AD-A081 873 APPLIED DECISION ANALYSIS INC MENLO PARK CA F/G 5/1
ANALYTIC PROCEDURES FOR EVALUATING DECISION AIDS.(U)
FEB 80 A C MILLER, P A MORRIS, R D SMALLWOOD N00014-78-C-0240
UNCLASSIFIED NL

1 OF 1
AD A
081873

END
DATE
FILMED
4-80
DTIC

(12) LEVEL II

ADA
APPLIED DECISION ANALYSIS, inc.
3000 Sand Hill Road, Menlo Park, CA 94025 (415) 854-7101

February 1980

PROGRESS REPORT

ANALYTIC PROCEDURES FOR EVALUATING DECISION AIDS

Prepared for:

Dr. Martin A. Tolcott
Code 455
Director, Engineering Psychology Program
Office of Naval Research
Arlington, Virginia

Prepared by:

Allen C. Miller
Peter A. Morris
Richard D. Smallwood
Ronald K. Langford
Robert S. Gibbons

Table of Contents

FIGURES

TABLES

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (*When Data Entered*)

REPORT DOCUMENTATION PAGE		**READ INSTRUCTIONS BEFORE COMPLETING FORM**
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (*and Subtitle*) Analytic Procedures for Evaluating Decision Aids		5. TYPE OF REPORT & PERIOD COVERED Progress Report 4/1/79 - 11/30/79
		6. PERFORMING ORG. REPORT NUMBER ADA Project #2017
7. AUTHOR(s) Allen C. Miller, III Ronald K. Langford Peter A. Morris Robert S. Gibbons Richard D. Smallwood		8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0240
9. PERFORMING ORGANIZATION NAME AND ADDRESS Applied Decision Analysis, Inc. 3000 Sand Hill Road Menlo Park, CA 94025		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		12. REPORT DATE February 1980
		13. NUMBER OF PAGES 61
14. MONITORING AGENCY NAME & ADDRESS(*if different from Controlling Office*) Defense Contract Administration Services Management Area-San Francisco San Bruno, California 94066		15. SECURITY CLASS (*of this report*) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (*of this Report*)

Unlimited

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

17. DISTRIBUTION STATEMENT (*of the abstract entered in Block 20, if different from Report*)

Unlimited

18. SUPPLEMENTARY NOTES

Reproduction in whole or in part is permitted for any purpose of the United States Government.

19. KEY WORDS (*Continue on reverse side if necessary and identify by block number*)

Automated Decision Aids
Decision Characteristics
Tactical Command and Control Decisions
Decision Taxonomy

20. ABSTRACT (*Continue on reverse side if necessary and identify by block number*)

This report describes research conducted by Applied Decision Analysis (ADA) for the Office of Naval Research (ONR) during the period from April to November 1979. The purpose of this research is to extend the analytical characterization of Navy tactical command and control decisions developed by ADA during the previous phase of this research, develop a practical method for evaluating and improving existing and proposed decision aids, and (continued)

apply it to the identification and design of aids for anti-air warfare (AAW) decisions at the unit and task-force levels. In addition, several analytical procedures have been developed to measure formally some of the characteristics and capabilities of decision aids that are outlined in the taxonomies. In particular, a method of measuring the credibility classes of decision aids is described in this report.

☒
☐
☐

Codes
and/or
Special
A

1. SUMMARY

This report describes research conducted by Applied Decision Analysis (ADA) for the Office of Naval Research (ONR) during the period from April to November 1979. The purpose of this research is to extend the analytical characterization of Navy tactical command and control decisions developed by ADA during the previous phase of this research, develop a practical method for evaluating and improving existing and proposed decision aids, and apply it to the identification and design of aids for anti-air warfare (AAW) decisions at the unit and task-force levels. In addition, several analytical procedures have been developed to measure formally some of the characteristics and capabilities of decision aids that are outlined in the taxonomies. In particular, a method of measuring the credibility classes of decision aids is described in this report.

The approach described in this report uses several taxonomies to describe: the characteristics of various command and control decisions, the characteristics of decision aids that could be used for command and control, and the information processing and analytic functions provided by people and decision aids to help a commander reach a decision. By comparing the functions needed to analyze a particular decision with those that can be provided by a typical decision-maker and his staff, we can determine the functional requirements for an appropriate decision aid. These requirements are used to evaluate existing aids or guide the development of new aids. Decision aids are viewed as sources of information that complement, rather than replace, intuition and judgment. This approach is being applied to the range of decisions encountered in anti-air warefare, and several areas have been identified where AAW decisions could benefit from improved aids. These areas are described briefly in this report and will be discussed in greater detail in a technical report that will be published at the completion of this research.

The research described here is in progress, and all of the taxonomies and analytical results discussed in this report are subject to revision. Many of the results are preliminary and incomplete. They build on the taxonomies developed during the previous phase of this research, and it is assumed that the reader is familiar with the technical report prepared by ADA at the end of that phase. See (7). Comments or suggestions concerning this research are welcome and should be sent to the authors.

2. THE USE OF TAXONOMIES AND ANALYTICAL PROCEDURES TO EVALUATE DECISION AIDS

One of the objectives of this research is to determine whether taxonomies can be used meaningfully to evaluate existing and proposed decision aids for Navy tactical command and control. In the previous phase of this research, we developed several related taxonomies to help someone think logically about the many aspects of decision situations and decision aids. This report describes an additional taxonomy of the functions required of an aid or decision-maker to analyze and make decisions, and a process for using the taxonomies to arrive at a logical framework for describing appropriate aids for various classes of decisions.

However, in the course of our research we have found the need for more precise analytical procedures than taxonomies to evaluate at least some of the characteristics of decision aids. While these procedures are more difficult to develop than the taxonomies, they are potentially better measures of the value of decision aids. Several of these analytical approaches are discussed in this report. At present, they are not fully developed, and deal with only a subset of the issues addressed in the taxonomies. However, they are already being used to clarify concepts used in the taxonomies, and they show promise for augmenting or replacing portions of the taxonomies.

Taxonomies represent a general conceptual approach for thinking about decision aids. They have several important advantages that guarantee that they will remain useful even after more precise evaluation procedures have been developed. Taxonomies provide a comprehensive checklist to ensure that all relevant aspects of an aid and the decisions it supports are considered. They can also organize the evaluation process into a series of logical steps by showing the sequence in which issues should be considered and compared. Taxonomies also provide a framework for decomposing general characteristics of decision situations and aids into basic elements that can be assessed or measured. This decomposition provides a guide for establishing the relative importance of desired characteristics. Finally, one of the major strengths of taxonomies is that they do not require precise, unambiguous assessments of the characteristics they contain, thus making it easy to use them.

However, the ambiguity permitted by a taxonomy is also one of its major weaknesses. It may not be possible to interpret correctly imprecise assessments of an aid's characteristics or capabilities, or to determine the implications of these assessments for the applicability of the aid in various decision situations. For instance, what does it mean for the usefulness of an aid if someone says an aid is "reliable" or a decision is "risky"? Taxonomies can provide a guide to evaluating aids, but not an operational measure of its value. We need an explicit, reproducible, and logically sound method for processing the information derived from a taxonomy to determine how much should be spent developing a particular aid.

The ambiguity inherent in taxonomies also means that they are not unique. Different taxonomies can be used to decompose and organize the same characteristics of a decision or an aid. The adequacy of a particular taxonomy is a matter of judgement.

While taxonomies provide a useful and flexible framework for conducting a general evaluation of decision aids, they should be augmented by more precise analytical procedures to guide specific resource allocations. Unfortunately, the analytical procedures and concepts needed to evaluate decision aids are not well developed. For instance, we do not have a basic theory to describe the value of decision aiding, analysis, and modeling, although several attempts have been made to develop one. (A survey of some of the work in this area is outlined in Chapter 5 of this report.) In addition, there is no formal measure of the credibility of decision aids, although it is clear that an aid must be credible to be useful to a decision-maker. (The issue of credibility is discussed in detail in Chapter 6 of this report.) Analytic l procedures such as these are at least as difficult to develop as taxonomies, but they provide an unambiguous and reproduceable way to evaluate at least some aspects of decision aids. Both taxonomies and more formal analytical procedures rely on subjective assessments of the properties of decision situations and aids, but taxonomies also require a subjective integration of these assessments to evaluate an aid. Analytical procedures rely on an explicit and testable set of logic rather than the intuition of the evaluator.

We anticipate that taxonomies and analytical procedures eventually will complement each other, and will be used together to evaluate decision aids. An evaluation procedure using both would probably start with taxonomies to make sure the important issues or features associated with an aid are considered and to identify the key tradeoffs that must be made. Then one or more rigorous analytical procedures would be used to produce formal measures of an aid's potential value. Finally, general evaluation procedures based on taxonomies would be employed to deal with aspects of an aid or decision situation that are poorly defined or for which no formal evaluation procedure exists. The question is not whether it is best to use taxonomies or more formal analytical procedures, but rather how we can integrate the results of both approaches. (See Chapter 6 of this report for the initial elements of a theory dealing with the use of formal analysis to update more intuitive logic.)

The following chapter describes an evaluation procedure based on taxonomies alone. This procedure will be revised and expanded as analytical procedures for evaluating decision aids become available.

3. A PROCEDURE FOR EVALUATING DECISION AIDS BASED ON A SET OF TAXONOMIES

Different taxonomies are needed to deal with various aspects of decision making and aiding. In particular, different taxonomies are appropriate for describing the features of decision situations, decision-makers, and decision aids. Additional taxonomies have been developed for the types of decision situations encountered in tactical command and control, and the analytical methods appropriate for different classes of decisions. See (4) and (8).

The previous phase of this research produced three closely related taxonomies. The first taxonomy groups Navy tactical command and control decisions into various categories. Representative decisions can be specified for each category, although the categories are sufficiently broad to include dissimilar decision situations. As a result, this taxonomy is useful mainly for identifying the range of decisions an aid may have to support. The second taxonomy contains a set of attributes or analytic measures that can be used to characterize a decision situation. These attributes are specified in both technical and intuitive terms, and a scale is defined for assessing each attribute. These attributes include: the number of decision strategies available to the decision-maker, the level of resources involved, the amount of information available and the time period within which the decision must be made. Table 1 contains a summary of the taxonomy of decision characteristics. The third taxonomy, which is not as fully developed as the others, specifies the characteristics of decision-aiding systems in terms that can be compared with the corresponding characteristics of the decisions for which they will be used. This taxonomy deals primarily with an aid's physical

TABLE 1: TAXONOMY OF DECISION CHARACTERISTICS

1. The Decision-maker's Resources

   Equivalent force level controlled by the decision-maker

2. The Importance of Decision to Decision-maker

   Importance of the decision relative to the other decisions made by the decision-maker over one year.

