

MEMORANDUM
RM-4823-ARPA
DECEMBER 1968

DECISIONMAKING AMONG
MULTIPLE-ATTRIBUTE ALTERNATIVES:
A SURVEY AND CONSOLIDATED APPROACH

K. R. MacCrimmon

PREPARED FOR:
ADVANCED RESEARCH PROJECTS AGENCY

The **RAND** *Corporation*
SANTA MONICA • CALIFORNIA

Copyright © 1968
THE RAND CORPORATION

This study is presented as a competent treatment of the subject, worthy of publication. The Rand Corporation vouches for the quality of the research, without necessarily endorsing the opinions and conclusions of the authors.

Published by The RAND Corporation

MEMORANDUM
RM-4823-ARPA
DECEMBER 1968

DECISIONMAKING AMONG
MULTIPLE-ATTRIBUTE ALTERNATIVES:
A SURVEY AND CONSOLIDATED APPROACH

K. R. MacCrimmon

This research is supported by the Advanced Research Projects Agency under Contract No. DAHCl5 67 C 0141. Views or conclusions contained in this study should not be interpreted as representing the official opinion or policy of ARPA.

DISTRIBUTION STATEMENT

This document has been approved for public release and sale; its distribution is unlimited.



PREFACE

This Memorandum was written as part of RAND's continuing research on the nature of the R&D decisionmaking process. The Memorandum examines several methods and techniques which have been advanced in the literature for making quantitative and qualitative evaluation between multiple attribute alternatives. It is planned eventually to be incorporated as a chapter in a forthcoming book on R&D management. It is published at this time to bring this type of information to the attention of interested individuals involved in the making of complex decisions.

The author is a consultant for the RAND Cost Analysis Department.

SUMMARY

In all types of decision situations the alternatives among which we must choose are characterized by multiple attributes (or properties). Jobs may be characterized by prestige, location, salary, and advancement opportunities, for example, while weapon systems may be characterized by vulnerability, reliability, cost, yield, and other such diversely measured attributes. How is a decisionmaker to choose from among complex alternatives? Clearly, decisionmakers do choose--decisions involving very complex alternatives are made all the time. This is not to say, though, that these decisions could not be improved. Most decisionmakers in such situations would like a method that would help them process the attribute-value information for each alternative.

Various methods have been proposed to help the decisionmaker with multiple-attribute decisionmaking. These range from techniques which consider all attributes at once to those which consider just single attributes, or proceed sequentially over single attributes. The various methods discussed include Dominance, Satisficing, Maximin, Maximax, Lexicography, Additive Weighting, Effectiveness Index, Utility Theory, Trade-offs, and Non-metric Scaling. The literature on these methods is fragmented and often discussed in contexts other than that of multiple-attribute decisionmaking. Therefore, in this Memorandum we critically review the assumptions underlying each approach and examine its information requirements. In addition, each method is described both in a general way and using a formal, abstract mathematical representation. Two examples, the choice of a weapon system and of a space suit, are used to illustrate the discussion.

Most theoretical discussions of decisionmaking methods assume that the exact values of the attributes are known--for example, the exact rate of job advancement, or the exact speed of the weapon system. In actual cases of decisionmaking, such information is seldom known with such precision. In this Memorandum we discuss situations where there is uncertainty about attribute values, presenting, in the concluding section, a combination of methods which may often be used as a more reasonable and valid approach than the choice of any single method.

These methods and the discussion presented here are applicable not only to the particular examples given but also to other military and nonmilitary decisions.

ACKNOWLEDGMENTS

The author would like to express his appreciation to various staff members and consultants of The RAND Corporation for their suggestions on various drafts of this paper dating back to 1965. F. S. Pardee emphasized the relevance of these techniques to military decisionmaking and provided assistance throughout the study. E. D. Harris, S. Mennine, B. Pinkel, and E. S. Quade gave useful comments and critical review. The detailed suggestions by F. S. Timson on the final draft were very helpful.

CONTENTS

PREFACE	iii
SUMMARY	v
ACKNOWLEDGMENTS	vii
Section	
I. INTRODUCTION	1
The Decisionmaking Process	2
Decisionmaking and Goals	2
Decisionmaking Complexity	3
Choosing Among Alternatives with Many Attributes	4
Scope and Organization of this Memorandum	5
II. MULTIPLE-ATTRIBUTE DECISIONS	7
Attributes, Goals, Criteria, and Dimensions	7
A Military Systems Example	8
Symbolic Representation	15
III. EXISTING APPROACHES TO MULTIPLE-ATTRIBUTE DECISIONMAKING ..	17
Full Dimensionality	17
Dominance	18
Satisficing	20
Single Dimensionality	22
Maximin	23
Maximax	25
Lexicography	26
Additive Weighting	28
Effectiveness Index	33
Utility Theory	35
Intermediate Dimensionality	37
Trade-offs	38
Non-metric Scaling	40
Summary of Procedures	43
IV. EXTENSIONS AND COMBINATIONS OF PROCEDURES FOR	
MULTIPLE-ATTRIBUTE DECISIONMAKING	45
Uncertainty in Attribute Values	45
A Subsystems Example	49
Eclectic Approaches to the Multiple-Attribute Problem ...	53
Combinations of Approaches	53
Current Uses of Multiple-Attribute Methods	56
BIBLIOGRAPHY	59



I. INTRODUCTION

Most important decisions in our modern society are made at many differing levels within organizations. The prime function of top level organizational participants is usually decisionmaking, and the content of their decisions pertains directly to the goals of the organizations. Even at the lower levels of the organization, however, decisions are made that may be quite similar in form, although not in content, to higher level decisions. These decisions are not made in isolation of each other; the output of one decision often becomes an input to another (at both higher and lower levels). We may conclude from this that a hierarchy of decisionmaking exists within organizations (Cyert and MacCrimmon, 1967).

