoff, Slovic and Lichtenstein (1979) suggest that people judging probabilities should perform many related judgements and compare them, to reduce bias. Huntington et alia (1982) describes the Stanford Energy Modeling Forum, which annually chooses a question of interest in the energy business, convenes a variety of experts, identifies existing models relevant to the topic, designs test scenarios to illuminate the models' behavior, and compares model results to identify their strengths and weaknesses. Bunn (1986), Hickman et alia (1987), and Morris et alia (1987) also suggest comparison of multiple computer models. Wack (1985) subsequently brought the scenario planning approach to the popular press and emphasized that scenario plans should be used to help management reorganize their mental model of reality. Holtzman (1985) observes that the decision analyst learns in early stages of the DA Cycle, and the domain expert learns from the analyst in later stages. Lane (1992) proposes an approach to consultancy whose goal is to accelerate the client's learning about the business by articulation and criticism of mental models.

Support human discussion

These systems support discussion or exchange of information among human participants in a decision process. Licker and Thompson (1985) and Ligeza (1988) retrieve decisions from previous similar cases. Eden et alia (1986) presents a scheme for the mapping of argumentation to synthesize the points of view of conference participants. Kettelhut (1989) describes a DSS that displays numeric rankings for important factors, to support further discussion. Shafer (1989) has observed that the use of decision networks in decision analysis helps achieve consensus about problem structure within an organization. Radford (1990) presents the following method of exploring solutions to complex decision situations in decision conferences: gather information, consider possible outcomes and participants' preferences for them, study courses of action that individual participants can employ, and then allow interaction between the participants.

2.1.4 Iterate

I treat prescriptions for iterative problem-solving processes here in three sections: search, the DA Cycle, and others.

Search

Besides being a model of human thought, the search, or generate-and-test, paradigm is an important tool in artificial intelligence work. Manheim (1966) describes a planning process that starts with a completely general action and explores sequences of operators that specify aspects of the plan. Newell and Simon (1972) set out and give evidence for the following theory of human problem solving: a task environment is represented in the human mind as a problem space, and problem solving takes place in that problem space. The structure of the task environment determines the possible structures of the problem space. The structure of the problem space determines the possible problem-solving approaches that can be used. Simon and Lea (1974) find that this problem space is searched for a solution that is satisfactory vis a vis dynamically adapting aspiration levels by generating new knowledge states and testing them. As with all of AI, this notion is intended to be both descriptive and prescriptive.

Decision Analysis

Raiffa and Schlaifer (1961) describe "statistical decision theory" and indicates how conjugate priors and explicit analysis of probabilities and utilities can be applied to decision problems. Howard (1965) applies statistical decision theory to a decision about contracting for engineering devices. Howard (1966a) christens the emerging discipline Decision Analysis. He discusses the role of the *decision analyst, gives a rule of thumb for sizing* the analytic effort, and discusses elicitation of risk aversion, time preference, and probabilities. Howard (1968) gives an early statement of the DA approach that sets out what was to be the dominant picture of the DA Cycle for many years: deterministic analysis for problem formulation, probabilistic analysis to determine a preliminary recommendation, and informational analysis to determine whether any information gathering activities are merited, or whether the recommendation should be acted upon immediately. If further information is gathered, analysis re-iterates through formulation and solution. Howard et alia (1975) identify two important principles underlying the DA Cycle: make every portion of the analysis meaningful to the DM, and quit if the answer becomes obvious. It also notes that a decision model will grow or shrink as important parts are elaborated or unimportant ones deemphasized. Howard (1983) recharacterized the informational phase of the Cycle as the basis appraisal phase and identified its role in refining both the model formulation and the decision-maker's intuitions. Thus the DA Cycle is responsive to the difficulties of unvarnished DT, in that it gives a chance to revisit crucial premises in the DT equivalence argument in light of their consequences. This article notes that the role of decision analysis is to reduce opaque situations to transparent ones where the decision theoretic axioms are obviously applicable. Holtzman (1985) formalized the role of evaluation of a formal model in fostering decision ownership in the decision-maker. This account subdivides each phase of the DA Cycle into (re)formulation of a model, evaluation of the model, and appraisal of differences between the formal model and the decision-maker's intuitions. Appraisal can either lead to improvement of the decision-maker's insight, thereby fostering decision ownership, or it can call for reformulation of the model, if the model is found to be deficient, or it can lead to confident choice if the model and the DM's intuitions agree.

The approach to decision models taken by DA authors is of interest here. Howard (1966a) develops a large deterministic decision model and then simplifies it to find a version that can be exercised probabilistically without undue computational cost, perhaps even without a computer.

To simplify the model, variables to which the utility function is not sensitive, as determined by Deterministic Sensitivity Analysis (DSA), are fixated at some reasonable value, allowing the rest to be analyzed probabilistically. In DSA, each input variable is varied from low to high values, holding all others fixed, and the effect on a target variable (normally the utility node) is noted. Howard (1968) set out the approximate 10th and 90th fractiles of a variable's prior distribution as the range of variable values to be employed in DSA. Howard (1983) describes tornado diagrams, in which the relative impact of agent variables on the target is depicted graphically. Howard (1968) recommends building increasingly detailed decision models during an analysis: pilot, prototype, and production models. The pilot model is constructed very quickly; its purpose is to help the analyst identify which features of the problem must be analyzed and which clearly do not have a substantial impact. Variables with high sensitivity should be elaborated on, creating a prototype model. The prototype model includes most or all of the interesting factors in a plausible way; it is well-enough tuned to the problem to support substantial critical appraisal. The production model includes all factors that are thought to be crucial to the analysis (especially uncertainty). Howard and Matheson (1976) set out a backward-recursive model formulation process: start with the utility node in a decision network model, and identify relevant predecessors backward until the relationship between the decision at hand and one's values is entirely specified. This can be effective when a vast welter of possibly relevant data makes data-based forward chaining an ineffective way to build models. McNamee and Celona (1987) give an example of expansion and contraction of a decision model in response to DSA.

Thus, typical changes to a model during the DA Cycle include: formulating an additional decision, adding an additional chance node (especially as a predecessor to a node whose distribution is not known sufficiently well), fixating an unimportant decision or chance node, identifying the distribution of an important variable more carefully, identifying deterministic relationships among variables, fixing "bugs" in the model, and investigating crucial parameters of the DM's preferences.

In addition to DA authors, others in the literature note that DSA can be used to focus attention in a model-building process. Shannon (1975) says DSA motivates subsequent modifications to the model. Eschenbach and McKeague (1989) note that sensitivity analysis can focus managerial attention, e.g., on needed refinements in data estimates. Rios Insua and French (1991) show how sensitivity analysis can help one focus on those judgemental inputs that are most important in determining choice and, therefore, need to be revised most carefully. Bankes (1993, p.445) says "One possible approach to [focus attention on aspects of modeling that are most critical for the question at hand] is selective resolution, where initial modeling is done with relatively aggregate models and the results of this preliminary analysis are used to guide the selective use of higherresolution models, with detail added only for the attributes that appear to have a large impact on the question of interest."

Other decision processes

Other processes for decision-making that have been set out in the literature include: Simon (1960) identified the three phases of decision-making as intelligence (declaring the decision), design (generating alternatives), and choice. Vari and Vecsenyi (1984) characterizes the decision process as planning (including problem structuring, information gathering and processing, alternative generation), choice, and implementation of the decision. Rasmussen (1985) gives a model of exception-based decision-making for process-control: when an unusual situation arises, identify the state of the system, forecast its consequences, set a goal, plan a sequence of actions to achieve it, and execute the plan. Kottemann (1986) distinguishes problem recognition, choice of technique, and primary decision-making. Weber and Konsynski (1987) define problem management as problem finding, problem representation, information surveillance, solution generation, and evaluation. To evaluate Decision Support Systems (DSSs), Adams, Courtney and Kasper (1990) advance the following model of DSS-aided decision making: problem identification, diagnosis, alternative development, and alternative selection.

2.2 Thinking and understanding

Having reviewed the literature regarding problem-solving, I now turn to thought and understanding. The purpose of this review is to identify ways of coming to know things that are roughly consistent with a justifiable analysis of a decision, and to design a way to summarize such analysis of a decision in comprehensible terms. I treat thinking first, then understanding.

2.2.1 Thinking

I review empirical or philosophical study of how people judge the outcome of the combination of a number of factors, insofar as this can form the basis for confident decision making. I find three basic approaches: direct judgement, decomposition, and recognition. This taxonomy corresponds roughly to that in Yu (1990), which identifies three major approaches to development of analytic structures: hierarchical and heterarchical decomposition, inductive reasoning, and deductive reasoning.

Boundedness

Before beginning a discussion of the internal structure of thought, it is important to acknowledge its boundaries. Much of the work of Herb Simon (e.g., 1955, 1969, 1982) is focused on the boundedness of human rationality; that fact that human thought takes place at a finite speed, and with a finite capacity for information. Other early study of this issue was in Bruner et alia (1956) and Miller (1956).

Decomposition

One important account of the way people think is that they first decompose the issue at hand, and then think or form judgements about the components. An important paradigm of this sort of mental decomposition is decomposition into a network of mental states related to one another by operators. If the states represent physical states of the world and the operators represent rules of causality, we call the decomposition simulation. If the states are knowledge states and the operators are rules of inference, we call it reasoning. If the space is likened to physical space and the operators resemble moving from one place to another, we call it search. In all three cases, the hypothesized method of thought is to repeatedly apply operators to identify new states until some criterion of completeness is achieved. For simulation or reasoning, normally this occurs when some attribute (e.g., profitability of a circumstance or truth value of a proposition) is specified. For search, a set of goal locations satisfies and terminates the process. I discuss simulation, reasoning, and search in this section.

Craik (1943) hypothesizes that human thought consists of translation of some external process into an internal representation (model) in words, numbers, or symbols, derivation of other symbols from them by an inferential reasoning process, and retranslation of these symbols into apprehensions or actions. Craik defines a model as a physical or chemical system that works in the same way as the process it parallels or imitates, but presumably more quickly or inexpensively. Johnson-Laird (1980) concludes that comprehension consists in first creating a propositional representation, and perhaps then also a mental model that can be used to simulate the situation in question. Kahneman and Tversky (1982) find that questions about causality and probability of events are often answered by an operation that resembles the running of a simulation model, where people construe the output of a simulation as an assessment of the ease with which variations of exogenous factors from their default values could produce different outcomes. Here the judgement of the likelihood of a scenario is decomposed into judgements of ease of variation of component events in a causal sequence. Isenberg (1986a, 1986b) finds that managers develop and test mental models reflecting their understanding of business situations, and that they reason by analogy to other better-known circumstances (which play the role of simulation models). Collins and Gentner (1987) find evidence that laymen construct generative models of system behavior by using analogy to map the rules of transition and interaction from known domains into unfamiliar ones.

As noted above, Johnson-Laird bases his account of reasoning on simulation; in (1988) he finds that, when reasoning, ordinary individuals imagine the states of affairs described in the premises, search for novel conclusions, and, *if* they are reasonably prudent, submit these conclusions to test by further search for counterexamples. Simon (1983) gives a similar account of the relationship of reasoning and search. Harman (1986) gives a theory of reasoning among defeasible all-or-nothing beliefs based on principles of mental conservatism (giving up as few relevant beliefs as possible) and avoidance of mental clutter (e.g., explicit recollection of the support of a belief).

Simon (1962) argues that complex natural systems will evolve in a hierarchically structured fashion, and suggests that such systems be decomposed hierarchically when solving problems. Manheim (1966) develops a model for planning or design processes; this model begins with a completely general plan, and applies a sequence of operators to this plan to specify one aspect after another of it. Planning is a search process in a network of partial plans, and the content of each node may be used to guide the remaining search. Although proposed as a normative model, the process is clearly based on observations of the actual planning process in the field of highway design. Newell and Simon (1972) give evidence that people represent a task environment as a problem space in the mind to enable problems solving; it specifies only that the structure of the task environment determines the possible structures of the problem space. Simon and Lea (1974) describe human problem solving as follows: 1) There is a problem space whose elements are knowledge states. 2) Generative processes (operators) take a knowledge state as input and produce a new knowledge state as output. 3) Test processes compare a knowledge state with the problem state or other knowledge states to identify differences. 4) Processes select which of these generators and tests to employ on the basis of the information contained in the knowledge states. Other articles anthologized in Simon (1979) characterize this process as means-ends search for a satisficing solution, guided by dynamically adapting aspiration levels. The solution of the entire problem comprises the serial application of the solutions of the sub-problems; in mathematical terms, the solution operator is the composition of the solution operators of the subproblems. Simon's (1979) General Problem Solver is a model of human cognition that represents a task environment into a means-ends hierarchy, and performs symbolic problem-solving there. Klir (1985) sets out an epistemological hierarchy describing the way people make sense of systems, in order to design a General Systems Problem Solver. Its lowest level is unanalyzed sensation; next, distinctions that identify data are defined, systems that generate these data are formulated, and analogies among these systems are identified.

The work on decomposition reveals more similarities than differences among simulation, reasoning and search. People are found to be comfortable reasoning with and about simulations, in particular by identifying aspects of a simulation that are invariant under plausible changes of the "inputs" of the simulation.

Direct Judgement

Articles summarized here describe human thought as relatively direct judgement. I discuss research into judgements of probability, preference, similarity, set membership, and impact.

Phillips and Edwards (1966) find conservatism (use of a sensible starting point, and adjustment of it in the right direction, but not enough) in a simple probability inference task. Many articles postulating and identifying linear additive models of human judgement are reviewed in Slovic and Lichtenstein (1971). Despite these articles' disagreement on the exact mathematical model of information incorporation, there is agreement among them that a human judge responds in a highly quantitatively predictable way to the information available to him. Tversky and Kahneman (1974) replicate the finding of anchoring and insufficient adjustment in probability judgements. Kahneman and Tversky (1979) develop the prospect theory of human decision making under risk, to systematize this finding. In prospect theory, probabilities are replaced by decision weights, which are generally a bit lower than the corresponding probabilities, except when the latter are below 5%. In a similar vein, Simon (1979) finds that the ability to extrapolate sequences to be a fundamental component of human cognition by modeling what the author believes to be the chief components of cognition and replicating human behavior in cognitive tests. (Lusk and Hammond 1991) is a recent work in Brunswick's lens paradigm, which characterizes human judgements in terms of linear additivity.

Tversky (1969) finds evidence that human choice is sometimes guided by a Lexicographic Semiorder of attributes, choosing according to the value of a "primary" attribute if a difference is noticeable, or choosing according to a less important attribute if there is no noticeable difference. For situations where this proves too restrictive a model, he also defines the more general additive difference preference model, in which the utility function is restricted to the sum of functions of differences along each attribute dimension. In prospect theory (Kahneman and Tversky 1979), value is again assigned to gains and losses rather than to final assets. The value function is normally concave for gains, convex for losses, and is generally steeper for losses than for gains. Korhonen et alia (1990) replicate these findings in a multiattribute choice situation.

Shepard (1957) develops a mathematical model to explain errors in stimulus-response situations as a confusion of similar stimuli: psychological inter-stimulus and inter-response distances in "psychological space" are postulated and solved for. Kruskal (1964) formulates the role of MultiDimensional Scaling (MDS) as the representation of objects geometrically by points in a psychological space of fixed dimensionality, so that the inter-point distances bear a monotonic relationship to experimental dissimilarities between objects. Recent applications of MDS, e.g., Coury (1987), attempt to identify how many as well as which attribute dimensions should be used to characterize subjects' judgements. Shepard (1987) gives a mathematical model of generalization (giving an old response to a new stimulus) based on psychological space and argues that these empirical regularities are derivable from principles of natural kinds and probabilistic geometry that are mandated by evolution. Nosofsky (1992) argues that any evaluation of how well a representation such as MDS accounts for similarity of data must occur within the framework of a formal process model that specifies what judgemental response will be given for a stimulus.

Saaty (1980) argues that comparison of the impact upon some criterion of a factor cannot be judged absolutely, but must be judged as being more or less than some other impact; and it argues that such comparisons are most naturally represented as a ratio. Saaty proposes a set of words for encoding ratio judgements, but, as with probabilities, subsequent research (e.g., Pöyhönen, Hämäläinen and Salo 1994) finds that the perceived meaning of the words varies substantially among subjects.

Zadeh (1965) defines a fuzzy set as one whose membership (characteristic) function, instead of mapping objects to true or false (1 or 0) as is done for a normal set, assigns to each object a grade of membership ranging between zero and one. Zadeh (1975) proposes to account for the meaning of linguistic terms as fuzzy subsets of a universe of discourse. Kochen (1975), McCloskey and Glucksberg (1978), Franksen (1979), Sticha et alia (1979), and Wallsten et alia (1986) employ this working hypothesis and find that human judgements regarding whether a given quantity satisfies size and probability terms, and whether specified items are included in natural category terms, can be represented using a continuum from true to false. For example, people judge 'A robin is a bird.' to be *more true* than 'A penguin is a bird.'.

Zadeh (1968) defines a fuzzy event as one whose occurrence takes intermediate truth values and defines a fuzzy logic where 'or' and 'and' are the max and min of the truth values in question. Zadeh (1975) conjectures that this mathematical formalism reflects use of certain natural language terms: that logical connectives follow the min/max rule, and that various linguistic hedges can be represented by mathematical operators on the fuzzy representation of the primary terms. Hersh and Caramazza (1976) studies the use of hedging terms, finding a reasonably successful fuzzy operator account for some phrases considered, but finding other phrases inconsistent with Zadeh's speculation. Oden (1977) finds evidence that Zadeh's max/min rule for logical connectives does not reflect natural language usage, but that people combine fuzzy truth values as if they represented independent probabilities. Franksen (1979) finds practically no empirical evidence to support the submission of fuzzy representations of psychophysical continua to arithmetic operations. Zwick and Wallsten (1989) investigate human probabilistic judgements regarding events described in fuzzy terms, and find that subjects are unable to carry out the complex calculations required by accounts that treat fuzzy events as a primitive; instead, they find that subjects essentially perform conventional probabilistic inference as if a clear event definition were made, and then attempt to combine inference results resulting from different hypothetical clear definitions.

In sum, people judge both preferences and probabilities by adjusting from an anchor, the adjustments are qualitatively appropriate, and preferences are formed by the sum of attribute-wise adjustments. Judgement of similarity is also found to behave as if responding to attribute-wise differences, but there is some evidence there that the combination may be by a distance norm, rather than by addition. Judgement of term applicability and set membership is on a continuum, and these fuzzy judgements are combined roughly like probabilities. The informal work on judgements of impact is clearly consistent with the base-case-and-adjustments notion of judgements, and it asserts adjustments to be intrinsically ratios, and thus to be additive on a log scale.

Recognition

The mode of thinking discussed here is to index into mental structures and retrieve (recognize) corresponding data structures (schemas) that contain the essence of relevant judgements or competence. These data had presumably been learned inductively.

Simon (1979) finds that man has a small short-term memory and an essentially unlimited long-term memory, which is indexed in chunks that permits rapid recognition of familiar stimuli and rapid access to stored information associated with them; it indicates that human knowledge is stored both in schemas and in productions. A production system used as an interpreter of a propositional network is a major portion of the theory of human cognition in Anderson (1983). In it, situation-specific productions that can be used to form judgements are sometimes available. Rumelhart (1985) characterizes schemas as active mental processes that contain knowledge and postulates a recognition device that evaluates schemas' goodness of fit to the data being processed. He suggests that we "understand" a problematic situation by encoding it in a set of schemas, and that these provide a reasoning mechanism and problem solving method for the problem at hand. Isenberg (1984, 1986b), Lord and Foti (1986), Klein and Calderwood (1987), find evidence that schemas or similar mental structures explain some of the behavior of managers.

2.2.2 Understanding

I supplement this examination of human thought with an examination of the literature regarding achievement of what is alternately called insight or understanding in the psychological and philosophical literature.

To understand something is to grasp its meaning. The meaning of a sentence is, very roughly, related to the speaker's intention to produce an effect in the audience. So the first thing to notice is that attempting to 'understand' the results of an analysis is slightly metaphoric, insofar as it treats the analysis as an artifact whose content was intended by someone, and attempts to recreate the intention. The metaphoric meaning we assign to 'understand a study' is to grasp what is important in it.

For this section, I take the word insight to mean a conclusion with reasonably broad policy implications. This definition is at least consistent with the remarks of Craik (1943), who said understanding allows control, and Lonergan (1957), who argues that insight 1) is synthetic and a priori, 2) comes suddenly and unexpectedly, releasing the tension of inquiry, 3) pivots between the concrete and the abstract, and 4) passes into the habitual texture of one's mind. In this section, I take the terms understanding and insight to be roughly synonymous.

There are a number of accounts of how to achieve insight or understanding in the philosophical literature. Lonergan (1957) argues that only consideration in a broader context can create insight. Friedman (1974) and Kitcher (1981) argue that reduction in the number of primitive laws or elements under consideration creates insight. Johnson-Laird (1980) gives some evidence that understanding often consists in creating a mental model. Railton (1978) argues we come to understand chance phenomena by subsuming them under probabilistic laws. Johnson-Laird (1980), Rumelhart (1985) and Muto (1988) argue that formalizing a circumstance creates understanding. Finally, explanation is widely identified as a source of understanding (Craik 1943, Friedman 1974, Kitcher 1981, Achinstein 1983, Salmon 1984, Harman 1986, Scriven 1988), so much so, that I turn briefly to analysis of explanation.

Two major strains in philosophical analyses of explanation are accounts of scientific explanation, which focus on subsumption of the phenomenon under a law or a causal account (Hempel and Oppenheim 1948, Sellars 1963, Salmon 1971, 1984, Friedman 1974, Railton 1978, Kitcher 1981), and speech act accounts of explanation, which focus on intentions and beliefs of speakers and audience (Wittgenstein 1945, Grice 1957, van Fraassen 1980, Achinstein 1983). In general, we may take an explanation to be something that gives an account that helps a person understand the meaning or cause of a thing. In our modeling context, we may take this to mean that an explanation of what's going on in a system (either a real-life system or a model) gives a causal account of its behavior that the user can understand.

Finally, it is worth noting what practitioners and computer modelers think creates understanding, because, as I suggested in the introduction, it can be expected that their comments are generally responsive to the circumstances that call for understanding, even if they do not give a formal account.

Howard (1968, 1983), Chao et alia (1985), Morris et alia (1987), Jimison (1988), Matheson (1990), Schutzelaars (1990), Taylor and Graves (1991), and Bankes (1993) say that model development guided by sensitivity analysis, value of information, or value of control give insight. Hogan (1978), Huntington et alia (1982), Bunn (1986), Morris et alia (1987), and Hickman et alia (1987) say that comparison and contrast of models creates understanding or insight. Howard and Matheson (1976), Eschenbach and McKeague (1990), Murphy and Weiss (1990), Geoffrion (1992), and Lee (1993) note that graphics create understanding. The importance of simplicity for comprehensibility is stressed in Schrattenholzer (1985), Molokova (1986), Moray (1987), and Bankes (1993). A wide variety of modelers cite the role of a causal account or model in creating understanding: Baskaran and Reddy (1984), Selfridge et alia (1985), Isenberg (1986b), Shachter and Heckerman (1987), Swartout and Smoliar (1987), Prior and Moscardini (1989), Sein and Bostrom (1990), and Augustine and Coovert (1991). A particular variant of this requirement is that a teleological account (one that specifies the purpose for which the model in question was created) is necessary for understanding: Schrattenholzer (1985), Neches et alia (1985), Gray and Borovits (1986), and Lind (1986).

It is difficult to summarize such a variety of viewpoints on understanding. Two points that will play a role later in this thesis are the importance of simplicity for understanding, and the importance of a causal account.

Chapter 3. Appraisal of the Literature

The previous chapter sets out the background against which this thesis should be viewed. This chapter contains an assessment of the problem-solving literature, and it sets out how I intend to make use of it. (I take stock of the literature on judgement in chapter 7, which shows how it motivates the design of a verbal summary facility.) This chapter is structured around my responses to the literature in the previous chapter: I affirm decision theory as the characterization of what one ought to optimize when solving problems; I affirm the construction, use, and reuse of simulation models in certain circumstances; and I affirm a variant of the decision analytic approach to use of decision theoretic computer models. Finally, I describe two outstanding problems in this overall approach, which will be addressed in subsequent chapters.

3.1 Decision Theory

The massive body of work axiomatizing decision theory sets it on a very firm footing. Howard (1992) states this axiomatization in a way that makes it intuitively appealing. He also defines the axiomatic composite as the axioms together with their consequences, and shows that the decision theoretic axioms (and none other) can meet these desiderata. This meta-theoretic approach is an important one. Addition of other desiderata, such as quick or low-cost analysis, might show that no method achieves them all, but my focus will be on high-stakes decisions, where the costs of analysis are small by comparison. Accordingly I view decision theory as an unproblematic component of my approach.

While a large body of literature uses the fuzzy calculus (or some other) for decisions under uncertainty, there are both conceptual and theoretical objections that appear to rule this out. Zadeh (1968) and Oden (1977) argue that imprecision is different from uncertainty, and defend the fuzzy calculus for the former only. Lindley (1982) gives a convincing argument that uncertainty in any decision that matters should be treated probabilistically.

Salo and Hämäläinen (1993) characterize AHP as being properly understood as a method of eliciting utility functions, thus implicitly repudiating the calculus of impacts as being a generally valid way to analyze decisions. This argument allows the eclectic analyst to capture Saaty's important insights into human judgements, and his associated computational machinery, and to use these in service of decision theoretic analysis, without confining oneself to multilinear models of system behavior, as is done in AHP.

3.1.1 DT requires a procedural context

But these remarks do not mean that we should simply instruct decision-makers to specify their alternatives, information, and preferences in a decision network model and solve it. These instructions, though conceptually correct, overlook pragmatic difficulties in the specification of the basis elements (as noted in Fischhoff 1986). More generally, the issue here is that human intellect is bounded by the activities that accompany thought, and any problem-solving formalism must be embedded in a process that addresses them. Communication facilities such as explanation or summary can give aid dealing with these difficulties, especially if embedded in a well-designed problem-solving process. Another way to say this is to note that the desiderata for a decision theory set out in Howard (1992) address perfect rationality, and they could be expanded to reflect the broader notion of adaptive rationality by adding to them the desideratum that there be a simple, quick and inexpensive procedure available to implement.

An important example of these difficulties is the fact that decision theory is helpful only if the DM believes the decision basis is good. To begin, she specifies conditional distributions that she must believe and preference-equivalencies that she must affirm as inputs. She must believe that these identify an equivalent simple prospect for each alternative. She should believe that the choice axiom is applicable to these. There are two kinds of premises in this argument that one particular alternative is best: procedural premises drawn from the DT axioms, and substantive premises drawn from her own specification. All of these premises must be believed if the conclusion (that one particular alternative is best) is to be believed. I call for any important decision making to be embedded in a procedure designed to give the decision maker comfort on these issues.

3.1.2 The DT Procedure

The decision theory axioms define a notion of perfect rationality to which we aspire. A straightforward but cumbersome problem-solving process may be fashioned from them. I describe the axioms and this process here as a point of departure for consideration of other more satisfactory processes.

Probability: Uncertainty about the occurrence of any event can be represented using conditional probabilities, which follow Kolmogorov's axioms and Bayes' rule. The former specify that each event has (unconditional) probability in [0,1], that 0 represents impossibility and 1 represents certainty, and that probabilities of exclusive events add. The latter defines the conditional probability of an event as the ratio of the joint probability of the event and its conditioning event to the unconditioned probability of the conditioning event.

Order: Prospects can be ordered according to preference if adequate distinctions are made.

Equivalence: For any three sufficiently well distinguished outcomes that one prefers at different levels, a probability can be assigned such that one is indifferent between receiving the middle outcome for sure, or that probability of the best outcome and one minus that probability of the worst outcome.

Substitution: The decision maker is to be indifferent to substitution of one prospect for another she has asserted to be equivalent.

Choice: If two alternatives each can lead to either of the same two possible outcomes, it is rational to choose the alternative that leads to the greater probability of the better outcome.

One of the fundamental results of DT is a constructive proof that any decision situation can be analyzed pursuant to the sketch given in the preceding paragraph. One form of this proof is the following set of instructions. Each instruction is directly justified by one of the DT axioms, and when all the steps are executed, a decision is reached.

1. Define enough distinctions among possible outcomes that aspects of the situation not being explicitly modeled do not prevent assignment of a preference measure to each outcome (to systematize complexity).

2. Assign a probability to each outcome (to represent uncertainty).

3. For each outcome, find a preference-equivalent lottery in terms of the best and worst outcomes (to represent preferences).

4. Substitute that lottery for the outcome.

5. Roll back the probabilities. Now each alternative is seen to be preference-equivalent to a lottery of the best vs. the worst.

6. Choose the alternative that is equivalent to the best probability of the best outcome

It should be noted that a constructive proof, such as the proof embodied in this DT Procedure, need only show that the right answer will ultimately be achieved; it need not argue that the procedure is an efficient one. The DT Procedure as stated here is hardly ever followed, because of the huge burden of assessment it entails. Without some focusing mechanism, a huge number of distinctions is required to ensure that all prospects are specified well enough to allow arbitrarily accurate specification of preference-equivalent lotteries. This problem is compounded by the fact that the number of assessments of probabilities and preferences can increase exponentially in the number of distinctions. In practice, simplifying assumptions are frequently used to reduce this assessment task, and computer models are used to encapsulate these assessments.

3.2 Reuse of Models

Creation of computer models can play an important role in creating knowledge of what the best action is, but this role is not as simple as "The model says X, so we must do X.". I argue here that model reuse can sometimes be of value, and identify some of the difficulties that cause suspicion of otherwise-efficient modeling processes.

Often, numerical structure can be found to simplify the specification of probabilities and preferences. Frequently there will be a number of events the DM is content to treat as being determined by other factors under consideration. Deterministic simulation models can be useful for this purpose. In addition, expertise about relationships among events (whether viewed deterministically or probabilistically) is often easier to express regarding small simple cases that comprise the overall situation. For instance, first principles of mechanics typically apply to small components of systems of interest. In such cases, computer models can be an effective way to handle bookkeeping detail generated by these assessments.

Building models for the nonce in ongoing enterprises may require duplicated effort at one or more of the following: re-defining distinctions, re-specifying and coding structural knowledge, re-assessing probabilistic knowledge, and re-assessing preferences. In addition, nonce models may overlook important distinctions that were previously identified.

Unfortunately, standing models tend to accumulate "baggage": sections that were relevant for previous decisions but are not currently useful. Baggage makes models and their data cumbersome to maintain and hard to interpret (Schrattenholzer 1985, Bankes 1993).

Reliance on model reuse is a gamble: you spend some time checking whether the world has changed enough to invalidate the old model and data, and trying to understand a model with extraneous distinctions. If all goes well, you avoid re-formulation, re-assessment and re-coding. The literature on tools for construction of composite models from preexisting modules reiterates all these points, but at a slightly smaller level of granularity -- reusable modules obviate effort at rewriting, but they create difficulties in maintenance and comprehension. If the problem of comprehension can be solved, this can be a fruitful proposition.

Production costing models at electric utilities exemplify all of these points: cost components of electric generation systems are often better known at a plant-wise level than in aggregate, building such models is time consuming, a substantial portion of such a model will typically be useful for subsequent decisions, but these models also typically have substantial amounts of data and logic that distract one's attention from the decision at hand.

3.3 Verbal Summary

The literature on communication-based problem solving seems especially relevant in a thesis oriented toward giving the DM knowledge of the best action. One important way to know something is to be told. But this is only sufficient if one has reason to believe what is told. There is little connection between the communication-based problem-solving literature and the psychological literature that explores what sort of judgements come naturally to people. This thesis sets out a procedure for summarizing analyses, and shows its place in an overall problem-solving process.

3.4 Decision Analysis

This subsection assesses the literature on problem-solving methods, finding differences of emphasis, but no obvious conflict, between the two major approaches, decision analysis and search.

3.4.1 Decision Analysis embodies many good ideas

Decision analysis has been said to embody teaching and learning, and the search paradigm ostensibly finds increasingly worthwhile knowledge states in the process, but the mechanism by which these improvements are supposed to take place is not always stated explicitly. I note that the DA Cycle, under Holtzman's interpretation, is a dialectical process. Hence it can be expected to give the continued improvement of ideas that this promises. Furthermore, the dialectic can be neatly embedded within a search process -- each iteration through the Cycle can be viewed as application of an operator to one's current model and state of knowledge, thereby producing a new model and state of knowledge. The remainder of this thesis explores what sort of information must be extracted from any given state in this search to make good this promise.

Although there are claims that DA Cycle modeling creates clarity of action, there is little explicit support for this claim in the literature. However, consideration of the DA Cycle in light of the problem-solving literature reviewed here suggests that its claims may be valid.

Heidegger's (1926, p. 190) idea of the hermeneutic circle proposes that understanding is an essentially cyclic process, always beginning with a pre-understanding and ending with an (improved) understanding. The idea here is directly applicable to problem-solving methods: if the method does not iteratively analyze and create an opportunity for improvement of understanding (or if it does not at least revisit one's judgement after some analysis), the method does not, in itself, offer any possibility for improved judgement.

The DA Cycle can be seen as combining various ideas from the problem-solving literature for improving DMs' judgements. It calls for iterated improvement of judgement, as described here, and in the generate-and-test literature (Newell and Simon 1972). The DA Cycle also generates, compares and improves two points of view: the DM's direct judgements regarding the situation, and the relatively more detailed modeling view, which is generally built from analysis into low-level judgements and re-combination of those judgements. This comparison of points of view picks up the benefits of the communication and comparison segment of the literature.

The following figure expands upon the treatment of the DA Cycle in Holtzman (1985). It shows that the two different forms of judgement, direct judgements and formal model, call for slightly different responses within the Cycle. The step of the Cycle previously called "appraise" shows up here in two parts: first, the interpretation of computer evaluation of the model, and second, the critical comparison of this interpretation to one's direct judgements.

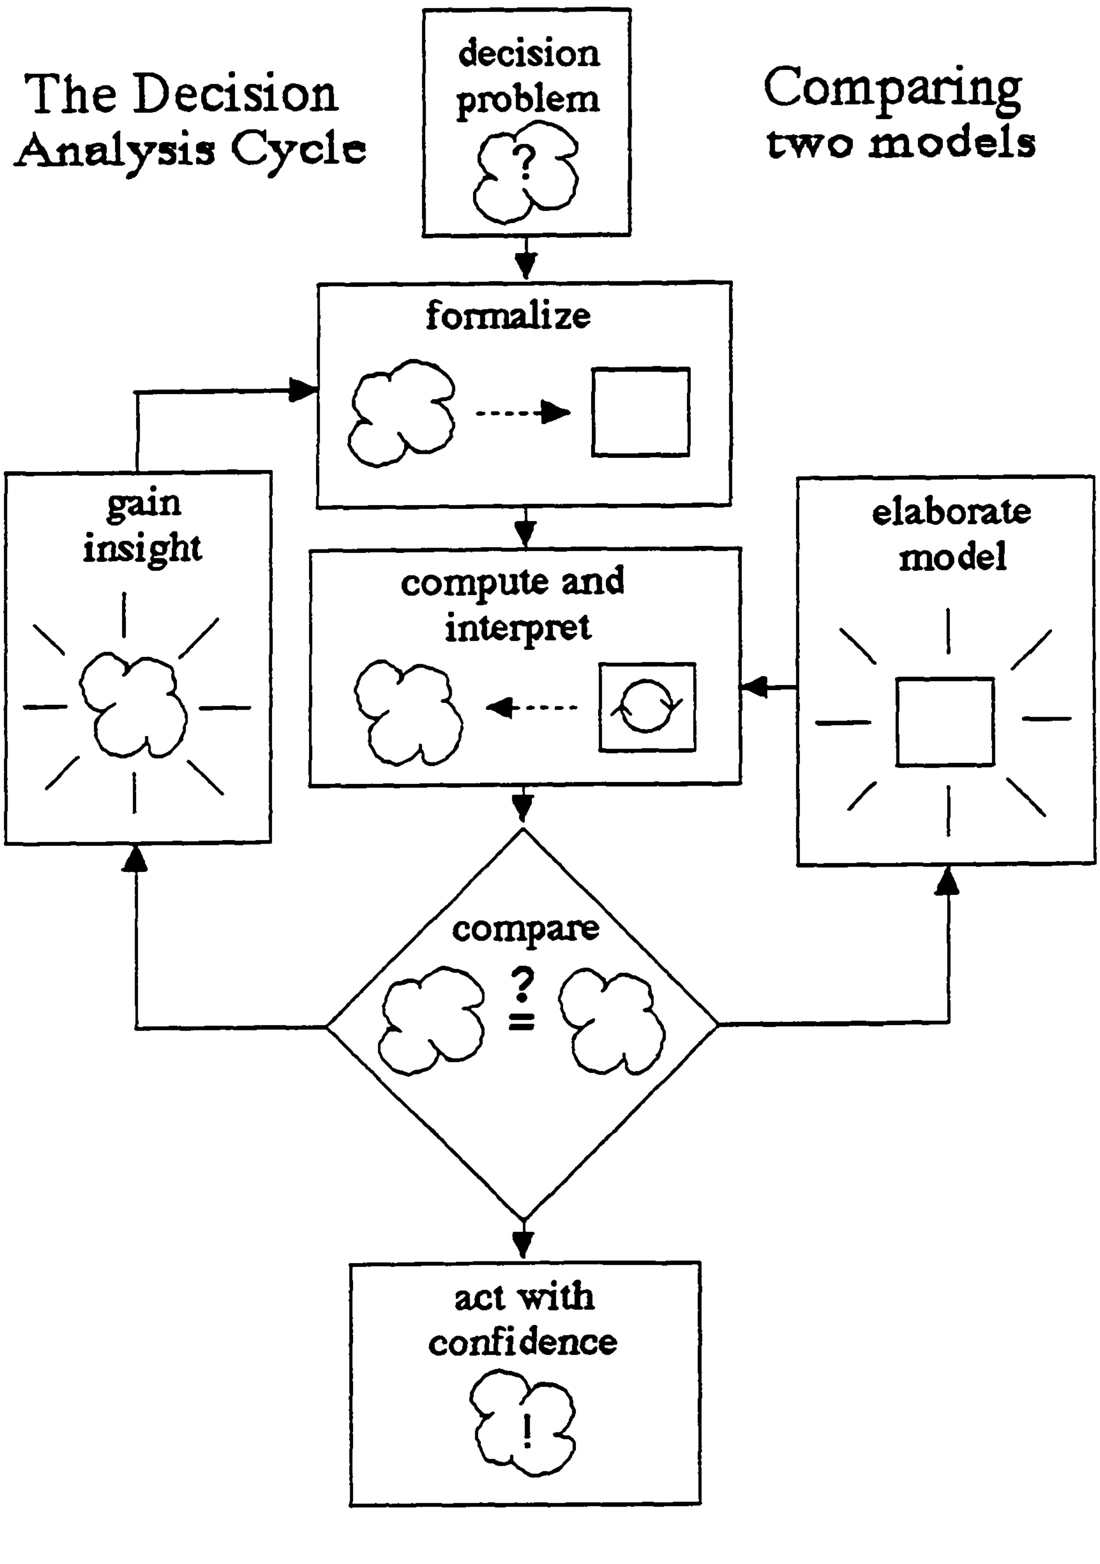
The Decision
Analysis Cycle
Comparing
two models
decision
problem
?
formalize
gain
insight
compute and
interpret
elaborate
model
compare
?
=
act with
confidence
!