3. The Number of Decision Strategies

   Number of alternatives available in the primary decision problem (i.e., first decision node in a decision tree)

4. The Number of Significant Factors

   Number of factors (i.e., state variables) that could have a significant impact on the outcome

5. The Number of Outcome Attributes

   Minimum number of outcome variables (i.e., attributes that must be considered to adequately represent outcomes

6. Outcome Measurability

   Percentage of outcome variables (attributes) requiring a subjective scale

7. Contingent Decisions

   The importance of contingent planning in this situation

8. Probabilistic Dependence

   Average or typical number of variables that have a significant and direct impact on each outcome variable

TABLE 1: TAXONOMY OF DECISION CHARACTERISTICS (Cont'd.)

9. The Degree of Risk

Importance of low-probability, high-consequence events

10. Review and Approval

Extent of required review and approval

11. Structural Uniqueness

Extent to which existing plans or procedures can be used to deal with the decision

12. The Quantity of Information

Number of messages related to decision received per day

13. The Variability of Information Value

Percentage of messages that are significantly more valuable than the average

14. The Reliability of Information Sources

Percentage of information sources considered reliable

15. The Time Available for the Decision

Time from recognition of a decision to the point where an action must be taken

16. The Frequency of Decision

Mean time between recurrence of the decision characteristics and capabilities, such as its cost, support requirements, and data processing and storage capabilities. Table 2 contains a preliminary version of the taxonomy of decision aid characteristics.

The taxonomies of decision characteristics and decision-aid characteristics are closely related. Each decision characteristic has implications for the features of a decision aid that would be appropriate for that decision. A preliminary mapping between the two taxonomies was developed and demonstrated during the previous phase of this research. This mapping is shown in Table 3. For example, the decision characteristic measuring the time available for a decision has an effect on four of the aid characteristics: cost of use, reliability, data processing capability, and complexity of the user interface. Working in the other direction, the decision aid characteristic measuring the cost of use is related to three decision characteristics: importance of the decision to the decision-maker, structural uniqueness, and the time available for a decision.

Table 3 can be used as a checklist to help someone explore all of the implications of a decision situation for the capabilities required of an aid to support that decision. An example of the procedure for evaluating a hypothetical aid using the taxonomies is contained in the previous report. However, the linkages specified in Table 3 are not well defined and the evaluation process requires a considerable amount of judgement on the part of the user. It is doubtful that two people attempting to evaluate an aid for the same decision situation would interpret all of the relationships specified in Table 3 as having the same meaning or importance.

To help overcome this problem, an additional taxonomy has been developed to describe the information processing and analytical functions that are provided by a decision aid and needed to examine

TABLE 2: A PRELIMINARY TAXONOMY OF THE CHARACTERISTICS OF DECISION AIDS

1. Unit cost of the aid, including a proportional share of the development costs.
2. Cost of using the aid, including the level of effort required to use or program it.
3. Support requirements, including data sources, other aids and equipment, and physical space.
4. Reliability, including redundancy, self-monitoring, and a capability for graceful degradation.
5. Data processing capability, as measured by the number of calculations and amount of data that can be processed per unit time.
6. Data storage capability, including both rapid-access and slow-access storage.
7. Capability to maintain data security, including access control and encoding.
8. Data verification capability, including error checking and cross checking data from multiple sources.
9. Communications capability, including data transmission rates and number of communication channels.
10. Ability to prioritize its own operations, including the ability to interrupt and restart a procedure.
11. Facilities for testing and updating algorithms, as measured by the ease with which procedures or analyses can be restructured.
12. Ability to monitor and update information, including information sorting and screening.
13. Complexity of the aid-user interface, including the level of training required of the user, the sophistication of the algorithm, and the extent to which the aid summarizes and supplies its outputs.
14. Compatibility with existing systems and procedures.

TABLE 3

RELATIONSHIP BETWEEN DECISION CHARACTERISTICS AND DECISION-AID CHARACTERISTICS

Aid Characteristics Decision Characteristics	1. Unit cost	2. Cost of use	3. Support requirements	4. Reliability	5. Data processing capability	6. Data storage capability	7. Data security	8. Data verification	9. Communications capability	10. Prioritized operations	11. Test and update algorithms	12. Monitor and update info.	13. User interface	14. Compatibility
1. DM's resources	✓		✓											
2. Importance of decision to DM		✓		✓			✓							✓
3. # of decision strategies					✓	✓								
4. # of significant factors					✓	✓								
5. # of outcome attributes													✓	
6. Outcome measurability													✓	
7. Contingent decisions					✓	✓				✓		✓		
8. Probabilistic dependence					✓	✓								
9. Degree of risk												✓		
10. Review and approval			✓			✓							✓	
11. Structural uniqueness		✓									✓		✓	
12. Quantity of information									✓			✓		
13. Variability of info. value								✓	✓	✓		✓		
14. Reliability of info.								✓	✓	✓		✓		
15. Time available		✓		✓	✓								✓	
16. Frequency of decision					✓									✓

a decision situation. Because this taxonomy specifies a common set of functional requirements for decision aids and decision situations, it defines the nature of the relationships shown in Table 3. In fact, the taxonomy of analytic functions can be used in place of Table 3 to relate the characteristics of decision aids to those of the decisions they support. It also can be used to describe the analytic functions provided by a decision-maker and his staff, which means that aids can be evaluated in terms of functions that do not duplicate those already available.

Although the taxonomy of analytic functions relies on the terminology of decision analysis and data processing, it is sufficiently general to describe intuitive and qualitative processes for dealing with a decision situation. The taxonomy contains information processing and analytical functions such as: sorting and reorganizing data, identifying patterns and relationships, using an analysis to predict outcomes, and identifying the most significant elements of a decision. This taxonomy is discussed in detail in the next chapter.

The manner in which the taxonomies can be used to evaluate a decision aid is outlined in Figure 1. The sequence of assessments and logic is shown as a flow chart, starting with a specific decision aid and the decisions it supports, and proceeding through a series of steps to an assessment of the aid's potential value. These steps make use of several taxonomies to describe and compare the features and capabilities of various aspects of the decision process. Some of the steps are optional; these are shown with dashed lines. These steps can be bypassed if the characteristics of the decision aid and its potential users are well understood, but doing so will make the evaluation process less explicit and more dependent on the intuition of the evaluator. If all of the steps in Figure 1 are carried out, the process makes use of four taxonomies: three that describe the characteristics of decision aids,

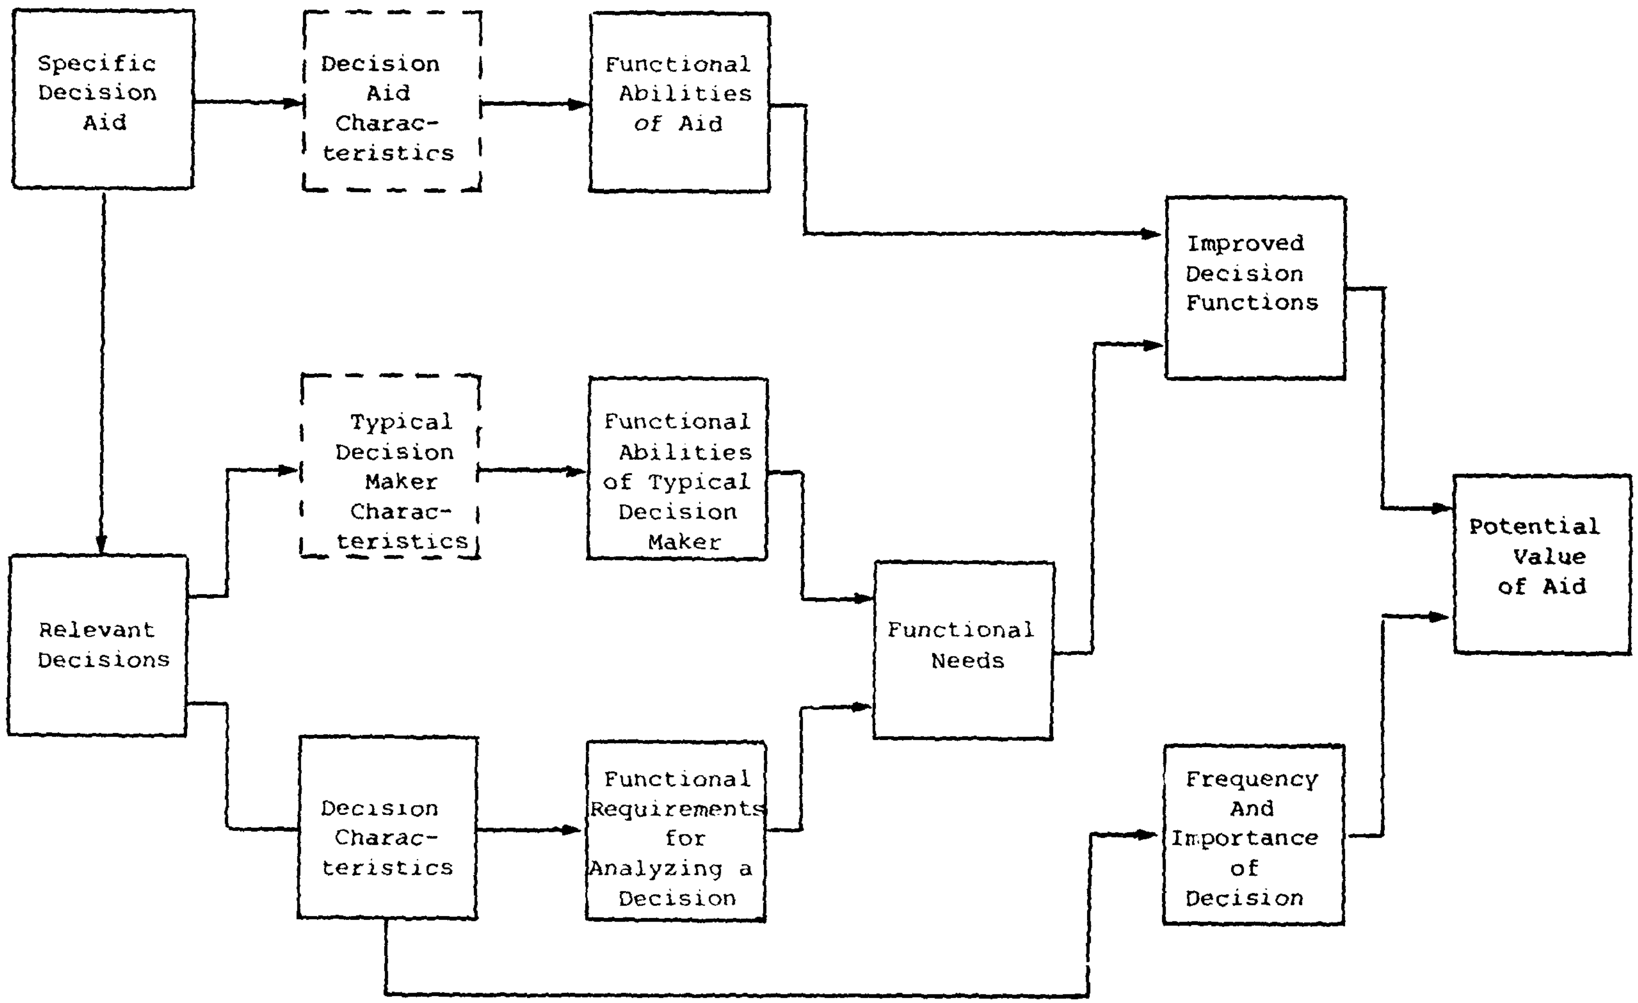

A PROCEDURE FOR EVALUATING DECISION AIDS

Figure 1

decision-makers, and decision situations; and one that specifies the information processing and analytic functions needed to study a decision problem. The last taxonomy is used to compare the requirements of decision situations with the capabilities of decision-makers and aids.