A particularly important example of organizational decisionmaking is found in the U.S. Government's decisions on military activities. On these decisions depend the lives of many people and the future and security of nations. In U.S. military affairs both the executive and the legislative branch of the government act in a decisionmaking capacity. There are many levels at which decisionmaking occurs. At the top level of the executive branch the President in his capacity as Commander in-Chief makes broad-scale decisions, often of a goal-setting or policy-making nature. He also performs the important function of relating military objectives to other national objectives. At the next level the Secretary of Defense engages in similar decisionmaking activities, primarily in the domain of military choices (although he may at times relate other objectives to these; see for example, McNamara, 1966). Below the Secretarial level the organization becomes more specialized, being partitioned into the various branches of the Armed Forces and some special cross-service areas. Decisionmaking at this level becomes particular to these sub-organizations. Such a chain of decisionmaking can be traced down to much lower levels where decisions might be made about minor procedural matters only.

THE DECISIONMAKING PROCESS

At each level we can observe a somewhat similar decisionmaking process--even though the content of the decisions will differ. Goals are being considered, information is being gathered and processed, evaluations are being made, and eventually a course of action is chosen. Although the process of decisionmaking centers around the element of choice, it should include all the activities leading up to the final choice, plus a post-choice stage of implementation and control (Dewey, 1933; Simon, 1957; Festinger, 1964). The representation in Fig. 1 gives a general idea of the process, where the feedback loops indicate the adaptive nature of the decisionmaking process.

In this paper we shall focus primarily on the evaluation and choice parts of this process. We assume that in the search activity various alternatives have been generated along with information about them, and that the decisionmaker's problem is to choose among these alternatives. To make his choice, however, the decisionmaker must refer to the goals explicated in the earlier parts of the decision process.

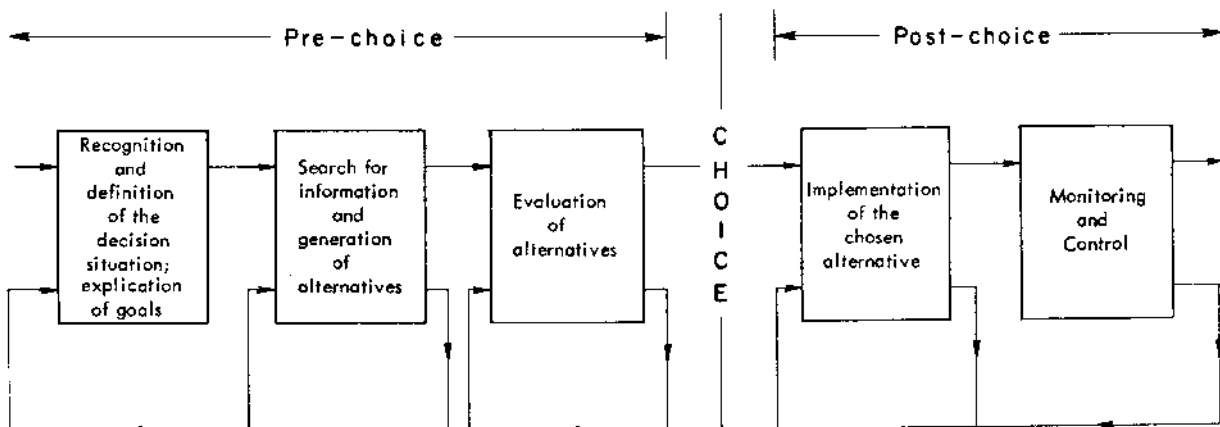


Fig. 1 -- The decisionmaking process

Decisionmaking and Goals

Decisionmaking is a response to a decision problem. Often this decision problem arises because of a discrepancy between the situation

at hand and the decisionmaker's goals (Newell and Simon, 1959). By a decisionmaker's goals we mean the organizational goals toward which a decisionmaker at a particular level is expected to direct his efforts. The goals will, of course, differ from organizational level to level, but the existence of a discrepancy between the goals at a particular level and the situation at hand is a common phenomenon. At the top level the goal may be a desired future military position with respect to other nations, while at a lower level it may be the more mundane matter of keeping maintenance costs under control. The methods to be discussed in this paper are applicable to decisionmaking at any level even though the examples used will pertain to decisionmaking at a particular level, and hence to the goals related to that level.

Decisionmaking Complexity

As noted above, we shall focus on the core of the decisionmaking process: evaluation and choice. Choices among alternatives are made difficult by two factors: (a) uncertainty and (b) constraints upon the information-processing capacity.

Often we do not know what the outcome of choosing any particular alternative will be. That is, the outcome and its associated desirability or utility may depend upon events outside the control of the decisionmaker. At times the primary events may be controlled by other organizations, as for example, in international conflict; at other times the events may result from some aspect of nature, as for example, the weather. In either case, the decisionmaker faces a situation of uncertainty.

Uncertainty is usually handled by game theory or decision theory. If the uncertainty is due to events controlled by other decisionmakers, the method has been to formulate game theoretic models (von Neumann and Morgenstern, 1947). When the uncertainty is not due to events controlled by a decisionmaker, the method used has been to formulate decision theoretic models (Savage, 1954).

Uncertainty is to some degree always present in decisionmaking, but often a more critical element in a decision situation concerns the

processing capacity of the decisionmaker--that is, his capacity to deal with the complexity of the problem at hand. Human beings have a limited information-processing capacity (Miller, 1956; Simon, 1957), and even when supplemented with computers, the capacity of organizational decisionmakers has very definite constraints.* Decisionmakers are forced to construct simplified representations because they can process only a limited amount of information about a situation (Cyert and March, 1963). Thus, even for decisions in which a relatively high degree of certainty is present, a decisionmaker may be faced with a very difficult situation. In this paper, we shall focus on models relevant to the constraints of the processing capacity of decisionmakers.

Choosing Among Alternatives with Many Attributes

A common type of processing-capacity complexity occurs when the alternatives are characterized by numerous attributes--attributes relevant to the decision situation at hand. A weapon system, for example, may be characterized in terms of performance, cost, availability date, and such factors. The attributes of these characteristics may be considered at various levels of aggregation; for example, performance may be broken out into range, delivery time, yield, vulnerability, accuracy, etc. The specific attributes to be considered (and their level of aggregation) will depend greatly on the goals or objectives toward which the decision is directed. The fact that there are usually multiple goals in any decision problem suggests that each of the alternatives directed toward the problem may have multiple attributes to be considered.