3.4.2 DA and Adaptive rationality

For the most part, the issues being addressed by sensitivity analysis are issues of adaptive bounded rationality (Simon 1955): we are explicitly considering the costs of identifying the rational action. Four costs we are trading among are assessment, computation, attention, and infidelity. Assessment of probabilities takes time and effort from both the domain expert and the decision analyst. Computation uses computer resources and analyst time. Requiring attention from the decision-maker is a cost, because decision-makers, like the rest of us, can only think of a few things at one time. I give the name infidelity to the divergence between one's decision model and the decision model that could be created if assessment, computation and attention were free. We can envision a four-dimensional space of decision modeling processes in which a zone about the origin is infeasible. We wish to choose from the feasible processes the one that achieves the best tradeoff among the dimensions.[2]

Three of the roles of DSA summarized above can be viewed in this framework. Model elaboration expends more assessment, computation and attention to reduce infidelity. Fixating variables and eliminating alternatives saves computation and attention costs, but at the cost of fidelity. Each of these roles can be thought of means to the end of a formal decision model that is as faithful to the decision-maker's initial mental model of the situation as possible given the costs involved.

3.5 Remaining Problems

The general approach to knowing the best action recommended in this thesis is to support an enhanced version of the DA Cycle approach to modeling. This recommendation entails that the DM or her organization have a basic familiarity with decision theory and modeling of situations. The purpose of the dissertation is to solve certain problems that would normally have to be solved

---

[2]These aspects of decision process imperfection may be compared to the dimensions of problem difficulty given in [Howard 1968]: complexity, uncertainty, and time dependence; and the dimensions of computer program costs in [Schoppers and Linden 1990]: response time, processing power required, program space, and inattentiveness to problem detail. In all three cases, the purpose of the taxonomy is to support choosing a solution method for problems; the difference is which aspect of a solution method is being chosen among. Howard's taxonomy is oriented toward the choice of a formal solution method, hence its dimensions do not address the cost of creating & using a formal method. Schoppers & Linden's taxonomy has in mind choice of computer hardware and software; hence three of its dimensions can be viewed as an elaboration of the computation dimension, and their "inattention" corresponds roughly to my "infidelity".

by DA expertise, reducing the weight of this requirement. As such, it can be viewed as both a useful tool for current decision analysts, and as a step toward future research that may be able to further reduce the expertise required to levels that are commonly available in the business community.

The remainder of this chapter sets out the specific problems within this approach that my dissertation addresses. The first is the relationship of DA Cycle modeling to model reuse, and the second is the need to ensure that the decision-maker justifiably believes a decision analytic recommendation.

3.5.1 DA and Model Reuse

The traditional approach to this issue in decision analytic practice has been the "SWAT team" approach: the high-powered DA consultants bring in their distinction- and parameter-elicitation apparatus, embed the DM in a DA-Cycle-based process that builds a decision model from scratch (often building and discarding one or two models along the way, Howard 1968) and convinces her what the best option is, and go on their way. Reusability (both before and after) is sacrificed in the name of focusing the analysis and using appropriate distinctions. Previously existing computer models are not used, and no consideration is given to use of the DA model for subsequent decisions. This cavalier attitude toward reusability is motivated by the decision analyst's conviction that every interesting decision is unique, and that reuse of an inappropriate set of distinctions (i.e., those which were well suited to some other decision) can lead to unhelpful results.

However, decision analytic procedures about use of models can be taken to show how to handle the problems associated with model reuse. Routine use of simple preliminary models may help distinguish cases where a standing model would be helpful from those where it would be superfluous, and analysis in the style of the DA Cycle may help to focus attention on crucial model outputs, helping the DM to understand model results.

Indeed, to support critical comparison, inputs and outputs for both sensitivity analyses and probabilistic optimization should be given in terms meaningful to the DM, not those used in some standing model's calculations. As her viewpoint evolves, the variables the DM considers crucial, and their level of detail, change. However, certain aspects of situations faced by an ongoing business often do not change. This juxtaposition of change and stasis presents a challenge: it seems that effort put into modeling situations of repeated interest should be reused, but repeated modification during an analysis of situation-specific variables in a reused model is troublesome and time-consuming. Thus one is led to inquire how such standing models can be incorporated into a decision analysis without sacrificing the appropriateness of the analysis or its comprehensibility to the decision-maker. While my remarks above sketch part of the reconciliation, I will address this further by setting out a computer architecture in support of standing models that can support DA Cycle modeling.

3.5.2 Understanding decision model results is hard

The other major issue to be treated in this work is the difficulty of understanding the results of a decision model.

As yet, there is no theoretical basis for the claim that DA-Cycle-based modification of the distinctions or parameters of a decision model <u>will</u> create knowledge. Psychological research shows that people can comprehend only a small number of 'chunks' of information at once, so the intrinsic clarity of a decision theoretic analysis may be lost if its demonstration requires too many transformation steps. Advocates and practitioners of the DA Cycle claim that it fosters "clarity of action", which corresponds roughly to the knowledge of the best action I pursue here, but no clear arguments in support of this have been adduced. In subsequent chapters, I will set out and defend an account of how the DA Cycle can produce knowledge of the right answer.

The second contribution of this dissertation is a natural language summary facility that highlights the important results from an analysis and summarizes them in English.

Chapter 4. Illustrative Example

This chapter gives an example of the early stages of a decision analysis using the Deft software developed as a part of this thesis. It does so to make concrete the ideas described so far, and to show the target toward which the subsequent theoretical developments are aimed. I set out different experts' views of the decision, describe a scenario describing how it could have been formulated using Deft, and reflect on important features of the illustration.

4.1 Decision: Synfuels commercialization

The decision used here is the decision by the U.S. government in the late 1970s whether to pursue rapid commercialization of synthetic fuels (synfuels) in the U.S. Persian Gulf instability and the appearance of substantial energy-related tax revenues from the "windfall profits tax" prompted the U.S. government to consider appropriating billions of dollars to spur domestic synthetic fuel production projects. This decision is of interest because a sizable amount of resources hung in the balance. In addition, the many authors writing about it (Synfuels Interagency Task Force 1975, Tani 1978, Congressional Research Service 1980, Harlan 1982) make it easier to ascertain in retrospect how important experts would respond to queries posed by a person or system aimed at formulating the problem, and allows me to demonstrate how a problem where experts have substantially different points of view about what is important can be structured and solved effectively.

It can be argued that this decision is on morally inappropriate ground, insofar as it relies upon the coercive power of U.S. government to tax for something besides defense of individual liberties. The author is generally sympathetic with this argument. In real decisions, this issue must be addressed, and unethical alternatives must be weeded out before analysis begins. Once this is done, there is no remaining need for ethical considerations in decision analysis, hence an example that makes no reference to ethical considerations will address all the issues of a typical decision analysis. For the purposes of this example, then, let us stipulate that all options considered here have been carefully considered and judged to be ethically appropriate.

There is more to the synfuels debate than can be captured in a few pages here, and some of these authors may not have been consulted by the U.S. decision maker in 1975. Nonetheless, for expository purposes, I choose a format that is abbreviated and personalized, treating the authors as experts being consulted at the time, and treating certain salient results of their work as if they were presented by that author at the time of the decision.

4.1.1 Tani

In January 1975, President Ford called for a program of Federal incentives for commercial production of one million barrels a day (MBD) of synfuels by 1985. Then he asked for a study to show that this was a wise course of action. Both the Synfuels Interagency Task Force (1975) and Tani (1978) describe this study. The former describes it in four volumes; the latter summarizes crucial points in an article.

Benefits to be studied included technological learning, development of industrial infrastructure, increase in domestic production to insure against oil price shocks or embargoes, and improvement of the U.S. international bargaining position. These were quantified as the sum of producer surplus, consumer surplus, cost of any disruptions of the petroleum market, and costs imputed to air pollution and boom towns. Study horizon was twenty years. Costs and benefits were discounted at 10% annually. The following variables were treated probabilistically: cost of synfuels, state of the Middle East cartel, oil price, and domestic energy surplus. The 1975 U.S. decision (between programs of 0, 0.35, 1, or 1.7 MBD capacity) was considered to be followed by a decision by private industry in 1985 whether to add 0 to 5 MBD additional synfuels capacity, in light of 1985 values of the uncertain variables, and subject to further changes of the uncertain variables.

The expected values for the four alternatives, in $billions, were 0, -1.65, -5.4, -11. Strength of the cartel, which drives oil price, was a dominant factor in the analysis. There was a large negative producer surplus associated with each program, due to the high anticipated cost of the synfuels. The positive consumer surplus and protection against embargo were smaller, as was the environmental impact. Citing failure to quantify U.S. international bargaining position, the Administration chose the "small" 0.35 MBD program instead of no program.

4.1.2 Zuckert

Congressional Research Service (1980) reports a debate on the merits of a crash program to commercialize synfuels. Remarks of four speakers pro and four speakers con are recorded. I focus on one argument from the pro side: Zuckert favored rapid commercialization, arguing that it would allow us to reduce foreign aid expenditures.

4.1.3 Harlan

Harlan (1982) asks whether a federal synfuels program is warranted and how large a program should be implemented. His basic approach is cost-benefit analysis, together with examination of the sensitivity of these results to variations of uncertain parameters. His key conclusions: a moderate program (0.5 MBD) is a robust choice across many uncertainties. It provides basic information and learning as well as some infrastructure development. The following factors make a larger program more attractive: low synfuel costs relative to oil price, high probability of abrupt changes in the market price or security costs of imported oil, low factor cost inflation (which, however, is made less likely by a large program), and high imputed value to development of infrastructure.

4.2 Development scenario

Deft supports an approach to modeling a decision with such a wide variety of issues by making the simplest possible model that comprehends them all, and performing sensitivity analysis on it. This can avoid extensive research or detailed modeling of variables that have little or no impact. We illustrate two passes through an iterative problem formulation process using Deft.

To begin, the decision maker, key experts, and a decision facilitator meet in a room with a large computer screen visible to all. For this case, the decision maker is President Ford, and the experts present are Tani, Harlan, and Zuckert. Ford asks Tani to present the point of view from the Interagency Task Force.

Tani identifies OPEC strength, oil price, cost of synfuels, amount of private investment, and net U.S. petroleum balance as crucial variables. He advocates a two-stage analysis, where each of the uncertain quantities evolves for a period, is observed by the private investors before investing, and then evolves further. U.S. costs, for both routine events and market disruptions, should be calculated using a deterministic market-clearing model.

The decision facilitator wishes to balance the competing demands of capturing all participants' insights completely, and building a model that is simple enough to allow it to be modified to capture different points of view as they arise. Feeling that the duplication of variables to capture just the information known to the private investors may not be necessary if the outcome of the private investment decision is not important to the U.S. policy decision, the facilitator creates nodes for only one instance of these factors. She creates a node by clicking in a vacant area of the diagram and responding to a dialog box, e.g.:

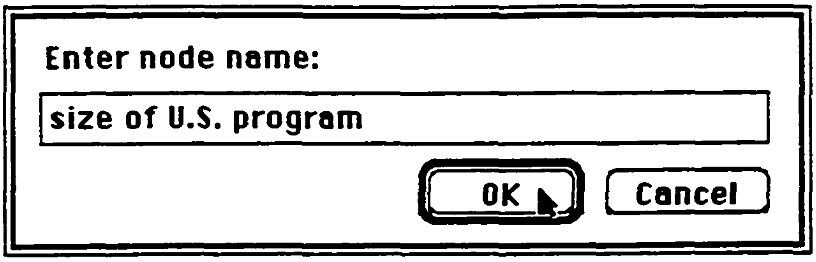

Figure 1. Dialog box to create a node representing size of the U.S. synfuels program

When she has typed in all the nodes, she draws a box around the primary decision, and a hexagon around the utility function by clicking on the "node" button, choosing decision and target, respectively, and then clicking on the appropriate node. This creates the following diagram:

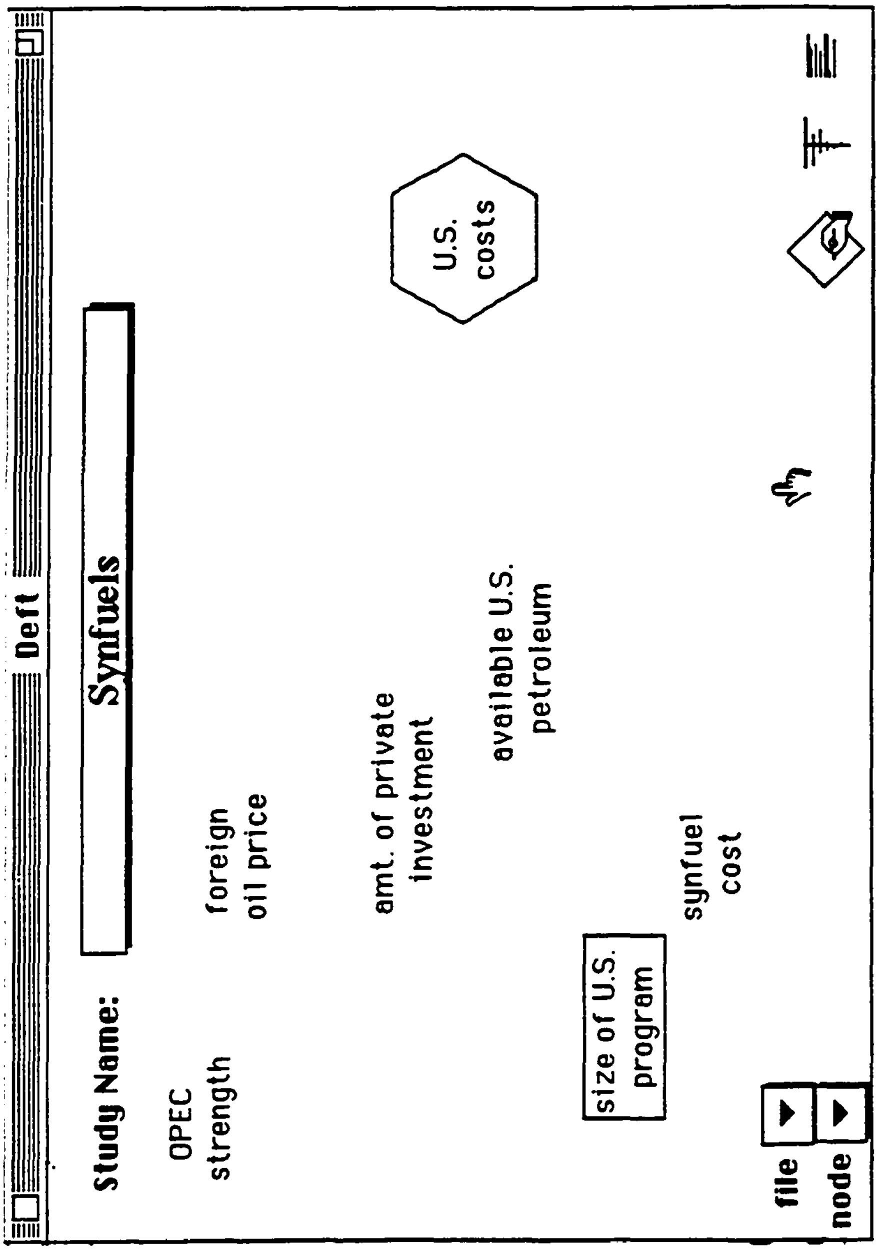

Figure 2. Synfuels network, before arrows are added

She then asks which variables are used to calculate U.S. costs in the model. Tani replies that all of them except OPEC strength are used; its influence is only by way of its impact on oil price. Accordingly she draws arrows from each node except OPEC to U.S. costs. She asks whether the remaining nodes are independent, or whether he feels any should be specified conditional on some other. Tani replies that it depends on how the nodes are defined - are the oil price and synfuel costs to be measured before or after the private investment decision is made?. In a later stage of the analysis (when assessing probabilities) it will be important to ensure that all nodes are defined so clearly that their value in any specified condition of the world can be ascertained without use of judgement. For now, a relatively general definition can be used. The facilitator and Tani agree that for the initial formulation, these should be average values over a twenty year study horizon. Tani then says that he does want arrows from these nodes to private investment, but that even with specification of these averages, he has some residual uncertainty how the market will respond. The only change to the network required here is addition of the two arrows; we are not yet specifying whether the market node is decision, chance, or deterministic. Tani adds that the size of the U.S. program and the amount of private investment in synfuels will help him specify the U.S. petroleum balance, so the facilitator draws in these arrows.

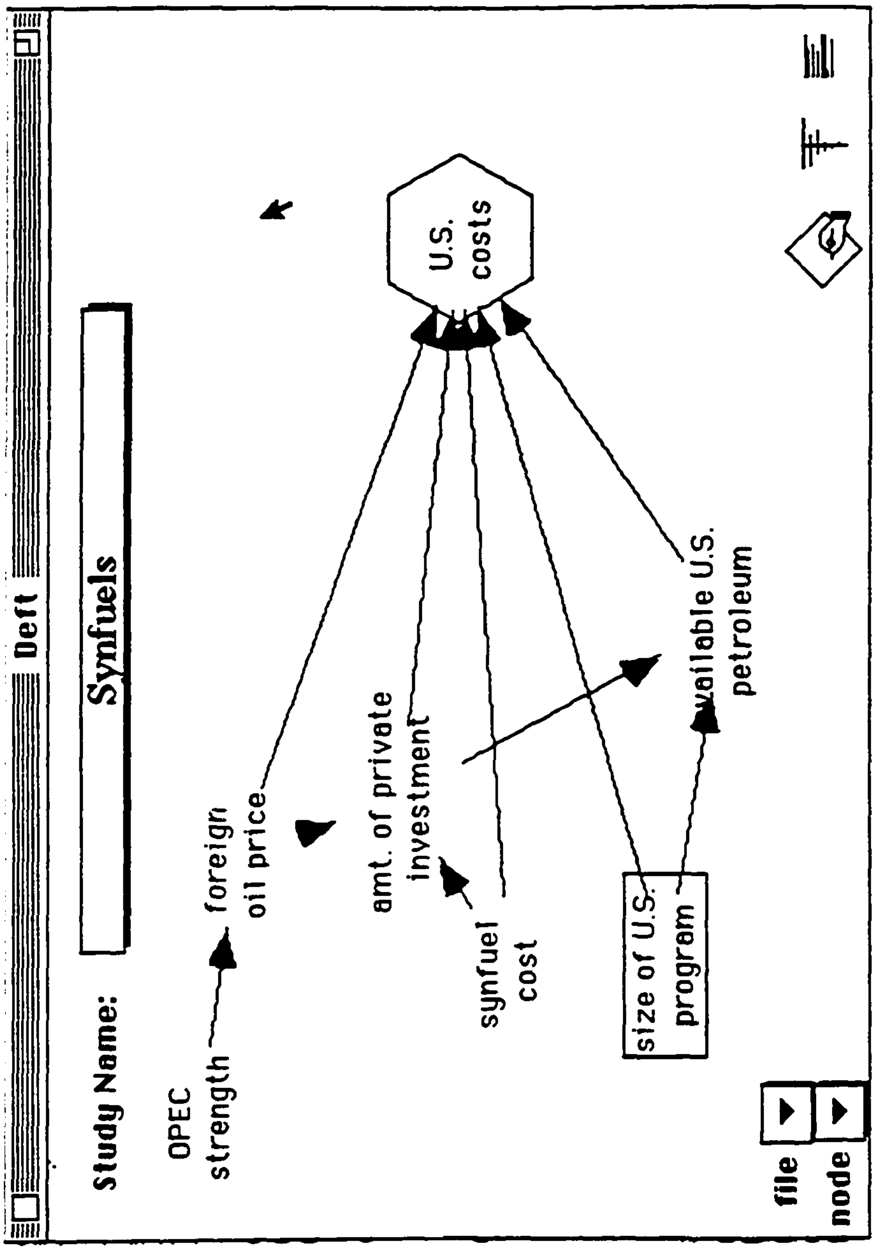

Figure 3. Initial synfuels decision network

At this point Gerry says "Thanks, Steve. That gives us a good start, do you have anything to add to this, Jim?". James Harlan says, "Well the way I had it figured, there are three crucial factors here: the difference between the cost of synfuels and the price of oil, the probability of market disruption, and the value we impute to the development of infrastructure. You already have the cost and price issues in your model. But Tani's model treats market disruptions as deterministic; to my best knowledge, they cannot be predicted, so I suggest treating them probabilistically. My last point can be viewed as a call to acknowledge that construction of infrastructure will build intangible assets by lowering the fixed cost for subsequent synfuel production."

The decision facilitator draws a node for frequency of market disruptions, and a related node for the cost of those disruptions, and draws an arrow from the U.S. petroleum balance node to the latter, and erases the arrow directly to costs. She says she does this because the only impact the availability of petroleum has on costs is by way of disruption costs. Tani and Harlan agree. She draws nodes for components of synfuel costs, and draws in the arrows to synfuel cost. The following diagram results:

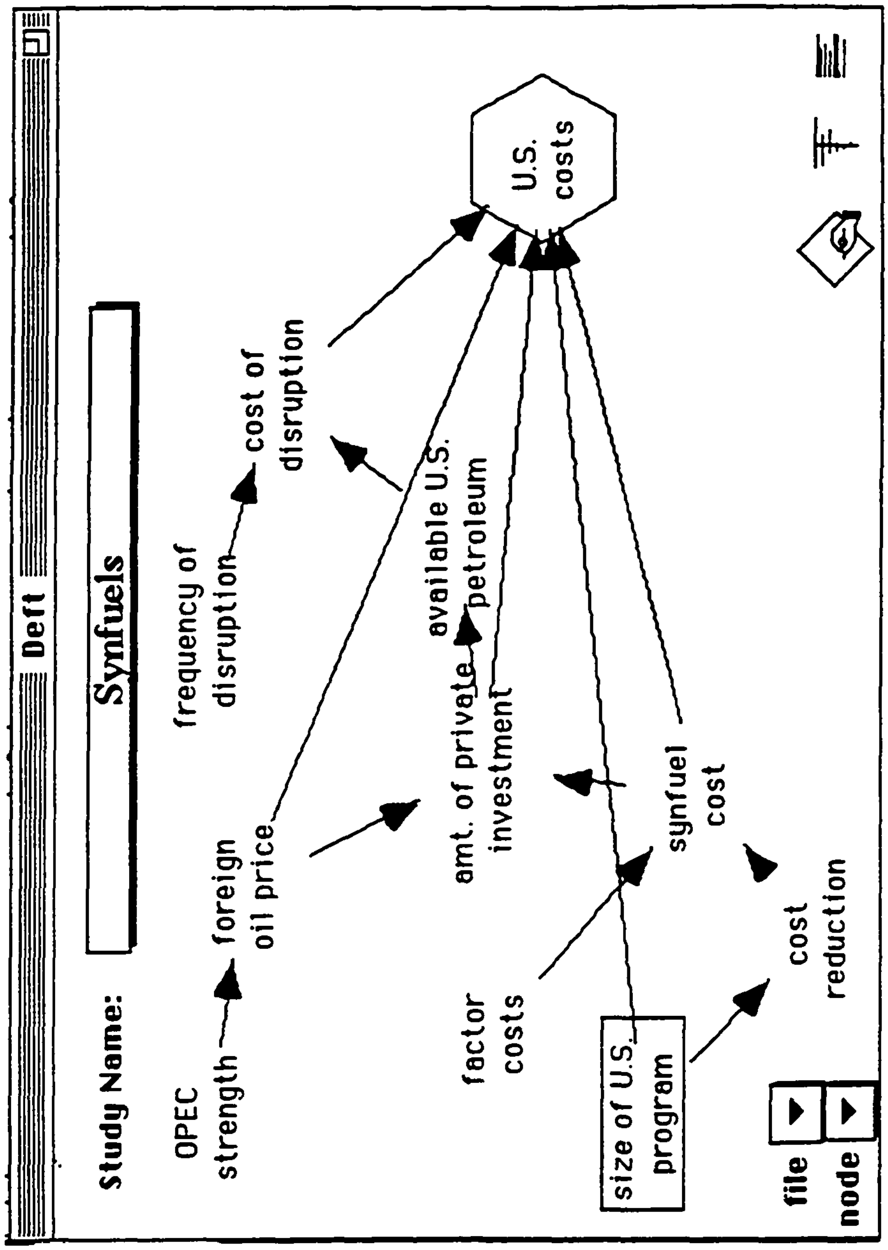

Figure 4. Synfuels network with disruption costs

The decision facilitator inquires whether there are other important factors that should be included in the first-cut model. Tani says that three other factors may be of interest: environmental costs, which apply equally to U.S. and private investment in capacity, the existence and size of a U.S. "Strategic Petroleum Reserve" (SPR), and the discount rate used. The decision facilitator adds a node, "synfuels capacity", to be the sum of US and private investment, and notes that this node and the size of the SPR have a bearing on the amount of petroleum available in the U.S. in the event of a market disruption. She adds appropriate nodes for these. President Ford notes that the choice of discount rate is one requiring careful consideration, and so the facilitator suggests that this issue be treated as an uncertainty pending a policy determination.

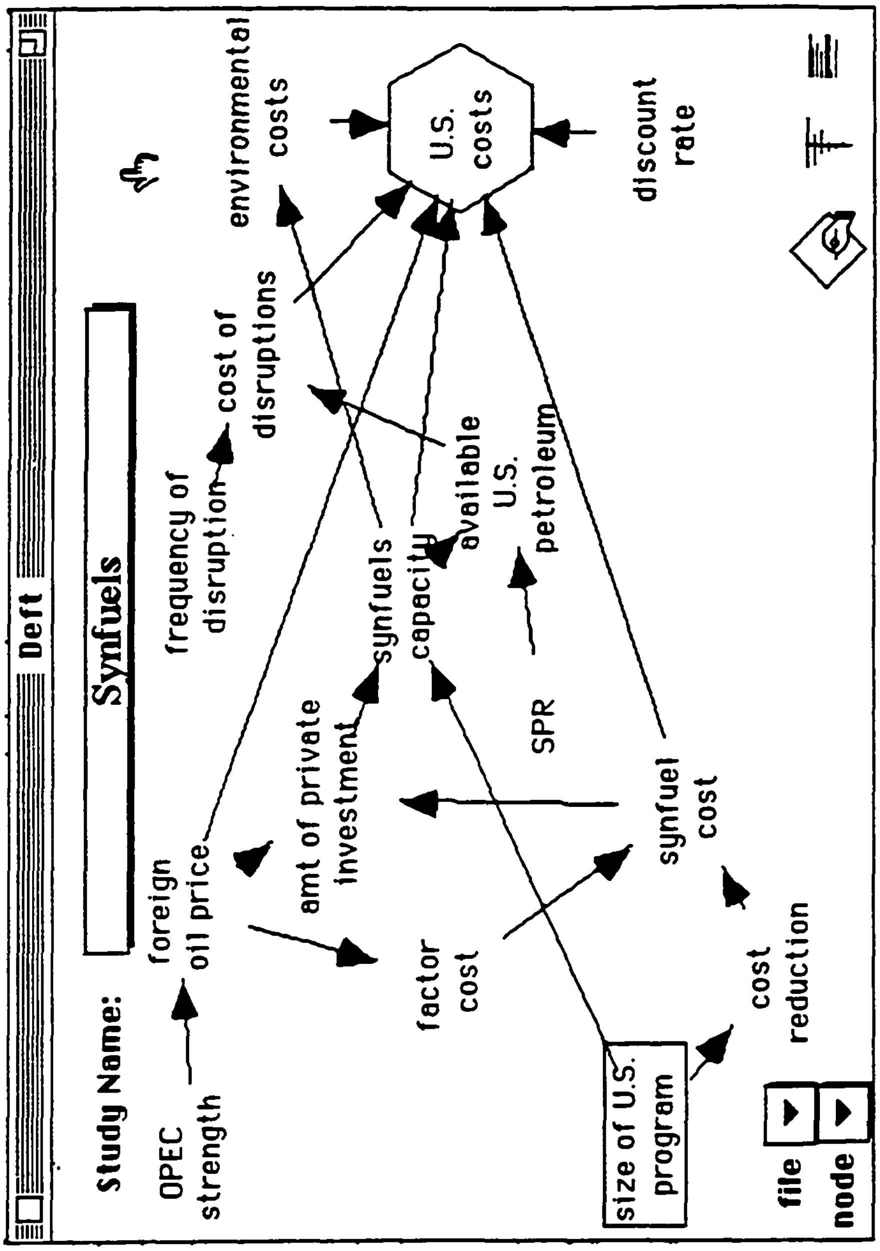

Figure 5. Synfuels network with private market model

At this point, the decision facilitator suggests, and all agree, to take a look at the basic behavior of this model, to make sure that the model so far can capture the diversity of opinion expressed. She clicks on the "tornado" icon near the lower right corner of the screen, at which point the Deft software uses a sequence of dialog boxes to prompt the assembled experts with a sequence of questions about the values of the nodes, and the response of the model to perturbations of these variables. These results will ultimately be displayed as a "tornado diagram". The first dialog box asks which kind of tornado, and the decision facilitator chooses a "Values" tornado, the default.

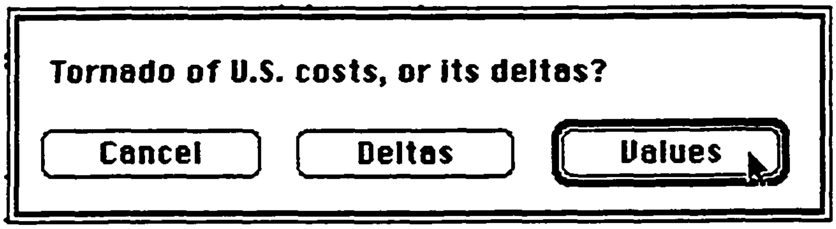

Figure 6. Decision facilitator chooses to create a basic "values" tornado diagram.

Then Deft presents a sequence of dialog boxes like the following, asking for a "base case" value for each node.

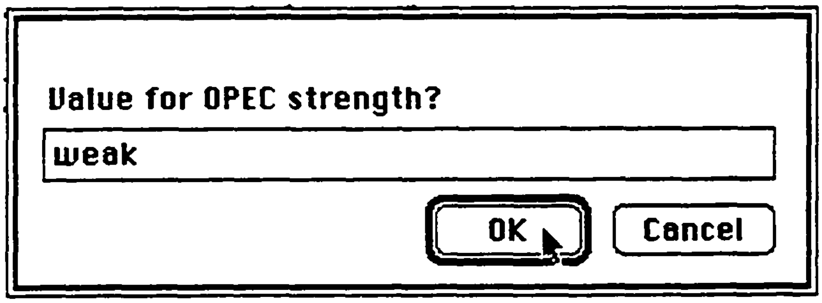

Figure 7. Facilitator specifies "weak" in the dialog box eliciting base case value of OPEC strength.

After some discussion, the experts feel it is more likely that OPEC will be weak than strong, but they are not comfortable ruling out the possibility that OPEC will be strong. The decision facilitator lets them know that these concerns can be reflected in answer to the next question, which maps out a range of values for this node. The purpose of the base case, she explains, is to set up the background conditions against which the impact of <u>other</u> variables is measured, so it merely needs to be a sensible, representative value.

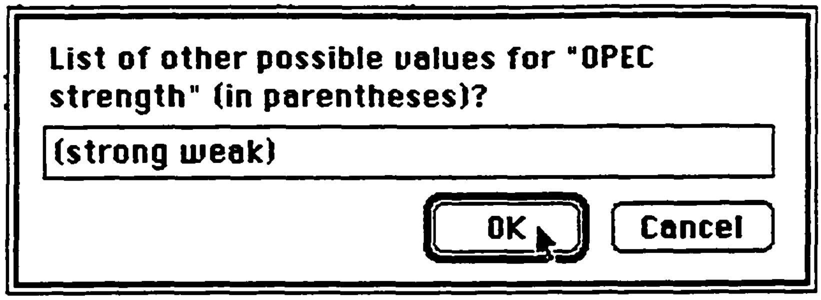

Figure 8. Facilitator specifies two verbal values in the dialog box eliciting a range of values for OPEC strength

Deft then handles the next node in the network, foreign oil price. This node has a predecessor, so the questions it asks are asked conditional on the predecessor. The experts' sense that oil price could rise rapidly causes them to call for a range of values of oil price that has a large upside.

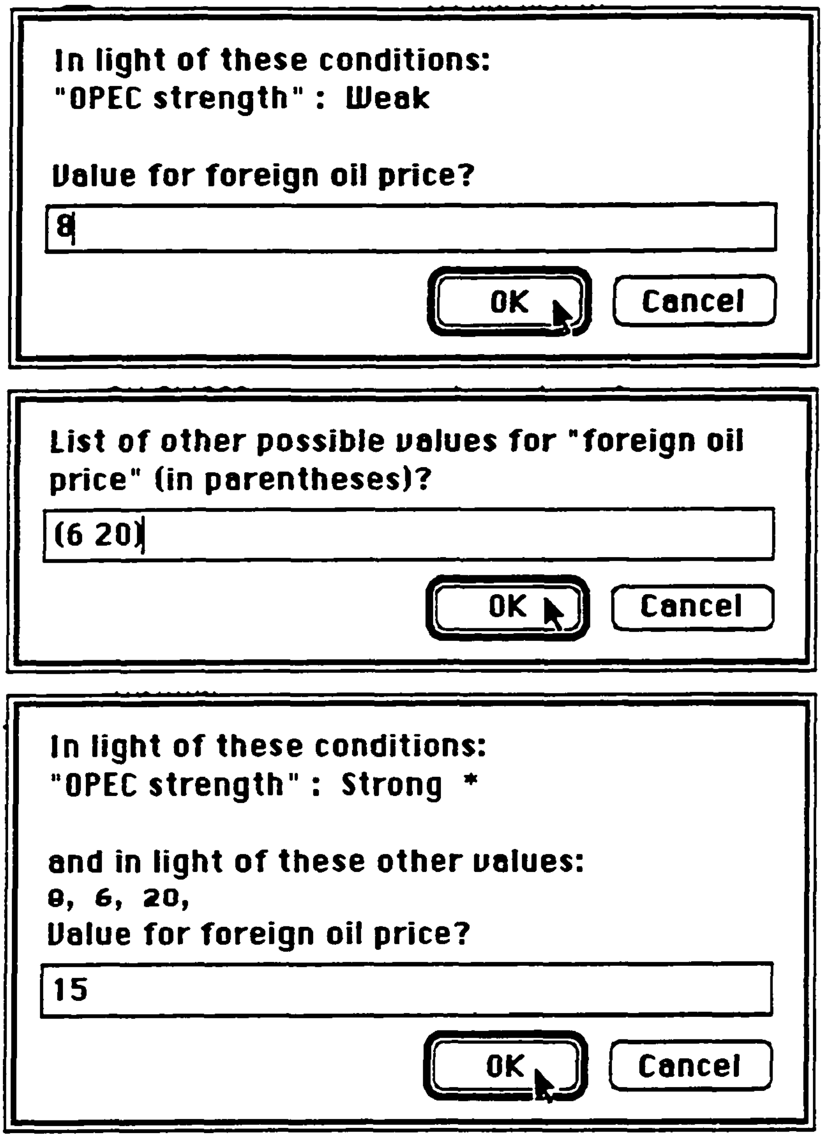

Figure 9. Facilitator specifies numeric values for foreign oil price when prompted by Deft.