The flow chart in Figure 1 starts with a specific decision aid, although it would be possible to start with a set of decisions and use the same logic to determine the analytical functions that an aid should perform to assist a typical decision-maker who must make them. If the evaluation process starts with decisions rather than aids, someone must design the aids that will provide the needed functions. The taxonomies specify the goals of the design process, but they cannot tell us exactly how an aid should be designed.

For simplicity, the following discussion of the evaluation procedure assumes that an aid has already been designed.

The first step in evaluating a specific decision aid is to identify the decisions for which it can be used. The biggest problem that may occur at this step is that an aid will be undervalued because its potential for aiding certain classes of decisions is not recognized.

In addition to identifying the decisions relevant to an aid, one should specify its characteristics in as much detail as possible. The taxonomy shown in Table 2 can be used to accomplish this task, although someone familiar with an aid may not need this kind of checklist to specify the aid's characteristics.

The same is true of specifying the characteristics of decision-makers who might use the aid. Although a taxonomy of decision-maker characteristics could be developed, it is probably not necessary to use this taxonomy to describe the officers who make

Navy tactical command and control decisions. Part of the problem with trying to specify the characteristics of decision-makers in much detail is that Navy officers differ in such characteristics as their level of training in analytical methods and prior experience with decision aids. However, other characteristics, such as the organizational structure within which the decision-maker must operate, are essentially the same for each Navy officer responsible for the same class of decisions. For the purpose of evaluating decision aids for tactical command and control, one must assume that they will be used in a manner specified by Navy procedures and that the individuals using them will have a level of training and skill typical of Navy officers.

Determining the characteristics of a particular decision does not involve the ambiguities inherent in describing an unknown decision-maker. The taxonomy outlined in Table 1 can be used to characterize decision situations. The complete version of this taxonomy is described in the report published at the end of the previous phase of this research. It includes assessment measures for each major decision characteristic, and the feasibility of the assessments has been demonstrated by applying the taxonomy to a set of representative command and control decisions. The characteristics dealing with the frequency and importance of a decision are especially important since they are used later in the evaluation process to determine the potential value of decision functions supplied or improved by an aid.

Once the characteristics of decision aids, decision-makers, and decision situations have been assessed, either intuitively or with the help of relevant taxonomies, they are translated into a common set of functional abilities and requirements. At present, this translation process is done intuitively, although mappings like those in Table 3 could be used to make the assessments more explicit. In each case the assessment of functional abilities or requirements is done for only one element of the process: aids, decision-makers, or decision situations. All comparisons among these elements are done after the assessments are made. Since the decision-maker is specified only as being a typical Navy officer, his functional abilities probably can be assessed once for each level of command. Only the functional abilities of aids and the functional requirements of decision situations need to be reassessed as new aids are evaluated.

Once functional abilities and requirements are assessed, the comparison process is relatively straightforward. Comparing the functional abilities of a typical decision-maker to the requirements of a particular decision situation produces a set of functional needs or deficiencies that a decision aid may be able to provide. If the evaluation process starts with a decision situation rather than a decision aid, the list of functional needs can be used as the starting point for designing an appropriate decision aid. However, if the evaluation process is applied to an existing or proposed aid, the functional needs are compared to the aid's functional abilities to determine the extent to which the aid can improve the decision process. The improvements in information processing and analysis are evaluated in terms of the frequency and importance of the decision situation. The resulting estimate of an aid's potential value would be realized if the aid is used whenever the decision situation under consideration is encountered. It may be necessary to discount this potential value if it is believed that not all decision-makers are willing to use the aid or if the aid is not available every time a relevant decision situation arises.

The procedure described in this chapter, including the taxonomy discussed in the next chapter, currently is being applied to the range of decisions encountered in anti-air warfare (AAW) at the unit and task force levels of command. These decisions include determining whether or not sensor data indicates the presence of an air threat, deciding what type of threat is present from imperfect and incomplete information, deciding the significance and danger of each threat, selecting the priorities for dealing with threats, deciding which defensive weapons are capable of destroying a threat, assigning defensive weapons to threats, determining when to fire a defensive weapon, and reallocating or reorganizing AAW forces in response to new information or combat results.

If all of these decisions were made by one person, many of them could be considered part of a single decision problem to allocate and use AAW weapons to counter air threats. However, AAW operations are very complex and involve numerous decision-makers, each concerned with a portion of the overall problem. Each of the decisions listed above could be made by a different individual in an actual combat situation, so it is appropriate to consider each decision separately and explore the information processing and analytic functions needed to support it. However, since the various AAW decisions are interdependent (e.g., the selection of an appropriate defensive weapon depends on an accurate determination of the type of threat, which depends, in turn, on recognizing the presence of a threat), consideration must also be given to the manner in which the results and logic of each decision are communicated to others.

Since we are starting with a set of command and control decisions rather than a specific decision aid, the emphasis in our current research is on the steps shown in the bottom two-thirds of Figure 1. Preliminary assessments have been made for the characteristics of each of the AAW decisions, the functional requirements for analyzing the decisions, and the extent to which these functions can be provided by typical decision-makers. The relative importance of the resulting functional needs is used as a basis for identifying areas where existing decision support systems can be improved or new aids developed.

The investigation to date indicates several areas where AAW decisions could benefit from improved aids. These areas include: a capability within AAW decision aids to efficiently communicate the uncertainty and credibility associated with previous decisions; an ability to summarize the development of an air battle over time to guide decision making at higher levels of command; automatic or semi-automatic generation of contact reports, damage reports, etc.; the flexibility in AAW decision aids to accept judgemental estimates from the user and adapt the aid's processing accordingly (e.g., a detection aid that could be told to look for expected enemy flight characteristics or tactics); and an aid's ability to consider changing levels of defensive resources and the uncertainties inherent in prior decisions when recommending allocations of weapons. The importance and feasibility of these (and other) potential improvements to AAW decision decision aids are currently under investigation. Details of these decision aids, or improvements to existing aids, will be specified in the final report for this research project.

4. A TAXONOMY OF INFORMATION PROCESSING AND DECISION ANALYTIC FUNCTIONS

This chapter describes the current version of a taxonomy of information processing and decision analytic functions. As discussed in the previous chapter, this taxonomy is used to relate the characteristics of decision aids to those of decision situations and decision-makers. The taxonomy is intended to be sufficiently general to describe the abilities and requirements of aids, decision-makers, and decisions in common terms. Research is in progress to test the completeness and applicability of this taxonomy, and this work probably will result in revisions to the version discussed here.

The taxonomy is based on the assumption that some attempt is made by the decision-maker to use relevant information to understand the nature and consequences of a decision. In other words, it is assumed that some information processing or analysis (perhaps at an intuitive level) is needed to make the decision. If this assumption is not valid, then there is little point in developing a decision aid.

Although rather perfunctory information processing tasks are included in the taxonomy, they are part of the process of reaching a decision. The taxomomy is not an attempt to detail every aspect of information processing, but only those functions that are directly relevant to decision making. Also, the taxonomy does not specify how the functions are carried out, but only what they accomplish.

An outline of the taxonomy is contained in Table 4. The taxonomy starts with the task of acquiring information and proceeds through a series of more complex functions until an analysis of the

TABLE 4: A TAXONOMY OF INFORMATION PROCESSING AND ANALYTIC FUNCTIONS

1. Acquire information

   A. Objective or preprocessed information

      (1) Observe and measure physical phenomena (e.g., radar signals)
      (2) Monitor information channels (e.g., communications from another command)

   B. Subjective information

      (1) Assess uncertain quantities (e.g., estimate of enemy intentions)
      (2) Estimate aggregate descriptors (e.g., a measure of combat readiness)

2. Store and recall information (i.e., no change or reorganization of data)

3. Organize information

   A. Aggregate or collect information into pre-established categories (e.g., important vs. routine)

   B. Sort or arrange data in a specified order and format (e.g., rank air threats in terms of proximity to task force)

   C. Compare and filter data (i.e., eliminate portions of information)

   D. Recognize pre-established patterns of information (e.g., look for flight characteristics of cruise missile using a prestructed model or template)

4. Analyze decision situation (i.e., decompose decision into basic elements)

   A. Identify and define significant elements of situation; search for new elements

      (1) Alternatives
      (2) Factors and issues affecting the decision
      (3) Uncertainties and information sources
      (4) Outcomes and preferences

TABLE 4: A TAXONOMY OF INFORMATION PROCESSING AND ANALYTIC FUNCTIONS (Cont'd.)

B. Identify patterns and relationships

  (1) Recognize trends or patterns in data
  (2) Define causal links (i.e., dependencies) among elements of situation
  (3) Fit functional relationships to data (correlate data); update relationships based on new data

C. Search for alternative representations of decision situation (i.e., Is there an easier or more insightful way to look at th problem?)

5. Determine implications of analysis

  A. Calculate summary data (e.g., compute total number of aircraft available)

  B. Use logic from an analysis to predict outcomes and extrapolate data

  C. Determine optimum alternatives

  D. Calculate sensitivity of outcomes and decisions to changes in data or assumptions

  E. Determine value of additional information and cost of delay

6. Integrate the results of an analysis with existing knowledge and intuition

  A. Identify the most important elements and relationships

  B. Develop a simple explanation of analytic results (that can be internalized by decision-maker); interpret and consolidate results of analysis

  C. Update decision-maker's internal model of decision situation

  D. Select an alternative

7. Transmit information and decisions (i.e., instructions)

  A. Verbal

  B. Alpha-numeric/digital

  C. Graphical

decision is integrated with existing knowledge, a choice is made, and the conclusions drawn from the analysis are transmitted to other individuals or organizations. The order in which the information processing and analytic functions are specified in the taxonomy often does not correspond to the order in which these functions are performed. In fact, a typical decision process can skip back and forth among several of the functions. The sequence of tasks contained in the taxonomy is only intended to suggest the general flow of analytical activity as a decision is reached.