At this point we should note that decision problems with multiple attributes are quite common, and that many of the methods to be discussed were originally presented in a nonmilitary context. Some examples of alternatives with multiple attributes considered elsewhere are: computing equipment evaluation (Miller, 1966), job selection (Fishburn, 1965a), product design (Terry, 1963), wife selection (May, 1954), and leisure time allocation (Papandreu and others, 1957).

* Games of "perfect information"--such as chess--are clear examples of this type of complexity (von Neumann and Morgenstern, 1947, sec. 15).

SCOPE AND ORGANIZATION OF THIS MEMORANDUM

In multiple-attribute decision situations it is obvious that some alternatives are preferable when particular attributes are considered (those associated with particular goals), while other alternatives are preferable when different attributes are examined. As the number of relevant attributes and alternatives rises, the ability of the decision-maker to handle the problem decreases. The information processing requirements may rapidly exceed the decisionmaker's processing capacity.*

How, then, are such choices actually made? How should they be made? These are the questions that will be at the core of this study. We shall deal with decision problems where there are a number of known alternatives characterized by attributes common to each, and the problem is to make a choice among these alternatives (assuming a future environment that is relatively well specified). We shall examine various methods that have been suggested for dealing with the multiple-attribute decision problem. Most of these approaches have been developed in the course of theoretical studies by various researchers. In general, they are proposals specifying how such decisions should be made: that is, they are normative. The descriptive aspect of such decisions or how they usually are made, will not concern us here; this is a separate discussion that has been treated elsewhere by Shepard (1964), Klahr (1967), and MacCrimmon (1967) among others.

In the next section we shall introduce a weapon system-selection problem. The problem has, by necessity, been considerably simplified to present some common approaches to multiple-attribute choice. We hope that the use of a concrete example (rather than an abstract discussion) will give the reader a clearer understanding of the methods considered, so that he will be able to generalize to other decision problems. This example is used to demonstrate the similarities and differences in the various approaches--a matter that seems particularly ambiguous in the current literature. In the process we shall emphasize the information requirements of each method. In order to add precision

*The reader should bear in mind that "decisionmaker" refers both to individuals and to whole organizations, and that it is quite possible for a decision problem to exceed the processing capacity of the total organization.

to the verbal discussion, each decision approach concludes with an abstract representation of the method given in symbolic terms.

Our discussion of the existing approaches to the multiple-attribute decision problem suggests possible extensions and variants of the methods considered, some of which are explored in the concluding section. This concluding section is presented in the context of a subsystem example drawn from an exploratory or advanced development project; its generalization follows directly and should be obvious. Uncertainty is introduced into the decisionmaking situation by considering the case where the attribute values are not uniquely known but are random variables.

II. MULTIPLE-ATTRIBUTE DECISIONS

ATTRIBUTES, GOALS, CRITERIA, AND DIMENSIONS

Choosing among alternatives characterized by multiple attributes has been called a multi-goal, multi-criteria, or multidimensional decision problem. Throughout this Memorandum it will be called the "multiple-attribute decision problem." Multiple attributes could be expected when the decisionmaker has multiple goals, and as a consequence uses multiple criteria. Although the correspondences among these various terms should be clear, we prefer the term "multiple-attribute" because of its direct reference to the characteristics of the alternatives themselves, the objects among which the decisionmaker is choosing. The terms "goals" and "criteria" seem to be more general, and must be inferred behaviorally from the choices among alternatives (real or fictitious). (For a corresponding rationale in decision theory, see Savage (1954), sec. 2.6.) We shall use the term "attributes" throughout the Memorandum, but other terms with essentially the same meaning are "performance parameters," "components," "factors," "characteristics," and "properties."

Although the term "multidimensional" is commonly used to characterize the types of decision problems considered here, we shall reserve the term "dimensionality" for the psychological space in which the problem is characterized, that is, the spatial representation of this problem in the decisionmaker's mind. In the initial formulation of a problem, then, the number of attributes and the number of dimensions will be equivalent, but as we shall see, the most common approach to these decision problems is to attempt to reduce the dimensionality of the problem. This reduction commonly results in alternatives that, while described by the original number of attributes, are evaluated in a space of smaller dimensionality (Klahr, 1967). Some of the methods we shall discuss attempt to reduce the problem to one dimension (usually a numerical value), while others try to reduce it, if not to one, at least to no more than a few of the original number of dimensions. Examples of the former are the various weighting procedures (e.g., Churchman and

Ackoff, 1954). The new non-metric scaling methods of Shepard (1962) are examples of the latter. There are a few methods that retain the original dimensionality. All of these methods are discussed in Sec. III of this Memorandum.

A MILITARY SYSTEMS EXAMPLE

The similarities and differences of the various approaches to be discussed may stand out more clearly in the context of an example. Although many possible examples could be used, we have chosen one at the military systems level, and thus necessarily must make highly aggregative and simplifying assumptions.

Suppose for a particular anticipated military requirement, say, within the general war mission, we must make a choice among designs for a future weapon system.* Let us consider three possible types of systems--call them X, Y, and Z. They could be as diverse as aircraft, satellites, and missiles, but perhaps more frequently would be various configurations of aircraft, for example. We shall assume that each would be expected to have a different force size to fulfill the requirement, say 250, 1000, 500 units respectively; and we shall assume that \$10 billion is expected to be available to acquire and maintain the chosen system.**

The relevant attributes in this decision problem would be generated by careful political-military consideration of this particular requirement within the overall mission and possibly also future uses of the proposed system. Perhaps in this case it might be permissible to consider only a single aggregate attribute, for example, "expected target destruction" (per dollar). It might be more meaningful, however, to

* See also Hitch (1964), Quade (1964), and Fisher (1966), for discussions of similar decision problems.

** All values are completely fictitious. The purpose of the numerical values is solely to indicate what type of quantitative information might be required.

break this up into "probability of reaching the target" and "expected destruction if target is reached," and to also include some indication of "future possibilities." These attributes, in turn, could be decomposed still further with the first involving speed, distance, delivery time, vulnerability, accuracy, etc., while the second would involve such things as yield, reliability, etc. Future uses might be characterized in terms of range, payload flexibility, and so forth. Each of these attributes could, of course, be further broken down.