The different answers given for oil price under strong and weak cartel capture what is viewed as a medium-level value of oil price under the given conditions. These values will be used later in the calculation of the model's sensitivity to OPEC strength. The facilitator reminds participants that these judgements are for scoping purposes only, to get order-of-magnitude sensitivity measurements, not as final predictions.

More dialog boxes of this nature are presented. If all are answered as they come up, roughly fifty questions will be asked, calling for discussion of a number of interesting scenarios.

Some of the judgements, however, are mechanical and uninteresting. For instance judgements of synfuels capacity are formed simply by summing the capacity built by government and non-government entities. To save the effort of answering repeated uninteresting question, the facilitator cancels instead of answering one of the questions put by Deft. This returns her to the decision network, where she double-clicks on the synfuels capacity node. It "opens up", and she types a formula into the formula field. She explains that Deft will calculate the value of synfuels capacity using this formula whenever a value is needed, instead of calling for direct judgements from the assembled experts.

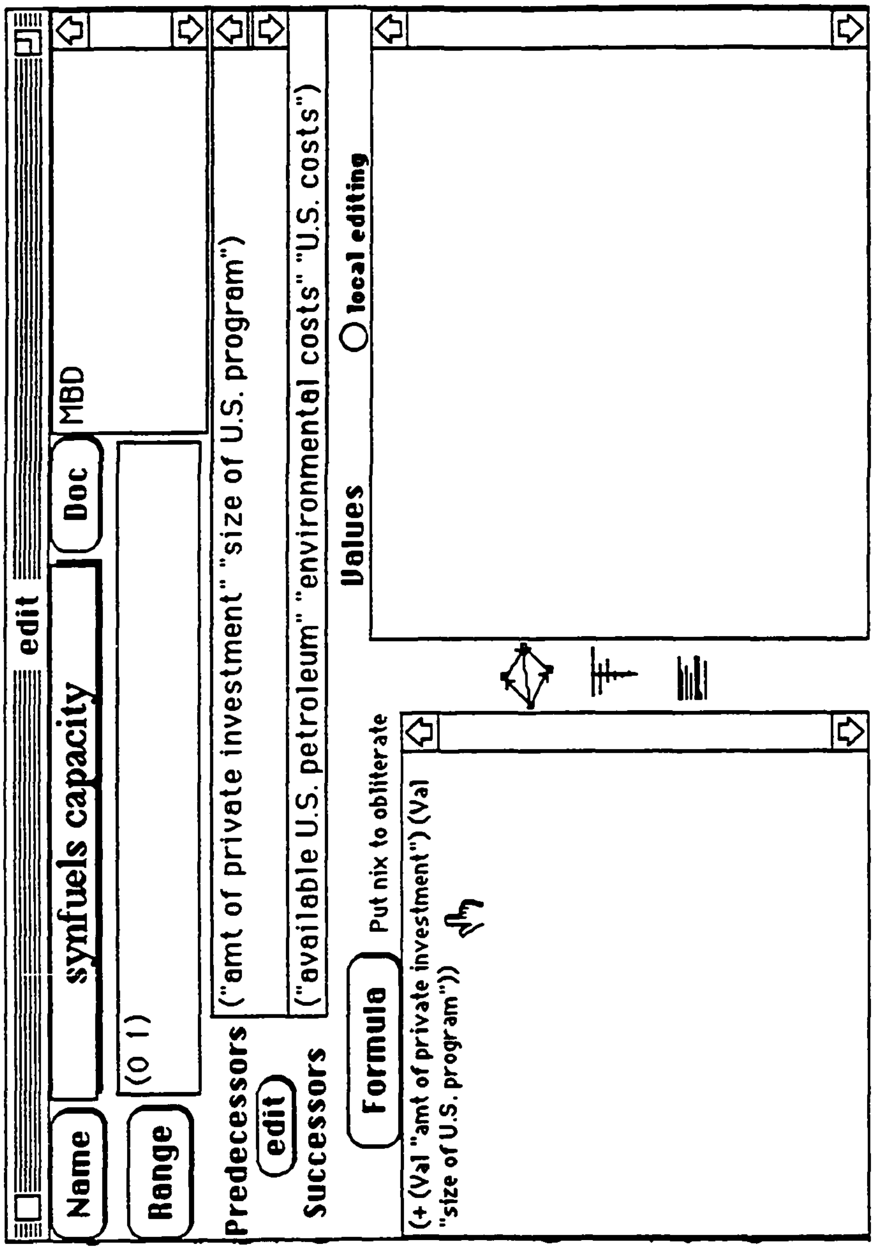

Figure 10. Contents of the synfuels capacity node, including the formula that will be used to specify its value in terms of private investment and U.S. program, as needed

After some discussion, it is agreed that "cost of disruptions" is a difficult problem. In order to more forward, a simple formula based on price response to past supply disruptions and the amount of oil available in the U.S. is crafted. Participants agree that a wide range of values should be given to this variable, to reflect substantial remaining uncertainty about its value.

The formula is based on the estimation that a 2MBD supply interruption in 1974 induced a $20/bbl price change. Taking a simple linear approximation, the change in price is $10/bbl per MBD of net reduction (an assumed 2 MBD reduction, less domestic synfuels production and gradual drawdown of the SPR, as reflected in "available U.S. petroleum"). This price, applied to an assumed U.S. consumption of 20 MBD, and multiplied by the fraction of the time in which this disruption is in effect, gives a simple estimate of the costs of oil disruptions.

The judgements of U.S. costs present a bit of a challenge, because the multiple predecessors give rise to multiple pieces of conditioning information:

In light of these conditions:

"cost of disruptions" : 10.55

"discount rate" : 8

"environmental costs" : 0

"foreign oil price" : 8

"synfuel cost" : 45

"synfuels capacity" : 0

and in light of these other values:

545271, 0, 1000000,

Value for U.S. costs?

Harlan says, "Well what we mean here is that the costs to the U.S. will be the sum of one-time fixed cost for construction of the capacity, and a stream of costs to represent the environmental problems and oil market disruptions, and the net difference between the cost of synfuels and the price of a comparable amount of oil. Do we have to sit down and figure this out each time it asks?" The facilitator replies that they do not, that she can use the node-formula capability to call a function written in the lisp programming language to perform these calculations. The logic is straightforward; she suggests they all break for a cup of coffee, and she is able to write and debug such a routine in a few minutes.

Having made these preparations, the facilitator again clicks on the tornado icon, the group answers more questions like the ones we have seen previously, and Deft then displays a tornado diagram:

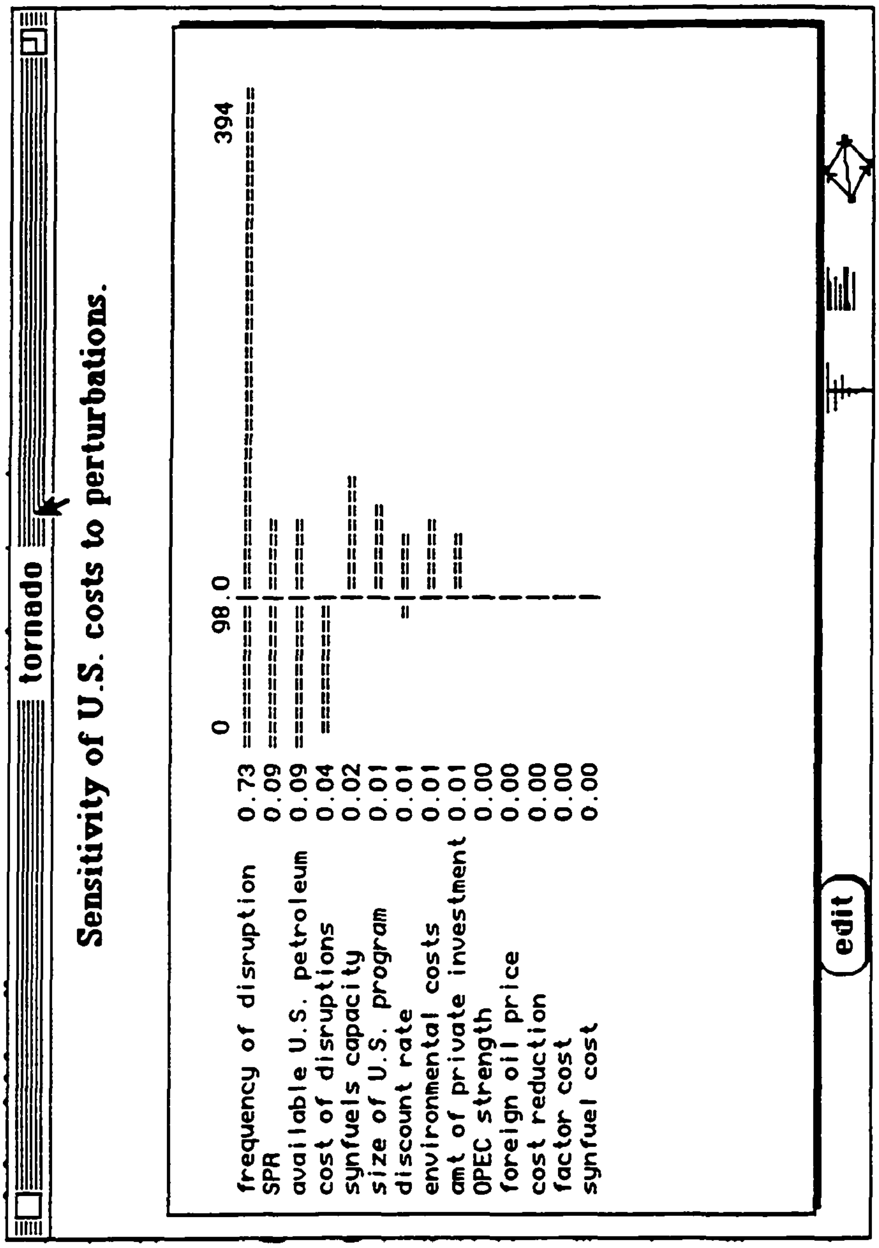

Figure 11. Initial tornado diagram shows that U.S. costs without the synfuels program are strongly affected by frequency of market disruption.

This diagram shows the response of the model to perturbations of each variable, individually. The amount of U.S. costs is graphed horizontally, with a scale in the header. Each bar shows the range of costs that would result from varying each variable through its plausible range and allowing all downstream nodes to take reasonable values. The vertical line shows that the base case, where no variable is perturbed, is $98 billions of costs attributable to the variables in this model.

This diagram shows that variables related to market disruption drive the U.S. fortunes. Upon reflection, the experts agree that it is appropriate for this to be much bigger than the impact of the synfuel program - the program may help, but its impact is not of the magnitude of a sizable disruption of the oil market.

The facilitator notes that this tornado diagram can be useful for finding obvious bugs in the model (there don't seem to be any), but that it could be that some of these important variables affect all options the same way, hence negating their importance for the decision at hand. She suggests running a tornado of deltas, which performs a similar perturbation analysis, but calculates for each perturbation case the difference of U.S. costs between two chosen options. She suggests using the "no program" option and the "1 MBD" option.

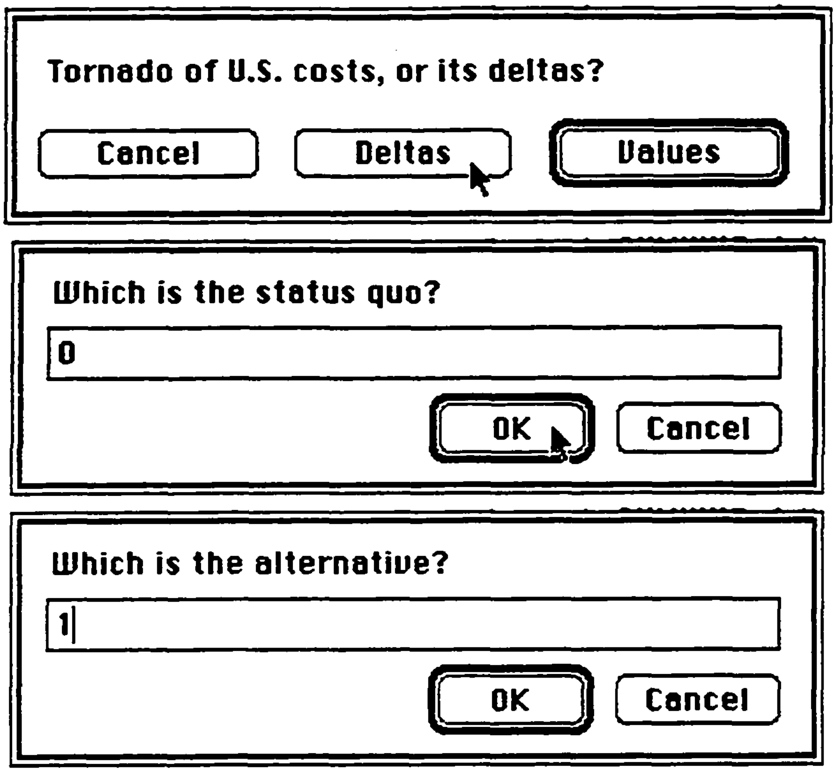

Figure 12. Facilitator specifies the alternative project sizes whose "deltas" are of interest.

The results are as follows:

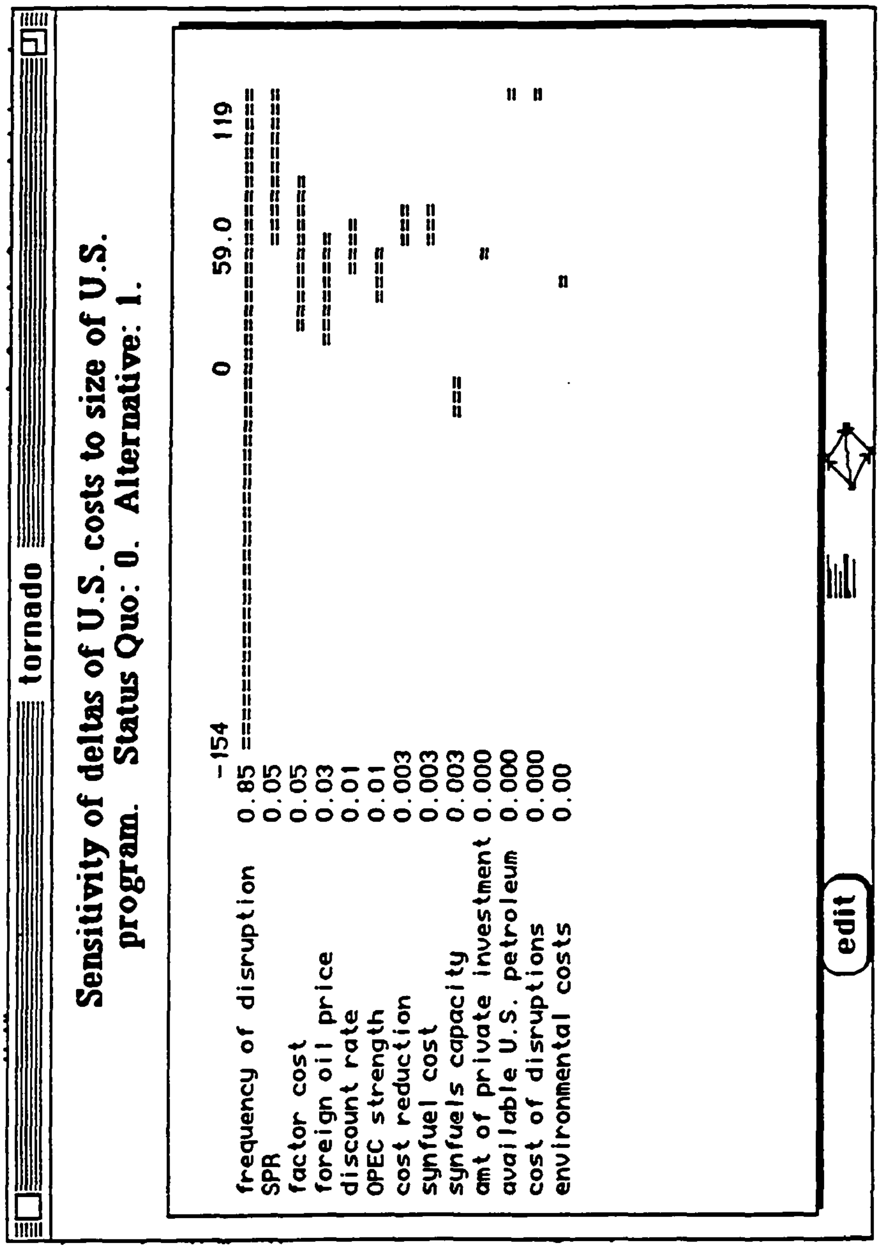

Figure 13. Tornado of deltas of the synfuels model shows a 1 MBD project costing roughly $50 billion more than no project, with delta-costs strongly influenced by market disruptions.

One point to note is that the amount of private investment does not have a large impact in either diagram, hence, at least for now, duplication of all variables to fully capture the logic of that decision, as accomplished in the Task Force analysis, does not seem merited.

The diagram indicates that frequency of disruption is driving the decision. This prompts Harlan to argue that if a disruption occurs, it may take years for the price to return to normal, hence the fraction of time under the effects of a disruption should be larger than the 0.1 the group had previously assessed. The group agrees to raise this to 0.3, partially because the results are surprising - under the base case assumptions, the project stands to lose about $60 billion. Nobody in the room expected it to lose that much. The facilitator explores the values assessed for other crucial variables, and the experts see that they had not been sufficiently careful with the definition of variables. In particular, the oil price variable needs to be a study-period average, but the value and range being assessed was for near term conditions. All present felt the price of oil would escalate in real terms. To fix this, the facilitator clicked on the oil price node to "open" it. The prices of 8 and 15 were there. The following diagram shows the node's contents after the price range has been respecified upward, but before the oil price values under weak and strong OPEC have been re-specified.

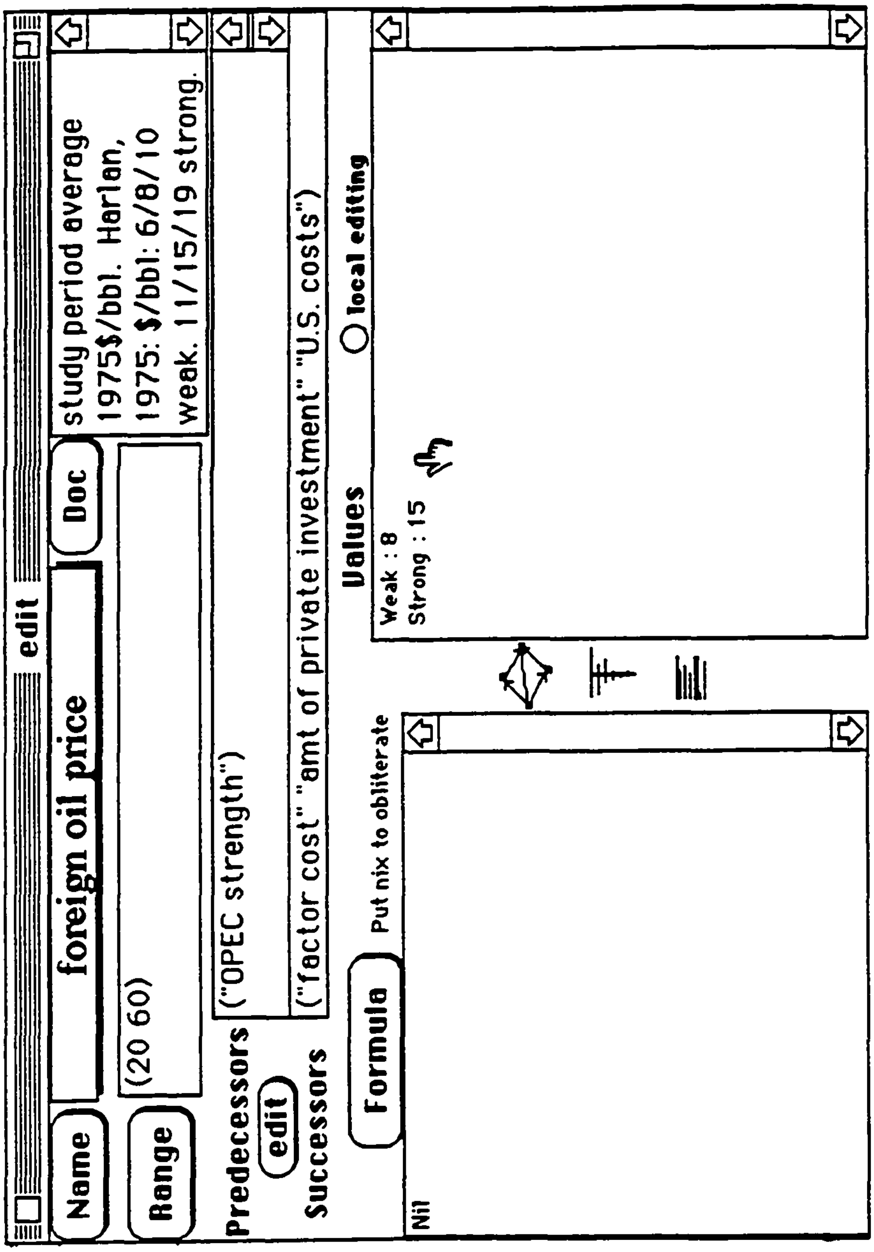

Figure 14. Contents of the foreign oil price node after modification of its range, but before modification of its values

To understand the direction of the impact of each of the variables and help the experts assess the reasonableness of the model to this point, the facilitator clicks on the icon for textual summary. The result is given as follows:

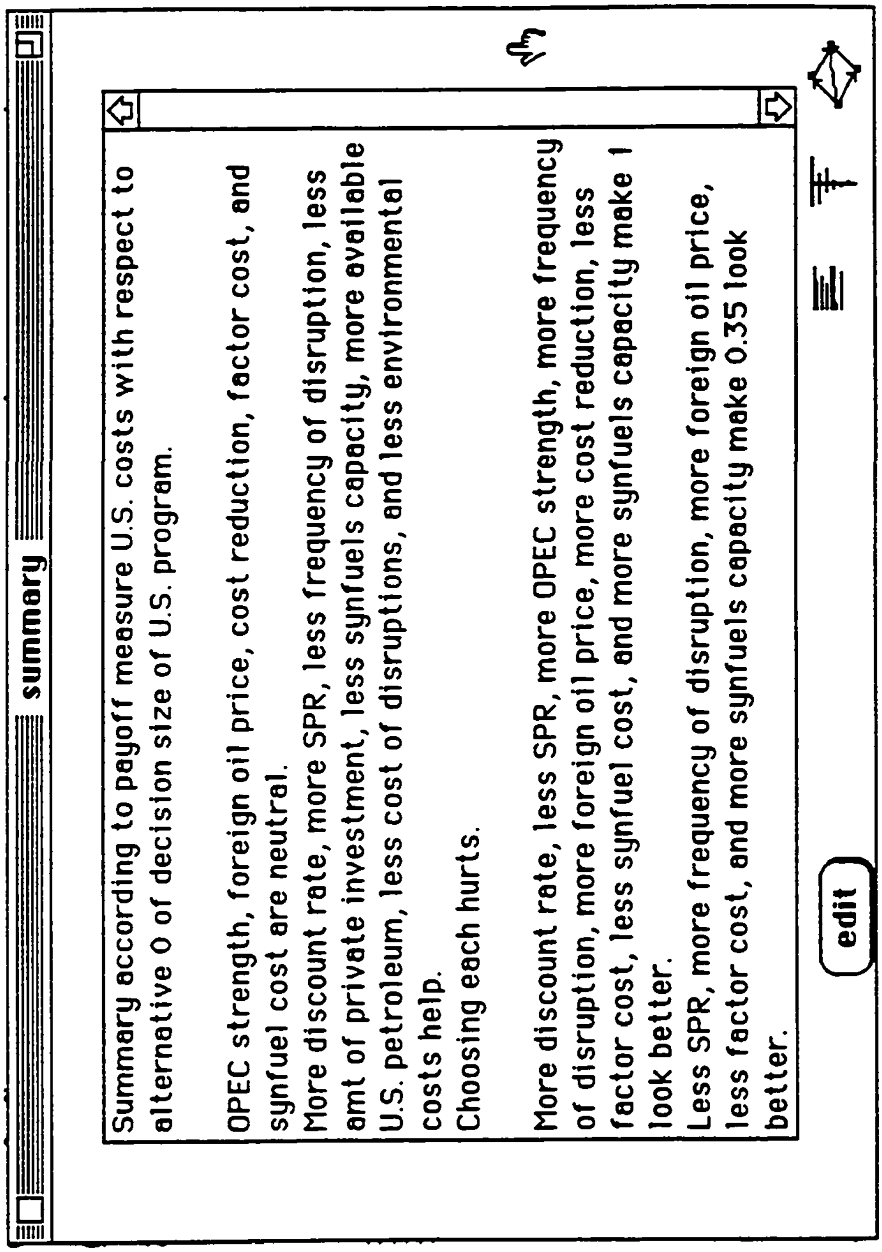

Figure 15. Verbal summary of the behavior of the synfuels decision model

The verbal summary shows the direction of impact of each of the factors. The second paragraph describes the relationships in the first tornado, the sensitivities of model results to perturbations of each variable. The third shows the results of the deltas, the impact of each perturbation on the favorableness of each alternative. The only "surprise" here is that oil price and OPEC strength are neutral in the basic costs - one would expect the U.S.'s costs to be strongly contingent on these. Investigation reveals that the simple cost function multiplied the difference of costs of oil and synfuels by the synfuels capacity, which only counts part of the petroleum needs of the U.S., and which could give incorrect results when comparing cases of different synfuels capacity. The facilitator makes an appropriate change to the cost function written in lisp.

"So, are we missing anything now?", the facilitator asks. Zuckert mentions his concerns about foreign aid, saying that a large program will allow foreign aid to be reduced from its $60 billion level. "How much?", the President asks. "That would take additional study to determine," Zuckert hedges. The facilitator says "Well, give us an upper bound.", and Zuckert says it could be up to $5 billion a year difference.

Harlan recalls that another important dynamic in his study of the issue is that substantial private investment in synfuel capacity would drive up the variable cost of factors of production. The decision facilitator notes that addition of an arrow from private investment to variable cost would create a directed cycle in the diagram, which can cause difficulty in probabilistic formulations. The problem is that the use of study-period-averages in the definitions of the terms used here allows early values of, for example, variable cost, to influence later values contributing indirectly to overall average variable cost; probabilistic models specified in this circular way run the risk of precluding any solution, or of allowing multiple solutions. To make a formulation that is smoothly extensible to a probabilistic decision model, one of two choices must be made - either the nodes must be combined into one multiple-valued node whose different components represent the variables of interest (in this case, private investment and variable and total synfuel costs) or this must be judgementally simulated in a diagram where one of the arrows in the directed cycle, presumably the "weakest", is omitted.

For the moment, the latter choice is made - to omit the arrow from synfuel cost to private investment, and specify the latter based only on foreign oil price. The facilitator edits the private investment node, and finds that the group's judgements stored in its values field show only one place where there would be difficulty specifying the amount of private investment in this way: in cases where foreign oil price is 50, private investment was judged to be either 0 or 5 MBD, depending on whether synfuel costs were 45 or 50. The experts agree that discarding the synfuel cost information and setting private investment to some intermediate level in the one borderline case is reasonable for the current purposes.

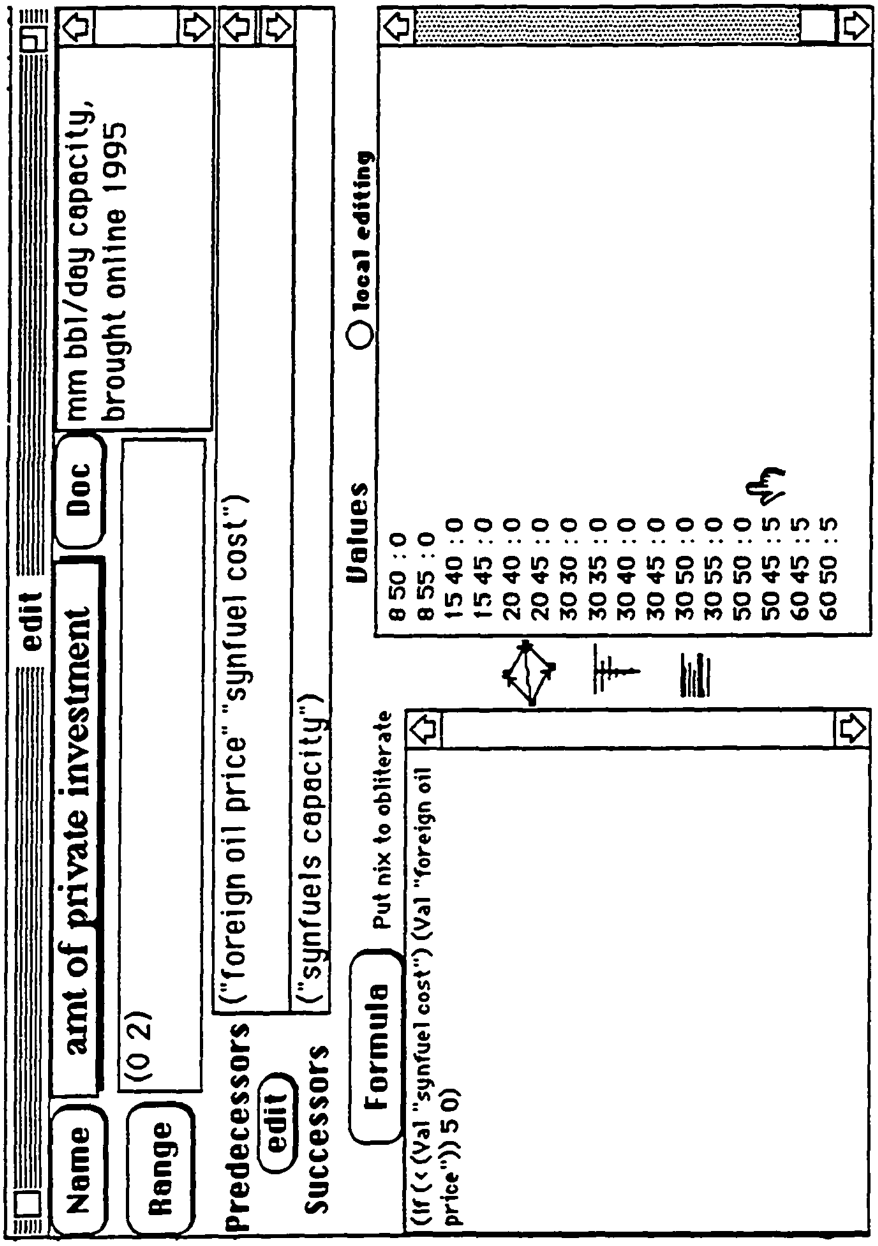

Figure 16. The second predecessor, synfuel price, changes the value assigned to private investment (from 0 to 5) only when the first predecessor, foreign oil price, is 50.

The facilitator redraws the arrows in the network accordingly and re-edits the private investment node. The synfuels cost information that had been present is no longer; only the foreign oil price is shown to the left of the colon. She changes the value for foreign oil price = 50 to 1 MBD private investment in capacity.

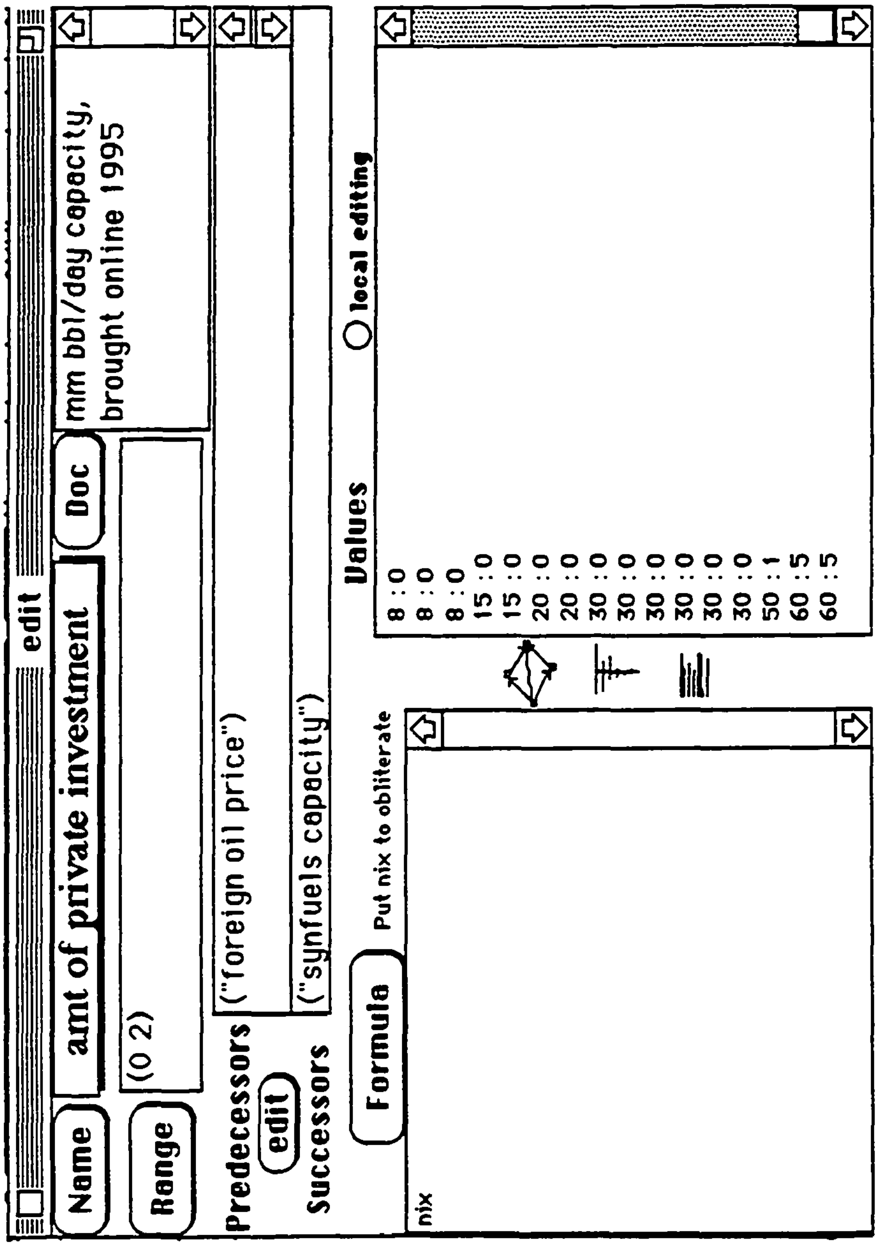

Figure 17. Contents of the private investment node after removal of arrow from synfuels price and modification of projected investment when foreign oil price is 50.

Now she can add an arrow from private investment to variable cost, allowing Harlan's theory of cost escalation to be explored. She clicks again on the tornado icon. Most values required for the sensitivity analysis are now cached, or are computable from nodes' formulae, so only three additional questions, all regarding variable costs when it is influenced by unusual amounts of private investment, are asked. Figure 18 shows the resulting tornado diagram.