The first function is acquiring information. The nature of this task depends on whether the information is objective or subjective, and whether its source is a data processing or storage system (e.g., another command center or a manual), a physical measurement, or an estimate by the decision-maker or his staff. If the information has been preprocessed and stored or transmitted by another individual or system, the process of acquiring the information requires only that is be successfully received and decoded. However, direct acquisition of physical data or subjective estimates requires some sort of measurement procedure. This procedure relies on sensors and displays for physical measurements, and on methods for expressing, testing, and interpreting subjective estimates (e.g., debiasing procedures for subjective probability assessments).

The second function is the storage and retrieval of information. This function does not involve any change or reorganization of the data, other than placing the information in a form that can be stored or interpreted after retrieval. For example, a computer may translate a message into a series of bits that can be stored on a disk, but when the message is recalled it will be essentially the same as when it was stored. Any reorganization of the data during the storage and retrieval process is considered part of the next function.

The third function includes the many ways that information can be organized and placed in a form that is more meaningful to the decision-maker. These procedures generally involve comparing the data to a pre-established pattern or criterion, or using the data to recognize when one exists. For instance, pieces of information can be grouped into a set of categories based on a rule for discriminating between categories (e.g., incoming messages can be classified as urgent or routine). Alternatively, information can be reorganized by sorting it according to a particular sequence or format (e.g., ranking air threats according to their proximity to a task force). Both of these operations change the order in which data is presented to the decision-maker without changing the total amount of information. In contrast, information can be filtered to eliminate data that does not meet an established criterion. A more sophisticated method of classifying information occurs when data is scanned to determine whether or not a designated pattern exists (e.g., examing the flight characteristics of an air threat to see if it is a cruise missile). The criteria used to reorganize data is fixed in each of these procedures. Developing a new criterion (i.e., a new model of the relationships among data elements) is part of the next function.

The fourth function is analyzing a decision situation. The process of recognizing and specifying the basic components of a decision is probably the most difficult and creative part of the decision process. Unfortunately, there are relatively few decision aids to help someone carry out this task. The outcome of a decision situation often hinges on how well the individuals involved have identified the options available to them and the factors that have a significant bearing on the decision. Some of these are likely to be uncertain, so it may be necessary to identify various options for gathering information to reduce the uncertainty. Any decision-making activity requires some consideration (either implicit or explicit) of the outcomes that could result from alternative courses of action, and the relative value of each outcome.

In addition to recognizing the elements of the decision problem, it is usually necessary to identify the relationships among them. This typically is done by recognizing trends or patterns in the data, and assuming that these patterns can be explained by dependencies among the elements of the decision. These causal links are often the most important and least recognized assumptions in an analysis. They can be established by fitting parametric relationships to the data, and updating the relationships as new data is received. One of the powerful ways to analyze a decision situation is to search for alternative ways to think about the problem. Comparing different models of the same situation often provides insight about the most significant elements of the problem.

Once a decision situation has been decomposed into a set of elements and relationships, logic can be used to determine the implications of the analysis. This is the fifth analytic function. The simplest form of logic is to calculate some summary information, such as the total amount of fuel in storage. Simple calculations like this require very little modeling beyond an assumption that the quantity being calculated is relevant to the decision. More sophisticated logic usually is required to determine the implications of an analysis. Often it includes some sort of extrapolation procedure to predict the outcomes resulting from a decision, and it may include calculations of the uncertainty associated with these outcomes. In addition, the logic used to predict outcomes can be imbedded in an optimization procedure to determine the course of action that produces the most desirable outcomes. Since the recommended course of action depends on all of the assumptions contained in the analysis, additional calculations often are used to test the sensitivity of the predicted outcomes and recommended decisions to changes in data and assumptions. If the analysis is probabilistic or dynamic, further logic can be used to determine the value of additional information or the cost of delay.

Numerous computer programs have been written to calculate the implications of an analysis. Few of these are sufficiently general to be useful for more than a few specific decisions. General purpose algorithms have been developed for broad classes of problems (e.g., linear programming), but they often must be adapted or tailored to fit a specific decision. In addition, they usually deal with only part of the problem, so they become part of a larger analysis.

The sixth function is often done intuitively and with little support, although it is one of the most important. It is the process of integrating the results of an analysis with existing knowledge and intuition. Analyses are rarely accepted without reconciling their results with the decision-maker's current understanding of the situation. The extent to which a decision-maker's current view of the problem is updated or replaced depends on the credibility and limitations of the analysis. (A detailed discussion of the concept of credibility is contained in Chapter 6). Often the purpose of an analysis is to give the decision-maker a few basic insights that then become the basis for a decision. This can be accomplished by identifying the most important elements and relationships in an analysis, and using them to form a relatively simple model of the decision situation that can be internalized by the decision-maker.

The last function is transmitting the results of the decision process to other individuals and organizations. These results can be instructions or decisions, or simply information that results from data processing and analysis. The form of the communications can be verbal, alpha-numeric, or graphical depending on the nature of information being transmitted. Typically a decision-maker has a variety of aids to support this function.

5. A REVIEW OF RESEARCH ON LIMITS TO HUMAN INFORMATION PROCESSING AND DECISION ANALYTIC ABILITIES

The procedure for evaluating decision aids using a taxonomy of decision functions is based on the idea that humans have limited abilities to perform certain functions, and aids may be able to overcome these functional deficiencies. In order to understand how an aid can help the decision-making process, we must first understand how humans make decisions without an aid. This chapter reviews some of the research that has been done on the limits to human rationality, and the extent to which decision aids and analytic procedures can overcome these limits.

The organization of this chapter is as follows. We first review several alternative theories of how humans make decisions. These theories are based on an interpretation of alternate forms of decision-making behavior as manifestations of different bounds on rationality. These bounds may be viewed as scarce resource constraints that influence the extent to which a decision-maker can act rationally in a given situation. We then discuss the degree to which analytical methods and decision aids can or should reduce these bounds to rationality, and the extent to which these methods introduce undesirable attributes of their own.

There is a great deal of evidence that unaided humans often make very poor decision-makers by any measure. Empirical research by psychologists, organizational behaviorists, decision analysts, and others has revealed that in many situations humans tend to be poor information processors, poor at computation, biased evaluators, and inconsistent at making choices, although some individuals seem to be naturally good decision-makers. Over the last twenty five years a great deal of progress has been made in understanding why people behave the way they do. Although a complete review of the literature is beyond the scope of this report, the work of three authors -- Simon, March, and Radner -- provides a good framework for this discussion.

Simon starts from the idea that all intendedly rational behavior is behavior within constraints (15,16). In his view the most important of these constraints reflect limitations in human computational capabilities and the organization and utilization of human memory. Simon makes two relevant points for this discussion: first, that the state of information of a decision-maker is an important characteristic of his decision process; and second, that a decision-maker may deliberately introduce simplification into a model of a situation in order to bring the model within the range of computational capacity. Simon concludes that a theory of rational behavior must be as concerned with rational actors coping with uncertainty and cognitive complexity as with the objective environment in which they make their decisions.

Simon distinguishes between two different types of rationality: <u>substantive rationality</u>, the extent to which appropriate courses of action are taken; and <u>procedural rationality</u>, the effectiveness, of the procedures used to choose these actions in light of a decision-maker's cognitive powers and limitations. Conceptually, there is no reason to treat the substantive and procedural problems as though they are independent of each other: The global optimization problem facing a decision-maker is to find the least cost or best return decision, including computational cost. In practice, however, this is rarely the way humans proceed. Not only are there cost constraints on the computation of solutions of problems, there are also constraints on time, memory, imagination and a variety of other phenomena that are difficult to overcome and even more difficult to take into account in a systematic way.

Radner also deals with the problems of what he calls "limited rationality". He characterizes a number of different types of behavior that seem to follow from a combination of different types of bounded rationality: (1) constant proportions, in which the allocation of effort is constant over time; (2) putting out fires, in which all effort is allocated to those activities that have the worst performance at any given time, and (3) staying with a winner, in which all effort is allocated to those activities that have the best performance at a given time. Radner emphasizes that such behavior is not necessarily the result of costly information and analysis, but is also influenced by the "limited capacities of humans and machines for imagination and computation." Since formal methods will never take the place of imagination, it is clear that we will always have to live with some form of bounded rationality.

March provides another more detailed characterization of types of rationality. He distinguishes between calculated and systemic rationalities (5). Action, he says, is presumed to follow either from explicit calculation of its consequences in terms of objectives, or from rules of behavior that have evolved through processes that are sensible but which obscure from present knowledge full information on the rational justification of any specific rule.

The distinction between calculated and systemic rationalities is a good starting point for discussing the role of decision aids. Four of the calculated rationalities identified by March are described below (the definitions are taken from (5)). Identifying the implicit resource constraints associated with these rationalities provides a natural framework for evaluating the potential impacts of analysis.

Limited rationality emphasizes the extent to which individuals or groups simplify a decision problem because of the difficulties of considering all the alternatives and information. This type of rationality manifests itself in oversimplified decision procedures such as heuristic search rules, incremental thinking, muddling through, and uncertainty avoidance. A scarce resource here is the ability to deal with complexity.

Contextual rationality emphasizes the extent to which choice behavior is embedded in a complex of other claims on the attention of decision-makers. The scarce resource that is emphasized is time: When the decision-maker cannot spend much time on a given decision, it is likely to look "irrational" when compared to other more carefully considered decisions.

Process rationality emphasizes the extent to which decisions find their sense in attributes of the decision process rather than in attributes of the decision outcomes. Such behavior is typical in large organizations in which, for purposes of control, many decision procedures are standardized. Decision-makers are rewarded not for good outcomes, but for adherence to the "rules". Flexibility and control may be viewed as scarce resources here.

Game rationality emphasizes the extent to which organizations and other social institutions consist of individuals who pursue individual objectives by means of calculations of self interest. The scarce resource here is cooperation.

For each of these types of rationalities a case can be made that decision aids can be of help in removing the bounds, or at least in providing a way of dealing with them. For example, an aid based on decision analysis could help someone deal explicitly with complexity and uncertainty, and enable them to do quick computations, even in very intricate situations. This type of aid is focused on reducing the bounds inherent in limited and contextual rationality. Similarly, an aid to modeling organizational behavior would focus on the bounds of process rationality by providing a framework for taking procedural or organizational problems explicitly into account in the decision process. An aid employing game theory addresses game rationality by offering a structure for identifying and analyzing ways in which cooperation can help the individuals in a group all achieve better outcomes. Thus, decision aids certainly have the potential to influence positively the factors inherent in the calculated rationalities. The scarce resources are relatively easy to identify, and also relatively easy to influence.