The decisionmaker would initially generate a reasonably exhaustive list of relevant attributes at each level. The extent of the subdividing of attributes would depend on the levels necessary to capture the essence of the problem and to attain some reasonable prospect of measurability. Not all the attributes generated need be used, since some may be redundant while others would not serve to discriminate between alternatives. An example of the first would be the redundancy contained in the attributes of speed, distance, and delivery time, while an example of the second would be the availability date if the same date applied to all systems.

Suppose such a logical analysis has taken place and the attributes remaining are range, delivery time, total yield, accuracy, vulnerability, and payload delivery flexibility.* The characterization of each system in terms of these attributes is given in Table 1. Force size and cost, although not shown in the table, remain, of course, attributes of the system.

The form of this table implies that we can characterize each system uniquely by each set of attributes. Yet we know that in most decision problems this information, even if available, would be uncertain. Such uncertainty would be especially likely in the research and development decision domain. At this point in our discussion, however, we shall assume that the attribute values are known uniquely, because this is the assumption most of the models make. We shall not take up

*The particular attributes used are for illustrative purposes only. If the reader feels other attributes would be more important in such a decision problem, he is free to add or substitute them. See also Fisher (1966), p. 2.

Table 1
A WEAPON SYSTEM DECISION PROBLEM

Attributes	Alternative Systems		
	X	Y	Z
Range (n mi)	10,000	8,000	5,000
Delivery time (hr)	5.0	0.5	1.0
Total yield (MT)	100	50	80
Accuracy (high-low)	average	low	high
Vulnerability (high-low)	average	high	very low
Payload delivery flexibility (high-low)	high	low	average

the extra complication of uncertainty about attribute values until the concluding section of this Memorandum.

We shall also assume that an alternative can be defined in terms of the attributes it possesses, with one value associated with each of the attributes, While this may seem an indirect way to define an alternative, a little reflection should convince the reader that he can describe missile systems in terms of attributes, just as he defines alternative jobs in terms of attributes such as salary, organization, work functions, reporting level, and other considerations.

Note that both qualitative and quantitative attribute values are represented in Table 1. Range, delivery time, and total yield are naturally expressed in numerical terms. Accuracy would also be considered numerical, in terms of circular error probable (c.e.p.) for single units of a system, but when considering the accuracy of the whole system it may be difficult to get a meaningful single number. Similarly, attributes such as vulnerability and payload delivery flexibility, while perhaps capable of numerical expression for a particular system alternative, are difficult to express quantitatively for cross-system comparisons.

Assuming for the moment that we are constrained to the information given in Table 1, which system should we choose? The choice, of course, will depend greatly on the characteristics of the mission, including the expected enemy response to the system. If we want the system with the lowest vulnerability we would choose system Z. For the system with the quickest delivery time we would choose system Y. But if we want the system with the greatest range we would choose system X. By thus focusing on single attributes only we could justify a choice of any of the systems. But remember that we (i.e., the decisionmaker) thought that each of the six attributes was relevant to the mission. Our task is to make a decision considering all six attributes. Inspection of Table 1 shows that part of the information is qualitative, part is quantitative, and that some of the numerical values differ by orders of magnitude. Most common approaches to decisionmaking would first place all attribute-value information on a comparable numerical scale.

Let us examine the question of quantifying the qualitative attribute values.* The attributes of accuracy, vulnerability, and payload delivery flexibility all have non-numerical values. One common quantification procedure,** given information of this kind, is to construct a scale associating the qualitative terms with numbers on the scale. For example, we might choose a 10-point scale and calibrate it in one of a number of possible ways. We could start with the end points, giving 10 points to the maximum attribute value that is practically realizable (or alternatively, physically realizable) and 0 points to the minimum attribute value that is practically realizable (or physically realizable). The midpoint would also be a basis for calibration since it would be the breakpoint between values that are favorable (or better than average) and values that are unfavorable (or worse than average).*** The type of scaling used could be of major importance in the final outcome, but since our principal interest is not in scaling procedures but rather in the methods that follow scaling (and also, as we

*For another approach, see Schelling, 1964.

**This type of scaling is used both in normative models of this kind (Fishburn, 1965; Sigford and Parvin, 1965) and also in descriptive psychological studies (Stevens, 1959; Gulliksen and Messick, 1960).

***Average is used in the sense of the intermediate anticipated value considering the prospective state of the art, for the proposed initial operational capability date.

shall see, in the methods that do not require scaling), we shall assume that one scaling procedure is selected and is used consistently. Thus, we can proceed by assigning attribute values more favorable than "average," a score of more than 5.0 points, while values less favorable would be assigned less than 5.0 points. Note that for the attributes of accuracy and payload delivery flexibility values of "high" are more favorable but that for the attribute of vulnerability a value of "low" is more favorable than "average."

Taking the attribute of payload delivery flexibility for the moment, how many points should we assign to the value "high"? Any value on the scale between 5.1 and 10.0 will satisfy the above constraints; it will be on the scale and greater than the value assigned to "average."

Commonly, values close to 10.0 will be reserved for extremely favorable characteristics; thus, for instance, "very high" might be assigned the value 9.0. This in turn constrains "high" to the interval 5.10-8.9. We might assign it the scale value 7.0. On the low end of the scale, "very low" might be associated with the value 1.0 and "low" with the value 3.0. A similar scale might also be used for the attribute of accuracy. Note, though, that such a scale must be reversed for the attribute vulnerability because low values for this attribute are more desirable than high values. These scale values are diagrammed in Fig. 2.

It should be obvious that a numerical assignment such as that given above is highly arbitrary. Many other scales are possible. Sometimes attempts are made to provide some consistency checks (Fishburn, 1965a). Such checks are desirable but make the scaling procedure a very involved activity posing many hypothetical questions to the decisionmaker.