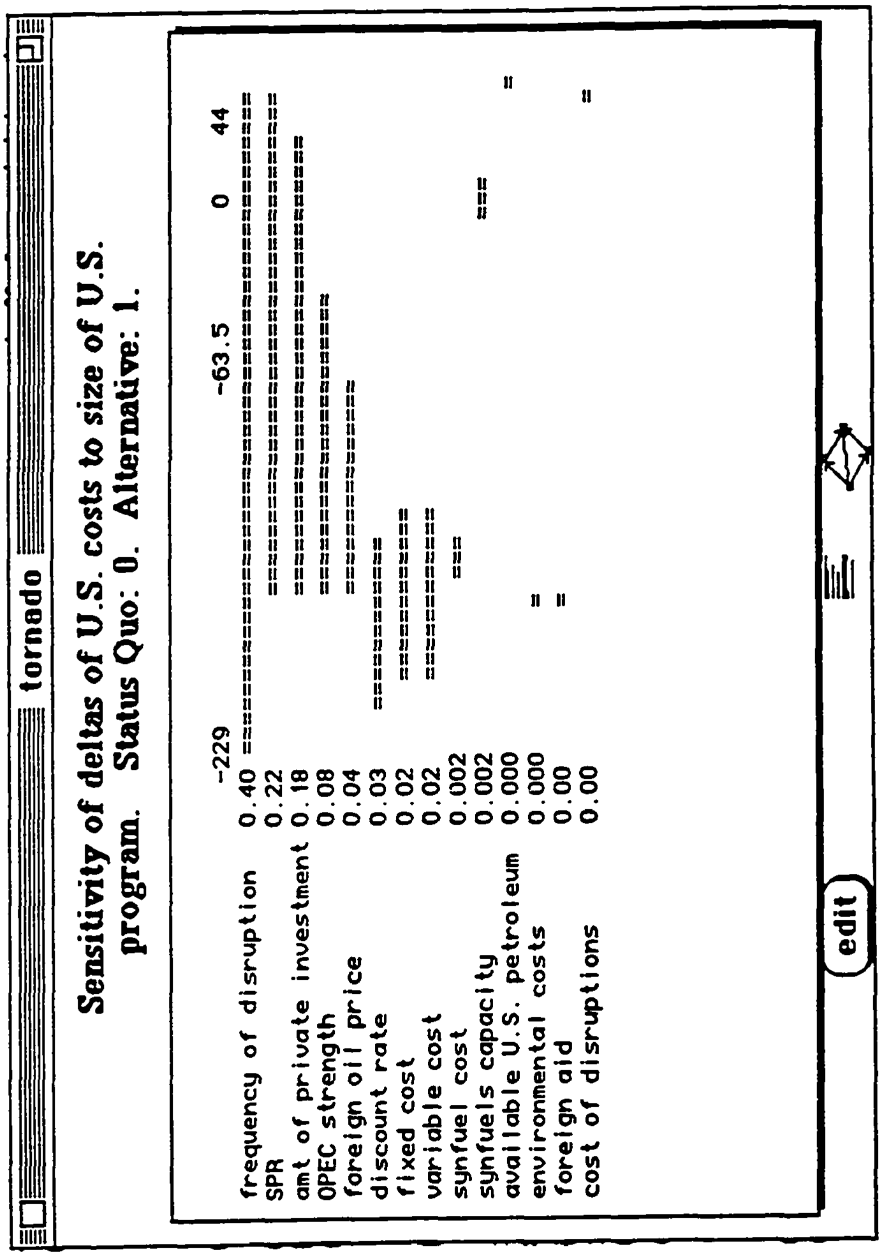

Figure 18. The synfuels program looks more attractive in this revised tornado diagram.

The tornado diagram shows the program in a much more favorable light, due to the reexamination of oil price forecasts and frequency of disruption. The facilitator calls for a verbal summary. No additional questions are required, and after a few seconds this summary is produced:

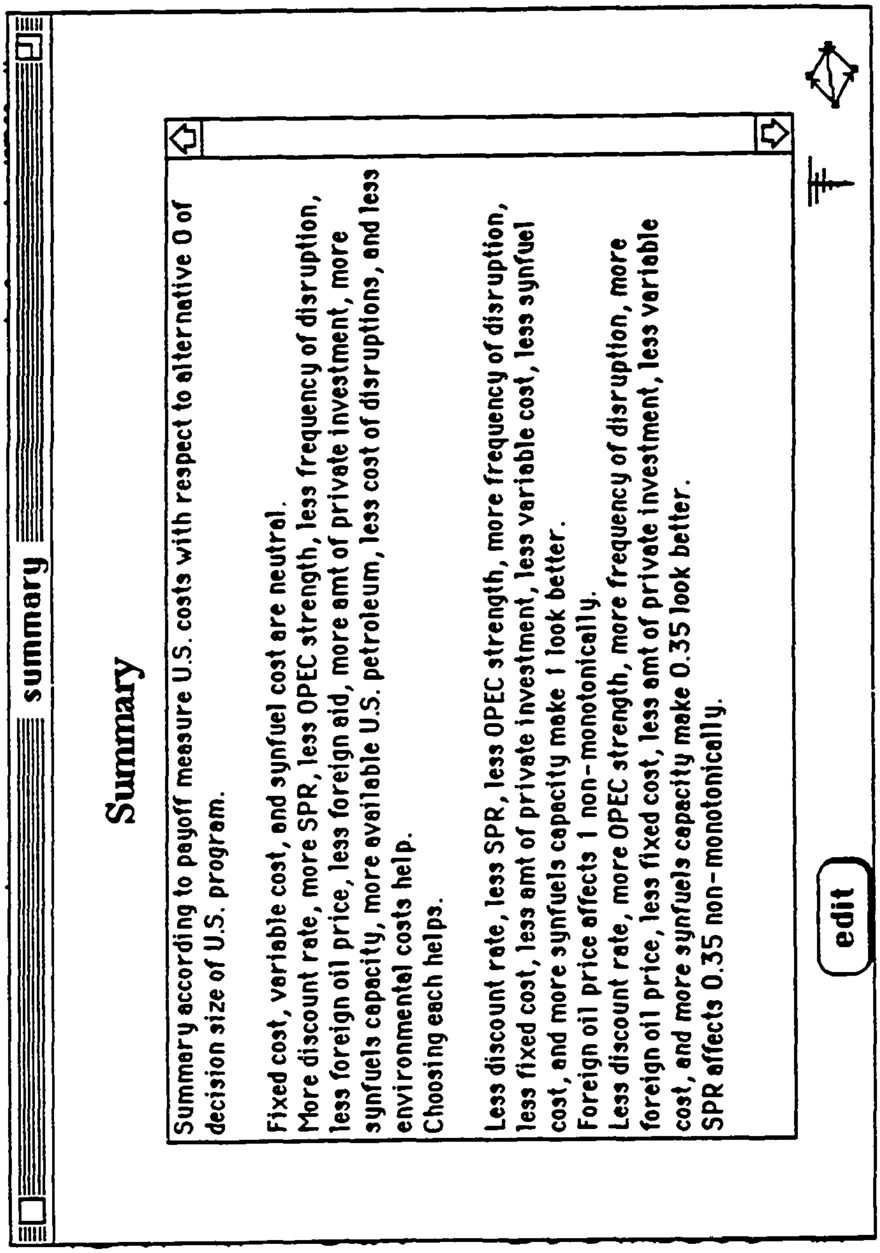

Figure 19. Automated verbal summary of the updated behavior of the synfuels model

This summary manifests a number of puzzles. Some are easily resolved. For example, the insensitivity of model results to cost components merely reflects the fact that the base case does not involve production of any synfuels. Other puzzles, such as the nonmonotonic impacts of foreign oil price (on the 1 MBD option) and size of SPR (on the 0.35 MBD option), and the counterintuitive impact of OPEC strength on the attractiveness of the 1 MBD option, call for careful scrutiny of the model.

Investigation of the impact of oil price on favorableness of a large (1 MBD) program shows that at low and medium levels, increasing oil price makes the program look more favorable, as one would expect. The facilitator reminds them that, "due to the simple private investment model, high oil price causes substantial private investment, which causes substantial factor cost inflation". Harlan says "Yes, but not to that extent - that's economically irrational market behavior." They find that the same phenomenon underlies the counterintuitive effect of OPEC strength. The facilitator suggests that this underscores the need for a sub-model of market response, "This is a portion of the problem where we have enough expertise to merit careful modeling.", she says.

A priori, we would expect the SPR to undermine a synfuels program - the reserves tend to soften the impact of disruptions, lessening the need for a potentially expensive synfuel program. However, the simple disruption cost model indicates that disruption costs would not change from low to medium levels of SPR here. The facilitator asks whether they are comfortable with this. In light of the tremendous sensitivity of the market disruption nodes, they agree that this entire issue should be looked at more closely - that perhaps a precursor, such as the occurrence of a middle east war should be added to the model, and that this might be related to OPEC strength. Furthermore, all the experts agree that environmental costs, in view of their small impact on the problem, should be modeled in the cost model, but should not be the focus of attention. The following diagram captures the new consensual problem structure.

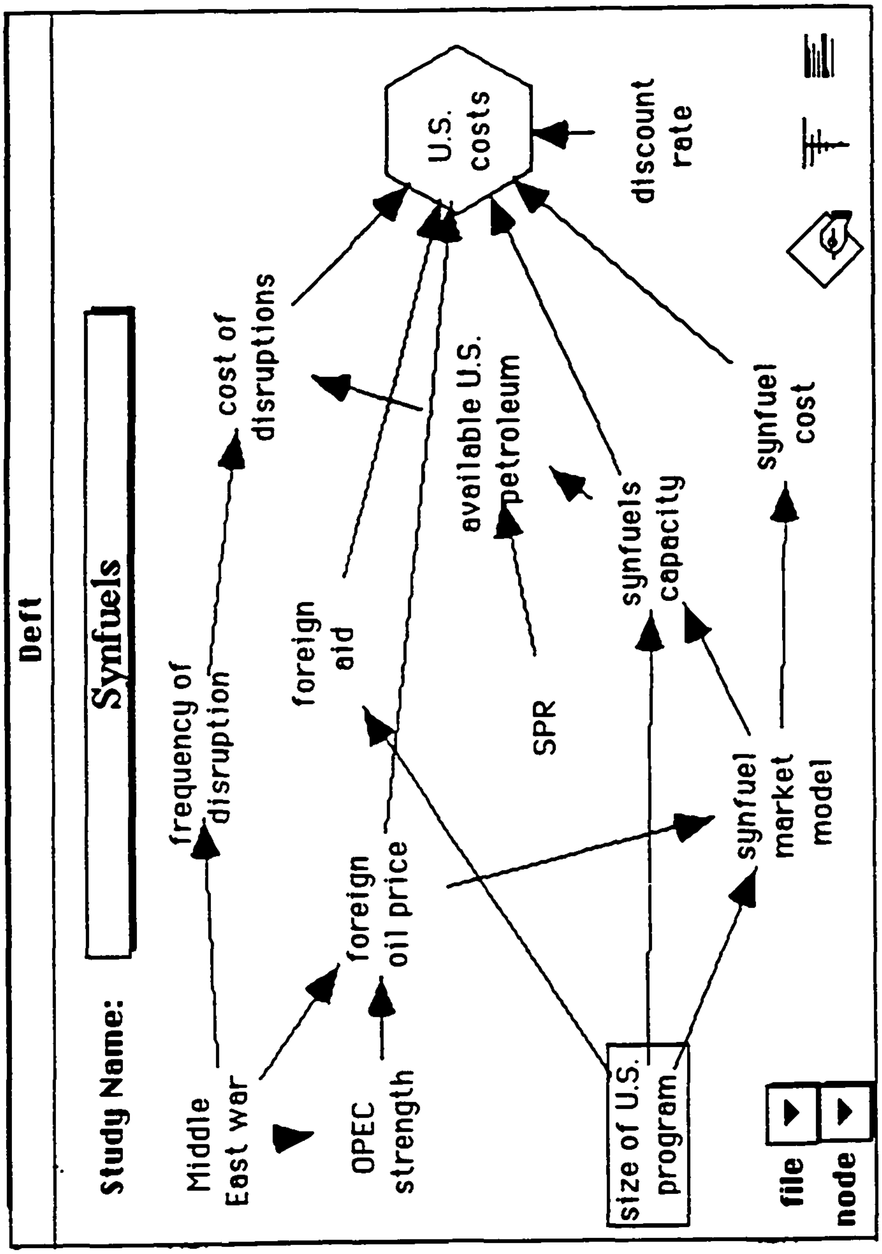

Figure 20. The synfuels decision network after modifications suggested by the verbal summary

This hypothetical conversation could be carried further, but enough has been said to allow some reflections on how Deft has been useful.

4.3 Conclusions from the illustrative case

This case illustrates a number of features of the Deft software, and of the approach to problem formulation on which it is based. I treat these features in three parts: sensitivity-guided modeling, verbal summary of deltas, and the role of decision analytic expertise.

4.3.1 Sensitivity-guided modeling

I discuss three topics here: Sensitivity-guided modeling helps achieve buy-in from participants, and it focuses modeling effort effectively. It does so by ensuring that their attention is focused on issues that make a difference to the decision. The Deft software described here fully supports this approach.

The addition of the foreign aid and environmental costs nodes shows the place of these considerations in the analysis. Seeing that the model is responding appropriately to one's issue gives one confidence in the model, and a feeling of ownership in it. This can help convince advocates (e.g., Zuckert) that their issue has been honored, and still allow attention to be focused on other variables that drive the decision.

Sensitivity-based modeling helps find the right use of computer models in an analysis. It avoids extensive modeling of issues that don't drive the decision (e.g., environmental costs, or detailed modeling of production costs of synfuels, or duplication of all uncertainties to reflect the amount of information available at the time of the private industries' choice of amount of investment). When a model is called for, this emerges naturally, as in the cases of the U.S. costs model, which is a routine calculation, and the synfuels market model. The boundaries of the latter emerged by noting nodes among which we wanted to put very many arrows - this is the set of nodes whose interrelationship we can specify closely. The input and output requirements of the model (inputs: size of U.S. program, foreign oil price; outputs: private investment in synfuels capacity and cost) emerge naturally as we plug in a node representing the model, and submerge nodes that would be represented in the model but not of interest to the decision at hand. Use of computer models to capture systematic expertise about situations that are roughly deterministic can be a valuable tool, if the model is well framed, and if it contributes to a well framed decision. In this illustration, the initial formulation of the problem differed substantially from the final formulation. This is typical. It is important not to formulate and delineate computer models until a good formulation of the decision makes clear the appropriate bounds of the model. The flexible implementation of node formulae in Deft allows direct judgements, simple formulae, and references to complex computer models to be seamlessly integrated into one analysis.

Although it may not seem remarkable, an unusual feature of this example is that it measures the sensitivity of modeling results to all the variables in the network. The network exhibits nodes that are not direct predecessors of the utility node (OPEC strength), and others that have direct predecessors themselves (cost of disruptions). As will be discussed later, most approaches to sensitivity analysis have trouble with one or the other of these classes of nodes.

4.3.2 Verbal summary

The verbal summary is succinct, clear and useful. It is succinct enough to fit on one page because it chooses a limited set of features of the model to address, and because it uses a "zero threshold" to screen out second order sensitivities that have no impact. It is clear because it juxtaposes favorable factors, thereby making counterintuitive results, such as the impact of OPEC strength on the attractiveness of the 0.35 option, stand out clearly. This focuses discussion and allows participants to identify places where the model needs revision easily. In addition, the verbal summary of second-order deltas points to insight, e.g., that a large SPR tends to obviate a synfuels program. Nonmonotonic behavior of the model is usually either a bug or an insight, hence close attention to such issues, e.g., the impact of the price of oil on the 1 MBD option, is warranted, even if it's not immediately clear whether the computer model or one's intuitions stands to be improved.

4.3.3 Decision analytic expertise

This scenario refers to a decision facilitator. This person symbolizes the requirement that the decision maker's organization have access to a certain level of expertise in decision analysis. I argue that if you want to know what the best action is, the best way is to invest in some decision analytic expertise, but I will also sketch what can be accomplished in its absence.

In this scenario, the facilitator did three kinds of tasks: she operated the software (drew nodes and arrows, edited nodes, initiated sensitivity analyses and summaries, and knew how to use nodes' ranges to capture diversity of opinion), and demonstrated expertise at computer modeling (choose study horizon, create synfuels capacity node to simplify the model, debug) and decision analysis (choose whether to type in all of Tani's nodes, remove a superfluous arrow, deal with cycles).

Tasks that I have characterized as merely operating the software already capture some of the activities that previously required decision analytic expertise. The construction of Deft reduces these activities to mechanical ones that could easily be performed by anyone who acquaints herself with the software. It can be argued that some of the tasks I characterize as requiring decision analytic expertise could be taught by a good software user's manual.

I do not consider it a disadvantage that my software calls on some modeling capabilities, especially since these need not be exhibited by the same person who does the decision facilitation task; most large companies already have people on staff with modeling expertise.

The third set of activities do call on some understanding of decision analytic modeling. These skills are, I submit, of fairly moderate nature; and they could be taught in a user's manual. Furthermore, a Deft-governed analysis would degrade gracefully in the absence of these skills. If no one had the expertise to judge whether to exclude Tani's extra set of nodes, they would all simply be included. This would call for more questions from the experts, and then many of those variables would, presumably, be found not to have much impact. They would then be removed, tending to converge toward the final formulation given in this example. The facilitator in this example removed an arrow when it was rendered superfluous; if this were not done, more questions might be asked of the experts, but the model would behave in essentially the same way as this example's model. If the facilitator did not have any expertise at breaking cycles artfully, the mere persistence of the Deft software not allowing the creation of directed cycles would force the group into a workaround similar to the one shown here. In all cases, lack of decision analytic expertise would lead to a more time-consuming process, but ultimately to roughly the same formulation of the problem. The skills of a decision analyst to assess probabilities for the resulting model would then be required for the DM to know the best action.

Decision analysis is a useful skill, so it makes sense that a company that has some will do better. However, due to the research of this dissertation leading to the creation of the Deft software, companies or individual users without can achieve one of its primary benefits, a well focused problem formulation, without decision analytic expertise, but at the cost of additional effort. Once a good formulation is at hand, probabilities must be assessed (Spetzler and Staël von Holstein, 1975) and the problem as formulated must be solved (Shachter 1986).

4.4 Chapter summary

This chapter has given an illustration of the use of the Deft software. Points to remember from the illustration are as follows:

Sensitivity-guided modeling

- Responsiveness to edits can earn participants' buy-in.
- Having a model as part of a decision network can be good.

- Sensitivity-based focusing avoids extensive modeling and data work that may not have much impact.
- Sensitivity-guided problem structuring shows where a model is needed, and what its boundaries should be.
- Sensitivity-guided problem structuring shapes up a problem quickly.
- It is easier to link in an existing model than to perform the extensive and difficult assessment of spot market benefits.
- *Measuring sensitivity to <u>all</u> the variables in the network* realizes the benefits of sensitivity analysis more broadly.
- Sensitivity-based modeling supports integration of multiple points of view into a single coherent analysis.

Verbal summary

- Verbal summary is clear.
- Juxtaposition of favorable factors helps when comparing to one's intuitions, and focuses discussion on important issues.
- Verbal summary of second-order deltas points to insight.
- Zero-threshold in summary makes it succinct.
- Nonmonotonicity improves understanding and points to insight or identifies areas of possible improvement.

Decision facilitation

- Deft integrates direct judgements and computer models smoothly and coherently.
- Deft is most effective if the user has a moderate understanding decision network models, but it will tend to lead even novice users to an appropriate problem formulation.

Experience with these summaries in actual and simulated decisions suggests the following evaluation: Identification of variables whose impact is non-monotonic often leads to identification of a logic bug in the simulation, and when it does not, it often leads to identification of an interesting aspect of the situation. Verbal summaries are easier to compare to one's intuitions of how the system should respond than numerical joint sensitivity analysis tabulations, and such comparisons are almost impossible from conventional output tables, which typically report the behavior of one sector of the model in one scenario. Juxtaposition of alternatives and feature-values that have a qualitatively similar impact can help a domain expert identify bugs or understand the logic of a simulation quickly. Identifying when a difference is essentially zero makes summaries more succinct. Use of this verbal summary in conjunction with other kinds of reports appears promising.

Chapter 5. VTAM Architecture

A crucial aspect of the DA Cycle is that both the computer model and the DM's direct judgements are subject to repeated change, until they come to agree. For computer models, this presents a difficulty that can be addressed by using a model architecture I call VTAM (Variable Translation and Management). This section describes the components of the VTAM architecture in turn, and then describes two examples. Broadly speaking, the architecture inserts two layers of processing between the user and the computer model. The upper layer implements high-level user specifications of study characteristics and summarizes studies for the user. The Deft software discussed in the preceding chapter works at this level. The lower level translates between the variables being used in the model and those used by the DM. I discuss this translation layer first.

5.1 Input/Output translation

I begin by discussing the issues associated with DA Cycle modeling with standing computer models, to motivate the variable-translation component of my VTAM architecture.

5.1.1 Changing variables and standing models

I define a model as a piece of software that requires a set of inputs to be specified, and then generates a scenario, i.e., a sequence of events that is likely to ensue from them, as a basis for generating its outputs. I also allow that a model may generate a probability distribution over scenarios, rather than just one scenario. If a model embodies the best information available to a decision-maker, she may want to base her decisions on its outputs. I call a model a standing model if a model is maintained and re-applied to multiple analyses. For example, a standing model in the electric business might require a specification of an electric utility's loads, capacity expansion plans and forecast fuel prices, and, embodying the best expertise of the firm on production of electricity, it would predict total production costs over an extended period of time based upon creation of an extended scenario prediction. Such a model can be useful for capacity expansion and rate-setting decisions.

If other analysis besides the model is required for a specific decision, the model's inputs should be consistent with the other analyses, and its outputs must be used in conjunction with theirs. There are difficulties in achieving this goal, however. First, the variables that are explicitly handled external to the model will change from one analysis to the next. Second, an event may be treated in a very detailed way in one analysis, and in a very aggregate way the next, requiring some sort of translation or processing to ensure consistency. And third, besides the requirement of dealing with different inputs, the model may be required to produce different outputs from decision to decision, either because different variables are of interest, or because different preferences are relevant in this decision.

This changing information environment for a the standing model from one analysis to the next is recapitulated within a single analysis when that analysis is conducted under the DA Cycle: conditioning variables are redefined to ensure that they are meaningful to the DM in the context of the current decision, variables are made to model events in more detail, and new variables are added and others are fixated in response to sensitivity analysis.

5.1.2 Benefits of Input/Output Translation

I propose that these difficulties be addressed by inserting a layer of processing in front of a standing model, to insulate it from changes in its information environment. Figure 1 illustrates the structural role played by this translation routine in a decision network.

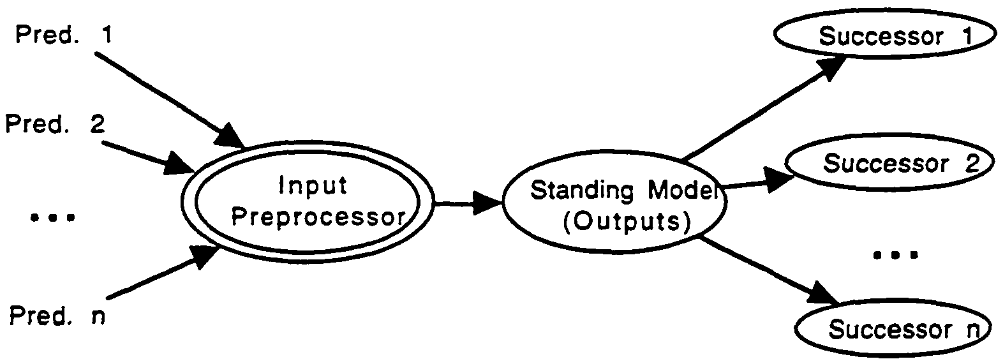

Figure 1. An input preprocessor insulates a standing model from changes to predecessors.

Figure 1 illustrates the basic idea most clearly on the inputs side. I define an input preprocessor node whose outcome space is all the possible input files for the standing model. I have not specified the nature of the predecessor nodes in this figure because it is not important for this discussion whether they are decision or event nodes, or whether there are arrows among them. The outcome space of the standing model is all the possible values of its output files. In this illustration, the translation to decision-specific terms is accomplished in each of the successor nodes. In many applications, the only successor to a standing model will be the utility node.

By making a commitment to a node that preprocesses values of nodes that are predecessors to the standing model into inputs the model can handle, the analysis is done primarily in terms meaningful to the DM, the model is insulated from changes to analysis variables, and the logic required for the translation is localized. These benefits are achieved as long as changes to the predecessors' meaning can be adequately captured by translating them into the model's inputs and using it to process these. Of course, there are times that, upon reflection, the model will be seen to be insufficient for its intended task. In these cases, it must be modified or rewritten, and the variable translation discipline can be used to insulate the rest of the decision network from these changes, if desired.

This variable translation discipline allows the analyst to achieve the benefits of model reuse while minimizing its costs, because, besides the benefits mentioned here, i/o translation is not computationally expensive. Very little of the runtime of a model employing this discipline will be spent on variable translation.

5.1.3 Translation by Parametrization

Typically, a large model has many variables relating to any given feature of the world, to permit use of different organizing principles and levels of detail. It is a helpful coding discipline to identify a minimal set of variables that fully specifies the state of the simulation, and to calculate all other variables from these, to guard against inconsistency among the variables representing the same feature of nature. I call this basic set of variables in the model its state variables. I shall contrast these variables with decision-making (DM) variables, which are tuned to a specific decision.

The variable-translation routine set out above is most easy to implement if one makes the fairly reasonable assumption that the state variables of a model are fairly disaggregated. Treating DM variables as parametrizations of these state variables and localizing the implementation of the parametrization forms a useful basis for the i/o translators in VTAM: The input translator takes specifications of input conditions and translates these into the disaggregated generic terms employed in the model. The output translator translates the disaggregated data from the model into terms amenable to human interpretation. Often this translation can be done on a variable-by-variable basis, but the localization of i/o translation is also responsive to circumstances where specification of some model state variable requires consideration of the interaction of DM variables

5.2 Condition Manager

This subsection describes the component of the VTAM architecture that implements the user's overall plan of investigation using the model, the Condition Manager.

It should be expected that the model will be run under multiple conditions, when conditions are not known with certainty, to explore the model's responses to them. Two major ways to generate sets of input conditions are: one-at-a-time variation of variables from a base case, and full factorial analysis for a chosen set of factors. The former may be used to support DSA; while the latter supports decision theoretic analysis. In addition, perhaps fractional factorial analyses or random sampling techniques may also be employed. I shall refer to these techniques collectively as condition generation. A condition manager localizes the logic for condition generation, thus allowing the condition-generation logic to be used for many models, and allowing different such techniques to be used for a given model, as needed.

The levels of the distinctions to be employed in the DSA and the probabilistic analysis must be chosen to be meaningful to the decision-maker to generate appropriate feedback from her and to persuade her. Accordingly, the routine that generates the input conditions must operate with DM variables and funnel its scenarios through the input translator. Note that this pushes the task of interpreting and implementing interactions of input factors to the level of a detail to be worked out by a technician writing the translator, not the decision maker.

These comments require that the Condition Manager and Summarizer be "outside" the i/o translators, motivating the dataflow diagram in Figure 2.

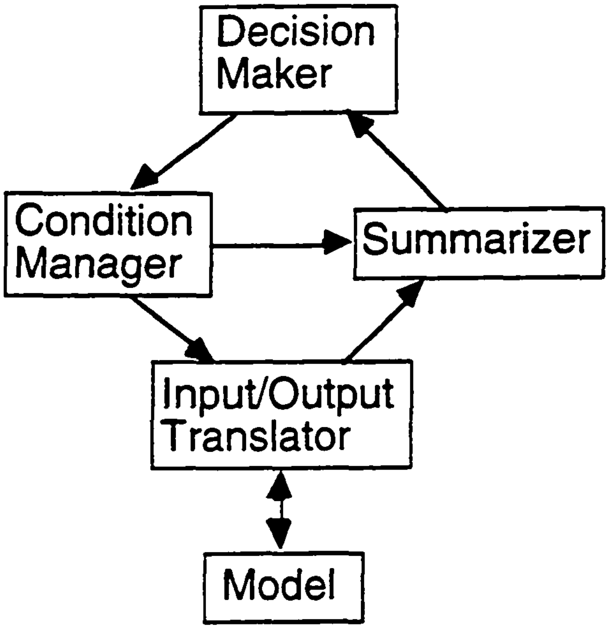

Figure 2. Data flow among modules of the VTAM architecture.

5.3 Summarizer

This subsection treats the component of the VTAM architecture that summarizes model results for the DM.

Many computer models used for real decision-making have been modified to create a large amount of output so that their outputs in various sectors can be verified by experts, and so that political constituencies can assure themselves that their issues are being modeled. Adding a condition manager and translation routines to a standing model is likely to exacerbate this problem, insofar as it allows more extensive analyses to be specified more easily. To avoid the information overload that this could cause, an output management routine may be employed to process results and support output summaries. One way to summarize a full factorial or DT analysis is to coerce it into the form of two-way sensitivity analysis and summarize it in corresponding fashion. This topic will be discussed in the next chapter.

A summarizer can also keep track of probabilities and roll back decision trees for decisions, insofar as a full factorial analysis gives access to all the required scenarios.

5.4 WSDM

This subsection describes the implementation of VTAM in a domain-specific model called WSDM.

5.4.1 WSDM itself

My first implementation of the i/o translation facility described here was in a computer model called WSDM (Western States Dispatch Model), which was used by Portland General Electric (PGE) to forecast the spot market for electric power among electric utilities in the western U.S. Like many other electric business models, WSDM is written in Fortran.

The realization of the VTAM architecture in WSDM is in two levels. The top level is the condition manager/summarizer level. In it, the conditions are checked, then the i/o translators and model are invoked once per condition, and finally the model outputs are summarized and reported. More will be said about this level later. The intermediate level is the translator level. At this level, the base-case data is read from disk, DM variables are reset according to the current conditions, the resulting variables are translated to model state variables, the simulation model (the lowest level of generality) is invoked to simulate the scenario, and its results are translated back into DM output variables. As originally written, WSDM took forecasts of demand characteristics and supply infrastructure and produced a point forecast of spot market behavior; it was simply a model. Subsequently high-level control routines comprising the VTAM architecture were added to it.

5.4.2 WSDM's Input translation

In the WSDM implementation, identification of the state variable changes required to implement DM variables is called cooking the raw inputs. In accord with the parametrization paradigm, multiple state variables typically must be reset in response to specification of a DM variable. For example, WSDM's load forecast inputs may be specified in terms of annual load growth rate, annual seasonal load shape, and circadian hourly shape, whereas the simulation itself needs the actual load in a given hour. Another example is the existence or absence of a Long-Term Intertie Access Policy (LTIAP) at Bonneville Power Administration (BPA) or of a major transmission facility called the Third PNW-PSW Intertie. These may be discussed as single events in policy forums, but to properly reflect them in the code, many changes to the numbers representing various parties' amount of access to transmission of power must be made.

In WSDM, inputs are cooked in an essentially lexical way: a set of text strings and associated data characterizing the required changes to model state variables is identified. I call elementary changes to data elements tweaks. For each tweak, these text strings specify the name of the variable or array to be changed, any required indices, a change-value, and a change-mode. The change-mode specifies whether the change-value is to replace the base-case value of the variable in question, or whether it is to be added to or multiplied by the base-case value.

Due to the lack of self-referential capabilities in Fortran, the facility that implements the lexical characterization of required changes in the actual data structures must be implemented as a long if/elseif/.../else construction, where each case handles one DM variable. Although initial coding was tedious, addition of a new DM variable was straightforward. I isolated this logic, along with associated validity checking, in a subroutine in WSDM.

5.4.3 WSDM's Output translation

In addition to the features and values in the condition-specification input file, the output measure (or measures) to be summarized must also be specified. Normally this will be an objective function, but it does not have to be. A formalism analogous to tweaks for specifying output variables could be created, but I have found that there are comparatively few called for, even from differing decisions. In the implementation discussed here, I simply calculate these output measures directly in the output translator and report whichever is called for.

5.4.4 WSDM's Condition manager

The WSDM condition manager processes files in a specification format tailored to creation of a full factorial analysis of input variables. The most important organizational feature of this format is that it allows specification of multiple features (i.e., the distinctions being highlighted in the study), and for each feature, it allows specification of one or more values. Note that it is not important to distinguish whether a feature is an alternative or an uncertain state of nature for generation of conditions for decision theoretic analysis, because all combinations of alternatives and states of nature must be considered, just as in full factorial analysis. For either sort of analysis, the condition manager invokes the model (with help from the i/o translators) for each combination of feature-values.

It may be noted that this implementation of full factorial analysis will also support one-way sensitivity analysis (DSA) of numerous variables, by treating individual perturbations of all the variables as different levels of one distinction. In addition, specifying a feature with only one value gives an easy way to supersede a datum that is read in from disk without disturbing the disk dataset, if this is desired.

WSDM allows features and values to be specified in an input file called a condition file using essentially the same formalism described above for tweaks: variable or array name, indices, change-value and change-mode. In addition, the user is allowed to give a name to the feature and its values, for use in outputs. Specification of values with both DM variables and state variables is allowed.

Since the condition manager's input specifications pass through the input translator, the ability to specify state variables in a condition file is not necessary. However, it may be more convenient than modifying the i/o translator for distinctions that are not likely to be used in subsequent studies.

It may be noted that the model need not be told whether aspects of the conditions are states of nature or policy options, as long as this is handled in the summarizer. So today we may treat the existence of a LTIAP at BPA as an exogenous uncertain factor, and tomorrow a BPA decision maker may use the same model with a different condition specification to explore the decision of whether to institute the LTIAP.

Here are some examples of possible enhancements to implementations that illustrate the potential of the VTAM architecture. The WSDM tweak formalism has proved satisfactory for many different kinds of analyses, but extensions (such as use of variables or wildcards for array indexing) would be helpful. WSDM treats the input data translator as specifying "initial" input conditions only, and an implementation built on this assumption has proved quite useful. However, this could be difficult to implement if the standing model does not read all of its input data into memory before beginning simulation. There does not seem to be any strong reason to enforce the regularity of changes happening "at the beginning of time", so one could imagine the formalism for specifying tweaks allowing a specification of the simulation-time at which the change to the data is supposed to occur.

5.4.5 WSDM's Summarizer

The summarizer of WSDM takes the outputs from a full factorial analysis and summarizes them in a numerical table reminiscent of that of two-way sensitivity analysis. A post-processor takes this summary and produces a succinct summary in English. The next chapter motivates and describes these capabilities in detail.

5.5 Deft

A more recent implementation of the top level of the VTAM architecture, written in conjunction with this dissertation, is a decision formulation tool called Deft. The preceding chapter gives and example of its use. This section describes Deft and shows how it manifests features of the VTAM architecture.

Deft is oriented toward formulation of decision network models. The user can add and delete nodes and arrows. This allows the problem formulation to evolve smoothly, without gathering clutter of unimportant distinctions. In accord with the decision analytic approach to problem formulation, Deft supports deterministic sensitivity analysis, to aid the analyst in focusing modeling attention. In particular, the program elicits assessments of the values of nodes that are required to specify deterministic sensitivity in complex decision networks, as set out in the next chapter. In this approach, nodes between the uncertain factors and the utility node are treated as deterministic, thus requiring incidental assessments. Deft elicits these as needed. In addition to eliciting direct assessments, Deft allows specification of an assessment function to be used instead of direct assessments. This serves two purposes: it allows the user to avoid repetitive systematic assessments, and it allows Deft to be used as an investigatory harness for a preexisting deterministic model.

Deft allows assessed values to be edited, in case reflection or sensitivity results suggest changes to assessments. It allows assessment sessions to be saved to file and restored later. It provides help text to the user as needed. The current implementation of Deft provides tornado diagram and verbal summaries (as described in the next chapter). A contemplated enhancement is graphical display of the decision network being constructed. Currently this is displayed only textually.

Note that Deft provides explicit support for the evolution of variables, which is central to the decision analytic pursuit of a well-framed problem. It also supports sensitivity analysis at an early stage in problem formulation, to guide this development.

Its central algorithm, which calls for assessments that support sensitivity analysis, is a condition management feature. Its basic output features, tornado diagram and English language summary, constitute the sort of high-level outputs toward which the summarizer function is directed. Its ability to support an assessment function may be used to support a computer model in the utility node. The impetus for building the model with an i/o translator on top is very substantial if the user makes full use of the problem-formulation-editing features of Deft, but it does not explicitly provide any such facility.

Deft has been used as a harness for a version of WSDM, for a probabilistic inference engine, and for operational, strategic and medical decision making.

5.6 Two instances of the VTAM architecture

Here is a tabular summary of the two instances of the VTAM architecture described in this chapter.

VTAM COMPONENT	WESTERN STATES DISPATCH MODEL (WSDM)	DECISION FORMULATION TOOL (DEFT)
Condition Manager	User defines features and values for full factorial analysis in an input file.	User specifies decision network node values for sensitivity analysis interactively, by an assessment function, or by subsequent editing.
Summarizer	numerical table in two-way sensitivity analysis format, verbal summary	tornado diagram, verbal summary
Input/Output Translator	Coder specifies feature-values as variable (or array cell), change-mode, and change-value.	not provided
Model	bulk power market model	many

Chapter 6. Conditional Deterministic Sensitivity Analysis[3]

Up to this point, I have reviewed the problem solving literature, chosen an approach that employs models in the service of decision analysis, illustrated its use, and made a very high-level characterization of the sort of computer system that can contribute to the decision analysis process. An important feature of such a process is that the model's sensitivity is used by the decision-maker to guide her further development of the model. The chapter makes two contributions: it formulates a variant of deterministic sensitivity analysis that is applicable to decision network models, and it focuses on use in the preliminary model-formulation phase of problem solving by ensuring that the assessment demands of its approach are minimal.