On the other hand, the case for decision aids is somewhat harder to make for the systemic rationalities. Here the scarce resources at play are somewhat harder to identify and also more difficult to influence. March identifies three types:

Adaptive rationality emphasizes experiential learning by individuals or organizations that permits the efficient management of considerable experiential information, but tends to separate current reasons from current actions. This type of rationality describes the intuitive decision-maker who has a hard time justifying his choices to others. The scarce resource here seems to be the ability to understand and articulate the reasoning process underlying ones' intuitive actions.

Selected rationality emphasizes the process of selection of decision procedures through survival or growth; choice is dominated by rules that have survived and evolved. Contrary to the decision analysis dictum of rewarding good decisions rather than good outcomes, this type of rationality stems from a system that rewards good outcomes (an individual who takes a poor risk would be rewarded if he happens to have a lucky outcome). In this case accountability appears to be a scarce resource; it is often much easier to measure the outcomes of a decision than to review the process by which the decision was made.

Posterior rationality emphasizes the discovery of intentions as an interpretation of action rather than as a prior position; thus, decisions are antecedent to goals and define preferences rather than follow them. The scarce resource here seems to be foresight or clear thinking, since this type of rationality implies taking actions before fully evaluating the consequences.

In principle, decision aids can reduce the bounds inherent in both selected and posterior rationality. However, it seems clear that no aid can hope to completely eliminate these bounds. For example, in the context of selected rationality, a decision analytic aid could provide a consistent framework for explaining the elements that went into making a given decision, and help a user appraise the quality of the decision process. On the other hand, it is not feasible to fully analyze every decision, even if every decision is made using the tools of decision analysis. Therefore, it seems inevitable that some form of selected rationality must remain in spite of the application of an analytic aid. Similarly, although an analytic aid can make it easier to choose actions that are based on one's preferences, it is likley that some sort of posterior rationality will always remain, particularly for minor decisions.

The case for decision aiding is most controversial in the context of adaptive rationality. It is this type of rationality that critics like Tribe (19), and Dreyfus and Dreyfus (2), defend to the exclusion of formal methods. Hoos (3) also discusses this problem. They feel that the quantification process required by many analytical aids can be inherently bad. Even if a formal analyis could lower the bounds inherent in adaptive rationality (a point they do not really concede), they do not believe that lowering the bounds is necessarily a good thing. They argue that the "intangible" elements that can't be quantified are often the most important parts of the problem, and that the process of formalizing a decision-maker's intuitive process tends to ignore these factors.

If we take the traditional view of either accepting or rejecting the conclusions derived from an analytical decision aid, then the critics' points seem well taken. It is certainly not reasonable to assert that adaptive rationality is always wrong or naive; most of us know persons that seem to be naturally good decision-makers. There is no logic that guarantees that an aid will be better than the decision-maker's intuition every time. On the other hand, most would agree that some of the benefits of decision aiding -- the ability to explore the logic of decisions, the ability to handle uncertainty and complexity, the expanded capability for communicating the results to others -- are very attractive. The challenge is how to take advantage of the benefits of decision aids without throwing away the rich body of intuitive knowledge and judgment that often cannot be formalized and quantified.

To summarize, there are a number of alternate forms of descriptive rationality. Each of these rationalities can be interpreted as the result of some bound or scarce resource constraint on the ability to formulate, think about, and solve decision problems. Formal decision-aiding methods have the potential to overcome these bounds in many cases. However, when decisions aids are used to the exclusion of unaided judgment and intuition, there is a valid question as to whether the aids do more harm than good.

The view that a conclusion derived from a decision aid must be completely accepted or rejected is perhaps the result of the cultural assimilation of the "scientific method" into our thinking. In the scientific method, a hypothesis is proposed and then either accepted or rejected on the basis of some experimental evidence. This type of thinking, which has served so well in the field of scientific discovery, is wholly inappropriate in the field of decision-making.

A broader and more useful view of a decision aid is that it provides a decision-maker with additional data on which to base his choices. The result of decision-aiding is information for the decision-maker, not a substitute for intuition, good judgment or clear thinking. Viewing analytical aids as complements to, rather than substitutes for, unaided intuition frees us from the unnecessarily rigid vision of formal analysis as a competitor of the intuitive process.

In science the only place for intuition is the design of an experiment; in the unavoidably personalistic realm of decision-making, intuition must play a central role. When faced with a one-of-a-kind decision problem, a decision-maker can do no better than act based on what he has (alternatives and resources), what he wants (values and preferences), and what he knows (information). Aids can help provide information and clarify choices, but it is unrealistic and unnecessary to expect them to wholly replace the decision-maker.

In effect we are claiming that an aid should be viewed as a special kind of expert. No one would ever argue that all expert opinions should be believed, or that all expert advice should be followed to the letter. It makes just as little sense to treat the conclusions derived from an analysis as inviolate. However, it is equally shortsighted to completely ignore an aid if it can provide insights, but fails to provide a perfectly credible answer. The information produced by the aid, just like the advice of the expert, must be tempered by the decision-maker's feelings about the credibility of the analysis that produced it. The formal development of these intuitive ideas has been started by Morris (9), and Nickerson and Boyd (12).

In summary, we have seen that the actual process by which humans make decisions is quite complex. We have attempted to argue that because of, not in spite of, the limited rationality with which the average person can hope to apply to a decision, the prospects for helpful decision aiding are good. On the other hand, the naive use of an analytical aid may well do more harm than good, and its user must clearly understand its proper role in order to ensure that it makes a positive contribution.

The following chapter will explore in more detail the idea of using an aid as an expert and will use this approach to derive a measure of the aid's credibility.

6. THE CREDIBILITY AND VALUE OF DECISION AIDS

Introduction

An important feature of a decision aid is its credibility. An aid that a decision-maker does not believe has little value, no matter how sophisticated it may be. An inexpensive aid that provides limited but credible information may be more worthwhile than a costly aid that provides detailed information that is not as credible.

This chapter reports the results of our research on decision aid credibility. The research has three main objectives:

i. Define precisely the concept of credibility

ii. Determine a quantitative measure of credibility that can be used to evaluate different decision aids

iii. Develop an operational procedure for assisting decision-makers in utilizing decision aids.

We have attempted to construct a conceptual structure for thinking about the issues of credibility, and an analytic procedure to help the Navy evaluate existing and proposed decision aids.

This chapter is divided into three main sections. The first two sections examine two general classes of decision aids. We first address the class of sophisticated decision aids that not only provide inputs to a decision-maker, but actually help in the process of formulating and making a decision. For such aids the issue of credibility is very complex. It is not obvious how to measure the credibility of such aids, and even less obvious how the value of the aid is related to the degree of credibility.

The second section starts at the other end of the spectrum and addresses the issue of credibility for a conceptually simple decision aid such as one that displays the location of objects detected by sonar or radar. Although such devices are not normally considered decision aids, they represent one end of the spectrum of the information processing functions described in Chapter 4. For these physical measurement devices the credibility may be related directly to the measurement error. We present an analysis of how the value of such aids is related to their credibility.

The third section presents a general conceptual framework for evaluating aids. We develop a concept named "stochastic decision trees" which appears to be a completely general way to think about the value of an aid that can be trum of decision problems. Although this work ly on the issues of credibility, it is our hope tha n trees can be used to unify the concepts developed in tions.

Credibility of Decision Tree Aids

There are many different types of aids that could help a decision-maker structure and make a decision. To bound the problem we focused on one representative aid whose purpose is to help the decision-maker structure a decision tree, and then evaluate the decision tree to recommend the best alternative. The aid would interact with the decision-maker to elicit both the structure of the decision and the probabilities and values that characterize the tree. A sophisticated version would provide additional assistance in helping the decision-maker provide these inputs (through a structured inquiry system or possibly through some additional modeling).

As discussed in Chapter 5, a large body of descriptive decision-making literature supports the conclusion that the credibility of a decision aid (or analysis, or model) influences the way it will be used. This existing work indicates the importance of the issue. However, there is little theoretical work on the subject of how the credibility of an aid should influence the way it is used. The closest related work has been in the area of model evaluation. This work is relevant for the decision aids we are examining, aids that produce models to assist in the decision-making process. We have included a list of references at the end of this report.

Approach

Our approach is best understood in the context of the traditional role of decision models. We begin with a brief review of the classical decision analysis paradigm.

In the traditional view, there are two solutions to a given decision problem. These are depicted in Figure 2. The first solution is the decision-maker's intuitive decision; i.e., the decision he would make without any formal analysis or model. The second solution is the model solution. We assume that a decision tree model has been constructed that structures the alternatives and outcomes and, based on a set of probability and value judgments, determines an "optimal" decision. This decision is called the "model decision".

In the traditional paradigm, the decision-maker must decide to accept or reject the model. If he rejects the model, he makes the intuitive decision; if he accepts the model, he chooses the model decision. There are at least two problems with this. First, there are no explicit criteria for determining whether to accept or reject the model. The decision-maker is often placed in the position of having to judge a model whose technical details he doesn't really understand. Second, and perhaps more important, in rejecting either the model or his intuition, the decision-maker is throwing away potentially useful information. The danger of neglecting intuition is the danger of neglecting important insights just because they are