The procedures that derive these numerical values use addition and multiplication operations* across attributes; from this, we note several implications. This type of scaling assumes that a scale value of 9.0 is three times as favorable as a scale value of 3.0. Thus, for example, for accuracy an attribute value of "high" is three times as

*Thus the scale values are necessarily cardinal, not just ordinal, even though the scaling procedure seems ordinal in nature.

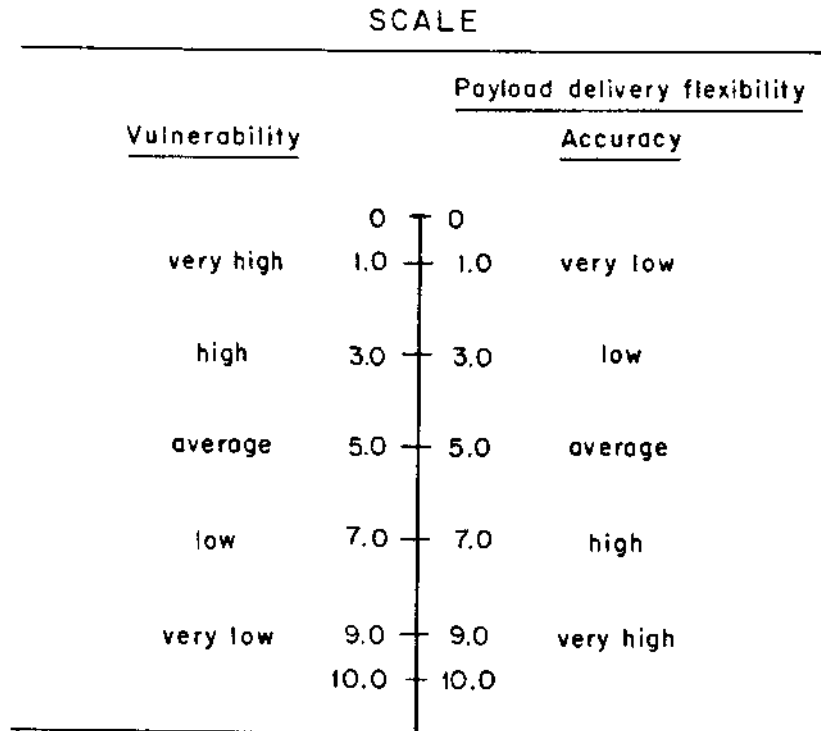


Fig. 2 -- Assignment of values on a numerical scale

favorable as one of "low." In addition, it assumes that the difference between "high" and "low" is the same as the difference between "very low" and "average" (4 scale points). Further, the combination of values across attributes implies that the difference between any two specific values (say, "high" and "low") is the same for each attribute; that is to say, that the difference between "high" and "low" accuracy is the same as the difference between "high" and "low" payload delivery flexibility.*

The reader should be aware of the arbitrariness of such a method. It should be obvious that various scaling assumptions can lead to quite different results. Because the 9, 7, 5, 3, 1 assignment is probably as reasonable (or unreasonable) as any other for our illustrative purposes here, we shall use the scale of Fig. 2. When these numerical

*This can be at least partially accounted for in the weights (to be discussed later) attached to each attribute, but this uses up some of the degrees of freedom in assigning weights and makes the process especially difficult.

values are substituted for the corresponding qualitative ones in Table 1, we obtain the values shown below in Table 2.

Table 2
ASSIGNING NUMERICAL VALUES TO ALL WEAPON SYSTEM ATTRIBUTES

Attributes	Alternative Systems		
	X	Y	Z
Range (n mi)	10,000	8,000	5,000
Delivery time (hr)	5	0.5	1
Total yield (MT)	100	50	80
Accuracy	5.0	3.0	7.0
Vulnerability	5.0	3.0	9.0
Payload delivery flexibility	7.0	3.0	5.0

Even after the qualitative attribute values have been quantified, we must still make all the numerical values comparable. Note that even for initially comparable values we can arbitrarily choose the units of expression. Total yield, for example, was expressed in megatons in our example, while range was given in nautical miles. Yield could have been expressed in kilotons and range in thousands of nautical miles. But in order to make these values comparable with each other, and with the numerical assignments of the values that were originally qualitative, we converted all values to the same scale, in our case the 10 point scale of Fig. 2, as follows: range (n mi x 10^3); total yield (MT x 10^{-1}); accuracy, vulnerability, and payload delivery flexibility in the original scale factors (1-10). Since lower values of delivery time are preferable to higher values, we must invert delivery times in order to have the final comparable values of the table in a unified format (high numbers are preferred to low ones). We therefore invert delivery times by dividing the delivery time in hours into the constant "5," that is, (1/hr x 5). The table of attribute values under such an assignment would appear as in Table 3*.

* Note, for example, that the highest scale value for each attribute does not have to get 10 points; rather, the extent to which the objectives of the mission are satisfied will determine the position of the values on the scale.

Table 3

COMPARABLE NUMERICAL VALUES FOR THE SAME PROBLEM

Attributes	Alternative Systems		
	X	Y	Z
Range (n mi x 10 ³)	10.0	8.0	5.0
Delivery time (1/hr x 5)	1.0	10.0	5.0
Total yield (MT x 10 ⁻¹)	10.0	5.0	8.0
Accuracy	5.0	3.0	7.0
Vulnerability	5.0	3.0	9.0
Payload delivery flexibility	7.0	3.0	5.0

SYMBOLIC REPRESENTATION

At this point let us make a more abstract representation of the multiple-attribute alternative problem. First, let us distinguish between particular attributes.

Let us call a particular attribute A_i , where $i = 1, 2, \dots, n$ if there are n attributes. Thus A_i is the name (e.g., vulnerability) of attribute i . Let a_i be the random variable representing the values A_i can take; it will also be used for particular values of the random variable such as "average vulnerability." Note that the attribute values are not necessarily numerical.