6.0. Background

A case reported by Staël von Holstein (1971) apparently employed a procedure to assess the sensitivity of variables taking account of probabilistic dependence, but the details are not given there. Rothenberg, et alia (1990) report an approach that calculates sensitivity of a computer model quickly by modifying each of its subroutines to cache sensitivities as they are calculated, and to reuse these when appropriate. This approach calls on numerical properties of the variables to determine reusability of cached sensitivities, and it generates a local measure of sensitivity, hence it is not directly applicable to our modeling context where variables may not be numeric, and where global sensitivity is more important. Neither of these considers the procedure in relation to decision networks. McNamee and Celona (1987) recommend that the model's sensitivity to scenarios where a set of dependent variables covary be measured. Korsan (1990) gives a procedure that calculates sensitivity of utility to individual variables, their squares, and their products. This procedure calls for assessment of marginal distributions for unconditioned nodes

[3]The basic idea of this chapter, to design a sensitivity analysis approach by reference to what I call structural decomposition, was communicated to me by Dr. Sam Holtzman, who had already been associated with an implementation of ideas along these lines at the Strategic Decisions Group. The development and analysis of the elicitation algorithm given here are my own work. Dr. Holtzman developed an initial variant of the example used here, but I have developed it substantially to make it serve my purposes here.

and assessment of conditional mean functions for nodes with predecessors. None of these articles discusses sensitivity analysis where there are future decisions.

This chapter proposes an extended definition of DSA for use with probabilistic and informational conditioning, and a simple assessment procedure to compute it. Section 6.1 describes DSA. In DSA, a range of values is specified for each variable, and the responses of the decision model to these perturbations are noted. DSA helps a decision-maker understand the effects of uncertainty in a model that is being formulated, to guide its further development. Section 6.2 shows that probabilistic conditioning is a widespread and important feature of decision models, and notes that DSA is not defined for conditioning variables. It sets out a procedure I call structural decomposition, in which auxiliary variables are created whose distributions can be stated unconditionally, and from which the original variables can be reconstructed. I then extend the definition of sensitivity and compare it to value of information and value of control by reference to a structurally decomposed decision network. This extended definition of DSA specifies how to handle variables that are directly or indirectly conditioned on the variable being analyzed. In Section 6.3, I note that structural decomposition leads to pragmatic difficulties, and I propose a simple assessment procedure called Conditional DSA (CDSA) that remedies these and obviates the auxiliary nodes. In CDSA, conditional base cases, as well as a base case and range, are assessed for each node that has conditioning predecessors. I show that the assessment procedure implements the extended definition of DSA. Section 6.4 compares the informativeness and assessment burden of CDSA to those of other proposals in the literature. Finally, Section 6.5 extends the CDSA algorithm for use in two-way sensitivity analysis.

6.1. Decision Analysis and the role of DSA

I begin by describing the conventional calculation and use of Deterministic Sensitivity Analysis (DSA) in decision analysis.

6.1.1. The DA Cycle guides basis development

The function of the Decision Analysis (DA) Cycle (Howard 1983, Figure 1) is to allow confident action by identifying an appropriate *basis*[4] for the decision: the specification of the alternatives, structural information, probabilistic information, and preferences that constitute the decision-maker's understanding of a decision problem. In the DA Cycle, the basis is alternately specified and appraised in light of the results of analysis based upon it. If the basis is found wanting, it is

---

[4]Here and throughout the paper, I italicize terms where I introduce them.

refined and modeling is repeated; otherwise, the recommendation emerging from the probabilistic analysis may be followed with confidence.

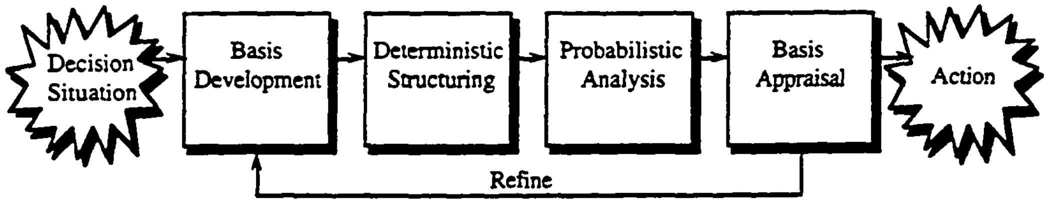

Figure 1. The DA Cycle, as given in Howard (1983)

In sensitivity analysis, the sensitivity of a variable (which I call the *target*) to variations of a given input variable (which I call the *factor*) through a range of values is measured. In DSA, sensitivity is measured for one factor at a time with all other variables fixed. Variables, or decisions, to which the model is insensitive may be fixated at some reasonable value, while the rest are treated probabilistically, or optimized. If a variable shows especial impact, a way may be sought to elaborate on its treatment, to model it more accurately, or to identify new alternatives that can affect it. Note that DSA helps refine a model under development, not a fully assessed model.

Normally the utility function, or a surrogate, is the target, and each chance variable is treated as a factor in DSA. In business problems, where utility is often a roughly linear function of a firm's profits, it may be more helpful to use profits, rather than utility, as the target. In addition, if there are only two alternatives, the difference between profits that would be realized under the two alternatives is often an enlightening target.

6.1.2. How to perform DSA

The previous subsection gives the motivations for performing DSA. This subsection recounts the traditional account of how to do so. Howard (1968) defines DSA as follows:

"The analysis begins by assigning each state variable a nominal value and a range that might correspond to the 10- and 90-percent point on its marginal cumulative probability distribution. Decision variables would also be assigned nominal values and ranges to reflect initial feelings about what the best decision might be. ... With all variables but one set to their nominal values, that one variable would be swept across its range to determine the effect on the value reading."

These nominal values are also known as the *base cases* of the variables. The definition of deterministic sensitivity is, using our terminology,

Definition 1. The deterministic sensitivity of a target to a factor is the set of values induced in the target by sweeping the factor across its range, while holding all other variables fixed at their base case values.

The following procedure for DSA is implied:

Original DSA Assessment Procedure.

1. For each factor,
   2. Specify a range of values for the factor, and
   3. For each range value,
      4. Assess (or compute) the target's value conditioned on the factor taking the range value and all others their base cases, and
   5. Note the highest and lowest target values identified in step 4 for this factor.

Factors and range values may be visited in any order in the two "For" loops. In step 2, the range consists of the base case and values corresponding to roughly the 10th and 90th percentiles of its conditional distribution, without the trouble of a careful assessment. If a factor's values have no intrinsic order, or if its impact on the target is nonmonotonic, it may be helpful to specify more than three values in the range. If it is binary, the two values may be used. If a decision is to be treated as a factor, a range of alternatives that captures the variety of alternatives available should be identified, or if there are not very many, all should be included in the decision's range. The value assessed in step 4 should be roughly the median, or if there is no intrinsic order to the variable, the most likely or reasonable value should be specified. The difference between the high and low values in step 5 is called the *swing* induced by the factor.

Terms like "roughly the median" are used here to suggest the semantics that makes this procedure reasonable, but to explicitly avoid demanding a careful assessment, which would defeat the purpose of this analysis - to ascertain where to perform careful assessments.

6.2. DSA and conditioning

Having given the traditional account of DSA, I note in this section that dependent variables cause trouble for it and develop an approach that works around this limitation.

6.2.1. The importance of conditioning

I begin by highlighting the growing importance of conditioning in decision modeling.

The problem of specifying a joint distribution over variables of interest (as required by decision theory) can be decomposed into simpler problems by conditioning: the marginal distribution of one or more variables is specified unconditionally, and others are specified conditional on these variables. This approach can be made relatively tractable if combined with *discretization* of each continuous distribution, approximation by a discrete distribution. Use of discrete or discretized variables allows distributions conditioned on them to be specified on a casewise basis. Decision trees explicitly represent conditioning of each probability, but this extensive treatment makes revision of the decision basis difficult. The development of decision networks allows probabilistic conditioning among variables to be specified without tracing the conditioning of each probability explicitly. Thus decision networks make it easier to repeatedly revise the set of variables under consideration, as called for in the DA Cycle.

Originally, formulation of a deterministic model was viewed as a crucial step in the DA Cycle. This reflected the modeling paradigm at the time, in which a complicated deterministic model specified the value of an objective function for any set of input variable values and decisions. However, a recognition of the importance of probabilistic modeling for decision-making under uncertainty led to a questioning of this premise. A new account of the DA Cycle, in Holtzman (1985), removes deterministic structuring from the Cycle altogether, reflecting the practicality of formulating problems with uncertainty in a decision network without ever formulating a deterministic model (see Figure 2). This underscores the need for a focusing technique like DSA that can be used with decision networks.

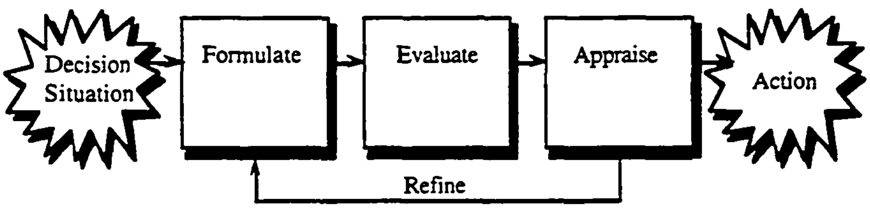

Figure 2. New account of the DA Cycle, as given in Holtzman (1989)

6.2.2. The Licensing Decision

I illustrate the modern approach to problem formulation with reference to a hypothetical Tire Manufacturing Process Licensing Decision. In this example, a tire manufacturer is considering a licensing agreement to use a new process. The process makes tires more biodegradable; hence it may make them more attractive to environmentally sensitive customers. However, this sensitivity could also call upon the firm to undertake more expensive methods to clean up effluents from the process. Demand for driving, and hence for the tires, is thought to depend on the price of gasoline. Land would have to be purchased for the facility to implement the process, hence the availability of land may affect the profitability of the endeavor. The availability of land might in turn be affected by the quantity of plots for sale in the area and by possible changes to local zoning law. Furthermore, there is some uncertainty regarding the cost of raw materials required by the project, and regarding the royalties that would have to be paid to license the process. Maximization of profits is the sole criterion for this decision. These considerations are illustrated in Figure 3.

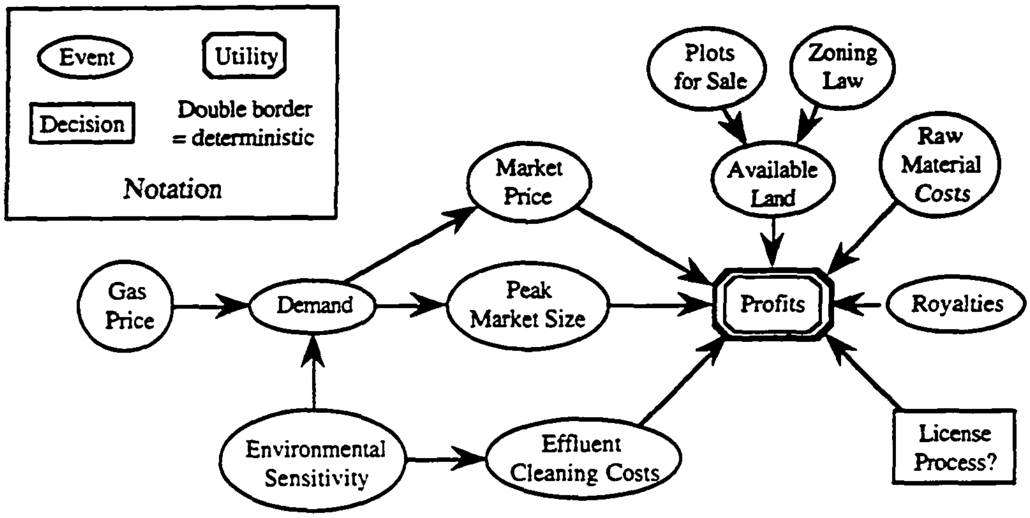

Figure 3. Decision network for the Licensing Decision

To determine whether this diagram shows an appropriate set of variables to consider in detail, the analyst should identify the variables whose variability induces the bulk of the variability of profit. An heuristic procedure commonly used for this purpose, described in McNamee and Celona (1987), is as follows.

The standard deviation of a variable, percentiles of its distribution, and its range values are all measured in units of the variable itself; whereas variance is in units of the variable squared. Assuming that a specified percentile of the factor generates that same percentile of the target's conditional distribution allows us to estimate the relative sizes of the standard deviation attributable to each factor by the relative sizes of their swings. These estimates of relative standard deviations must be squared to estimate relative variances. Further assuming that the variables are independent allows us to estimate the total variance of the target as the sum of the relative variances from the factors. Table 1 tabulates swing for as many variables as possible in this example, and shows the calculation of their contribution to the variance of Profit under these assumptions. If we felt the unanalyzed variables contributed nothing to the variance, we would fixate all variables except Peak

Market Size and Average Market Price, because these two apparently contribute over 90% of the variance.

	----FACTOR----			----TARGET----				CONTRIB.	% OF
FACTOR NAME	LOW	MED	HI	LOW	MED	HI	SWING	TO VAR.	VAR.
Available Land	50	120	280	-30	125	135	165	27225	.03
Demand	Small	Med.	Large	?	125	?	?		
Environmental Sens.	Low	Med.	High	?	125	?	?		
Effluent Cleaning Costs	4	11	13	145	125	100	45	225	.00
Gas Price	75	150	225	?	125	?	?		
Market Price	34	83	115	-130	125	270	400	160000	.20
Peak Market Size	1400	1500	1550	-320	125	460	780	608400	.76
Plots for Sale	Small	Large	All	?	125	?	?		
Raw Material Costs	1080	1320	1530	127	125	122	5	25	.00
Royalties	37	43	48	155	125	95	60	3600	.00
Zoning Law	Mod't	None	Extr.	?	125	?	?		
Total								799475	1.00

Table 1. Sensitivity of Profits to various factors

The traditional format for displaying DSA results is a *tornado diagram.* In it, the set of target values induced by the range values of each factor are plotted and connected into a horizontal bar, and the bars are arranged in descending order of swing. It may be noted that the target value corresponding to all the variables being at their base case will be in all of these bars, and that unimportant variables will have a very narrow bar about this central target value. Figure 4 shows a partial tornado diagram for the example.

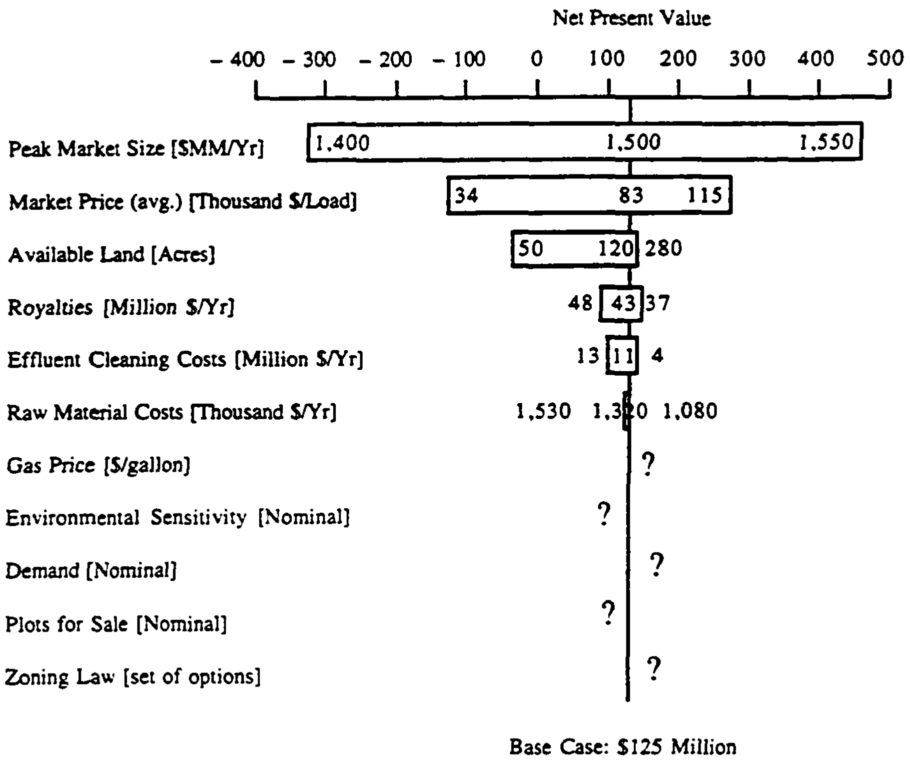

Figure 4. Tornado Diagram for the Licensing Decision

The most obvious feature of this analysis is the undefined sensitivity of a number of variables. If we interpret Howard's definition strictly, the model's sensitivity to these variables is zero, because all other variables, including all those that determine profit, are held fixed in the analysis of their impact. But this clashes with our notion of what DSA should determine. Consideration of the sensitivity data we have suggests that it may be worthwhile to build a careful model of the two most sensitive variables. In fact, our preliminary decision network (Figure 3) does so, but is it the right one? It could be that these variables are largely determined by Demand, or even by Gas Price and Environmental Sensitivity. We would like some way to investigate these precursor nodes with DSA.

6.2.3. Defining DSA with structural conditioning

This difficulty arises frequently in decision analysis. Analysis of the sensitivity of a target variable to a factor with all others held fixed is not meaningful if we expect a change in the factor to be accompanied by substantial changes of other variables. Thus DSA, in its simplest description, does not represent cases where one wants to assess nodes conditionally. However, if we consider that the Profits node of traditional models was often a complicated computer model with many incidental variables, we realize that these incidental variables are not held fixed in DSA in practice. They are allowed to vary according to their nature. This motivates the following extension of the definition of DSA:

Definition 2. The deterministic sensitivity of a target to a factor is the set of values induced in the target by sweeping the factor across its range, allowing its deterministic successor nodes to respond, and holding all other variables fixed at their base case values.

6.2.4. Changing probabilistic conditioning to structural conditioning

The previous subsection extends the definition of DSA to handle structural, but not probabilistic, conditioning. In this subsection, I work out an approach to problem reformulation that transforms a decision network node with probabilistic conditioning into an equivalent network with only structural conditioning, thus allowing Definition 2 of DSA to be employed. I call the process structural decomposition.

Consider an earlier stage of this analysis. Upon initial consideration of the issue, it is felt that the level of demand for tires is an important determinant of the profits from licensing the process. If the amount of profit is taken to be completely determined by demand and by whether the process is licensed, the decision network exhibits only structural conditioning. However, if it is felt that demand depends on the price of gasoline, probabilistic conditioning is introduced into the decision network, as shown in Figure 5.

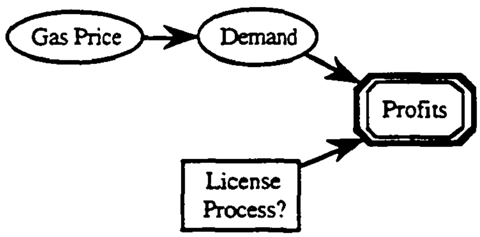

Figure 5. Decision network for Licensing Decision with conditional Demand

If we can reformulate the problem to replace its probabilistic conditioning with structural conditioning, we can apply Definition 2 of DSA to all chance nodes. This is what is normally attempted in systems analysis. In my simplified example, variables that "explain the remaining variance" in Demand would be sought. For instance, it may be thought that the environmental sensitivity of consumers would affect the level of demand. If this were thought to explain all the variance of Demand (i.e., if we felt we could assess it deterministically), the latter could be considered an incidental variable in the calculation of Profits, yielding a star-shaped decision network exhibiting only structural conditioning, as illustrated in Figure 6.

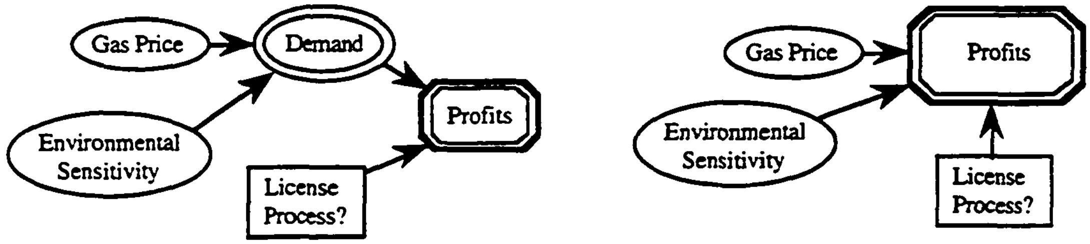

Figure 6. Fortuitous reformulation to eliminate probabilistic conditioning

This reformulation allows the sensitivity of profits to Gas Price and Environmental Sensitivity to be measured. But there is no reason to believe that variables allowing this sort of reformulation can always be found, or that they would necessarily explain all the remaining variance in a node such as Demand.

By making reference to the formal contents of decision network nodes, I derive an approach that always allows this sort of reformulation. Figure 7 shows possible contents of a Demand node conditioned solely on Gas Price, according to the view set out in Smith et alia (1993): each node contains a *distribution tree* consisting of a conditioning tree with an atomic distribution at the end of each of its branches. A *conditioning tree* is like a decision tree, listing all possible outcome of each conditioning variable, in turn. (This is possible only for discrete-valued nodes, to which I confine my attention.) It differs in that each juncture is marked with a simple dot; no distinction is made between event nodes or decisions in a conditioning tree. An *atomic distribution* is a probability distribution over all possible outcomes of the node under the given conditions; its juncture is designated by a circle.

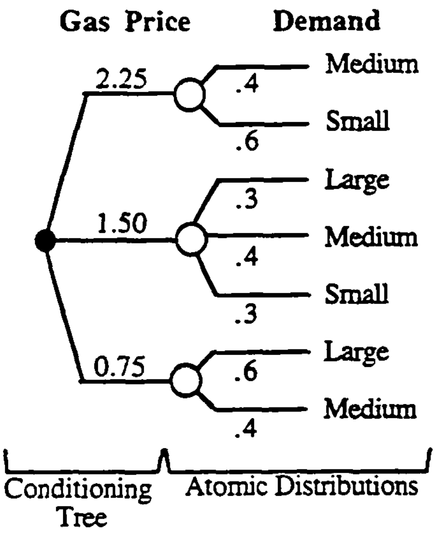

Figure 7. Distribution Tree for Demand

Any node may be *structurally decomposed* by replacing it with a deterministic *log* node that handles the logic of its conditioning tree and one *dist* node for each atomic distribution in the node. All the dist nodes, as well as the original predecessors, become predecessors of the log node. As suggested by the conditioning tree notation, it is immaterial for purposes of structural decomposition whether a predecessor is a decision or event node. Note that this process requires that the values into which Gas Price is discretized be specified to structurally decompose the Demand node, because they are used in the definition of the new dist nodes. Figure 8 shows two decision networks that represent the same state of information as Figures 5 and 7 by structurally decomposing the Demand node. On the left, the log node Demand is now deterministic; it simply chooses the value of the appropriate dist node predecessor, based on the value of Gas Price. On the right, it is considered incidental to calculation of profits. The probabilities that were in the Demand node before are now present only in the dist nodes.

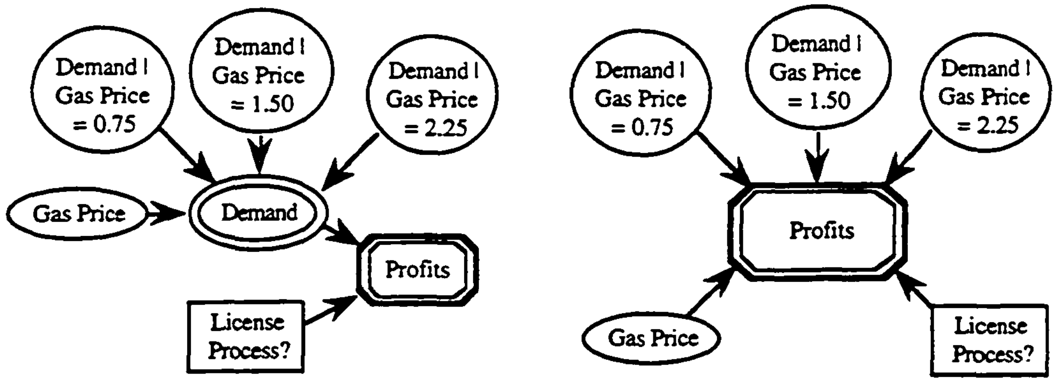

Figure 8. Structurally decomposing Demand, and an equivalent decision network

Having created the dist nodes, the analyst is obliged to consider whether they should be specified conditional upon one another, to maintain the integrity of the decision network. If this creates additional probabilistic conditioning, the dist nodes themselves must be structurally decomposed. Since I confine my attention to variables that take a finite number of discrete values, the number of nontrivial dist nodes that may be created in this way is bounded by the number of degrees of freedom in the joint distribution; hence this process must eventually terminate.

Even when informational predecessors of future decisions prevent formation of a star-shaped decision network like the one shown here, structural decomposition creates a new decision network in which the effects of uncertainty can be seen clearly, because the elementary sources of uncertainty and the mechanism of its effects are isolated for examination in the dist and log nodes, respectively.

6.2.5. VOI, VOC and sensitivity under structural decomposition

Having shown how to structurally decompose a decision network, I now show that structural decomposition offers a clear way to think about value of information and value of control, and I show the relationship of DSA to these measures.

The Value of Information (VOI) of a chance variable for a given decision is defined as the most one should be willing to pay an omniscient clairvoyant to reveal its value before the decision. As is shown in Howard (1990), there must be no arrow from the decision node to a chance node (an *influence* arrow) if its VOI is to be calculated from the decision network. He defines what has come to be called Howard Canonical Form, in which there are no influence arrows. Matheson (1990) describes a process whereby any decision network may be put into Howard Canonical

Form: each chance node bearing an influence arrow from a decision is replaced by a constellation of chance nodes without influences, one for each alternative in the decision node, together with a deterministic node that bears the influence arrow, as well as incoming arrows from each of the new nodes. This allows the addition of informational arrows from the new nodes to the decision (as is required in the calculation of the joint VOI of these nodes) without creating a directed cycle in the diagram. The VOI of an influenced node is the joint VOI of the new nodes in its Howard Canonical Form constellation. The analogy between this constellation of nodes and the dist nodes of a structurally decomposed decision network allows us to note that the Howard Canonical Form process amounts to structurally decomposing all chance nodes, but only with respect to decision predecessors, not to chance node predecessors. Hence every structurally decomposed decision network is in Howard Canonical Form.

The Value of Control (VOC) of a chance node is the most one should be willing to pay to an omnipotent wizard to specify its value (or its relationship to its predecessors). The motivation of the concept is to determine whether alternatives should be sought that could change the event in question. Matheson (1990) describes the calculation of VOC of a node with no predecessors and states that further analytic machinery is required to define the VOC of a node with predecessors. Structural decomposition provides such machinery. The VOC of a conditioned node can be defined as the joint VOC of its dist nodes in the structurally decomposed decision network, allowing the mechanism of the downstream log nodes to be realized.

These comments may be summarized by reference to Figure 9. In this network, E has been structurally decomposed into E' etc. Value of information is the maximal price one would pay for the arrows from E' etc. to D. Value of Control is the maximal price one would pay for arrows from D' etc., which specify how one would like them to come out, to E' etc. (Arrows among the decisions or from the decisions to the utility node U are suppressed to avoid clutter.)

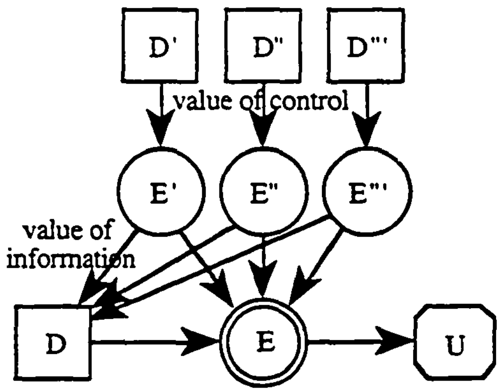

Figure 9. Decision network for calculating value of information and value of control

The sensitivity of a target to a factor is the amount of change induced in the target by changes to the factor and corresponding changes to its deterministic descendants, but not to its ancestors. Thus the sensitivity of a profit node to a factor can be viewed as conditional VOC: the VOC of that node, conditioned on all others being held at their base case values.

It has been argued that VOI and VOC are "better" measures of sensitivity than DSA, because they remove the limitations imposed by this conditioning. Although this is true, DSA is still the appropriate tool for use to guide early stages of problem formulation in the DA Cycle, because this "limiting" conditioning also substantially reduces the amount of assessment required. Just how far this assessment burden can be reduced will be discussed later in this chapter. First I must extend DSA to apply to all decision network models.

6.2.6. Defining DSA with probabilistic and informational conditioning

This subsection extends the definition of DSA to cases with probabilistic and informational conditioning.

The intent of DSA is to measure the effect of the intrinsic variability of a node, i.e., the uncertainty about that variable that would remain even if its predecessors were known, under base case conditions. The dist node of a factor node corresponding to all others at their base case represents this intrinsic variability, and the log node descendants of the factor specify the natural effect of that variability. Accordingly, I extend the definition of DSA as follows:

Definition 3. The deterministic sensitivity of a target to a factor node in a decision network is the set of values induced in the target by structurally decomposing the factor's event node descendants, sweeping the factor's base case dist node across its range, allowing its log node successors to respond, and holding all other variables fixed at their base case values.

Definition 2 handles probabilistic conditioning, and Definition 3 handles probabilistic conditioning, but neither definition specifies how informational conditioning is to be handled. Three approaches suggest themselves: one overall base choice could be specified for a future decision regardless of conditioning, a reasonable choice for each condition could be assessed, or the decision could be deterministically optimized for each condition. The former, unconditional choice, corresponds to open loop sensitivity analysis, as described in Howard (1971). This is what was intended in Howard's 1968 definition of DSA. The latter, conditional optimization, corresponds roughly to closed loop sensitivity analysis in Howard (1971), and is also most consistent with the notion that DSA is conditional VOC. However, it requires assessing variable values *corresponding to each* alternative for all successors of the decision. The intermediate approach, conditional assessment, captures the ability of future decisions to compensate or respond to events, but without imposing this extra assessment burden. Hence it seems most consistent with the spirit of DSA, which is to reflect a reasonable judgement of the effect of uncertainty, but without a combinatorial proliferation of assessments. Accordingly I reach a final definition of DSA:

Definition 4. The deterministic sensitivity of a target to a factor node in a decision network is the set of values induced in the target by structurally decomposing the factor's event node descendants, sweeping the factor's base case dist node across its range, allowing its event node descendants to respond, making a reasonable choice for each future decision considering its information state, and holding all other variables fixed at their base case values.

Note that models in the traditional paradigm are specified without conditioning, so each factor serves as its own (sole) dist node, and no node has successors. Hence this definition contains the traditional notion of DSA as a special case.

6.3. Simplified DSA assessment procedure

In this section, I identify some disadvantages of structural decomposition and propose an assessment procedure that remedies them by assessing a miniature scenario for each range value under consideration.

6.3.1. Disadvantages of structural decomposition

Structural decomposition requires careful analysis of the problem merely to define nodes: both the directions of arrows and the discretization of conditioning nodes are wired into the definition of dist nodes. Discretization is generally considered to require a time-consuming full assessment of the conditioning node's distribution. This makes the frequent modifications of the decision network required in the DA Cycle difficult.

The number of assessments required in structural decomposition is large. The number of dist nodes one must assess in a structurally decomposed decision network, although limited by the degrees of freedom in the joint distribution to be assessed, can still be large. In addition, many of these nodes are of no intrinsic interest, insofar as their definitions refer to discretizations or variables that may be subsequently modified or removed from the decision network. Furthermore, many assessments required for these nodes are not needed for DSA. This point will be illustrated in Section 4.

6.3.2. CDSA assessment procedure

In this subsection, I give an assessment procedure that implements my extended definition of DSA, but without formulating unnecessary nodes or assessing unnecessary values. In this procedure, one specifies conditional ranges for factor nodes and Conditional Base Cases (CBCs) for their descendants in the original decision network.

CDSA Assessment Procedure:

1. For each factor under consideration,
    2. Specify a range of values for the factor, and
    3. For each value in the factor's range,
        4. For each of the factor's descendants (including the target) in order,
            5. Assess the descendant's CBC in light of the factor's value and any relevant CBCs upstream, with all other variables at their base case, and
    6. Note the highest and lowest target values identified in step 5 for this factor.

Factors and range values may be visited in any order in steps 1 and 3. In step 2, the range that would have been assessed for the factor's base case dist node is employed as the range of the factor. In step 4, the factor's descendants may be visited in any order that honors arrow direction. Such an ordering is guaranteed to exist for acyclic networks, and it is not hard to identify. In step 5, each *Conditional Base Case (CBC)* for event nodes should be assessed as the base case value of the dist node whose definition refers to the factor and any relevant upstream nodes taking the values currently under consideration. The *Conditional Base Choice (also abbreviated CBC)* for future decisions is chosen in a reasonable fashion for the state of information currently under consideration.

Under this procedure, only descendants of the factor being analyzed are assigned some value other than their base case, just as in DSA in a structurally decomposed decision network. In step 5, these nodes are assigned the value that the corresponding log node would have been taken in a structurally decomposed network. Accordingly, it identifies the same values of the target variable, and step 6 identifies the sensitivity as defined in Definition 4.

6.4. Relationship to other work

This section compares the number of assessments required and the amount of information generated by CDSA, structural decomposition, and other approaches to sensitivity analysis in the literature, by reference to a section of my example decision network.

Consider a small network with only the three uncertain nodes considered before: Gas Price, Environmental Sensitivity, and Demand (Figure 10), and assume that Profits can be calculated from Demand without additional assessments. The comparison would be more tedious but no more enlightening if the whole network were considered.

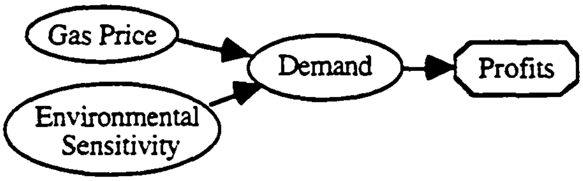

Figure 10. Fragment of the Licensing Decision network

In traditional DSA, only three assessments (of Demand) would be required, but the sensitivity of Profits to Gas Price and Environmental Sensitivity could not be measured.

McNamee and Celona (1987) require the analyst to specify joint variation of the variables. This is difficult here because it requires us to assert that Gas Price and Environmental Sensitivity, as well as Demand, covary. The intention expressed in the decision network is that Gas Price and Environmental Sensitivity are marginally independent. So nine assessments, one for each of the three variables under low, medium, and high conditions of the co-variation, would be required to measure the model's sensitivity to the co-variation, and these assessments would require specification of some of co-variation of nodes that is not intended. Presumably two other assessments for Demand could be employed to measure its impact.

Korsan (1990) requires full marginal distributions to be assessed for Environmental Sensitivity and Gas Price; and requires a conditional mean function for Demand. For purposes of comparison, I assume that ten assessments are sufficient to map out a marginal distribution. So Korsan's approach requires 20 assessments for the two marginals. Korsan does not discuss assessment of conditional mean functions, commenting only that assessment of a mean function conditioned on only one variable is of equal difficulty to assessment of a marginal. For our case with two predecessors, this suggests that ten slices with ten assessments each, i.e., a grid of 100 points, must be assessed. This large number invites speculation to flesh out Korsan's method: perhaps a less dense grid would suffice, and in some cases it may be possible to identify a class of functions to which the conditional mean function belongs, and to assess its parameters. Here I simply report the total of 120 assessments as an upper bound of the number of assessments called for by Korsan's method.

Under structural decomposition, the two predecessors would be discretized to three levels each, generating nine new dist node predecessors of Demand. Each of these eleven nodes would require three assessments, and Demand would require two assessments for each of them, plus one for them all at their base case, for a total of 56 assessments.

Under CDSA, the base case scenario requires one assessment for each node -- 3 assessments. Assessing the impact of Demand requires assessments of two range values, high and low. Assessing the impact of Gas Price requires assessments of high and low values for it, and corresponding values for Demand -- 4 assessments. Similarly the assessments for Environmental Sensitivity call for four assessments. Put another way, this is three assessments each of Gas Price and Environmental Sensitivity, and seven for Demand (base case, four CBCs, and two range values). The total is 13 assessments.

Rothenberg, et alia (1990) would call for evaluations of subroutines that correspond to the assessments required by CDSA if its focus is widened to the global measure of sensitivity pursued here, and if its notion of when cached sensitivities can be reused is confined to equality of all inputs, thus allowing for non-numeric variables.

METHOD	ASSESSMENTS	NUMBER OF NODES WHOSE IMPACT IS MEASURED
Howard 1968	3	1
McNamee and Celona 1987	3*3 +2 = 11	2 (Demand, and the co-variation)
Korsan 1990	2*10 +10*10 = 120 (upper bound)	3+ (the three nodes' impacts are reported, and first- and second-order effects are distinguished)
Structural Decomposition	11*3 +11*2 +1 = 56	3
CDSA	3 +2 +4 +4 = 13	3

Table 2. Comparison of sensitivity analysis methods

6.5. Extension to two-way sensitivity analysis

This algorithm can be extended for sensitivity analysis where two factors vary in concert. A common use of this kind of procedure is for decision analyses, where it is frequently of interest to know the variation of one's utility function under variation of an uncertain variable and different choices of alternatives. The algorithm can be stated for any pair of factors, regardless of node type.