DM Solution

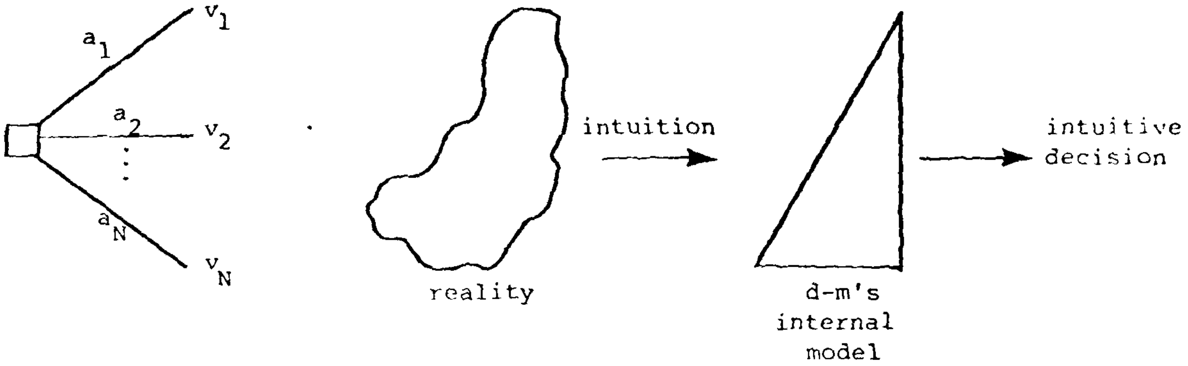

Model Solution

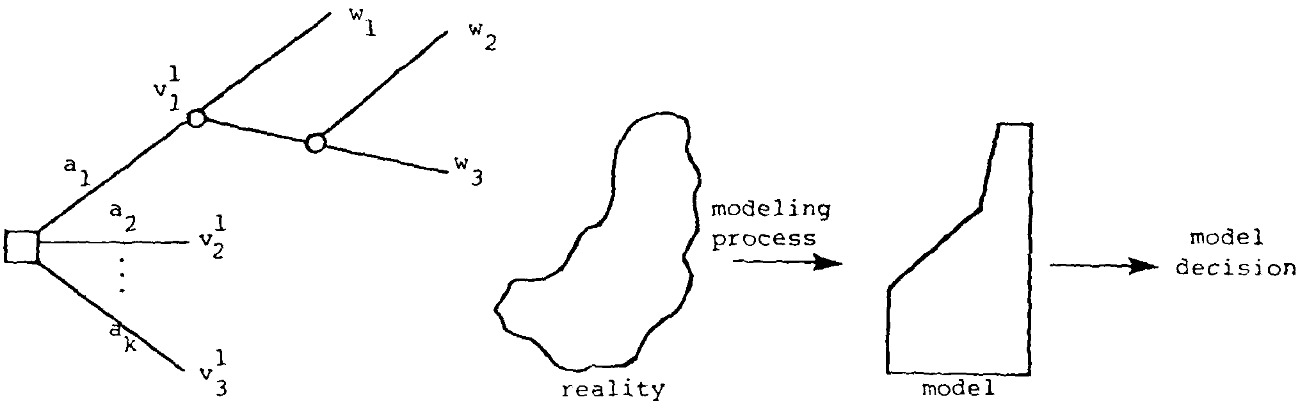

THE CLASSICAL PARADIGM

Figure 2

hard to model; the danger of neglecting the model is the danger of neglecting valid information and logic just because it is hard to understand.

A more general approach that we shall take is to view the model as providing new information about the value of the decision alternatives. In this view, the question of accepting or rejecting the model is meaningless. The model simply provides a source of information that the decision-maker can use to update his intuitive judgment about the value of each alternative. The updated values, based on the information supplied by the model, are then used to select the best alternative. This broader formulation treats the model as a complement rather than a substitute for intuition.

A feature of the analysis is that it provides a framework for determining the value of further analysis. For decision aids that help the decision-maker structure and think about a problem, this is especially important. For these aids there is a continuous decision of whether to stop using the aid and make the final choice, or continue using the aid to structure the problem further.

Credibility of a Model

It is easy to speak and think of credibility as being a physical attribute of a model. However, this is not a very useful way of thinking since two people may have very different feelings about the credibility of the same model. A more productive view is that the credibility of a model is related to how accurately it represents one's perceptions of reality. Since representation is by nature a subjective process, so must be its evaluation.

Another fallacy that goes to the heart of evaluating the credibility of a model is that there is only one correct model. It is easy to fall into the conceptual trap of believing that if we just had enough time we could achieve a perfect representation of reality. The logical flaw in this is that no matter how detailed a model is, we could always add additional structure. In fact, the best we can do is accurately represent our perception of reality, which is always limited. The following simple example will amplify these ideas and provide a useful reference for the discussion that follows.

Example: The Cuban Missile Crisis

Consider the simple decision tree illustrated in Figure 3 that depicts President Kennedy's decision in the Cuban missile crisis. One of his options was to blockade Russian ships. If he chose a blockade, he had to be concerned with the possibility of precipitating a nuclear war. Let us assume there was some probability p of a war occurring. In this simple model, there are three possible outcomes labeled w_1, w_2, and w_3 which are the values the President would attach to each outcome.

Figure 3 also shows a more detailed model of the same situation. Model II represents more of the problem structure, including the possibility that the Russians might try to run the blockade with some probability q, resulting in two different conditional probabilities of nuclear war, r and s. The model also includes the option of economic reprisal in the case of no blockade. In this augmented decision tree there are six basic outcomes on which the President would have to assign values u_1 through u_6.

Most observers would agree that Model II is in some sense more credible than Model I. It is more detailed and includes more of the actual problem structure. However, it is easy to show using elementary probability theory that Model II could be reduced to a

Two Decision Models:

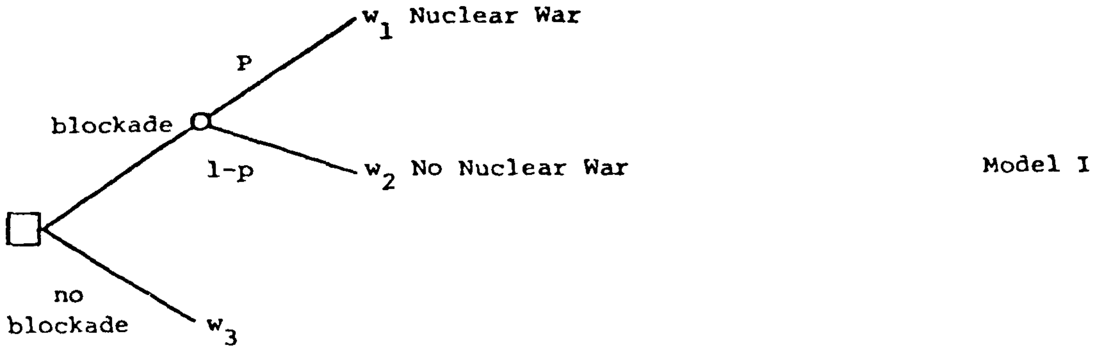

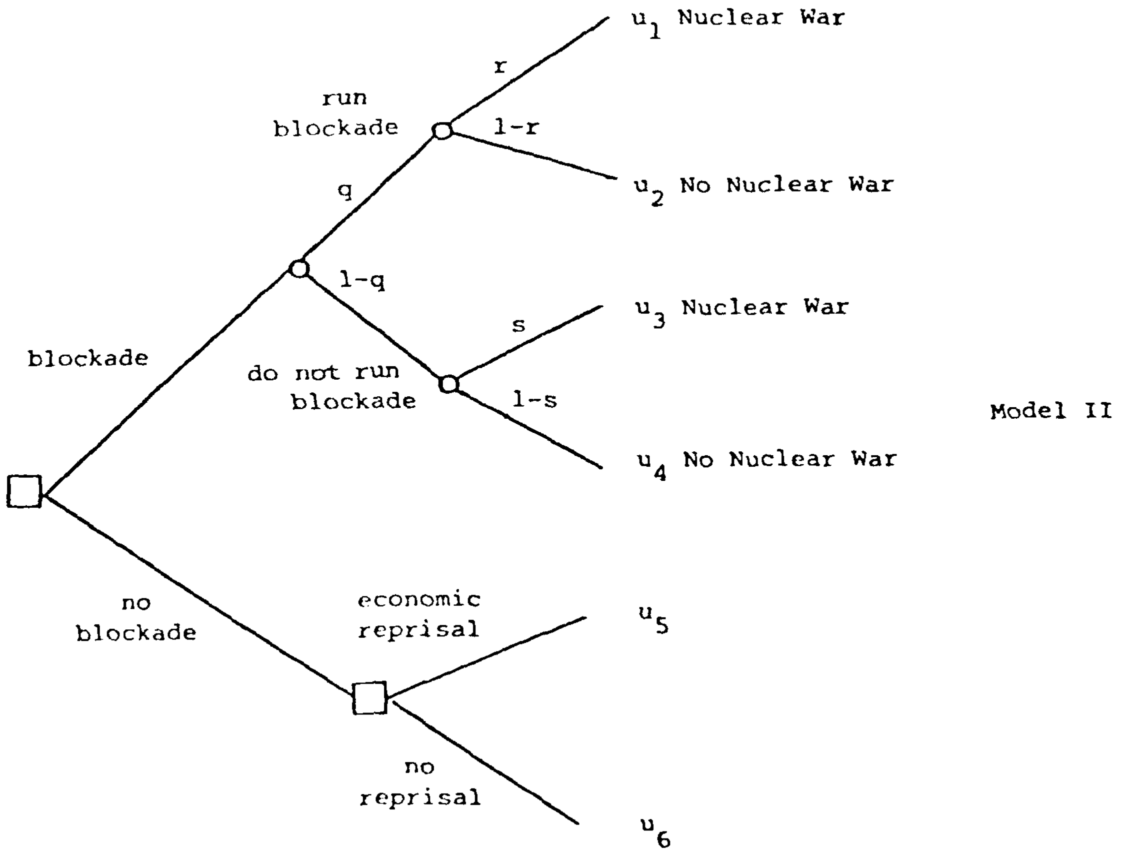

EXAMPLE - THE CUBAN MISSILE CRISIS

Figure 3

form exactly equivalent to Model I (for example, the probability of nuclear war would be $(qr+(1-q)s)$. The only possible difference would be in the numerical values of the probabilities and values.

Since the models are logically consistent, it makes no sense to think of either structure as being more or less correct. It is also apparent that an assessment of credibility must rely in part on information not available by measuring characteristics of the model such as number of variables, level of detail, etc. For example, it is easy to construct examples of more detailed decision models of the Cuban missile crisis that are less credible than those in Figure 4.

Perfect Credibility

A more relevant question to ask about the two models is how might the analysis change with additional modeling or information gathering. Consider two cases, corresponding to different feelings a decision-maker might have about a given model:

Case 1: There is a good chance that further modeling will significantly change his assessment of the value of each alternative.

Case 2: There is no chance that further modeling or information would lead to any new beliefs.

Most would agree that the credibility of the model in Case 2 is higher than in Case 1. Intuitively, if we know that our modeling assumptions are absolutely valid in the sense that no information we might observe, or modeling we might perform, prior to the decision would change the answer, then the model is as credible as possible. Conversely, if our assumptions are likely to change with more modeling or feasible information gathering, then the current model is not completely credible.

Consider again Model I, the simple model of the Cuban decision. Why is this not a very credible model? One reason is that the model, in effect, forces the decision-maker to specify probabilities of very complex outcomes. Thus, we might expect the decision-maker to feel that the probability assessment in Model I is in some sense relatively inaccurate, and likely to change with more thinking and modeling. Similarly, if we ask him to assign a value to the consequences of no blockade, he might be quite uncomfortable doing so since it depends on a number of unstated factors (some of which are included in Model II). Considering these factors explicitly may well affect the assessed value of this alternative.