Let us define an alternative in terms of the attributes it possesses; there will be one (and only one) attribute value for each of the n attributes. Formally, then, an alternative is an element of the Cartesian product $A_1 \times A_2 \times \dots \times A_n$. The generic form of a particular alternative $(a_1, a_2, \dots, a_n)^j$ will be denoted as \mathcal{A}^j . The set of all alternatives is obviously a subset of the Cartesian product of attributes, i.e., $\{\mathcal{A}^j\} \subseteq \prod A_i$. The i^{th} attribute of the j^{th} alternative will be denoted by a_i^j . When the attribute values have been scaled and they are all on a comparable, numerical (cardinal) scale, the point value assigned to the i^{th} attribute of the j^{th} alternative will be denoted by \bar{a}_i^j .

We assume that a simple ordering relation of preference, \succ , read as "is preferred to," is defined on each attribute, A_i , separately. Thus, for example, the decisionmaker can specify whether he prefers the attribute of range to be 8000 n mi or 5000 n mi, other things being equal. There is no need for the preference ordering in the basic attribute values to be monotonic. The term $a_i^1 \succ a_i^2$, for example, means that the value of attribute i for alternative 1 is preferred by the decisionmaker to the value of attribute i for alternative 2.*

This notation will be used to characterize the various methods described in Sec. III. These formal definitions will be given in addition to, and not in lieu of, informal verbal descriptions of each method in order to add rigor and clarity to the discussion. The reader who does not find the formal descriptions helpful is encouraged to skip them.

In the next section we shall discuss the various methods for resolving the multiple-attribute decision problems, by examining methods which treat the problem as presented in Table 1 format as well as those requiring the more demanding and arbitrary format of Table 3.

* This can also be written as $a_i^2 < a_i^1$.

III. EXISTING APPROACHES TO MULTIPLE-ATTRIBUTE DECISIONMAKING

Even though some of the approaches we discuss have been intended only as descriptive models of such decisionmaking, we are most interested in their normative implications. Our focus will be on the methods available to a decisionmaker faced with a multiple-attribute decision problem. In the process of describing the methods we shall highlight the assumptions and information requirements of each.

We shall organize our discussion into general categories representing the dimensionality of the problem that the method handles. If we call the original dimensionality "n," meaning that n attributes have been selected to characterize each alternative, then a method may either deal with all n attributes or reduce the problem to some lesser number. It would be desirable to consider all n attributes without omitting any from consideration and without imposing any arbitrary assumptions in order to collapse them. Such methods, however, either do not help reduce the complexity of the original problem or are relatively weak because they often do not give a unique solution. As we shall see, such methods are most useful in reducing the set of alternatives for final consideration.

At the other extreme are methods that reduce the problem of n attributes to a single dimension. Such methods operate either by imposing conditions that remove n-1 dimensions from consideration or by imposing assumptions that allow the n dimensions to be mapped (or reduced or combined) into a one-dimensional space.

Intermediate between these extremes lie some methods that either consider two or more attributes at a time or that reduce the space to a dimensionality less than n but greater than 1. We shall consider these categories in that order.

FULL DIMENSIONALITY

The two principal methods that treat multiple-attribute problems in their full dimensionality are dominance and satisficing. By

treatment in the full dimensionality we mean that each dimension or attribute is considered separately and independently. It must, by itself, meet certain requirements. A disadvantage or unfavorable value in one attribute cannot be offset by an advantage or favorable value in some other attribute. Each attribute must stand on its own.

Dominance

When comparing all alternatives, if some one alternative has higher attribute values for all attributes, we say that this alternative "dominates" the others. We can weaken this notion somewhat and say that if one alternative is at least as good as the other alternatives on all attributes, and is actually better in at least one of them, then this can still be considered the dominant alternative. Conversely, if one alternative is worse than some other alternative for at least one attribute, and is no better than equivalent for all other attributes, then we can say the former alternative is dominated by the latter.

In the missile example of the previous section (see Table 1) none of the three alternatives is dominated in this sense. We can say, therefore, that each of them is "admissible." If, however, system Y and system X had the same delivery time, then we could say that system X dominates system Y. That is, when each of the six attributes is considered separately, system X is at least as good as system Y, and for some of the attributes, such as range, system X is actually better.

If we had to consider a modified system Y (modified because it is like Y except for a delivery time identical to system X), we could remove it from consideration because this system would now be distinctly inferior to system X. In this manner, dominance can be a useful method for reducing the size of a decision problem. Dominance is not a very powerful method for making a final decision, however, because there are usually a number of alternatives remaining after the method is applied (Wohlstetter, 1964). In the unmodified example, all three systems, X, Y, and Z, remain. In the modified example, X and Z both remain because although X dominates the modified Y it does not dominate Z. To say that a number of alternatives remains means that, when comparing any

pair, each alternative has a higher attribute value than at least one of the other attributes considered.

In the dominance procedure the decisionmaker decides solely upon the basis of the question whether one attribute value is more preferred than another. Thus, as long as we can state that "high" accuracy, for example, is preferred to "low" accuracy, we have enough information to establish possible dominances. Numerical information about attribute values is unnecessary but can, of course, be used just as easily if the information happens to come in numerical form. A straightforward application of the dominance procedure, then, implies no assumption about the decisionmaker's degree of preference for particular attribute values; he does not have to establish how much more he prefers "high" accuracy to "low" accuracy. Nor does dominance require the decisionmaker to assess the relative importance of each of the attributes. Further, because each of the attributes is considered separately, the decisionmaker does not have to be concerned with such questions as, for example, how much he prefers "low" vulnerability to "high" accuracy.

Dominance is one of the most easily applied and commonly accepted decisionmaking procedures. The formal concept has been around for a long time, and has undoubtedly been applied intuitively for even longer. In economics its formal use was pioneered by Pareto (1848-1923) and hence non-dominated alternatives are often called "Pareto-optimal." Heavy reliance was placed upon the notion of dominance by Wald in his development of statistical decision theory. In this field it is commonly known as "admissibility" (Wald, 1950; Blackwell and Girschick, 1954).

Even after removing the dominated alternatives, the decisionmaker is often left with a complicated problem. Although the dominance concept seems to have been used mainly in normative decisionmaking, perhaps because of its obviousness, the dominance concept seems to be quite descriptive of the procedures followed in actual decisionmaking (see, for example, MacCrimmon, 1967).