Two-way Sensitivity Analysis Assessment Procedure:

1. For each node,
    2. Reset the context
    3. If it is a factor, then
        4. Specify its range of values
        5. For each value in the factor's range
            6. For each subsequent node,
                7. If it is the other factor, then
                    8. Specify its conditional range
                    9. For each value in the second factor's range,
                        10. For each subsequent node
                            11. Specify its CBC
                        12. Note the value of the target for this pair of factor values
                13. Else specify its CBC
    14. Else specify its CBC

Steps in preceding algorithms make reference to assessments in light of previously assessed values. This logic is captured here by use of a construct called "context". When reset, it contains no variable-value bindings, and as each CBC is specified and each range value is instantiated, its value is noted in context. For simplicity, an ordering of the nodes consistent with arrow directions can be identified ahead of time, and the iterations under steps 1, 6, and 10 should be performed according to this order. The conditional range in step 8 is assessed similarly to a normal range, with the specification that the approximate 10th and 90th percentiles are to be judged in light of the context.

6.6. Chapter Summary

In a structurally decomposed decision network, each conditioned node is replaced by one log node that reconstructs its behavior and a set of dist nodes that can, hopefully, be specified unconditionally. In such a network, the sensitivity of a node is the sensitivity of its base case dist node, and its values of information and control are the joint value of information or control of its dist nodes. In my Conditional DSA (CDSA) assessment procedure, Conditional Base Cases (CBCs), as well as a base case and range, are assessed for each node that has conditioning predecessors in the decision network. By employing these CBCs to "transmit the influence" of predecessor nodes, CDSA replicates the sensitivity analyses that would result from structural decomposition: Sensitivity to uncertain variables reflects their direct effect, their indirect effects as transmitted by downstream variables, and the mitigation of their effect by subsequent decisions. Sensitivity to decisions reflects their direct effect and their indirect effects as enhanced by subsequent decisions or transmitted through downstream variables influenced by them. Unlike previous work, CDSA is able to measure the sensitivity of a model's results to all its variables, and its assessment burden is light. The procedure can be extended in a natural way for assessment of two-way sensitivity analysis.

Chapter 7. Verbal Summary

This chapter gives an account for how the DA Cycle can help identify the best action, discusses what kind of summary will contribute to this, and shows how to generate such a summary automatically. Its fundamental premise is that a summary organized in the way human judgements are formed will be more easily comprehensible than other kinds, and that this will support more well focused criticism and elaboration by the decision maker and domain experts.

7.1 Literature

This section gives a brief summary of the salient conclusions from the empirical psychological literature reviewed in chapter 2. These conclusions will be explored in more depth in the body of this chapter.

Boundedness of human comprehension, while obvious, will nonetheless be a crucial finding for this work. Decision analyses tend to be complex, but people cannot comprehend complex things all at once. To engender belief in a recommendation, one must summarize it briefly.

The tendency to form judgements as adjustments from a base case suggests that this may form the basis of a clear way to present results of an analysis to a decision maker. There is some empirical support to the existence of linear additive mental models in some contexts, but these results conflict with the more widespread finding of conservatism - adjustment of a judgemental response that is quantitatively insufficient. While the conservatism results are of concern here if the presentation is numeric, these concerns are not present if only qualitative adjustments are presented, insofar as people handle these properly. However, quantitative information has the advantage of allowing us to simplify accounts of decision analyses by making arguments that the impact of a given factor is with certain bounds, and hence is small enough to be ignored.

The ease with which people conceive of causal scenarios in simulations suggests that enumeration of scenarios might also be a good way to summarize an analysis. This is roughly consistent with a decision analysis, which may be viewed as a process that identifies a probabilistic simulation model of the situation at hand and finds the optimal decision or control policy for that model.

Another possible use of entity-and-operator/decomposition thought paradigm is to treat the entire analytic process as a search for an appropriate knowledge state for the decision maker, and to formulate comments after each stage of the analysis that can lead to improvement of that state.

Recognition of a circumstance and retrieval of a pre-formed response is not consistent with the analysis of a unique high-stakes decision. I affirm recognition-based thought for unimportant or routine problems, or for decisions that must be made very quickly, but I offer it no direct support in my work, which is directed toward deliberative decision-making.

7.2 How the DA Cycle works

This section sets out and defends my account of the way the DA Cycle contributes to knowledge of the best choice.

7.2.1 The divide/compare/improve process

I propose that a process I refer to as "divide/compare/improve" can be used to rationally reconstruct successful operation of the DA Cycle process. This account combines both of the approaches to improvement of knowledge in the literature: generate and compare multiple viewpoints, and iteratively improve upon viewpoints found to be wanting. The viewpoints involved are the formal decision model and the DM's direct judgements. We divide these two accounts into parallel components, consider corresponding sub-accounts, and identify the superior one. Finally, we improve the deficient sub-account by replacing it with (or making it more like) the superior one. In a successful decomposition, improvement to one component will not degrade any other component.

When the two viewpoints' judgements regarding each of the factors agree, the DM's judgements constitute a high-level summary of a rigorous DT equivalence argument, thus giving her justified belief (knowledge) that her judgements indicate the best action. In addition, the now-unified viewpoint will be seen to have been repetitively improved, strengthening the DM's assent.

Two topics, the principle underlying this analysis[5] into parallel components, and the nature of "the model's outputs" being examined, are discussed in this chapter. For now, let us say that the logic inherent in a decision theoretic optimality argument is only one principle by which accounts could

[5]While the word 'analysis' normally refers to any systematic process leading to new conclusions, I shall use it more narrowly, to refer to only such processes that actually break the subject matter down into parts.

be analyzed, and that the value of the utility node is a credible operationalization of "important model outputs".

7.2.2 Divide/compare/improve is consistent with practitioners' accounts of DA

In this subsection, I argue that this "divide/compare/improve account" of the DA Cycle is consistent with existing accounts of it. I begin by reviewing the role of deterministic sensitivity analysis (DSA) in the DA Cycle, as this is frequently cited as the technical source of the DA Cycle's results.

Sensitivity analysis is said to contribute to five goals in the DA Cycle:

- identifying portions of the decision model that are broken or in need of elaboration,
- identifying portions of the decision-maker's mental model that are incorrect or inadequate,
- identifying decisions or variables whose impact is small enough that they can be treated deterministically (be "fixated"),
- eliminating impotent alternatives, and
- creating insight.

The first four items focus attention toward important variables and away from unimportant ones; they can be viewed as supplying a decomposition principle for analytic criticism and improvement. This decomposition principle is to view a decision model as the aggregation of effects of a set of important factors. The DA Cycle is characterized as a dialog (i.e., a structured encounter of two points of view) in Holtzman (1985) and Thomas and Samson (1986). This characterization is consistent with the second phase of my account: critical comparison of the two accounts.

7.2.3 Application to organizations

The work underlying this dissertation is oriented toward a single decision-maker. The psychological mechanisms I describe may be able to work if delegation is accompanied by trust. If this is to take place, DM should request a summary of an analysis at a high level such that she can evaluate the general reasonableness of each of its components, and her experts can verify the specific analyses supporting them. The DM's belief of the overall recommendation's optimality is as strong as its informational components, so if each of these is endorsed by a trusted subordinate with expertise in the area, the DM can come to believe the recommendation. This claim that organizations can make use of the same mechanisms as individuals may need further defense than can be offered in a work of this scope, e.g., it may require designing a new organizational structure and arguing that it can be implemented. This dissertation must be viewed alongside the rest of the literature on individual rationality, as being a subject of intrinsic interest to individuals and consensual groups, and a possibly worthwhile long-range target for heterogeneous organizations.

7.2.4 Value of the divide/compare/improve account

The arguments in this section are intended to set out my basic account for how the DA Cycle can create knowledge and give the reader some reason for optimism that a useful tool can be built upon this framework. The remainder of the chapter identifies an appropriate analytic principle for the "divide" phase. The ultimate purpose is to generate the verbal summary capability of the Deft software described in chapter 4.

Two caveats should be noted at this point. First, my account does not explicitly specify how the critical comparison and choice of a superior account is to be accomplished. However, I do take first steps in this direction, by suggesting that this problem can be approached by decomposition. Second, the divide/compare/improve account leaves the possibility that a modeling result judged superior to intuitions will be opaque; that its technical merit will be unarguable, but that the DM will not achieve an intuitive grasp of the point. In certain circumstances, the decomposition approach can solve this problem, but in others, it merely pushes the problem to a lower level of detail. I argue that this is as it should be. No DM can understand everything, hence DMs should make sure they understand at least the high-level issues and make do, if needed, with an "opaque" understanding of some details.

7.3 Content of the summary

A basic premise of this dissertation is that a perspicuous summary of a decision model being considered can draw the decision maker's attention to facets of it that call for revision or elaboration. The divide/compare/improve account can, if elaborated properly, serve as the basis for designing an approach to summarizing that performs this focusing. This subsection develops desiderata for the analytic principle of divide/compare/improve, and examines a few candidate principles in its light.

7.3.1 Desiderata

Desiderata for the analyzing principle may be drawn from the purposes it must serve: It must serve to divide both model-based analyses and direct judgements; these divided accounts must be compared; one or the other must be improved; and the DM must come to believe that her understanding of the situation identifies the best option. This gives rise to seven desiderata:

- Decomposition of natural thought in this way is easy.
- Decomposition of model results in this way is straightforward.
- Judging which of two components is better is possible.
- Changes to natural thought decomposed this way are easy to identify.
- Improvement of models decomposed this way is straightforward.
- These changes to one's direct judgements must maintain their believability.
- Natural thought divided thus should indicate which alternative is best.

It should be noted that the sixth desideratum stops short of saying that division according to the analytic principle fully represents one account or the other; it merely states that it represents enough of the formal account to command assent.

7.3.2 Candidates examined

I now turn my attention to consideration of the following principles of decomposition: a list of important distinctions, a set of DA basis elements, events in a simulation, and base outcome and adjustments in response to important factors. The latter will be considered in three variants: qualitative impacts, quantitative impacts, and conditional specification of impacts. The merits of each candidate will be discussed in turn, in light of the desiderata set out above.

7.3.2.1 Important Distinction List

The first candidate principle of decomposition for divide/compare/improve to be considered here is to list the important distinction's employed in the two viewpoints. As noted above, many authors think it's important to identify the right distinctions when solving problems or making decisions. A request to specify possibly relevant factors normally can elicit many responses from a decision-maker. Identification of important distinctions in models is straightforward: just do DSA on all distinctions and list the biggest ones. However, if some difference exists, there is often no clear way to judge which set of distinctions is better without deeper analysis. The purpose of the decomposition is to facilitate this comparison, so this must be seen as a significant disadvantage. Learning that a new distinction is important is a common sort of learning. It is not easy, but it is often done when learning a new field. Assimilating the fact that a variable is unimportant is easy. Processing such information is the basis of the phenomenon studied by Johnson-Laird (1980), that people are good at drawing conclusions based on the insensitivity of outcomes to certain factors. Adding a variable to a model is not terribly hard for most modeling systems, and fixating a variable is even easier. Mere opaque belief that a new computer-identified distinction is important is easy to come by. Mere identification of important distinctions does not indicate which action is best. In sum, there is no clear relationship between lists of important distinctions in an analysis and knowing the right action; this makes it difficult to critically evaluate distinctions, and leaves the DM without a recommendation, even if critical comparison and improvement is achieved.

7.3.2.2 Basis elements and DT argument

The next analytic principle I examine is "basis elements". The basis of a decision analysis is the alternatives, information, and preferences that the DM specifies. Together, these suffice to produce a decision theoretic recommendation of action. The elements of the basis are individual alternatives, and probabilities and utilities of individual events. This constitutes an elaboration of the distinction-list analytic principle, insofar as distinctions are tagged as volitional (alternatives), preferential (preferences), or neither (information).

Although the decision basis constitutes an excellent organizational principle for decision theory, it is not ideal as a principle for critical comparison of such an account to the DM's direct judgements. First, it is hard to break direct judgements down directly into basis elements. Decision analysis has developed elicitation procedures to achieve this end, but, as has been noted elsewhere, these are relatively expensive and time-consuming. Of course, a DT model is already specified in terms of basis elements, so this "translation" is trivially possible. If the two accounts are decomposed into basis elements, comparison of one basis element to another would not be prohibitive, and it is easy enough to add a new variable to a DT model. Insights from such a model may not, however, be readily assimilable, due to the extensive logic required to connect basis elements to an optimal action, hence they may not command assent from the DM. If believed, they would, of course, direct action. In sum, the essential problem of basis elements as an analytic principle is that it is too remote from the nature direct human judgements, making translation to and from them too difficult to support confident acceptance of a recommendation.

7.3.2.3 Simulation results

Enumerating the sequence of events in a projection/planning model can be a useful way of summarizing it, and one hears of "walking the DM through the scenario" as a way of achieving her buy-in. This is reasonable in light of Kahneman and Tversky's (1982) finding that people answer questions about events by "running a mental simulation model" and noting the ease with which variations of initial conditions from their default values produce different outcomes. Hence it will probably not be too difficult to encode the DM's judgements as a simulation by asking a sequence of questions like "And then what happens?". A simulation model is a special case (deterministic) of a decision network model, so making formal models along these lines is not difficult. If two projections from a set of conditions differ, it is not too difficult to identify the first step where they vary and decide which account of that step makes more sense. Improving a simulation model, should this prove necessary, is not that difficult. However, it may be difficult to update one's judgement in response to a finding that a more detailed simulation is more appropriate due to the cognitive difficulty of assimilating a large body of simulation results. However, if this is accomplished, it seems that it would command assent, and that it would indicate the best action to choose. This simulation-based analytic principle seems quite attractive, failing only insofar as the model's simulation is too extensive to be assimilated, especially if uncertainty is treated by creation of many possible simulation paths.

7.3.2.4 Base value and responses to important factors

The next analytic principle we will consider is to view both accounts as specifying a "base case" and responses to important factors. In light of Tversky and Kahneman's (1974, 1979) findings that people think using anchor and adjustments, it is natural to believe that elicitation of base case and sensitivities will be tractable. In limited experience with Deft, though not trivial, this typically can be accomplished in a couple of hours by an experienced analyst by asking, first, what is the most representative outcome of all the variables, and then, how much the payoff measure would vary if any one variable were varied. In models that can be represented as a star-shaped decision network, such questions are fairly easy to answer, because the inputs are independent - they do not interact. In the more general case, interactions of variables must be taken account to create a consistent scenario for the specification of the response of the model to variation of a variable. Chapter 6 specifies a consistent way to accomplish this, and chapter 4 describes a computer program that automates this elicitation process. The same logic may be used to perform DSA on an existing simulation model.

If a model's sensitivities are judged to be superior to the directly judged adjustments, it is easy to identify the required change to one's judgements -- just believe a variable has a different level (or direction) of impact. In my experience, this adaptation is frequently made, both by modelers and nontechnical personnel in existing practice. For instance, in the PNW energy business, it was easy for nontechnical managers to assimilate a remark like "Bad hydro increases our subsidy.", even if they didn't know how or why the quality of hydro was related to our subsidy. Similarly, they had little trouble assimilating and believing qualitative comparative statements like "Bad hydro increases our subsidy more than Trojan decommissioning.". I call this kind of understanding opaque understanding to acknowledge that it is not clear -- clear knowledge could be explained or defended in detail, while opaque knowledge is known only by faith that its source that has proved reliable when examined in other circumstances. While it is fairly easy for opaque knowledge of sensitivity to be accepted and believed, it seems somewhat more difficult for nontechnical domain experts to assimilate unexpected base cases. For instance, it was my experience in the PNW energy business that one particularly adept lawyer, named Marcus Wood, had a more useful understanding of policy-oriented models than many technical personnel, because he made a point to understand their sensitivities. Often mere opaque qualitative understanding proved quite useful to him.

If a model shows inappropriate sensitivity to a variable, one must investigate why the model gives this response. This can be done by tracing through the decision network to see where unanticipated response creeps in (after the fashion of Suermondt 1992), analyzing sensitivity of intermediate variables, if needed. This task is not trivial, but localizing the problem makes it more manageable, and the decomposition of the overall critical comparison into lower-level comparisons is a step closer to the granularity where the domain expert chose to assess the variables. The task of analyzing and fixing problems in a model's point forecast, also while nontrivial, may be assumed to be soluble, since deterministic modeling has been the dominant mode of analysis for many years in many industries.

Knowing the base case outcome and the sensitivity of the outcome to uncertain inputs often suggests a course of action to the decision-maker, especially if no change to a single variable can change which alternative is best.

7.3.2.5 Variants of base case and adjustment

Here I consider variant approaches to summarizing base case and sensitivity.

Response of which target

I have suggested that the target variable for sensitivity analysis, i.e., the variable that constitutes the crucial output of the model, should be utility, arguing that being able to characterize its behavior gives good insight into the behavior of the model. However, DSA may be performed on any deterministic state variable. Another target typically investigated is profit, which frequently bears a monotone near-linear relationship to utility. Other variables, especially variables that constitute the essential result of a substantial subsector of the model, can also profitably be explored.

Decompositional v. holistic explanation

If a decision problem is represented by a decision network or systems dynamics model, a natural propagation of base cases and sensitivities from node to node can be identified. Nodewise characterization of these BC&A's is thus a conceivable variant of BC&A. An approach similar to this is discussed by Suermondt (1992). In the early stages of model development where many nodes are under consideration, this can create a very large amount of information, making assimilation difficult.

Qualitative vs. quantitative BC&A

Although merely qualitative summary entails a loss of information, the information that remains is often sufficient for a DM to formulate improvements to insight or models, and occasionally it can indicate the best choice. Consideration of quantitative aspects of sensitivity allows an argument to be made that certain variables (those with little impact) are unimportant, allowing attention to be focused on the relatively few variables to which the objective function is sensitive. This eases assimilation when the cognitive burden would otherwise have been high. However, full assimilation of the quantities associated with the remaining variables is more difficult than assimilating their quality. Perhaps the easiest taxonomy to assimilate is a hybrid, which initially uses quantity to characterize certain features as unimportant, and then reports the qualitative aspect of remaining features.

Partial 2-way sensitivity analysis

An elaboration of the BC&A analytic principle may be considered if there is only one decision with no informational conditioning. This elaboration is to run two-way sensitivity analysis of each uncertain variable with the decision. I shall refer to this as partial two-way sensitivity analysis (P2SA).

The results of this exploration can be interpreted in two ways - to measure the sensitivity of one's utility to uncertainties conditioned on each alternative, or to identify whether the relative favorableness of one alternative versus another changes under different uncertain conditions. The former interpretation is sometimes of use, especially when a variable is essentially irrelevant under one alternative but has significant impact under some other alternative. The latter interpretation is what we will focus on here, however, because it contributes to a particularly important kind of insight for the DM. In this interpretation, not only will the DM be thinking about whether an alternative looks good, but also whether any particular uncertainty can make it look bad.

If one alternative is better than another under all circumstances considered in this analysis, this suggests that that alternative is better overall. This condition, though interesting, is not truly a form of dominance; it has been called tornado dominance[6].

The mathematics of optimal choice for one decision with no informational predecessors is simple: choose

$$d^* = \arg\max_d \sum_x u(x,d)P\{x|d\}$$

where d is the decision, x is the state of all other variables in the model, u is the utility assigned to a state of the world, and $P\{a|b\}$ refers to the probability assigned to a, given that b is the case.

To choose d*, it must be the case that

$$\sum_x u(x,d^*)\ Pr\{x|d^*\} > \sum_x u(x,d')\ P\{x|d'\} \quad \text{and}$$

$$\sum_x u(x,d^*)\ Pr\{x|d^*\} - \sum_x u(x,d')\ P\{x|d'\} > 0$$

for any other alternative d'. In general, the number of possible values of x in a decision network model will be large. DSA considers only a small number of these cases - those where most components of x are at their "base case" value. A heuristic basis for the interpretation of sensitivity analysis proposed here is to assume that virtually all of the probability mass associated with varying the i'th component of x is at the conditional base case for that perturbation (whose utility we denote $\underline{u}$). Further, we recall that the perturbations are specified as percentiles of the conditional distribution, making their probabilities independent of the decision. Hence we may calculate an approximate equality for the expression given above.

---

[6]This name was coined by David Lowell.

$$\sum_x u(x,d^*)\Pr\{x|d^*\} - \sum_x u(x,d')P\{x|d'\} \approx$$

$$\sum_{x_i} \underline{u}(x_i,d^*)\Pr\{x_i|d^*\} - \sum_{x_i} \underline{u}(x_i,d')P\{x_i|d'\} =$$

$$\sum_{x_i} \underline{u}(x_i,d^*)-\underline{u}(x_i,d')\Pr\{x_i\}$$

The term $\underline{u}(x_i,d^*)-\underline{u}(x_i,d')$ is the "delta" of alternative d* with respect to d'. This summation is a weighted average of deltas, according to the probabilities associated with the deviations of the i'th component of x. P2SA takes a look at these deltas, and if most or all are positive, it suggests concluding that d* is more attractive than d'. The conclusion is not rigorous, but it can be useful in determining the nature and amount of effort to put into addition elaborations of the decision model.

One important benefit of consideration of these deltas is that mere reliance on sensitivity of utility to perturbations could cause a substantial amount of effort to be spent on a variable that has a large but equal effect on the outcome of all alternatives, even though such a variable is really of little consequence for the decision. An example of this sort is sometimes encountered in the electric business: the price of oil affects the fortunes of the firm significantly, but may make little difference in capacity planning decisions because its impact is similar in all cases. In such cases, the delta for oil price would be zero, and consideration of P2SA would let the DM know that effort could profitably be redirected from oil price to other variables.

7.3.3 Overall assessment

I summarize the chief advantages and disadvantages of the analytic principles here with respect to two aspects that repeatedly were important in the discussion above: whether they support a comprehensible decomposition of models and intuitions, and whether the analysis thus decomposed is sufficiently consistent with a justifiable analysis of the problem at hand.

Merely distinguishing important distinctions is succinct and easy to judge, but gives insufficient information to make a decision. Identifying individual basis elements (probabilities and utilities) specifies the correct alternative, but is too extensive and too remote from natural modes of human thought to serve well. Consideration of steps in a simulation is congenial to human thought and to many kinds of models, but simulations required to understand a problem are often too extensive to command understanding and belief. BC&A presents difficulties in auditing base cases, but these have been studied extensively in the modeling literature. Base-case and qualitative adjustments is very succinct, and it is even more congenial to human judgements, but is somewhat less indicative of an optimal decision. Holistic BC&A is substantially more tractable than decomposed BC&A in a large decision network. BC&A conditioned on one decision gives somewhat better indications of the best choice, but is more extensive. Similar ends can be served a bit more succinctly by presenting sensitivity of "deltas" from a base case. None of these organizing principles emerges unambiguously superior. The approach chosen in this thesis, qualitative conditional holistic BC&A, attempts to capture the diagnosticity of P2SA with the brevity and intuitiveness of qualitative BC&A. The following diagrams give a simplified summary of the arguments in this section. Base case and sensitivity is the best combination of comprehensibility and justifiability, and qualitative base case and summary with a noise threshold is the best variant of base-case and adjustment.

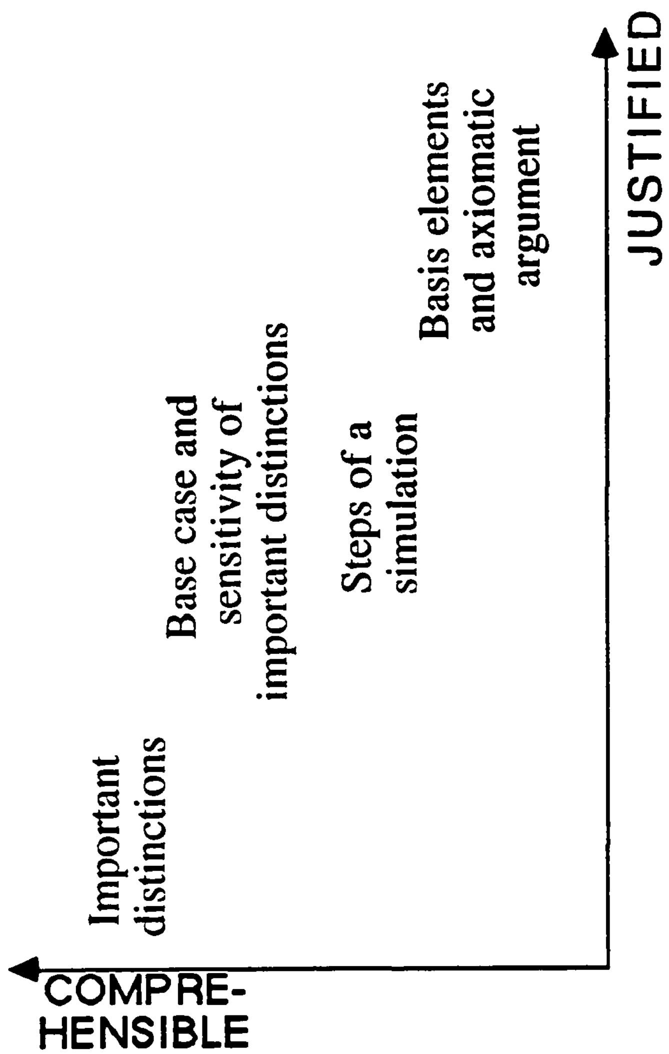

Figure 1. Reporting base case values and sensitivities (adjustments) to inputs is a comprehensible and justifiable approach to summarizing an analysis.

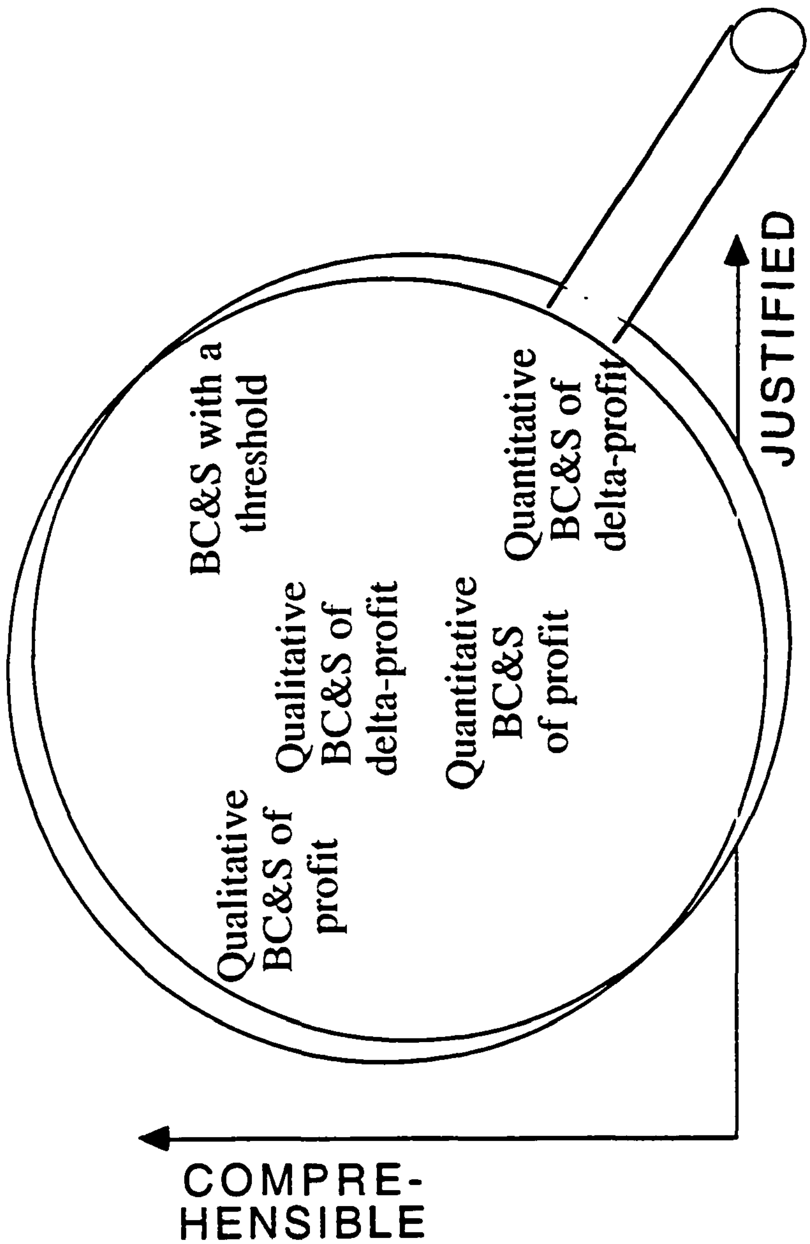

Figure 2. Reporting qualitative base case and sensitivities, with a zero threshold, is the most comprehensible and justifiable approach.

7.4 Presentation mode of the summary

In this section, I consider the merits of graphical and textual presentation of BC&A.

Tornado diagrams, the traditional presentation, show base case and relative size of impact graphically. Multiple tornado diagrams (from two-way sensitivity analysis) can be printed in different colors and intermingled to juxtapose the sensitivity bars of a variable under different conditions. In this way, the quantitative sensitivity of variables is clearly shown. Both the multidimensional scaling literature and the human factors literature suggest that representing quantities graphically can be meaningful to DMs. However, the direction of the model's response to variation of a variable may not be clear in a tornado diagram.

To support a comparison of component impacts, it is important that the qualitative impact of factors or interactions be clear. This is problematic in environments where both favorable and unfavorable figures of merit (costs and benefits) are routinely used as proxies for utility, and where enough variables are of interest that even the results of a sensitivity analysis present more information than can be reliably be interpreted by the DM.

As suggested before, reporting the qualitative impact of factor effects verbally puts sensitivity information into a form that is readily understandable by the DM. In addition, this allows generalizations, which can make such summaries quite brief. Brevity is said to be an aid to understanding in the cognitive systems engineering literature, experimental psychology (Miller 1956, Bruner et alia 1956, Waugh and Norman 1965), cognitive science (Simon 1955), and philosophical accounts of explanation as being fundamentally in service of simplification (Friedman 1974, Kitcher 1981).

Text gives the ability to collect similar effects and generalize over them in a way that can be reported naturally and succinctly in English. Philosophical literature suggests that subsumption of a point of view under a broader one creates understanding or insight. Hempel and Oppenheim (1948) and Railton (1978) conclude that the explanation of a phenomenon consists in its subsumption under laws or under a theory. Lonergan (1957) feels that only the consideration of a set of increasingly broad contexts can create an insight into insight. Kitcher (1981) argues that scientific explanation aims at a kind of insight that is achieved by exhibiting the phenomena as manifestations of common, underlying structures and processes. The summary system described in chapter 4 employs this sort of subsumption to summarize qualitative P2SA.

A textual-tabular hybrid presentation has also been suggested. Favorable factors and alternatives could be presented in a column of a table. This approach shares the benefit of juxtaposing ostensibly equivalent factor-impacts, and could be even more succinct if the meaning of the columns could be conveyed clearly by their headings. This tabular hybrid approach was encountered too late in this research to be considered more fully than this brief discussion.

Graphical and verbal presentation present different aspects of the situation well. Use of both seems to be indicated.

7.5 Implementation of the verbal summary

This section develops a specific approach to verbal summary, and describes the implementation of it in this thesis work.

Here I formally state the aspects of a decision situation I intend to summarize.

Consider a situation with only one decision. Let the variables $x_1,\ldots,x_n$ be all the features of interest whose values are not known with certainty. Let $u(a,x_1,\ldots,x_n)$ be the deterministic function that maps an alternative (the first argument) and a set of feature-values (the remaining arguments) to the corresponding value of a target variable. For simplicity, I shall refer to this as utility. Denote base-case values of variables by addition of an "o" to the subscript. Define the i'th projection of u as:

$$u_i(x_i) \equiv u(a_o,x_{1o},\ldots,\mathbf{x_i},\ldots,x_{no}),$$ [7]

and define the 0'th projection in terms of the alternative chosen:

$$u_o(a) \equiv u(\mathbf{a},x_{1o},\ldots,x_{no}).$$

Define the delta of alternative a as:

$$\partial(a) \equiv u_o(\mathbf{a}) - u_o(\mathbf{a_o}).$$

Finally define the conditional delta of alternative a to be:

$$\partial_i(a,x_i) \equiv u(\mathbf{a},x_{1o},\ldots,\mathbf{x_i},\ldots,x_{no}) - u(\mathbf{a_o},x_{1o},\ldots,\mathbf{x_i},\ldots,x_{no}).$$

These relationships can be characterized clearly in English. $\partial(a) > 0$ means that alternative a is beneficial in the base case. $u_i(x) > u_i(x_o)$ means that value x of feature i improves utility in the base case. $\partial_i(a,x) > \partial_i(a,x_o)$ means that value x of feature i makes alternative a more favorable

[7] Bold type is used for emphasis only here. All variables are considered to be single-valued here.

than it would have been in the base case. I call these relationships Elementary Qualitative Conclusions (EQCs).

A summary that stated all EQCs individually would create significant information overload. Generalization can present modeling results in a more succinct fashion. I have built an explainer that begins with a corpus of EQCs and performs quantification over variables to produce a few quantified statements that subsume several EQCs, allowing individual statements of them to be omitted. I have chosen to generalize in order to support the lines of reasoning discussed in sections 7.2.3.4 and 7.2.3.5, thinking of the model as a base case utility and adjustments therefrom. There are two design issues here: over which arguments to generalize, and in which order to perform the generalization. The following table lists patterns of EQCs that I summarize, and corresponding verbal characterizations.

RELATIONSHIP	VERBAL CHARACTERIZATION
$\forall a \in A \;\; \partial(a) > 0$	"Choosing alternatives <A> helps <u>."
$\forall i \in I$ $\neg((\forall x,y \;\; x<y \rightarrow u_i(x) < u_i(y))$ $\wedge (\forall x,y \;\; x>y \rightarrow u_i(x) < u_i(y)))$	"Features <I> have a non-monotonic impact on <u>."
$\forall (i,x) \in X \;\; x \neq y \rightarrow u_i(x) > u_i(y)$	"Feature-values <X> help <u>."
$\forall (i,x) \in X \;\; \partial_i(a,x) > \partial_i(a,x_o)$	"Feature-values <X> make <a> more attractive."

In this table, A is a set of alternatives, X is a set of feature-value pairs, I is a set of features, and u is a utility function. These sets are specified by enumeration or by universal quantification over all existing alternatives or feature-value pairs in the model, if appropriate, in the verbal summary.

7.6 Summary of other analyses as two-way sensitivity

As the DA cycle progresses, one builds and assesses a decision network model, and evaluates its recommended action. Here, again, the DA literature suggests reappraisal of the model in light of sensitivity analysis. This section shows how to use the machinery already in place for DSA to summarize the results of the full decision theoretic evaluation.

The accessibility of two-way sensitivity analyses suggests that aspects of a full factorial analysis or a decision analysis could be summarized by "boiling its data down" to pseudo-data in the form of P2SA.[8] This was done in the WSDM model discussed in the previous chapter. The method of "boiling down" the data employed there is to generate a set of summary numbers in the form of the P2SA results by reporting the probability-weighted average of all simulation games that agree with the feature-value and alternative for each cell in the table.

7.6.1 Numerical table

The following table summarizes a WSDM analysis in the format discussed here. In it, three alternative capacity contracts are analyzed. The table show utility projections in the leftmost column, alternative deltas across the top, and conditional deltas of alternatives in the body of the table.

Combined effect of important factors on value of a capacity contract

SCECOS is the variable being summarized.

alternative: uncertainty	no deal	1d cap	7d cap
Base case	2154.	1.9	1.0
hydro: good	2239.	2.6	2.7
hydro: bad	2069.	1.2	-0.7
oil price: $1.50	1396.	2.4	2.0
oil price: $5.00	2912.	1.4	-0.1
coal escl: 2%	2000.	2.5	2.0
coal escl: 8%	2308.	1.3	-0.1
year: 1989	1601.	2.6	2.5
year: 1997	2707.	1.2	-0.6

[8] The advantages of reporting studies as two-way sensitivity analysis of this form (sensitivity of all state variables under each alternative, deltas from the base case) was shown to this author by Ed Cazalet in professional work.

The utility function here is a portion of total system costs (in $ millions) to SCE. Since we are measuring costs, positive deltas tend to indicate an unfavorable option. 1d cap and 7d cap refer to alternative forms of a capacity contract under consideration in this decision: those for which the energy used to deliver the capacity must be delivered within one day, and those that allow it to be returned up to seven days later.

7.6.2 Decision tree

The following decision tree illustrates the calculation of the first few numbers in the summary table (those for variations in hydro). The numbers on the branches are the weighted averages of the utilities from all branches to their right in the tree, and they correspond to the numbers that appear in the first few lines of the summary table. The numbers in subsequent lines of the table would be evident in a decision tree with the variables re-ordered so that the corresponding uncertainties appear leftmost.