Model II on the other hand decomposes the problem into more basic elements. Presumably, the decision-maker would be more comfortable assessing the more basic elements and having the model perform logical computations than he would be if he must keep track of everything in his head. We would expect him to be more confident that further modeling would not change the results.

In summary, we have found that the notion of model credibility is related to the decision-maker's feelings about how future modeling and information could affect the model results. Furthermore, the stability of the model results is a function of the stability of the model inputs and assumptions. If all the assumptions were invariant to future information or modeling, then the model would be perfectly credible. We are thus led to a formal definition of a "perfectly credible model". This definition will be useful both as a conceptual thinking tool and for quantitative analysis.

<u>definition:</u> A perfectly credible model is a model whose results are estimated by the decision-maker to be invariant to additional feasible information gathering or modeling.

It is important to distinguish the notion of a "perfectly credible model" from the notion of a "perfect model" as proposed by Nickerson and Boyd (12). They define a perfect model as one that predicts with certainty the value of any alternative. There are two reasons why this is unacceptable as a definition of a perfectly credible model. First, such a model is unattainable in virtually all real decision problems. Second, many non-deterministic models are completely credible in the usual sense of the word. A classic example is modeling sequence of coin flips as a Bernoulli process with probability 0.5 of heads or tails. Intuitively, the Bernoulli model is completely credible because we believe no amount of information or modeling would change our conclusion that the probability is 0.5. Although we could imagine extreme situations in which we could make more refined observations (thumb velocity, height from the floor, resiliency of the coin, etc.), in most cases it is not possible to make such observations before the decision.

Our definition of a perfectly credible model takes into consideration the amount of information feasible to observe before the decision. This of course also limits the scope of the modeling that may be performed before the decision.

Sources of Model Credibility

It is useful to think about decision-tree models such as Model I and Model II as black boxes that have certain inputs and outputs. The inputs to each model are the probabilities describing the chances of uncertain outcomes, and the values describing the worth of those outcomes. The decision analysis methodology provides an explicit logical way to determine the best alternative given these inputs. This is the alternative with the highest expected value. (For a risk averse decision-maker, these values will be utilities which reflect a desire to minimize the uncertainty in the outcomes.)

It is important to recognize that there are really three basic types of inputs to the description of the decision problem. In addition to the probabilities and the values, there is the structure of the decision tree itself; i.e., the alternatives and events described by the tree. A useful way to view both the probability and value assignments is to regard them as summaries of situations that could always be structured in more detail. For example, the assignment of a value to the no-blockade option in Model I must take into account, at least implicitly, other possible actions that might be taken. The issue is whether to include this structure explicitly in the formal model, or implicitly in the assessment of the value.

It is always possible to interpret an event probability as an estimate of the unknown frequency of occurrence of similar events. If we are certain about the frequency, then the probability assignment won't change with further information or modeling. However, if the frequency is unknown, then additional modeling or information will potentially alter the probability. We are thus led to a natural definition of a "perfectly credible probability" as one on which there is no uncertainty concerning the underlying frequency. The probability assignment of 0.5 to a fair coin is perfectly credible because the frequency is virtually known and no modeling or information would change it. On the other hand, a probability assignment of 0.5 to the event that our next president will be Republican is not perfectly credible because additional modeling or information gathering could easily change it. In most circumstances a perfectly credible probability will be unattainable. However, the definition is useful because it represents the best we can do.

Similarly, we define a "perfectly credible value" as one that does not change with further modeling or information. Conceptually, we could obtain perfectly credible values by building a sufficiently detailed decision tree so that the outcomes are unambiguously defined and do not include implicit uncertain events or unstructured alternatives.

An expected value calculated by rolling back a part of a decision tree will be perfectly credible if and only if all the probabilities and values in that portion of the tree are themselves credible. Thus, a basic consistency condition is that a model is perfectly credible if and only if all its parameters are.

A useful analogy to the process of specifying probabilities and values is the physical measurement process. Just as there are measurement errors when making measurements of physical phenomena, so there will be assessment errors in specifying probabilities and values. Thus, we presume that both the intuitive value and the model value are, in general, not equal to the perfectly credible value; rather they both are estimates of this value that contain, in general, some error.

The decision-maker's prior uncertainty about the perfectly credible value of each alternative may be assessed directly as a probability distribution. This assessment, which we shall discuss later, is a measure of the decision-maker's confidence in his own intuition. We call it the "intuitive assessment".

The uncertainty in the perfectly credible value implied by the model may be calculated from an assessment of the uncertainty in the errors of the inputs. The decision-maker assesses the uncertainty in each parameter and then the model is used to calculate how the uncertainty propagates through the tree. (This calculation is a straightforward application of probability theory.) We call the resulting probability distribution the "model-based assessment." The model-based assessment is a measure of the decision-maker's feelings about the accuracy of the model. If the model were a human expert, the model-based assessment would be a measure of the expert's confidence in his advice.

It is useful and intuitive to think of the range of uncertainty represented by the model-based assessment as a measure of the credibility of the model. The more certain the model-based assessment, the more credible the model. The following definition of model credibility makes this notion precise.

> definition: The credibility of a model is equal to the precision of its probability distribution on the perfectly credible value.

The precision of a probability distribution is a well-defined technical term equal to the reciprocal of the variance. Infinite precision corresponds to complete certainty. Thus, a perfectly credible model is one whose credibility is infinity.

The evaluation of a given model depends on the relative credibility of the model and the decision-maker's intuitive assessment (his intuitive model). The next section presents a method for performing this evaluation.

A General Result

We are now in a position to determine how the decision-maker should update his intuitive assessment, based on reception of the model-based assessment. This is the central analytic result of our research.

We regard the model result as information. This information may be thought of as the result of an experiment, the modeling process. Viewed in this way, the updating problem is a classical inference problem, and Bayes' theorem may be applied to determine the posterior probability distribution on the perfectly credible value of each alternative, based on observation of the model-based assessment.

The posterior distribution depends on the likelihood function, which is an assessment of the relative likelihood of every possible model the aid might produce. Unfortunately, although it is possible in theory to assess the likelihood function directly, for most reasonably sophisticated decision aids, direct assessment is practically impossible.

One way to simplify the assessment is to *approximate* the likelihood function. A result of our research is that the likelihood function can be obtained directly from an assessment of the credibility of the model if the shape of the likelihood function is similar to (or can be approximated by) that of a Normal distribution and if two additional behavioral assumptions are satisfied.

The first assumption is that, if the decision-maker has no prior knowledge, or has done very little thinking about the decision problem, then he will adopt the state of information represented by the model as his own. This assumption would be accurate if the decision-maker does not feel that the model has any systematic bias. In fact, a good decision aid will be designed to help the decision-maker structure his decision problem to eliminate just such systematic biases.

The second assumption is that the result of the formal modeling process may be viewed as the outcome of an independent experiment. This means practically that our results will apply only to reasonably extensive modeling efforts that extend the decision-maker's thinking in a substantial way. Decision aids that result in very simple models (as in, say, a short-term crisis situation) would likely produce results highly dependent on the results of the decision-maker's intuitive model. Although we have some preliminary ideas or how to characterize such dependence, we have not addressed the subject in our current research.

For brevity we shall skip the mathematical details here and show the implications of these assumptions. Let v_d and c_d be the value (expected value) and credibility (the precision) of the decision-maker's intuitive assessment of a given alternative, and let v_m and c_m be the value and the *credibility* assigned by the model. We are interested in the updated value and credibility, based on observance of the model result. The result is that the new parameters are related to the decision-maker's *prior* evaluation and the model's assignment in a very simple way:

$$\text{updated value} = \frac{c_m v_m + c_d v_d}{c_m + c_d} \qquad (6)$$

$$\text{updated credibility} = c_m + c_d$$

Thus we see that both the model-based value v_m and the intuitive value v_d are weighted by their relative *credibilities*.

The above result allows us to determine the posterior distribution for the perfectly credible value for any decision alternative from a set of relatively simple assessments. The evaluation is based on an explicit measurement of the credibility of the model. The credibility measurement is derived from a simple set of assessments. We discuss the assessment task below, and conclude with a step-by-step summary of the methodology.

The Assessment Task

Apart from the task of building the model, the decision-maker must make three types of assessments in order to encode the credibility of a decision tree model. First, he assesses the possible variation in the intuitive values he assigns to each alternative before using the aid. Then he encodes his judgment about the possible variation in the probability inputs to the model, and his judgment about the possible variation in the values at the tips of the tree.

The possible variation in each parameter can be summarized by the credibility or precision of the probability distribution on that parameter. The easiest and most direct way to estimate the precision is to have the decision-maker provide a 70 percent probability range such that there is equal probability that the revealed value will fall on either side of the nominal value. Since the standard deviation of a probability distribution defines a range having approximately 68 percent probability (exactly 68.2 percent for the Normal distribution), we may use the 70 percent range as an estimate of the standard deviation. The precision, or credibility, is then calculated directly as the reciprocal of the square of the width of the range.

To summarize, the crediblity assessment task consists of specifying a single 70 percent range for each parameter that is not perfectly credible.

The Step-by-Step Procedure

We now summarize the methodology in terms of a procedure for utilizing a model produced by a decision aid:

Step 1 - Encode Intuitive Judgement

The decision-maker attaches a nominal value and a range to each alternative. The result is the best decision based on intuition and a measure of the credibility the decision-maker attaches to his intuition.

Step 2 - Construct the Model

Using the decision aid, the decision-maker constructs a decision tree and assigns nominal probabilities and values.

Step 3 - Encode Parameter Credibility

The decision-maker assigns a 70 percent range to each probability and value in the tree.

Step 4 - Calculate the Model Credibility

The model credibility is calculated from the parameter credibilities. As discussed earlier, this calculation is done using the model itself.

Step 5 - Calculate the Updated Expected Value of Each Alternative

The updated value is computed using the relative credibilities of the model and the decision-maker's intuition.

Step 6 - Calculate the Credibility of the Updated Model

The posterior credibility is computed as the sum of the intuitive and model credibilities. This is useful for evaluating further modeling efforts.

It is important to note that the decision-maker is only involved in the first three steps of the methodology. Steps 4, 5, and 6 are all logical operations that could be performed by the aid itself. Thus, the decision aid could be designed so that it assists in its own evaluation.

An important benefit of augmenting the aid to assist the decision-maker assess its credibility is that we can use the posterior credibility as a gauge of the value of refining the model. The new paradigm, which treats the model result as information allows us to utilize the powerful results of information-value theory. In particular, it is possible to calculate the value of a perfectly credible model. This provides an upper bound on the possible benefits of using an aid that helps the decision-maker produce a decision model.