A formal statement of dominance is as follows: Denote one alternative by $(a_1, a_2, \dots, a_n)^1$ and a second by $(a_1, a_2, \dots, a_n)^2$. Then

we say that the second alternative dominates the first if $a_i^1 \leq a_i^2$ for all i , and further $a_i^1 < a_i^2$ for some i .

Satisficing

A second method that treats the problem in its full dimensionality also has strong intuitive appeal. In order to apply the satisficing method described by Simon (1955), the decisionmaker supplies the minimal attribute values he will accept for each of the attributes. These are his minimal values, goals, or specifications. In the missile example suggested above suppose that the decisionmaker specified the following minimal requirements: range, at least 5000 n mi; delivery time, at least 5 hr or less; yield, at least 60 MT; accuracy, at least average; vulnerability, no worse than average; and payload delivery flexibility, at least average. Given these minimal acceptable values, both system X and system Z are acceptable; that is, they satisfy these requirements. System Y fails to satisfy them because its total yield is too low, its accuracy is too low, its vulnerability too high, and its payload flexibility too low.

As with the dominance procedure, often after applying the satisficing method we are still left with a number of feasible alternatives. In fact, we may assume that such procedures have been applied, if not formally at least intuitively, throughout the whole design process, and that the alternatives we are given have passed at least some such initial screening; for example, we can assume that all designs with a range less than 5000 n mi have been dropped from further consideration. With the satisficing method, in contrast to the dominance procedure, we can successively change the minimal requirements and hence successively reduce the feasible set.* For example, we may now decide that the accuracy should be high, and thus system X would be dropped, leaving us with only system Z satisfying all the requirements. When used in an iterative fashion we can sometimes narrow down the alternatives to a single choice. Since dominance cannot be used in this manner, satisficing is a more powerful decisionmaking tool.

*Alternatively, if none of the alternatives met the specified requirements, the requirements could be changed so as to increase the feasible set.

Neither satisficing nor dominance requires that the attribute information be in numerical form. We need only know which value is preferred, other things being equal. Further, with satisficing (as with dominance) we do not need information on the relative importance of the attributes. In a sense, though, the information requirements for satisficing are greater than for dominance because we need to have information on the minimal acceptable attribute values. It should be noted that if we simply use minimum cutoff values for each of the attributes, none of the alternative systems gets credited for especially good attribute values. Thus, system Y gets dropped even though it has the quickest delivery time of any of the systems. The attempts to credit alternatives with especially high values suggest other procedures to be discussed later.

Because of its strong intuitive appeal, satisficing has long been used. Simon, who has formalized it, uses it primarily as a descriptive model to show how people do make decisions. He asserts that people form and use minimal attribute values in the manner described above, choosing the first satisfactory alternative. If this is so, the decision process is sequential, and a final choice will depend on the order in which the alternatives are uncovered. Thus, in the example shown in Table 1, if the decisionmaker uses the minimal values previously specified, he will choose either system X or system Z, depending on which alternatives he discovered (or analyzed) first.

Since we are interested here in the normative implications of these procedures, we need not accept questionable descriptive implications. Instead, we can expect the decisionmaker to apply minimal requirements to all available alternatives as a first pass in reducing the set with which he must deal. He can then tighten some of the minimal attribute values and make a second pass in an attempt to reduce the set still more. This procedure can be repeated as often as the decisionmaker desires.

A formal statement of the (normative) satisficing procedure is given below:

Suppose a set of minimal attribute values (g_1, g_2, \dots, g_n) is defined on $A_1 \times A_2 \times \dots \times A_n$. An alternative j is satisfactory only if $g_i \leq a_i^j$ for all i . Any unsatisfactory alternative j' -- that is, an alternative for which $a_{i'}^{j'} < g_{i'}$, for some i' -- is dropped from consideration.

SINGLE DIMENSIONALITY

Dominance and satisficing are the two main procedures that treat multiple-attribute decision problems in their full dimensionality. From the discussion it should be apparent that both methods are particularly effective in reducing the set of alternatives to be evaluated in a final choice. They are relatively weak in the elimination of alternatives, although the satisficing procedure may be strengthened by applying it iteratively. Dominance utilizes an alternative-alternative comparison: that is, each of the available alternatives is compared with each of the other available alternatives. Satisficing, on the other hand, involves an alternative-goal approach in which each of the alternatives is compared with a minimal goal vector, although this vector can also be considered as a fictitious, minimally acceptable alternative. It is our opinion that both dominance and satisficing are most useful when combined with the procedures now to be discussed.

In each of the following procedures the n attributes characterizing an alternative are reduced to a single dimension. The first three methods retain, as the single dimension, one of the n attributes of the original problem. An integral part of these methods, then, is specifying which one of the attributes is to be considered. The other three procedures to be discussed attempt to map the information from the n dimensions to a numerical scale. In the mapping operation the focal point of interest will be the types of functions proposed.

Maximin

There is an old saying that "the chain is only as strong as its weakest link." In choosing among chains we would examine the weakest link in each and presumably select the chain with the strongest weakest link. This procedure can be used in decision-making situations. As decisionmakers we would examine the attribute values for each alternative, note the lowest value for each alternative, and then select the alternative with the most acceptable value in its lowest attribute. This is called selecting the maximum (across alternatives) of the minimum (across attributes) values, or the maximin.

Under this procedure a very high degree of comparability is required. We not only have to compare values across attributes for each single alternative, we also have to characterize each alternative by its lowest attribute value when comparing it with other alternatives. Yet if these lowest attribute values come from different attributes, as they sometimes do, we may be basing our final choice on single values of attributes that differ from alternative to alternative. Because such a high degree of comparability is necessary with the maximin, all attributes must be measured on a common scale -- which, however, need not be numerical. The maximin procedure reduces the characterization of an alternative to its value for a single attribute; all other (n-1) attributes for a particular alternative are ignored.

In applying this procedure to the missile example of the previous part of the Memorandum, we note that the lowest attribute value for system X is 1.0 (for delivery time); the lowest attribute value for system Y is 3.0 (for accuracy, vulnerability, and payload delivery flexibility); and the lowest attribute value for system Z is 5.0 (for range, delivery time, and flexibility). Since the largest of these three minimum values is 5.0, system Z would be selected because it has the maximum-minimum value.