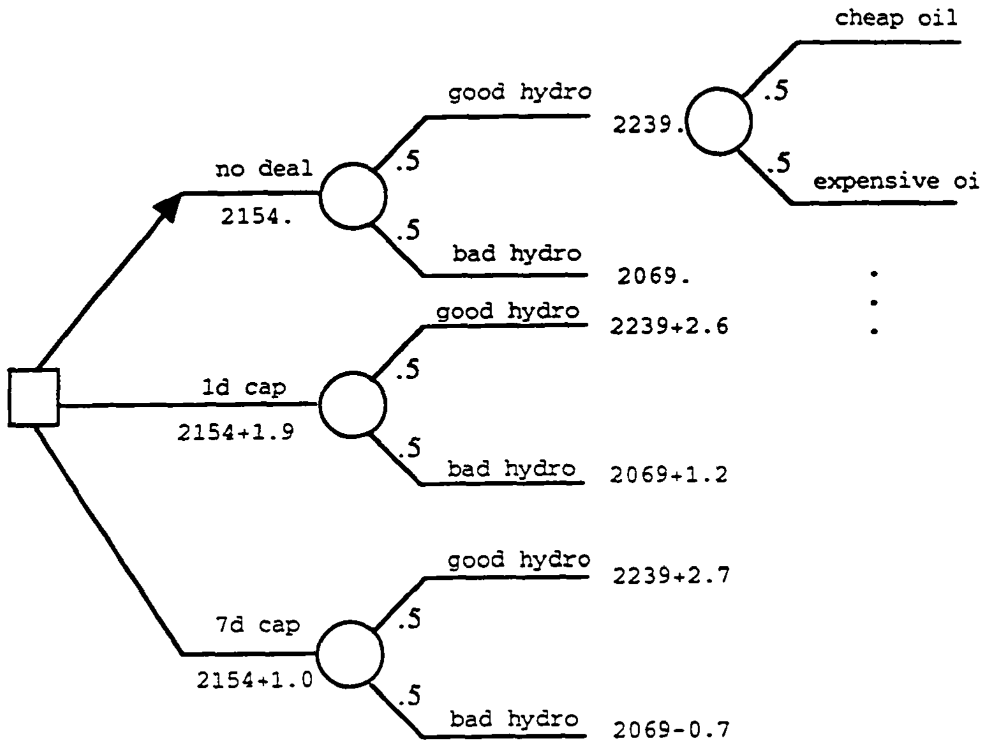

Figure 3. Decision tree for the capacity purchase decision

7.6.3 Tornado diagram

Here are "tornado diagrams" portraying the results of this study. Very different pictures appear, depending on whether total costs, or their deltas, are reported. Even with the numerical annotations, the direction of impact of the variables is not immediately evident here; it must be "figured out".

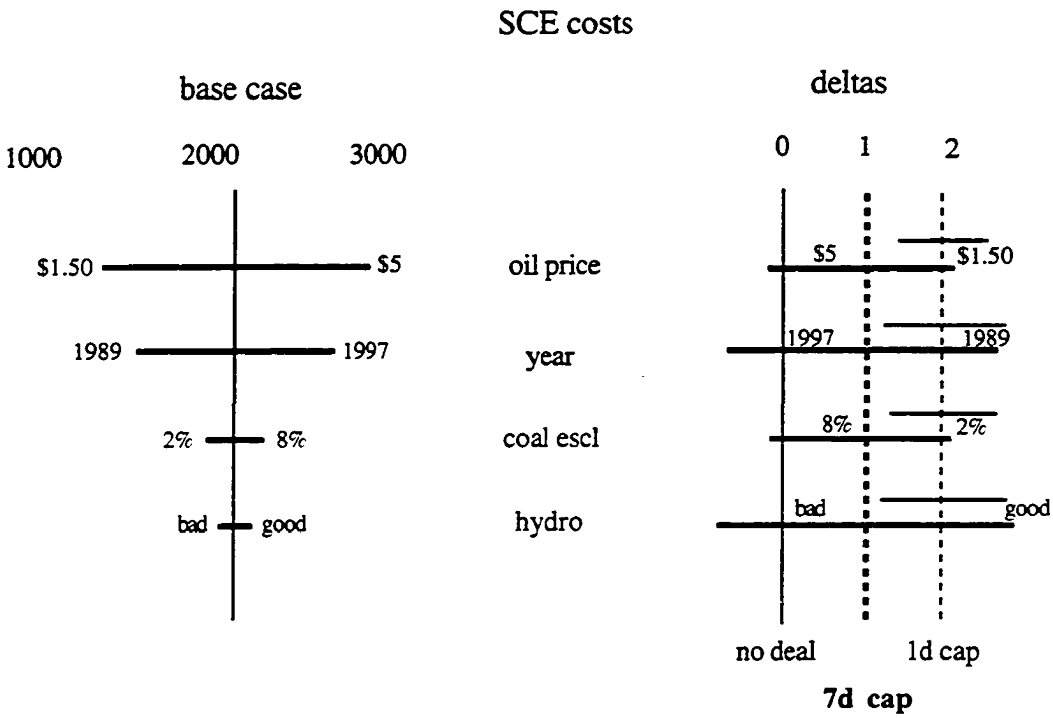

Figure 4. Hand-drawn tornado diagrams for the capacity purchase decision

7.6.4 Verbal summary

Here is the computer-generated summary produced for this power marketing study:

Summary according to measure SCECOS:

Choosing each option hurts SCECOS.

Feature values bad hydro, low oil price, low coal escl, early year, help SCECOS.

Feature values bad hydro, high oil price, high coal escl, late year, make 7d cap option more attractive.

Feature values bad hydro, high oil price, high coal escl, late year, make 1d cap option more attractive.

Feature value bad hydro makes each option more attractive.

Feature value high oil price makes each option more attractive.

Feature value high coal escl makes each option more attractive.

Feature value late year makes each option more attractive.

7.7 The chapter in review

This chapter presents a rational reconstruction of the DA Cycle as iterative division of computer models and intuitions into comparable parts, comparison of these, and improvement as appropriate. It argues, based on existing studies of judgement and decision making, in favor of dividing these accounts according to what their "base case" is and "how the model responds" to perturbations. This decomposition is tractable to human experts, hence it can support critical assessment of a model. The chapter shows how to generate a verbal summary of a computer model along these lines, and it shows other presentations of similar information to make clear what the meaning of the summary is.

Chapter 8. Looking back; looking forward

8.1 Summary

The ultimate goal toward which this dissertation aims is for decision-makers to know the best action to take. Knowledge is adequately justified belief. Decision theory gives reason (justification) to believe that a choice is optimal if the Decision Maker (DM) believes all the important aspects of the decision situation have been represented properly in the formulation that was solved. The thrust of this dissertation has been to motivate and present a tool whose use can help the DM formulate a decision situation well, and give her confidence that it captures all that is relevant about her situation. This, together with the use of decision theoretic reasoning about the resulting formulation, gives the DM justified belief that the recommended action is the best one.

The choice axiom of Decision Theory (DT) specifies a self-evident principle for choice in a very limited set of circumstances, and DT's other axioms specify ways to transform the prospects at hand into a set of equivalent prospects, to which the choice axiom applies. DT is an analysis of the word "best" based on the construction of equivalent circumstances that are easy to judge.

The DA Cycle is a problem-solving approach wherein the judgements forming the basis of a decision theoretic calculation are repeatedly appraised and revised in light of the results of decision theoretic calculations. These revisions often change the definition of one or more variables, leading to a statement of the problem that captures its essence more clearly. As the DM's viewpoint evolves, the variables she considers crucial, and their level of detail, change.

While the variables under consideration change from to decision to decision, and even during the course of a given decision, certain aspects of situations faced by an ongoing business often do not change. This juxtaposition of change and stasis presents a challenge: we would like the effort put into modeling situations of repeated interest to be reused, but modification of a generic reusable model to reflect specific circumstances can create difficulties. Often such a model operates at a high enough degree of detail that the DM's variables can be considered to be parametrizations of the model's detailed data. In this case, I propose that the model be insulated from these changes by wrapping a variable translation routine around it, localizing the code changes required by changing variables or parametrizations. Thus the selection of sets of exogenous input values and any desired summarization of outputs can take place in the set of variables considered important in the problem at hand without constantly changing the underlying model.

Many problem-solving approaches in the literature, including decision theory (DT), can be viewed as reducing a complex characterization of a situation to a simple one where the best choice is obvious, but few discuss how justified belief in the initial characterization is to be generated. For instance, in discounted cash flow analysis, alternative cashflow streams that are difficult to compare directly are converted to equivalent amounts of cash now, which are easy to compare; but DCF itself offers no procedure for verifying the cashflows. Two themes emerge from articles that do address enhancement of participants' points of view: generate and compare multiple viewpoints, and iteratively improve upon viewpoints found to be wanting. I argue that the DA Cycle can employ both of these approaches to generate understanding of DT arguments. The viewpoints involved here are formal DT equivalence arguments (as embodied in a model) and the DM's direct judgements. In the DA Cycle, these are iteratively compared, and one or the other is improved until they converge.

I argue that the comparison and improvement takes place using a different organizing principle than the complex DT argument. A common way for people to form judgements about a complex situation is to think of a base-case outcome and a set of adjustments to this outcome in response to exogenous variation of individual factors. These judgements are directly comparable to the sensitivity of a DT model to those factors. Comparing intuitions to model results in terms of base case and sensitivities breaks down the problem of comparing and reconciling the two overall accounts of the situation into a set of simpler task of reconciling their accounts of the response of the system to various exogenous factors. As the two accounts are reconciled, at least one of them improves, and they converge. When the DM's judgements agree with a high-level summary of a rigorous DT equivalence argument, the agreement gives her justified belief (knowledge) that her intuition is coherent and that the formulation captures her best expertise about her situation; and reflection on the nature of DT can convince her that the decision theoretic recommendation for that formulation is the best possible action.

To support this comparison of component impacts, the qualitative impact of factors or interactions must be clear. This is problematic in environments where both favorable and unfavorable figures of merit (benefits and costs) are routinely used, and where enough variables are of interest that even the results of a sensitivity analysis present more information than the DM can easily interpret. Reporting only the qualitative impact of factor effects puts sensitivity information into a form that is readily understandable by the DM. It also gives the ability to collect similar effects and generalize over them in a way that can be reported naturally and succinctly in English. Although this entails a loss of information, the information that remains is often sufficient for a DM to formulate improvements to insight or models. As an example, the system implemented for this research gives outputs like "The following factors <a list of factors> make the base case more favorable.", or "The following factors <list> make alternative x more attractive.".

The decision analytic approach dictates that a simple model be elicited from the DM and experts, and both DA and behavioral decision theory suggest that that sensitivity analysis be used to motivate its elaboration. DA's traditional notion of deterministic sensitivity analysis, one-at-a-time perturbation of variables to observe the response of the objective function, is consistent with this program, but implementations in the literature are not, insofar as they fail to measure the sensitivity of the objective function to some kinds of nodes in an evolving decision network model. Following the workaround typically employed in the face of such difficulties, we extend the definition of deterministic sensitivity to allow probabilistic descendants of the variable being perturbed to respond "according to their nature", rather than holding them fixed. An algorithm that accomplishes this with a small number of elicitation questions is developed here and used in the Deft software.

In an account based loosely on the United States' decision among synfuel commercialization alternatives, an implementation of these facilities guides model development by identifying bugs, allowing certain simplifying approximations to be employed, and suggesting more elaborate treatment of crucial variables. It gives insight by showing surprising results that are subsequently explained. It creates knowledge of the best action as participants see that the analysis is responsive to their concerns and expertise, and as the DM finds that it agrees with his improved apprehension of the situation.

8.2 Possibilities for Future Research

This section considers two lines of inquiry that follow naturally from this dissertation work: verifying its utility, and extending its functionality.

8.2.1 Verifying usefulness

This dissertation takes seriously an approach to problem formulation that is written about in the decision analysis literature, and it develops two technologies to support that approach - one to support elicitation of a simple model, and one to support verbal summary of the model's behavior. It argues that the former is parsimonious, and this argument appears strong. On the latter point, it offers an argument based on previous empirical research in the field of behavioral decision theory to suggest that its approach to verbal summary is effective. This argument is less convincing, because results in empirical psychology frequently surprise the researcher.

Thus it would be of value to test the effectiveness of this software and approach empirically. Such a test could either be holistic, testing whether a realization of the entire approach used in this dissertation achieves its aims, or it could be piecemeal, testing whether existing components contribute to the achievement of the milestones delineated in chapter 7, and whether achievement of those milestones contributes to the overall goal.

The ultimate aim of this dissertation is for the DM to know what the best action is, in a cost-effective way. Relying for now on the argument that the parsimony of the elicitation procedure keeps analytic costs reasonably low, this section addresses the simpler question of whether this work contributes to effective problem-solving processes. Only after establishing more clearly the benefits of these procedures can a good study be designed to compare them to its costs.

In the holistic approach, subjects would be asked to perform a decision-making task, some with the fruits of this dissertation and some without, and we would test whether they know the best action. Decision-making as discussed in this dissertation is an exploration of one's expertise, and that of one's experts. An important challenge for the experimenter is to design a decision situation and a way of presenting it that gives subjects enough (artificial) expertise to explore, that does not pre-think the analytic structure to be used for the decision, and that is nonetheless tractable within the time a subject is willing to contribute to an experiment. The claim of this dissertation is that Deft, in conjunction with the assessment and solution of a decision network model, creates knowledge of the best action. In a holistic test, these capabilities would have to be made available to the subject, either by integration of Deft with software supporting assessment and solution of decision networks, or by provision of decision analytic assistance to the subjects. Upon completion of their exercise, subjects could be asked to report, on a seven-point scale, their degree of agreement with the statement "I know what the best action is.".

In the piecemeal approach, milestones and mechanisms hypothesized in chapter 7 can be tested. To test the divide/compare/improve account, we can test whether subjects can perform the elementary judgements of impact into which an analysis breaks down a whole problem, whether these are easier or more reliable than forming a holistic judgement, and whether identification of counterintuitive model behavior suffices to suggest modifications to the model. Presumably these tasks will be easier with a directly assessed model than with a complicated reused one.

Furthermore, my conjecture that identification of counterintuitive model responses is eased by juxtaposition with other factors ostensibly having the same qualitative impact could be tested fairly easily in a laboratory setting.

One could also test whether subjects exposed to explanation approaches tailored to other ways of thinking (e.g., telling a "causal" story, or describing an additive model, or giving a graphical presentation thereof), or to no explanation at all, can formulate as many worthwhile elaborations to the model as subjects with Deft.

If a Bayesian perspective toward the analysis of data is taken, one study could achieve the benefits of both approaches. This hybrid methodology was employed by Chu and Fehling (1994). That work presented and evaluated a new approach to conflict resolution. In that case, as in this one, the cognitive task for subjects (taking part in a full-scale conflict resolution session) was difficult and time-consuming, so relatively few trials were performed, and data was taken from trials to measure both the outcome of interest (whether a high quality settlement agreement was reached), as well as the status of other variables related to the success of the new approach (efficiency of communication). In a Bayesian perspective, the experimenter knows something going into the experiment, and uses the experimental data to update beliefs about pertinent phenomena. In Chu and Fehling (1994), the experimenters had relatively diffuse priors on the quantities of interest, but, after reading and discussion with domain experts, they had beliefs of moderate strength that there would be correlations among quantities that they could measure. They created a belief net to represent these prior beliefs, and used the experimental data to update these beliefs. The data, if taken quantity by quantity was statistically insignificant, or barely significant; but each datum measured was favorable to the new approach, and when taken together, they amounted to strong evidence in favor of the new approach.

This Bayesian approach could be employed to integrate the information from an experiment that measured the variables implicated in both the holistic and piecemeal analyses. A belief network could be constructed representing the experimenter's relatively diffuse priors on users' difficulty of forming judgements of qualitative and quantitative impacts of factors, users' facility with identifying counterintuitive results, the number of elaborations proposed, and the degree of assent to "knowing the best choice". Measurements of these variables from experimental trials could be used to update these priors.

The fact that philosophers (Hegel, Feyerabend (1974) and Longino (1990)) and problem solvers (Mitroff and Betz (1972) and Hogan (1978)) call for critical comparison constitutes a basis for a prior that is slightly favorable to the Deft approach, but in view of the current distaste for non-data-oriented pursuit of science, the research results might be more convincing if one chose not to avail oneself of this opportunity.

8.2.2 Extending functionality

Another way to pursue this research is to extend the conjectural elaboration of Deft. One could do this either by reference to other findings in the judgement literature, or in ways that give better justification to the belief that the model is good and agrees with the DM's direct judgements.

Work done to follow up on Brunswick's lens theory and other linear models (e.g., in Slovic and Lichtenstein, 1971) suggests that judgements are not merely qualitatively additive, but quantitatively so; a summary that treated quantitative aspects and emphasized additivity could be considered. Evidence that people reason causally (Johnson-Laird 1980, Kahneman and Tversky 1982) suggests that a completely different approach to summary, one that gives an account of events in a causal sequence, might work well.

The generalizations given in Deft's summary are all implicitly qualified with the phrase "Under base-case conditions...". To release this assumption, the system could identify and search through additional sensitivity cases, to verify the robustness of these generalizations. Doing so would increase the number of assessments required, but would give increased assurance that the verbal summary is correctly reporting the behavior of the model being constructed, and perhaps increased confidence with the elaborations so elicited. Such analyses would be unhelpful in the early formulation phase of the DA Cycle, but they would seem more appropriate if the analysis is largely based on a preexisting model that is not too expensive to run, or in a later stage when more of the required assessments have been made. Sensitivities that could be explored include full two-way sensitivity analysis of uncertain variables with other uncertain variables, or full (multi-way) factorial analysis.

Bibliography

[Achinstein 1983] Achinstein, P., "The illocutionary theory of explanation", The Nature of Explanation, Peter Achinstein, Oxford University Press, 1983. Excerpted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.199.

[Adams et alia 1990] Adams, D.A., Courtney, J.F.Jr., Kasper, G.M., "A process-oriented method for the evaluation of decision support system generators", Information and Management, 19(4):213-25.

[Anderson 1983] Anderson, J.R., The Architecture of Cognition, Harvard University Press, Cambridge, Mass.

[Augustine and Coovert 1991] Augustine, M.A., Coovert, M.D., "Simulation and information order as influences in the development of mental models", SIGCHI Bulletin, 23(1):33-5.

[Bailey and Duban 1990] Bailey, J.J., Duban, S.L., "Explanation and learning in medicine", Computers & Education, 15(1-3):91-97.

[Banerjee and Basu 1990] Banerjee, S., Basu, A., "A knowledge based framework for selecting management science models", Hawaii International Conference on System Sciences 23(3):484-493.

[Bankes 1993] Bankes, S., "Exploratory modeling for policy analysis", Operations Research, 41(3):435-449.

[Baskaran and Reddy 1984] Baskaran, V., Reddy, Y.V., "An introspective environment for knowledge based simulation", Proceedings of the 1984 Winter Simulation Converence, ed S.Shepard et alia, IEEE, Dallas TX, p. 644-651.

[Bayes 1763] Bayes, T., "An Essay toward Solving a Problem in the Doctrine of Chance", Philosophical Transactions of the Royal Society, p.370-418. Reprinted in Facsimiles of Two Papers by Bayes, ed. W.Edwards Deming, 1963, Hafner Pub. Co., New York.

[Bernoulli 1730] Bernoulli, D., Exposition of a new theory on the measurement of risk, English translation of "Specimen theoriae novae de mensura sortis," Commentarii academiae scientiarum imperialis Petropolitanae, 1730 and 1731, 5, 175-192, 1738, by Louise Sommer, Econometrica, 22, 23-26, 1954.

[Blanning 1985] Blanning, R.W., "Determining the applicability of fuzzy logic to knowledge-based DSS", Hawaii International Conference on System Sciences 18(1):530-538.

[Blanning 1986] Blanning, R.W., "A relational framework for information management", Decision Support Systems, A Decade in Perspective, ed. E.R.McLean & H.G.Sol, Elsevier Science Publishers B.V., North-Holland.

[Blanning 1988] Blanning, R.W., "An entity-relationship framework for information resource management", Information and Management 15(2):113-119.

[Boettner 1985] Boettner, R., "Fuzzy approach for the selection from a limited number of (nondominated) alternatives", Annual Review in Automatic Programming v 12 pt 1. Publ by Pergamon Press, Oxford, Engl and New York, NY, USA p 259-262.

[Bosch and Weyant 1989] Bosch, D.K., Weyant, J.P., Review of initial experience with CPUC rules for production simulation model use and suggested modifications, Energy Modeling Forum, Stanford University, Stanford CA.

[Bradley and Clemence 1988] Bradley, G.H., Clemence, R.D.Jr., "Model integration with a typed executable modeling language", Hawaii International Conference on System Sciences 21(3):403-410.

[Bradshaw and Boose 1990] Bradshaw, J.M., Boose, J.H., "Decision analysis techniques for knowledge acquisition: combining information and preferences using Aquinas and Axotl", International Journal of Man-Machine Studies, 32(2):121-186.

[Bradshaw et alia 1989] Bradshaw, J.M., Boose, J.H., Covington, S.P., Russo, P.J., "How to do with grids what people say you can't: The application of decision analysis methods in Axotl and personal construct methods in Aquinas to design problems", Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-based Systems Workshops, Banff, Canada, November 1988.

[Bridges and Johannes 1988] Bridges, S., Johannes, J.D., "Explanation production by expert planners", Fourth Conference on Artificial Intelligence for Space Applications, NASA Conference Publication n 3013. Publ by NASA, Washington, DC, USA. p 323-330.

[Brown and Lewis 1989] Brown, C.E., Lewis, T.G., "HELM: Hierarchical Environment for Linear Modeling. I. The schema", Hawaii International Conference on System Sciences 22(3):449-58.

[Bruffaerts et alia 1989] Bruffaerts, A., Henin, E., Pirotte, A., "Sound basis for the generation of explanations in Expert Systems", Philips Technical Review, 44(8-10):287-295.

[Bruner et alia 1956] Bruner, J., Goodnow, J., Austin, G., A Study of Thinking, Wiley.

[Buckley 1987] Buckley, J.J., "Portfolio analysis using possibility distributions", Approximate Reasoning in Intelligent Systems, Decision and Control. Sanchez, E.; Zadeh, L.A. (editors), Pergamon, Oxford, UK. p. 69.

[Bunn 1986] Bunn, D.W., "Forecasting with multiple predictors: pluralistic methods in decision-support", IEE Colloquium (Digest) n 1986/119. Publ by IEE, London, Engl p 7.

[Castillo et alia 1991] Castillo, D.G., Dolk, D.R., Kridel, D.J., "GOST: an Active modeling system for costing and planning NASA space programs", Hawaii International Conference on System Sciences 24(3):396.

[Chandrasekaran et alia 1986] Chandrasekaran, B., Josephson, J., Keuneke, A., "Functional representations as a basis for generating explanations", Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics p 726-731.

[Chao et alia 1985] Chao, H.P., Peck, S.C., Wan, Y.H., "LOAD: an electric technology R&D planning model with contingent decisions", Resources and Energy v 7 n 2 Jun 1985 p 163-177.

[Chu et alia 1989] Chu, P.-.Y., Moskowitz, H., Wong, R.T., "Robust Interactive Decision-analysis (RID): Concepts, methodology, and system principles", Hawaii International Conference on System Sciences 22(3):255-261.

[Chu and Fehling 1994] Chu, P.Y.V., Fehling, M.R., "Cognitive Conflict Resolution: mediation analysis and strategies", IEEE International Conference on Systems, Man and Cybernetics, SMC-94, San Antonio, p.1553-1558.

[Collins and Gentner 1987] Collins, A., Gentner, D., "How people construct mental models", Cultural Models in Language and Thought, ed. D.Holland and N.Quinn, Cambridge University Press.

[Congressional Research Service 1980] Congressional Research Service, The Pros and Cons of a crash program to commercialize synfuels, U.S. G.P.O. Su.Docs 57-542 O.

[Cooper 1984] Cooper, G.F., NESTOR: a computer-based medical diagnostic aid that integrates causal and probabilistic knowledge, dissertation, Stanford Univ.

[Coté and St-Denis 1992] Coté, V., St-Denis, R., "Bridging the gap between CASE tools and project management through a decision support system based on metrics", Hawaii International Conference on System Sciences 25(3):300-309.

[Courtney, Paradice and Ata Mohammed 1987] Courtney, J.F.Jr., Paradice, D.B., Ata Mohammed, N.H., "A knowledge-based decision support system for managerial problem diagnosis", Decision Sciences, 18:373-99.

[Coury 1987] Coury, B.G., "Multidimensional scaling as a method for assessing internal conceptual models of inspection tasks", Ergonomics, 30(6):959-973.

[Craik 1943] Craik, K., The Nature of Explanation, University Press: Cambridge, reprinted 1952.

[D'Ambrosio et alia 1987] D'Ambrosio, B., Fehling, M.R., Forrest, S., Raulefs, P., Wilber, B.M., "Real-time process management for materials composition in chemical manufacturing", IEEE Expert, 2(2):80-93.

[Dempster and Ireland 1989] Dempster, M.A.H., Ireland, A.M., "Object-oriented model integration in a financial decision support system", Decision Support Systems 7:329-340.

[Dolk and Konsynski 1984] Dolk, D.R., Konsynski, B.R., "Knowledge representation for model management systems", IEEE Transactions on Software Engineering, SE-10(6):619-628.

[Dutta and Basu 1984] Dutta, A., Basu, A., "Artificial intelligence approach to model management in decision support systems", Computer, 17(9):89-97.

[Eck et alia 1990] Eck, R.D., Philippakis, A., Ramirez, R., "Solver representation for model management systems", Hawaii International Conference on System Sciences 23(3):474-483.

[Eden et alia 1986] Eden, C., Williams, H., Smithin, T., "Synthetic wisdom: the design of a mixed-mode modelling system for organizational decision making", Journal of the Operational Research Society, 37(3):233-241.

[Elkan 1993] Elkan, C., "The paradoxical success of fuzzy logic", Proceedings of AAAI-93 and IAAI-93, p.698-703.

[Elsaesser 1989] Elsaesser, C., "Explanation of probabilistic inference", Uncertainty in Artificial Intelligence 3, ed L.N.Kanal, T.S.Levitt, J.F.Lemmer, Elsevier Science Publishers B.V. (North-Holland), p.387-400.

[Eschenbach and McKeague 1989] Eschenbach, T.G., McKeague, L.S., "Exposition on using graphs for sensitivity", Engineering Economist, 34(4):315-33.

[Fehling and Breese 1988] Fehling, M.R., Breese, J.S., "A computational model for decision-theoretic control of problem-solving under uncertainty", Proceedings of the 4th AAAI Workshop on Uncertainty in Artificial Intelligence. UAI-4.

[Feyerabend 1974] Feyerabend, P., Against Method, New Left Books.

[Fischhoff 1986] Fischhoff, B., "Decision making in complex systems", in Intelligent Decision Support in Process Environments, ed. E.Hollnagel et alia, Springer-Verlag, Berlin, p.61-86.

[Fischhoff, Slovic and Lichtenstein 1979] Fischhoff, B., Slovic, P., Lichtenstein, S., "Subjective sensitivity analysis", Organizational Behavior and Human Performance, 23:339-359.

[Fjeldstad and Konsynski 1986] Fjeldstad, O.D., Konsynski, B.R., "Reapportionment of cognitive responsibilities in DSS dialogues", Decision Support Systems: a Decade in Perspective, ed. E.R.McLean and H.G.Sol, Elsevier Science Publishers B.V. (North-Holland).

[Franksen 1979] Franksen, O.I., "Fuzzy sets, subjective measurements, and utility", International Journal of Man-Machine Studies, 11(4):521-545.

[Friedman 1974] Friedman, M., "Explanation and scientific understanding", the Journal of Philosophy 71(1):5-19. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.188.

[Gaines 1987] Gaines, B.R., "An overview of knowledge-acquisition and transfer", International Journal of Man-Machine Studies, 26(4):453-472.

[Gass 1977] Gass, S.I., "Evaluation of complex models", Computers and Operations Research, 4:27-35.

[Gass 1986] Gass, S.I., "A process for determining priorities and weights for large-scale linear goal programmes", J. Opl. Res. Soc. 37(8):779-785.

[Geoffrion 1985] Geoffrion, A.M., "An introduction to Structured Modeling", Management Science, 33(5):547-588, 1987. A draft was available in 1985.

[Geoffrion 1989] Geoffrion, A.M., "Reusing structured models via model integration", Hawaii International Conference on System Sciences 22(3):601-611.

[Geoffrion 1991] Geoffrion, A.M., "FW/SM: a prototype Structured Modeling environment", Management Science, 37(12):1513-1538.

[Geoffrion 1992] Geoffrion, A.M., "The SML language for structured modeling" (in two parts), Operations Research, 40(1):38-57 and 40(1):58-75.

[Gheorghe and Stoica 1987] Gheorghe, A.V., Stoica, M.V., "PC and fuzzy dynamics-aided decision: a case for energy sectoral models", Foundations of Control Engineering, 12(4):153-65.

[Gheorghe et alia 1985] Gheorghe, A.V., Stoica, M.V., Barac, A., Barbantan, M., "Fuzzy Models for Energy Demand Management", Cybernetics and Systems, 16(2-3):181-189.

[Godo et alia 1989] Godo, L., Lopez de Mantaras, R., Sierra, C., "MILORD architecture and its management of linguistically expressed uncertainty", International Journal of Intelligent Systems, 4(4):471-501.

[Goguen 1986] Goguen, J.A., "Reusing and interconnecting software components", Computer 19(2):16-28.

[Goldman and Breese 1992] Goldman, R.P., Breese, J.S., "Integrating model construction and evaluation", Uncertainty in Artificial Intelligence, proceedings of the eighth conference, Morgan Kaufmann, p.104.

[Gray and Borovits 1986] Gray, P., Borovits, I., "Contrasting roles of Monte Carlo simulation and gaming in decision support systems", Simulation, 47(6):233-239.

[Grice 1957] Grice, H.P., "Meaning", Philosophical Review, 66:377-88.

[Hämäläinen et alia 1986] Hämäläinen, R.P., Seppaläinen, T., Ruusunen, J., "Microcomputer-based decision support tool and its application to a complex energy decision problem", Hawaii International Conference on System Sciences 19(1a):494-502.

[Harlan 1982] Harlan, J.K., Starting with Synfuels: benefits, costs and program design assessments, Ballinger Publishing Co., Cambridge MA.

[Harman 1986] Harman, G., Change in View: Principles of Reasoning, MIT Press, Cambridge MA.

[Hasling et alia 1984] Hasling, D.W., Clancey, W.J., Rennels, G., "Strategic explanations for a diagnostic consultation system", International Journal of Man-Machine Studies, 20(1):3-19.

Hawaii International Conference on System Sciences, proceedings available from IEEE Service Center, Piscataway, NJ.

[Hayes-Roth et alia 1989] Hayes-Roth, B., Washington, R., Hewett, R., Hewett, M., Seiver, A., "Intelligent monitoring and control (health care systems)", IJCAI-89, Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, ed. N.S.Sridharan, Morgan Kaufmann, Palo Alto, CA, vol 1, p.243-249.

[Heckerman and Horvitz 1990] Heckerman, D.E., Horvitz, E.J., "Problem formulation as the reduction of a decision model", Proceedings of the Sixth Conference on Uncertainty in AI, 1990. Also: Knowledge Systems Laboratory report KSL-89-67 #326, Stanford CA.

[Heidegger 1926] Heidegger, M., Being and Time, English translation 1962, Harper, New York.

[Heise 1975] Heise, D.R., Causal Analysis, John Wiley & Sons, New York.

[Helman and Bahuguna 1986] Helman, D.H., Bahuguna, A., "Explanation systems for computer simulations", Winter Simulation Conference Proceedings 1986. Available from IEEE Service Cent (Cat n 86CH2385-3), Piscataway, NJ, USA p 453-459.

[Hempel and Oppenheim 1948] Hempel, C.G., Oppenheim, P., "Studies in the logic of explanation", Philosophy of Science 15:567-579. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.9.

[Hersh and Caramazza 1976] Hersh, H.M., Caramazza, A., "A fuzzy set approach to modifiers and vagueness in natural language", Journal of Experimental Psychology: General, 105(3):254-276.

[Hickman, Huntington and Sweeney 1987] Hickman, B.G., Huntington, H.G., Sweeney, J.L., "Macroeconomic Impacts of Energy Shocks", volume 163 of the Contributions to Economic Analysis series, North-Holland, Amsterdam.

[Hogan 1978] Hogan, W.W., "Energy modeling: building understanding for better use", Presented at the 2d Lawrence Symposium on the Systems and Decision Sciences, Berkeley CA, Oct 3, 1978.

[Holtzman 1985] Holtzman, S., Intelligent Decision Systems, Addison-Wesley, Reading MA, 1989. This book was available as an unpublished dissertation in 1985.

[Hopcroft and Ullman 1979] Hopcroft, J.E., Ullman, J.D., Introduction to Automata Theory, Languages, and Computation, Addison-Wesley, Reading, MA.

[Horvitz and Breese 1990] Horvitz, E.J., Breese, J.S., "Ideal partition of resources for metareasoning", Report KSL-90-26 #309.

[Howard 1965] Howard, R.A., "Bayesian decision models for systems engineering", IEEE Trans. Systems, Science, Cybern., SSC-1(1):36-40. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.837-841.

[Howard 1966a] Howard, R.A., "Decision Analysis: applied decision theory", Proceedings of the Fourth International Conference on Operations Research, John Wiley and Sons, New York, p.55-71. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.95-113.

[Howard 1966b] Howard, R.A., "Information value theory", IEEE Transactions on Systems Science and Cybernetics, SSC-2(1):22-26. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.777-784.

[Howard 1968] Howard, R.A., "The foundations of decision analysis", IEEE Trans. Sys. Sci. Cybern. SSC-4(3):211-219. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.579-589.

[Howard 1971] Howard, R.A., "Proximal decision analysis", Management Science, 17(9):507-541. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.843-879.

[Howard 1977] Howard, R.A., "Risk Preference", Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.629-663. Available in the 1977 edition of Principles and Applications of Decision Analysis, ed. Howard and Matheson.

[Howard 1983] Howard, R.A., "The evolution of decision analysis", Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.5-16.

[Howard 1988] Howard, R.A., "Decision analysis: practice and promise", Management Science, 34(6):679-695.

[Howard 1990] Howard, R.A., "From Influence to Relevance to Knowledge", Influence Diagrams, Belief Nets and Decision Analysis, ed. R.M.Oliver and J.Q.Smith, John Wiley & Sons, 1990, p.3-23.

[Howard 1992] Howard, R.A., "In praise of the Old-Time Religion", in Utility Theories: Measurement and Applications, ed. W.Edwards, Kluwer, Hingham, MA, p.27-55.

[Howard and Matheson 1976] Howard, R.A., Matheson, J.E., "Influence diagrams", Principles and Applications of Decision Analysis, ed. Howard and Matheson, p. 719-762. Previously available in Development of automated aids for decision analysis, Miller et alia, SRI International Technical Report 3309.

[Howard et alia 1975] Howard, R.A., Merkhofer, M.W., Miller, A.C., Tani, S.N., A Preliminary Characterization of a Decision Structuring Process for the Task Force Commander and His Staff, SRI International Technical Report, MSD-4030. Menlo Park CA.

[Huff 1954] Huff, D., How to Lie with Statistics, W.W. Norton & Company, Inc., New York.

[Huntington, Weyant and Sweeney 1982] Huntington, H.G., Weyant, J.P., Sweeney, J.L., "Modeling for Insights, not Numbers: the experience of the Energy Modeling Forum", Omega, The Int. Journal of Mgmt. Sci., 10(5):449-462.

[Isenberg 1984] Isenberg, D.J., "How senior managers think", Harvard Business Review, Nov-Dec 1984, v 62, p.81-90.

[Isenberg 1986a] Isenberg, D.J., "The structure and process of understanding: Implications for managerial action", The Thinking Organization, ed H.Sims and D.Gioia, Jossey-Bass, San Francisco, p.238-262.

[Isenberg 1986b] Isenberg, D.J., "Thinking and Managing : a verbal protocol analysis of managerial problem solving", Academy of Management Journal, 29(4):775-788.

[Jacob et alia 1989] Jacob, V., Moore, J., Whinston, A., "Analysis of human and computer decision-making capabilities", Information & Management, 16(5):247-255.

[Jamieson 1989] Jamieson, P.W., "New paradigm for explanation in medical causal reasoning systems", Proceedings - Annual Symposium on Computer Applications in Medical Care. Available from IEEE Service Cent (cat n 89TH0286-5), Piscataway, NJ, USA. p 17-21.

[Jimison 1988] Jimison, H.B., "Generating explanations of decision models based on an augmented representation of uncertainty", Uncertainty in Artificial Intelligence 4, 1990, ed. R.D.Shachter, et alia, Elsevier Science Publishers B.V. (North-Holland), p.351-365.

[Johnson-Laird 1980] Johnson-Laird, P.N., "Mental models in cognitive science", Cognitive Science, 4:71-115.

[Johnson-Laird 1988] Johnson-Laird, P.N., "Reasoning and common sense", in IEE Colloquium on 'Inference' (Digest No.37). London, UK. p. 1.

[Kahneman and Tversky 1979] Kahneman, D., Tversky, A., "Prospect theory: an analysis of decision under risk", Econometrica, 47(2):263-291.

[Kahneman and Tversky 1982] Kahneman, D., Tversky, A., "The simulation heuristic", in Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press.