The ideas described above appear to be promising in terms of practical application to decision tree aids. Furthermore, the ideas can be extended to a broader class of aids (e.g., aids of known structure with a finite number of inputs). The main issue that needs additional attention is that of the dependence between the modeling process and the process the decision-maker uses when making unaided decisions. The results stated here assume that these processes are independent. Further research is necessary to explore more thoroughly the issues of probabilistic dependence, both in parameter estimation and between the formal model and the intuitive model.

The Value of Decision Aids for Assessing or Measuring Parameters

This section presents an analytic framework for evaluating the benefits of a class of simple decision aids. We examine decision aids that provide information to the decision-maker rather than help him in the process of making his decision. While the approach is compatible in concept with the decision tree methods described above, we are able to obtain stronger results by narrowing our focus. The hope is that this specialized methodology can be extended to other aids, and, with further research, the concepts of the previous section and the concept of stochastic decision trees outlined in the next section can be merged with these results to form a general procedure for evaluating decision aids.

We assume that the information produced by the aid allows the decision-maker to make a more accurate estimate of a variable or factor that is important to the decision he is facing. We characterize each potential decision situation for which the aid is relevant by three measures: the importance of the situation, the importance of having an accurate estimate, and the decision-maker's prior uncertainty about the uncertain factor. The methodology quantifies the value of an aid in terms of these three measures. By then specifying the frequency with which different types of decision situations occur, we can quantify the overall value of an aid. These ideas are explained in more detail below.

Framework

Assume for a given decision situation that there is a true state of the world. The aid provides information that leads to an estimate of the true state. For example, consider the problem of locating a submarine. The true state of the world is the actual location of the submarine. A decision aid that displays the location of objects detected by sonar helps to pinpoint a location of a submarine.

On the basis of the information supplied by the aid, the decision-maker makes a new estimate of the actual state of the world. This estimate may or may not be equal to the estimate the aid provides. In many cases the decision-maker may choose to incorporate his own intuition about the true state of world, as well as the information provided by the aid.

We shall posit that in a given decision situation there is an underlying cost function that describes the cost of an inaccurate estimate. The simplest such function would assign a fixed cost to incorrectly estimating the true state of the world (as, for example, in a "hit or miss" situation). A more general cost function would describe how the cost increases as the estimate diverges from the true state of the world.

The final ingredient in our formulation is an assessment of the decision-maker's uncertainty about the true state of the world prior to using the decision aid. If the decision-maker were absolutely certain about the state of the world, then the value of the aid would be zero. The more uncertain the decision-maker is about the true state of the world, the higher the value of a decision aid that provides information about the true state.

Thus, we describe a given decision situation in terms of three characteristics.

- The magnitude of the decision problem (i.e., a measure of its importance).
- A measure of the importance of accurately estimating the true state of the world in a given situation.
- The decision-maker's uncertainty about the true state of the world in the given situation.

For each specific setting of the above characteristics we can compute the value of an aid by looking at the value of the decision situation with and without using the aid.

Any decision aid probably will be used in a sequence of different decision situations. By characterizing the frequency with which each decision situation occurs, we can calculate the total value of the aid by multiplying the value in each situation by the number of times the situation occurred, and aggregating this value over all possible situations.

The above ideas provide what appears to be a natural framework for quantifying an evaluation of this class of decision aids. We have experimented with several mathematical models that attempt to quantify and build a rigorous analytic foundation for the framework described above. For brevity, we omit the detailed mathematical description here. The basic result is that it appears that we can quantify the characteristics of any given decision situation in terms that a decision-maker can understand and assess. With further research we believe that it will be possible to apply the framework to the evaluation of real decision aids.

Stochastic Decision Trees

One important dimension for classifying decision problems is the degree to which they are similar to other decisions faced by the same decision-maker. At one end of the spectrum we might have decisions that are virtually identical to a number of other decisions, while at the other end of the spectrum we have decisions that are completely different in structure and content. For example, the decisions associated with launching successive aircraft from a carrier are very similar. On the other hand, situations such as the Cuban missile confrontation represent decisions that a typical decision-maker is likely to face only once in a lifetime.

Somewhere between these two extremes are decision situations that are similar in structure but have differences in degree that can be represented by a set of parameters. For example, the detection of an unidentified aircraft approaching a task force creates a decision situation that can be typified by several parameters such as weather conditions, readiness of the fleet, trajectory of the aircraft, and the consequences of failing to identify correctly and intercept the aircraft.

This kind of situation can be characterized by a decision tree that identifies the initial actions, potential responses, secondary actions and ultimate resolution of the encounter. The part of the problem that will change from situation to situation is the parameters of the decision tree (i.e., the probabilities and values placed on the outcomes). In principal, one could imagine drawing a decision tree for each generic type of situation and then characterizing the possible situations likely to be faced by the decision-maker by a joint probability distribution over the probabilities and values in the decision tree. We have termed such an entity--the decision tree structure and the joint probability distribution over its parameters--a "stochastic decision tree". One can imagine constructing such a stochastic decision tree for each of the major decision situations faced by a decision-maker, or at least those decision situations that have similar structures.

From a practical point of view it is not likely that the joint probability distribution required in the specification of the stochastic decision tree could be constructed directly from known data or experience. Instead, we could build models that characterize various scenarios likely to be faced by the decision-maker and then use these models to generate the required joint probability distributions.

Once a stochastic decision tree has been specified it is possible to calculate the potential value of a decision aid to the class of decisions it represents. For example, if one were considering a decision aid that supplies additional information about one of the state variables, then it is possible to calculate the value of that additional information to the class of decisions represented by the stochastic decision tree. This calculation can be carried out in one of two ways: An analytical expression can be derived for the value of each alternative action as a function of the probability for the specific variable under consideration. With this calculation it is possible to compute the expected value of additional information about the unknown variable. If this calculation proves to be too difficult, the value of the additional information about the stochastic decision tree can be calculated using a Monte Carlo technique. In this technique, random samples are chosen from the joint probability distribution for the parameters of the stochastic decision tree and the value of the additional information for this particular setting of the parameters is calculated. Successive samples then yield the distribution over the value of information.

Although stochastic decision trees seem like a natural way to describe decision situations in which the structure is constant but the parameters are not, more research is required on this topic. This research should be directed toward two important topics: a) validation of the general approach as a useful descriptor of military decision situations, and b) development of methodologies for using the concept to calculate the value of decision aids under the assumptions of normative decision-making.

REFERENCES

1. B. De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources", Studies in Subjective Probability, Kyburg and Smokler, Ed., Wiley, New York, 1963.

2. H. Dreyfus and S. Dreyfus, "Uses and Abuses of Multi-Attribute and Multi-Aspect Models of Decision-Making", Unpublished Report.

3. I. Hoos, Systems Analysis in Public Policy, University of California Press, Berkeley, 1972.

4. R.A. Howard, "Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968).

5. J.G. March, "Bounded Rationality, Ambiguity, and the Engineering of Choice", Bell Journal of Economics, 1979.

6. J.E. Matheson, "The Economic Value of Analysis and Computation", The IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, 1968.

7. A. Miller, T. Rice, M. Metcalfe, "An Analytic Characterization of Navy Command and Control Decisions," March 1979, ADA Final Report, ONR Contract Number N00014-78-C-0240.

8. A.C. Miller, J.V. Rowney, J.R. Payne, "The Naval Task Force Decision Environment," September 1974, SRI Technical Report NWRC-TR-8, ONR Contract Number N00014-74-C-0357.

9. P.A. Morris, Bayesian Expert Resolution, Ph.D. Dissertation, Stanford University, 1971.

10. P.A. Morris, "Combining Expert Judgments: A Bayesian Approach", Management Science, Vol. 23, 1977.

11. P.A. Morris, "Decision Analysis Expert Use", Management Science, Vol. 20, 1974.

12. R.C. Nickerson and D.W. Boyd, "The Use and Value of Models", to appear in Operations Research, 1979.

13. R. Radner, "Satisficing", Journal of Mathematical Economics, Vol. 2, 1979.

14. T.R. Rice, "The Value of Data For a Quadratic Decision Problem", Management Science, Vol. 24, 1978.

15. H.A. Simon, "The Behavioral Model of Rational Choice", Quarterly Journal of Economics, Vol. 69, 1955.

REFERENCES (continued)

16. H.A. Simon, "On How to Decide What to Do", Bell Journal of Economics, 1979.

17. R.D. Smallwood, "A Decision Analysis of Model Selection", The IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, 1968.

18. S.N. Tani, "A Perspective on Modeling in Decision Analysis, Management Science, Vol. 24, 1978.

19. L. Tribe, "Policy Science: Analysis of Ideology?", Philosophy and Public Affairs, 1972.

DISTRIBUTION LIST

Director, Engineering Psychology
Programs (Code 455)
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Defense Documentation Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Craig Fields
Director, Cybernetics Technology
Office
Defence Advanced Research
Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Naval Analysis Programs
Code 431
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Statistics and Probability
Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Commanding Officer
ONR Branch Office
ATTN: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Branch Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Analysis and Support Division
Code 230
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Opertions Research Program
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Branch Office
ATTN: Dr. J. Lester
Bldg. 114, Section D
666 Summer Street
Boston, MA 02210

Dr. James McGrath
Navy Personnel Research and
Development Center
San Diego, CA 92152

Dr. Donald A. Topmiller
Chief, Systems Effectiveness
Branch
Human Engineering Division
Wright Patterson AFB, OH 45433

CDR P.M. Curran
Human Factors Engineering
Division
Code 604
Naval Air Development Center
Warminster, PA 18974

Dr. S.D. Epstein
Analytics
2500 Maryland Road
Willow Grove, PA 19090

Dr. G. Hurst
University of Pennsylvania
Wharton School
Philadelphia, PA 19174

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 OLN
ENGLAND

Dr. Rex Brown
Decision Science Consortium
Suite 421
7700 Leesburg Pike
Falls Church, VA 22043

Dr. Arthur Siegel
Applied Psychological Services
Science Center
404 E. Lancaster Street
Wayne, PA 19087

I.R. Mirman
Assistant for Special Projects
HQ AFSC-DL
Andrews AFB, MD 20334

Dr. Edgar Johnson
Organizations & Systems Research
Laboratory
U.S. Army Research Laboratory
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Amos Freedy
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. C. Kelly
Decisions and Designs, Inc.
8400 Westpark Drive, Suite 600
P.O. Box 907
McLean, VA 22101

Mr. Georg Pugh
Decision Science Applications, Inc.
1401 Wilson Boulevard
Arlington, VA 22209

Mr. David Walsh
Integrated Sciences Corporation
1640 Fifth Street
Santa Monica,CA 90401

Mr. Joseph Wohl
MITRE Corporation
Box 208
Bedford, MA 01730