[Kelly 1955] Kelly, G.A., The Psychology of Personal Constructs, 2 volumes. New York: Norton.

[Kettelhut 1989] Kettelhut, M.C., "Building consensus for funds allocation decisions through implementation of a decision support system", Hawaii International Conference on System Sciences 22(3):803-811.

[Kirkwood 1991] Kirkwood, C.W., "ADAM: an Algebraic Decision Analysis Modeling system", Technical Report DIS 90/91-6, Dept. of Decision and Information Systems, Arizona State University, Tempe AZ 85287-4206.

[Kitcher 1981] Kitcher, P., "Explanatory unification", Philosophy of Science 48:507-531. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.167.

[Klein and Shortliffe 1991] Klein, D.A., Shortliffe, E.H., "Interactive diagnosis and repair of decision-theoretic models", in Proceedings. Seventh IEEE Conference on Artificial Intelligence Applications (Cat. No.91CH2967-8). Los Alamitos, CA, USA: IEEE Comput. Soc. Press, 1991. p. 289-93.

[Klein and Calderwood 1987] Klein, G.A., Calderwood, R., "Decision models: some lessons from the field", Conference Date: 1987 Oct 20-23. IEEE, Systems, Man & Cybernetics Soc, New York, NY, USA. Available from IEEE Service Cent (Cat n 87CH2503-1), Piscataway, NJ. p 247-251.

[Klir 1985] Klir, G.J., Architecture of Systems Problem Solving, Plenum Press, New York.

[Kochen 1975] Kochen, M., "Applications of fuzzy sets in psychology", in Fuzzy Sets and Their Applications to Cognitive and Decision Processes, ed. L.A.Zadeh, et alia, Academic Press, New York, p. 395-408.

[Korhonen et alia 1990] Korhonen, P., Moskowitz, H., Wallenius, J., "Choice behavior in interactive multiple-criteria decision making", Annals of Operations Research, 23(1-4):161-79.

[Korsan 1990] Korsan, R.J., "Towards better assessment and sensitivity procedures", Influence Diagrams, Belief Nets and Decision Analysis, ed. R.M.Oliver and J.Q.Smith, John Wiley & Sons Ltd., p.427-455.

[Kosy and Wise 1984] Kosy, D.W., Wise, B.P., "Self-explanatory financial planning models", Proceedings of the National Conference on Artificial Intelligence, AAAI-84 Distributed by William Kaufmann Inc, Los Altos, CA, USA p 176-181.

[Kottemann 1986] Kottemann, J.E., "Some requirements and design aspects for the next generation of decision support systems", Hawaii International Conference on System Sciences 19(1a):339-344.

[Kottemann and Dolk 1988] Kottemann, J.E., Dolk, D.R., "Process-oriented model integration", Hawaii International Conference on System Sciences 21(3):396-402.

[Koussev et alia 1989] Koussev, T., Weiss, M.P., Reiss, K., "Graphic explanation environment for expert systems", Second International Conference on Software Engineering for Real Time Systems, Cirencester, Engl, IEE Conference Publication n 309, p 11-15.

[Krishnan 1989] Krishnan, R., "PDM: a knowledge-based tool for model construction", Hawaii International Conference on System Sciences 22(3):467-74.

[Kruskal 1964] Kruskal, J.B., "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis", Psychometrika, 29(1):1-27.

[Kukich 1985] Kukich, K., "The feasibility of automatic natural language report generation", Hawaii International Conference on System Sciences 18(1):546-556.

[Lagomasino and Sage 1985] Lagomasino, A., Sage, A.P., "Representation and interpretation of information for decision support with imperfect knowledge", Large Scale Systems, 9(3):169-191.

[Lakatos 1970] Lakatos, I., "Falsification and the methodology of scientific research programmes", in Criticism and the Growth of Knowledge, ed., I.Lakatos, A.Musgrave, Cambridge Univ. Press, p.91-197.

[Lane 1992] Lane, D.C., "Modelling as Learning: a consultancy methodology for enhancing learning in management teams", European Journal of Operational Research, 59(1):64-84.

[Langlotz 1989] Langlotz, C.P., A Decision-theoretic Approach to Heuristic Planning, dissertation, Stanford Univ.

[Langlotz et alia 1988] Langlotz, C.P., Shortliffe, E.H., Fagan, L.M., "A Methodology for Generating Computer-based Explanations of Decision-theoretic Advice", Medical Decision Making, 8(4):290-303.

[Laplace 1812] Laplace, P.S., "A Philosophical Essay on Probabilities", "Théorie analytique des probabilités", 1st ed 1812, 3d ed: Paris: Mme Ve Courcier, 1820. The 3d edition is translated as A Philosophical Essay on Probabilities, tr.: F.Truscott and F.Emory, Dover Publications, New York, 1952.

[Leal and Pearl 1977] Leal, A., Pearl, J., "An interactive program for conversational elicitation of decision structures", IEEE Transactions on Systems, Man, and Cybernetics, SMC-7(5):368-376.

[Lebailly et alia 1987] Lebailly, J., Martin-Clouaire, R., Prade, A., "Use of fuzzy logic in a rule-based system in petroleum geology", in Approximate Reasoning in Intelligent Systems, Decision and Control. Sanchez, E.; Zadeh, L.A. (editors) Oxford, UK: Pergamon,. p. 125-44.

[Lee 1993] Lee, K.C., "A cognitive map knowledge-based strategic planning simulation", Hawaii International Conference on System Sciences 26(3):249-267.

[Lenard 1986] Lenard, M.L., "Representing models as data", Hawaii International Conference on System Sciences 19(1a):389-396.

[Lenard 1987] Lenard, M.L., "An object-oriented approach to model management", Hawaii International Conference on System Sciences 20(1):509-515.

[Liang 1986] Liang, T., "A graph-based approach to model management", Proceedings of the 7th International Conference on Information Systems, San Diego, 428-429, 1986.

[Licker and Thompson 1985] Licker, P.S., Thompson, R.L., "Consulting systems: group decision support by one person", Hawaii International Conference on System Sciences 18(1):466-475.

[Ligeza 1988] Ligeza, A., "Expert systems approach to decision support", European Journal of Operational Research, 37(1):100-110.

[Lind 1986] Lind, M., "Decision Models and the design of knowledge based systems", Intelligent Decision Support in Process Environments, ed. E.Hollnagel et alia, Springer-Verlag, Berlin, p.197-210.

[Lindley 1982] Lindley, D.V., "Scoring rules and the inevitability of probability", International Statistical Review, 50(1):1-26.

[Liu and Tomsovic 1986] Liu, C.C., Tomsovic, K., "Expert system assisting decision-making of reactive power/voltage control", IEEE Trans Power Syst, PWRS-1(3):195-201.

[Lonergan 1957] Lonergan, B.J.F., Insight: a study of human understanding, Philosophical Library, New York.

[Longino 1990] Longino, H.E., Science as social knowledge: values and objectivity in scientific inquiry, Princeton University Press, Princeton, NJ.

[Lord and Foti 1986] Lord, R.G., Foti, R.J., "Schema theories, information processing, and organizational behavior", The Thinking Organization, ed. Henry P. Sims Jr. and Dennis A. Gioia, Josey-Bass, San Francisco, p.20.

[Luce and Raiffa 1957] Luce, R.D., Raiffa, H., Games and Decisions, Wiley.

[Luenberger 1979] Luenberger, D.G., Introduction to Dynamic Systems, Wiley, New York.

[Lusk and Hammond 1991] Lusk, C.M., Hammond, K.R., "Judgment in a dynamic task: Microburst forecasting", J. of Behavioral Decision Making, 4(1):55-73.

[Ma et alia 1989] Ma, P.C., Murphy, F.H., Stohr, E.A., "Representing knowledge about linear programming formulation", Annals of Operations Research, 21(1-4):149-72.

[Maeda and Murakami 1988] Maeda, H., Murakami, S., "A fuzzy decision-making method and its application to a company choice problem", Information Sciences, vol.45, no.2, p. 331-46.

[Malakooti 1988] Malakooti, B., "A decision support system and a heuristic interactive approach for solving discrete multiple criteria problems", IEEE Transactions on Systems, Man and Cybernetics, SMC-18(2):273-84.

[Manheim 1966] Manheim, M.L., Hierarchical Structure: a Model of Planning and Design Processes, MIT Press, Cambridge MA.

[Manheim et alia 1990] Manheim, M.L., Srivastava, S., Vlahos, N., Hsu, J., Jones, P., "A Symbiotic DSS for production planning and scheduling: issues and approaches", Hawaii International Conference on System Sciences 23(3):383-390.

[Manheim et alia 1991] Manheim, M.L., Srivastava, S., Vlahos, N., Tseng, C.P., "Working with an intelligent assistant: Experiments with a Symbiotic DSS for production planning and scheduling", Hawaii International Conference on System Sciences 24(3):386.

[Mannino, Greenberg and Hong 1990] Mannino, M., Greenberg, B., Hong, S.N., "Model libraries: knowledge representation and reasoning", ORSA Journal on Computing, 2(3):287-301.

[Matheson 1990] Matheson, J.E., "Using influence diagrams to value information and control", Influence Diagrams, Belief Nets and Decision Analysis, ed. R.M.Oliver and J.Q.Smith, John Wiley & Sons, 1990, p.3-23.

[Matheson and Howard 1968] Matheson, J.E., Howard, R.A., "An introduction to decision analysis", Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.17-55.

[McCloskey and Glucksberg 1978] McCloskey, M.E., Glucksberg, S., "Natural categories: well defined or fuzzy sets?", Memory and Cognition, 6(4):462-472.

[McCoy and Boys 1987] McCoy, M.S., Boys, R.M., "Human performance models applied to intelligent decision support systems", IEEE Proceedings of the National Aerospace and Electronics Conference 1987. Available from IEEE Service Cent (Cat n 87CH2450-5), Piscataway, NJ, USA p 1352-1359.

[McGovern et alia 1991] McGovern, J., Samson, D., Wirth, A., "Knowledge acquisition for intelligent decision systems", Decision Support Systems, 7:263-272.

[McNamee and Celona 1987] McNamee, P., Celona, J., Decision Analysis for the Professional with Supertree, The Scientific Press, Redwood City, CA.

[Mili 1988] Mili, F., "A framework for a Decision Critic and Advisor", Hawaii International Conference on System Sciences 21(3):381-386.

[Mili and Cioch 1990] Mili, F., Cioch, F.A., "Documenting decision models for informed and confident decisions", Hawaii International Conference on System Sciences 23(3), p.494-503.

[Mili and Szoke 1992] Mili, F., Szoke, I., "Assisted model selection, evaluation and comparison", Hawaii International Conference on System Sciences 25(3):485-93.

[Miller 1956] Miller, G., "The magical number seven, plus or minus two", Psychological Review, 63:81-97.

[Mitroff and Betz 1972] Mitroff, I.I., Betz, F., "Dialectical decision theory: a Meta-theory of decision making", Management Science, 19(1):11-24.

[Molokova 1986] Molokova, O.S., "Formation of an individual explanation in expert systems", Sov J Comput Syst Sci, 24(1):34-44.

[Moore et alia 1989] Moore, J.C., Richmond, W.B., Whinston, A.B., "An economic framework for computing", Hawaii International Conference on System Sciences 22(3):114-122.

[Moray 1987] Moray, N., "Intelligent aids, mental models, and the theory of machines", International Journal of Man-Machine Studies, 27(5-6):619-629.

[Morris et alia 1987] Morris, P., Sandling, M., Fancher, R., Kohn, M., Chao, H.P., Chapel, S., "Utility Fuel Inventory Model", Operations Research, 35(2):169-184.

[Muhanna and Pick 1988] Muhanna, W.A., Pick, R.A., "Composite Models in SYMMS", Hawaii International Conference on System Sciences 21(3):418-427.

[Murphy and Weiss 1990] Murphy, F.H., Weiss, H.J., "Approach to modeling electric utility capacity expansion planning", Naval Research Logistics v 37 n 6 Dec 1990 p 827-845.

[Mutalik et alia 1988] Mutalik, P., Fisher, P., Swett, H., Miller, P., "Structuring coherent explanation: The use of diagnostic strategies in an expert critiquing system", Proceedings - Annual Symposium on Computer Applications in Medical Care. Available from IEEE Service Cent (cat n 88CH2616-1), Piscataway, NJ, USA. p 26-31.

[Muto 1988] Muto, T., "Cognition and understanding: a cognitive science view", Journal of the Institute of Electronics, Information and Communication Engineers, 71(11):1123-5.

[Neches, Swartout and Moore 1985] Neches, R., Swartout, W., Moore, J., "Enhanced maintenance and explanation of expert systems through explicit models of their development", IEEE Transactions on Software Engineering, SE-11(11):1337-1351.

[Neustadter, Geoffrion et alia 1992] Neustadter, L., Geoffrion, A., Maturana, S., Tsai, Y., Vicuna, F., "The design and implementation of a prototype Structured Modeling environment", Annals of Operations Research, 38:453-484.

[Newell and Simon 1972] Newell, A., Simon, H.A., Human Problem Solving, Prentice-Hall, Englewood Cliffs NJ.

[Nosofsky 1992] Nosofsky, R.M., "Similarity scaling and cognitive process models", Annual Review of Psychology, 43:25-53.

[Novak 1987] Novak, V., "Automatic generation of verbal comments on results of mathematical modelling", Approximate Reasoning in Intelligent Systems, Decision and Control. Sanchez, E.; Zadeh, L.A. (editors), Pergamon, Oxford, UK. p. 55-68 .

[Oden 1977] Oden, G.C., "Integration of fuzzy logical information", Journal of Experimental Psychology: Human Perception and Performance, 3(4):565-575.

[Oden 1979] Oden, G.C., "Fuzzy propositional approach to psycholinguistic problems: an application of fuzzy set theory in cognitive science", in Advances in Fuzzy Set Theory and Applications, ed. M.M.Gupta, et alia, Elsevier North-Holland, New York.

[Ogasawara and Russell 1993] Ogasawara, G.H., Russell, S.J., "Planning using multiple execution architectures", Proceedings of the 13th International Joint Conference on AI, Chambery, France, 1993.

[Owen 1978] Owen, D.L., "The use of influence diagrams in structuring complex decision problems", Proceedings, Second Lawrence Symposium on Systems and Decision Sciences. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p. 763-771.

[Paradice and Courtney 1986] Paradice, D.B., Courtney, J.F., "A Kantian based system for business problem formulation", Hawaii International Conference on System Sciences 19(1):371-378.

[Paradice and Courtney 1988] Paradice, D.B., Courtney, J.F., "Dynamic construction of statistical models in managerial DSS", Annals of Operations Research, 12(1-4):321-36.

[Phillips and Edwards 1966] Phillips, L.D., Edwards, W., "Conservatism in a simple probability inference task", Journal of Experimental Psychology, 72:346-57.

[Pöyhönen et alia 1994] Pöyhönen, M.A., Hämäläinen, R., Salo, A., "An experiment on the numerical modeling of verbal ratio statements", research report A50, Systems Analysis Lab, Helsinki University of Technology, Espoo, Finland.

[Pracht and Courtney 1986] Pracht, W.E., Courtney, J.F., "Visual user interface for capturing mental models in model management systems", Hawaii International Conference on System Sciences 19(1a):535-545.

[Pratt, Raiffa and Schlaifer 1964] Pratt, J.W., Raiffa, H., Schlaifer, R., "The foundations of decision under uncertainty: an elementary exposition", J. Am. Statist. Assoc., 59:353-375.

[Prieto-Diaz and Freeman 1987] Prieto-Diaz, R., Freeman, P., "Classifying software for reusability", IEEE Software, 4(1):6-16.

[Prior and Moscardini 1989] Prior, D.E., Moscardini, A.O., "Mathematical insight through system dynamics", Transactions of the Institute of Measurement and Control, 11(4):196-207.

[Radford 1990] Radford, K.J., "Strategic/tactical model for resolution of complex decision situations", Information and Decision Technologies, 16(4):333-346.

[Raghavan 1990] Raghavan, S.A., "BIRBAL: a computer-based Devil's Advocate", Hawaii International Conference on System Sciences 23(3):391-402.

[Raghavan 1991] Raghavan, S.A., "KNOB: A Mutation-based approach for generating stereotype behaviors in active decision support systems", Hawaii International Conference on System Sciences 24(3):417.

[Raiffa 1968] Raiffa, H., <u>Decision Analysis: Introductory Lectures on Choices under Uncertainty</u>, Addison-Wesley, Reading MA.

[Raiffa and Schlaifer 1961] Raiffa, H., Schlaifer, R., <u>Applied Statistical Decision Theory</u>, MIT Press, Cambridge MA.

[Railton 1978] Railton, P., "A deductive-nomological model of probabilistic explanation", Philosophy of Science 45:206-226. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.75.

[Rasmussen 1985] Rasmussen, J., "Role of hierarchical knowledge representation in decisionmaking and system management", IEEE Transactions on Systems, Man and Cybernetics, SMC-15(2):234-243.

[Rennels et alia 1987] Rennels, G.D., Shortliffe, E.H., Miller, P.L., "Choice and explanation in medical management: a multiattribute model of artificial intelligence approaches", Medical Decision Making 7:22-31.

[Rios Insua and French 1991] Rios Insua, D., French, S., "A framework for sensitivity analysis in discrete multi-objective decision-making", European Journal of Operational Research, 54(2):176-90.

[Rothenberg et alia 1990] Rothenberg, J., Shapiro, N.Z., Hefley, C., "A 'Propagative' approach to sensitivity analysis", Proceedings of the AI, Simulation & Planning in High Autonomy Systems conference, ed. B.Zeigler and J.Rozenblit, IEEE Computer Society Press, Los Alamitos CA, 10-16.

[Rotmans and Vrieze 1990] Rotmans, J., Vrieze, O.J., "Metamodelling and experimental design: case study of the greenhouse effect", European Journal of Operational Research, 47(3):317-29.

[Rumelhart 1985] Rumelhart, D.E., "Schemata and the cognitive system", Handbook of Social Cognition, ed. R.S.Wyer and J.K.Srull, Erlbaum Assoc., Hillsdale NJ, 161-188.

[Ruokangas 1988] Ruokangas, C.C., "Real-time control for manufacturing Space Shuttle main engines: work in progress", in Fourth Conference on Artificial Intelligence for Space Applications (NASA Conf. Publ. 3013). Washington, DC. p. 5-18.

[Saaty 1980] Saaty, T.L., The Analytic Hierarchy Process, McGraw-Hill, New York.

[Saaty 1982] Saaty, T.L., Decision Making for Leaders, Lifetime Learning publications, a division of Wadsworth, Inc., Belmont, Calif.

[Saaty 1986] Saaty, T.L., "Axiomatic foundation of the analytic hierarchy process", Management Science, 32(7):841-855.

[Sakawa and Yano 1989] Sakawa, M., Yano, H., "Interactive decision making for multiobjective nonlinear programming problems with fuzzy parameters", Fuzzy Sets and Systems (24 Feb. 1989) vol.29, no.3, p. 315-26.

[Salmon 1971] Salmon, W.C., "Statistical explanation and causality", from Statistical Explanation and Statistical Relevance, W.Salmon, U.of Pittsburgh Press, 1971. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.75.

[Salmon 1984] Salmon, W.C., Scientific Explanation and the Causal Structure of the World, Princeton University Press, Princeton NJ.

[Salo and Hämäläinen 1993] Salo, A.A., Hämäläinen, R.P., "On the measurement of preferences in the Analytic Hierarchy Process", research report A47, Systems Analysis Laboratory, Helsinki University of Technology, Espoo, Finland.

[Savage 1954] Savage, L.J., The Foundations of Statistics, John Wiley. 2d ed, Dover, 1972.

[Schmucker 1984] Schmucker, K.J., Fuzzy Sets, Natural Language Computations, and Risk Analysis, Computer Science Press, Rockville MD.

[Schoppers 1991] Schoppers, M., 1991, "Editorial introduction: Special Section on real time knowledge based control systems", CACM, August 1991.

[Schrattenholzer 1985] Schrattenholzer, L., "Experience with the operation of an energy model set", Bridge Between Control Science and Technology, Proceedings of the Ninth Triennial World Congress of IFAC 1984. IFAC Proceedings Series 1985, Pergamon Press, Oxford, Engl. n 6 p 3265-3268.

[Schutzelaars 1990] Schutzelaars, A.A.J.H., "SECDS: a Simulation Environment for Complex Dynamic Systems", in Proceedings of the 1990 Summer Computer Simulation Conference. Svrcek, B.; McRae, J. (editors) San Diego, CA, USA. p. 860-6.

[Scriven 1988] Scriven, M., "Explanations, predictions, and laws", Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.51.

[Sein and Bostrom 1990] Sein, M.K., Bostrom, R.P., "An experimental investigation of the role and nature of mental models in the learning of desktop systems", in Desktop Information Technology. Proceedings of the IFIP WG 8.2 Working Conference on Desktop Information Technology and Organizational Worklife in the 1990's. Kaiser, K.M.; Oppelland, H.J. (editors) Amsterdam, Netherlands: North-Holland. p. 253-76.

[Selfridge et alia 1985] Selfridge, M., Daniell, J., Simmons, D., "Learning causal models by understanding real-world natural language explanations", Second Conference on Artificial Intelligence Applications: The Engineering of Knowledge-Based Systems. Available from IEEE Service Cent (Cat n 85CH2215-2), Piscataway, NJ, USA p 378-383.

[Sellars 1963] Sellars, W., "Theoretical explanation", Philosophy of Science: the Delaware Seminar 1963, vol 2, University of Delaware. Reprinted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.156.

[Sember and Zukerman 1989] Sember, P., Zukerman, I., "Strategies for generating micro explanations for Bayesian belief networks", Proceedings of the Fifth Workshop on Uncertainty in Artificial Intelligence, p.295.

[Shachter 1986] Shachter, R.D., "Evaluating influence diagrams", Operations Research, 34(6):871-882.

[Shachter and Heckerman 1987] Shachter, Ross D.; Heckerman, D.E., "Thinking backward for knowledge acquisition", AI Magazine, 8(3):55-61.

[Shafer 1989] "Beyond Subjective Expected Utility", presented to the Decision Analysis Colloquium, Decision and Ethics Center, Stanford University, Stanford CA, unpublished.

[Shannon 1975] Shannon, R.E., Systems Simulation, Prentice-Hall, Englewood Cliffs, NJ.

[Shepard 1957] Shepard, R.N., "Stimulus and response generalization: a stochastic model relating generalization to distance in psychological space", Psychometrika, 22:325-45.

[Shepard 1987] Shepard, R.N., "Toward a universal law of generalization for psychological science", Science 237:1317-23.

[Shoval 1986] Shoval, P., "Comparison of decision support strategies in expert consultation systems", International Journal of Man-Machine Studies, 24(2):125-139.

[Silverman 1991] Silverman, B.G., "Expert critics: operationalizing the judgement/decisionmaking literature as a theory of 'bugs' and repair strategies", Knowledge Acquisition, 3(2):175-214.

[Simon 1955] Simon, H.A., "A Behavioral Model of Rational Choice", Quarterly J. of Economics, vol.66, pp.99-118.

[Simon 1960] Simon, H.A., The New Science of Management Decision, Prentice-Hall, Englewood Cliffs NJ. 3d edition, 1977.

[Simon 1962] Simon, H.A., "The architecture of complexity", Proceedings of the American Philosophical Society, 106:467-482.

[Simon 1969] Simon, H.A., The Sciences of the Artificial, MIT Press, Cambridge Mass., 2d ed. 1981.

[Simon 1979] Simon, H.A., Models of Thought, vol 1, Yale University Press, New Haven.

[Simon 1982] Simon, H.A., Models of Bounded Rationality: Behavioral Economics and Business Organization, MIT Press, Cambridge Mass. This is volume 2 of 2.

[Simon 1983] Simon, H.A., "Search and reasoning in problem solving", Artificial Intelligence, 21:7-29.

[Simon and Lea 1974] Simon, H.A., Lea, G., "Problem Solving and Rule Induction", in Models of Thought, ed. H.A.Simon, Yale University Press, New Haven, 1979.

[Sinha et alia 1984] Sinha, A., Caplan, L., Desai, B., Evens, M., Hier, D., Hill, H., "Making decision support systems self explanatory", Proceedings - IEEE Computer Society's Eighth International Computer Software & Application's Conference, COMPSAC 84. Available from IEEE Service Cent (Cat n 84CH2096-6), Piscataway, NJ, USA p 358-362.

[Slovic and Lichtenstein 1971] Slovic, P., Lichtenstein, S., "Comparison of Bayesian and regression approaches to the study of information processing in judgment", Organizational Behavior and Human Performance, 6:649-744.

[Smets 1988] Smets, P., "Belief functions versus probability functions", Uncertainty and Intelligent Systems. 2nd International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems: IPMU '88. ed. B.Bouchon et alia, Springer-Verlag, Berlin, 1988. p. 17-24.

[Smith, Holtzman and Matheson 1993] Smith, J.E., Holtzman, S., Matheson, J.E., "Structuring conditional relationships in influence diagrams", Operations Research, 41(2):280-297.

[Smyth 1993] Smyth, P., "Neural Networks", Normative Systems conference, University of Southern California, unpublished, 1993.

[Spetzler and Staël von Holstein 1975] Spetzler, C.S., Staël von Holstein, C.A.S., "Probability encoding in decision analysis", Management Science, 22:340-358. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.601-626.

[Srinivas et alia 1989] Srinivas, S., Russell, S., Agogino, A., "Automated construction of sparse Bayesian networks from unstructured probabilistic models and domain information", Uncertainty in Artificial Intelligence 5, ed. M.Henrion, et alia, North-Holland, p.295-308.

[Staël von Holstein 1971] Staël von Holstein, C.A.S., "A tutorial in decision analysis", Presented at the Third Research Conference on Subjective Probability, Utility, and Decision Making, London, Sept. 1971. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson p.129-157.

[Sticha et alia 1979] Sticha, P.J., Weiss, J.J., Donnell, M.L., "Evaluation and integration of imprecise information", Technical report, Decisions and Designs, Inc., McLean, Virginia, August 1979.

[Suermondt 1992] Suermondt, H.J., Explanation in Bayesian belief networks, dissertation, Stanford Univ.

[Swartout 1985] Swartout, W.R., "Knowledge needed for expert system explanation", AFIPS Conference Proceedings v 54: 1985 National Computer Conference. Publ by AFIPS Press, Reston, VA, USA p 93-98.

[Swartout and Smoliar 1987] Swartout, W.R., Smoliar, S.W., "Explaining the link between causal reasoning and expert behavior", Proceedings - Annual Symposium on Computer Applications in Medical Care 11th. Available from IEEE Service Cent (Cat n 87CH2446-3), Piscataway, NJ, USA p 37-42.

[Synfuels Interagency Task Force 1975] Synfuels Interagency Task Force, Recommendations for a Synthetic fuels Commercialization Program, U.S. G.P.O. Su.Docs stock number 041-001-00112-3, 041-001-00111-3, 041-001-00114-0, 040-000-00340-1.

[Tani 1978] Tani, S.N., "Decision Analysis of the synthetic fuels commercialization program", National Computer Conference, AFIPS v 47, p.23-29. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, 435-443.

[Tanniru and Murray 1987] Tanniru, M.R., Murray, T.J., "Model validation in a DSS environment", Hawaii International Conference on System Sciences 20(1):655-664.

[Taylor and Graves 1991] Taylor, G.D., Graves, R.J., "Integrated decision-making in a flexible assembly system: sensitivity analysis and extended testing", Production Planning and Control, 2(4):335-46.

[Thomas and Samson 1986] Thomas, H., Samson, D., "Subjective aspects of the art of decision analysis: exploring the role of decision analysis in decision structuring, decision support and policy dialogue", Journal of the Operational Research Society, 37(3):249-265.

[Tse and Fehling 1989] Tse, E., Fehling, M., "An intelligent control system", Proceedings, IEEE International Symposium on Intelligent Control, IEEE Comput. Soc. Press, p.194-199.

[Tversky 1969] Tversky, A., "Intransitivity of preferences", Psychological Review, 76(1):31-48.

[Tversky and Kahneman 1974] Tversky, A., Kahneman, D., "Judgment under Uncertainty: heuristics and biases", Science, 185:1124-1131. Reprinted in Principles and Applications of Decision Analysis, ed. Howard and Matheson, p.901-910.

[van Fraassen 1980] van Fraassen, B.C., "The pragmatic theory of explanation", The Scientific Image, by Bas van Fraassen, Oxford University Press, 1980. Excerpted in Theories of Explanation, ed. J.C.Pitt, Oxford University Press, New York, 1988, p.136.

[Vari and Vecsenyi 1984] Vari, A.; Vecsenyi, J., "Selecting decision support methods in organisations", 1984 Journal of Applied Systems Analysis, 11:23-36.

[von Neumann and Morgenstern 1944] von Neumann, J., Morgenstern, O., Theory of Games and Economic Behavior, Princeton University Press, Princeton NJ, second edition, 1947.

[Wack 1985] Wack, P., "Scenarios: Uncharted waters ahead", Harvard Business Review, Sept-Oct 1985, p.73-89, and also its companion article "Scenarios: Shooting the rapids", Harvard Business Review, Nov-Dec 1985, p.139-150.

[Wagner and Wechsung 1986] Wagner, K., Wechsung, G., Computational Complexity, D. Reidel Publishing Co, Dordrecht.

[Waldspurger et alia 1992] Waldspurger, C., Hogg, T., Huberman, B., Kephart, J., Stornetta, W., "Spawn: a distributed computational economy", IEEE Transactions on Software Engineering, 18(2):103-117.

[Wallsten et alia 1986] Wallsten, T., Budescu, D., Rapoport, A., Zwick, R., Forsythe, B., "Measuring the vague meanings of probability terms", J. Experimental Psychology: General, 115(4):348-365.

[Wang et alia 1988] Wang, S.P., Xia, A.B., Zhao, Z.X., "Model building and management system using predicate calculus and relational framework", Proc 1988 IEEE Int Conf Syst Man Cybern v1 (of 2) p 219-222.

[Watkins et alia 1992] Watkins, P.R., Lin, T.W., O'Leary, D.E., "AI integration for enhanced decision support", Hawaii International Conference on System Sciences 25(3):133-144.

[Waugh and Norman 1965] Waugh, N., Norman, D., "Primary memory", Psychological Review, 72:89-104.

[Weber and Konsynski 1987] Weber, E.S., Konsynski, B.R., "Problem management: neglected elements in DSS", Hawaii International Conference on System Sciences 20(1):774-781.

[Wellman et alia 1989] Wellman, M., Eckman, M., Fleming, C., Marshall, S., Sonnenberg, F., Pauker, S.G., "Automated critiquing of medical decision trees", Medical Decision Making, 9:272-284.

[Wellman 1990] Wellman, M.P., Formulation of tradeoffs in planning under uncertainty, Morgan Kaufmann, San Mateo, CA.

[Weyant 1990] Weyant, J.P., "Application of models in the process of legislation", Energy, 15(3/4):187-201.

[Wick and Slagle 1989] Wick, M.R., Slagle, J.R., "Explanation facility for today's expert systems", IEEE Expert, 4(1):26-36.

[Wick and Thompson 1989] Wick, M.R., Thompson, W.B., "Reconstructive expert system explanation", "Reconstructive expert system explanation", Artificial Intelligence, 54 (1992) 33-70.

[Will 1975] Will, H.J., "Model management systems", Information Systems and Organizational Structure, p. 467-482, E.Grochla and N.Szyperski (eds). Walter de Gruyter, Berlin.

[Wittgenstein 1945] Wittgenstein, L., Philosophical Investigations, translated to English by G.E.M.Anscombe and R.Rhees, 3d ed., 1958, Macmillan, New York.

[Yu 1990] Yu, O.S., "Interface between mental and computer models", Energy, 15(7/8):621-629.

[Zadeh 1965] Zadeh, L.A., "Fuzzy sets", Information and Control, 8(3):338-353.

[Zadeh 1968] Zadeh, L.A., "Probability measures of fuzzy events", Journal of Mathematical Analysis and Applications, 23(421-427).

[Zadeh 1975] Zadeh, L.A., "The concept of a linguistic variable and its application to approximate reasoning", Information Sciences, 8(3):199-249, 8(4):301-357, 9(1):43-80.

[Zwick and Wallsten 1989] Zwick, R., Wallsten, T.S., "Combining stochastic uncertainty and linguistic inexactness: theory and experimental evaluation of four fuzzy probability models", International Journal of Man-Machine Studies, 30(1):69-111.

Glossary

ancestor: predecessor of a node or its ancestors

atomic distribution: a probability distribution over all possible outcomes of a node under the given conditions

base case value: a value of a variable that is representative of all the variable's possible values under the circumstances; frequently taken to be the conditional median. Sometimes called conditional base case.

belief network: a directed graph whose nodes represent uncontrollable variables, and whose arrows indicate probabilistic conditioning. Also called relevance diagram or Bayes network.

chance node: event node that is not deterministic

conditional base case: base case

conditional base choice: choice for a decision thought to be reasonable under given conditions

conditional deterministic sensitivity analysis: a procedure defined in this dissertation, which accomplishes a similar function to deterministic sensitivity analysis, but makes explicit the fact that some variables in the analysis are to be treated as auxiliary to the calculation of utility and allowed to vary

conditioning tree: a graphical structure like a decision tree, listing all possible outcome of each conditioning variable, in turn, without distinguishing decisions from uncertainties or specifying probabilities

decision network: a directed graph whose nodes represent both controllable and uncontrollable variables, and whose arrows indicate informational and probabilistic conditioning, respectively. Also called decision diagram.

decision node: decision network node that contains a set of alternatives decision-making (DM) variables: variables of interest for a specific decision

delta: the difference of a target variable (typically profit or utility) under two specified conditions (typically the status quo and some other alternative)

descendent: successor of a node or its descendents

deterministic node: event node whose distributions are all degenerate (i.e., the value of the corresponding variable is completely determined in each condition)

discretization: approximation of a continuous distribution by a discrete one

dist node: node with no predecessors, containing one atomic distribution from a node being structurally decomposed

distinction: a division of experience into a set of degrees of instantiation of a feature

distribution tree: a conditioning tree with an atomic distribution at the end of each of its branches

event node: decision network node that contains the decision-maker's conditional probability distributions over the outcomes of an event

factor: input variable whose effects on a target variable of a system is to be measured

Howard canonical form: having no probabilistic conditioning of an uncertain node (i.e., having only structural and informational arrows)

influence arrow: arrow from a decision to an event node

influence diagram: collective term for belief and decision networks.

informational conditioning: the condition where the outcome of an event will be known when a decision is to be made, thus allowing the optimum decision-making policy to be specified conditional on the event's outcome

input variable: a variable conceived of as specifying the conditions that control the behavior of a system

log node: node with degenerate atomic distributiosn, whose purpose is to capture only the logic of a node's distribution tree lottery: a prospect about which one's uncertainty regarding specific unresolved events is formally stipulated. This has also been called a "deal" by authors attempting to avoid the unsavory connotations of gambling.

nominal value: base case value

outcome: the state of the world, understood and characterized according to a specific set of distinctions.

output variable: a variable conceived of as being the result of the behavior of a system

predecessor: a variable upon which the probabilities of a variable's distribution are specified. All predecessors and successors referred to in this paper are "direct" predecessors or successors. What might be called indirect predecessors or successors are called ancestors and descendents in this paper.

probabilistic conditioning: specification of the distribution of a variable by reference to another, conditioning, variable.

prospect: an outcome. This word is used to emphasize that an outcome may contain within it substantial uncertainty regarding distinctions not currently under discussion.

scenario: a sequence of events that is likely to ensue from a set of initial conditions

simulation model: a piece of software that realizes a scenario (or a probability distribution over scenarios) from a set of inputs as a basis for generating its outputs. Under this definition, portions of a computer program that do not realize the scenario, such as condition-specification logic, or output writers, are not part of the model.

standing model: a simulation model that is maintained and re-applied to multiple analyses

state variables: a minimal set of variables that fully specifies the state of a simulation

structural conditioning: a variant of probabilistic conditioning in which all conditional probabilities are zero or one

structurally decompose (a node): replace a node with a deterministic log node that handles the logic of its conditioning tree and one dist node for each atomic distribution in the node

successor: a variable whose distribution of outcomes is specified conditionally swing: the amount of change induced in the target variable by changes to the factor in question

target: an output variable whose response to perturbations of input variables is to be mesaured

target: the variable whose responses is being measured in sensitivity analysis

tornado diagram: a figure displaying the sensitivity of a target variable to individual variations of various factors in the model

utility node: decision network node that specifies the decision maker's utility function

utility: an output variable that the DM wishes to optimize. Some authors call this the objective function, or the value node of a decision network.

value of control (of a chance node): the most one should be willing to pay to an omnipotent wizard to specify its value as a function of its predecessors

value of information (of a chance node): the most one should be willing to pay an omniscient clairvoyant to reveal its value before a given decision is made

value: (of a variable) the realized number calculated from the underlying situation and represented by a variable. Also called a degree of a distinction or a level of a feature.

variable: a (typically numerical) representation of the degree to which a distinction is instantiated. Sometimes also called a random variable, or measure