It should be clear that this procedure utilizes only a small part of the available information in making a final choice--only one attribute per alternative to be exact. So even if an alternative is outstanding

in all but one attribute, another alternative that is only average on all attributes would be chosen over it. The maximin procedure, then, has some obvious shortcomings in decision situations. What is appropriate for chains is not necessarily appropriate for general decisionmaking. The procedure is reasonable for chains because all the links are used at the same time and they are essentially interchangeable, in that no one link is worth more, or performs a different function, than some other single link.

The maximin and its reverse, the minimax procedure is widely used in game theory. In fact, the fundamental theorem of game theory involves the equivalence between the maximin and the minimax (von Neumann and Morgenstern, 1947, p. 153), but we shall not discuss this further here. Maximin is a reasonable procedure in game theory because it is assumed that one is operating against a self-optimizing opponent (or opponents). Thus, instead of the dimensionality of the problem arising from attribute values, it arises from the choice of alternatives available to the other player (see, for example, Schelling 1964). It could be said that the opponent is choosing an "attribute" such that his loss is lowest (assuming a constant sum game). Given this action of the opponent, we can take, for each of the alternatives available to us, the lowest "attribute value" as our expectation of his choice. We are then essentially restricted to these values as our range of payoffs, and thus we maximize among these values.

We can see from this example that applications of exceptionally ingenious procedures in game situations can lead to ludicrous results in individual decisionmaking. In a general decisionmaking situation a maximin method would be reasonable only if the decisionmaker assumed that some malevolent nature was trying to inflict the worst possible outcome on him. Any decisionmaker who believes this and makes his choices accordingly deserves the outcomes he will receive.

The maximin* procedure may be abstractly stated as:**

$$A^{\dagger} = \max_{\{A^j\}} \text{er} \min \bar{a}_i^j .$$

Maximax

We can reverse the procedure described above by characterizing an alternative by its best attribute value rather than its worst attribute value. In this case we would identify the highest attribute value for each alternative, then compare these maximum values to select the alternative with the largest such value, the maximax procedure.

Under the maximax procedure the highest attribute value for system X is now 10 for range and total yield, the highest value for system Y is 10 for delivery time, and the highest value for system Z is 9 for vulnerability. Maximizing these maximum values would lead then to a choice of either system X or system Y.

Note then that this procedure, as with the maximin procedure, may evaluate different attributes in a final choice among alternatives. Thus, we may be evaluating the range and total yield of X against the delivery time of Y and both of these against the vulnerability of Z. Furthermore, the maximax method has the same incompleteness as the maximin method. No attention is paid to the n-1 lower-valued attributes of an alternative. Even if system Z had all 9's while system X had only a single 10 and all the rest of X's attribute values were very low -- say, around 2 or 3 -- the maximax procedure would lead to a choice of X.

*The useful, but infrequently used, expression "max'er" (maximizer) identifies a member of the index set rather than a value in the domain of maximization. Thus, the expression above picks out that A^j that yields the $\max[\min \bar{a}_i^j]$.

**The notation \bar{a}_i^j could be weakened here to mean "comparable but not necessarily cardinal numbers," as discussed above.

The comparability assumptions and incompleteness properties of maximax, then, do not make it a very useful technique for general decisionmaking. However, just as maximin may have a domain in which it is quite reasonable, such as choosing among chain-like alternatives, maximax may also be reasonable in some specific decision situations. An implicit assumption of this method is that, in some sense, we are weighting the purposes (or attributes) equally; that is, we do not care which single purpose is the one pursued.

The maximax criteria can be abstractly stated as:

$$A^{\dagger} = \max_{\{A^j\}} \max_i \bar{a}_i^j .$$

Lexicography

Lexicography can be called a single-dimensional technique only in the sense that one dimension at a time is considered. The specifications for this method include the procedure for comparison across attributes, albeit one at a time. In some decision situations a single attribute will seem to predominate. By this we mean that it is obviously the most important one to that decisionmaker. In the same way, in choosing a job, the factor of salary will override all other factors for some people. One way of treating this situation is to compare the alternatives on this one attribute. If one alternative has a higher attribute value (e.g., the highest paying job) than any of the other alternatives, it is chosen and the decision process ends. If no single alternative is predominant in this, the primary attribute, then the non-maximal alternatives (that is, those with less than this common maximum value) are dropped from further consideration. The remaining alternatives are then compared in the next most important attribute. The process continues in this lexicographic fashion either until a single alternative is chosen or until all attributes have been considered.

In the missile example, suppose that accuracy is by far the most important attribute. Since system Z has the highest accuracy, it is the preferred alternative. System X and system Y both have a lower accuracy than Z, but if we use the lexicographic procedure, there is no need to consider the other attributes. Because of this, we see that the lexicographic procedure suffers from the same type of incompleteness as do maximin and maximax. Lexicography is somewhat more demanding of information than these two, because it requires a ranking of the importance of the attributes, whereas maximax and maximin do not. Since it does not require comparability across attributes as did maximin and maximax, information of the most basic kind (i.e., Table 1) can be used. None of these three unidimensional procedures, maximin, maximax, and lexicography, requires numerical information.

The term "lexicographic" comes from the correspondence between this procedure and the arrangement of words in a dictionary. These words are first ordered alphabetically by their first letter--this tells us that "battle" should be listed before "war." However, it does not tell us in which order to list "battle" and "bomb." In such a case we consider the second letter, and can thus determine that "a" comes before "o" in the English alphabet, so that "battle" should be listed first. This, of course, requires some sort of "preference" order and the order of proceeding from left to right alphabetically; when applied to general decisionmaking, the lexicographic method requires information on the preference among attribute values and the order in which attributes should be considered. In both of these cases, though, we only need ordering or ranking information and not (necessarily) numerical values.

Because of its limited information requirements, lexicography has received serious consideration as a decision technique in a number of areas. It has been considered as an alternative to a numerical measure of the von Neumann-Morgenstern utility (see next section) by Hausner (1954), Thrall (1954), Debreu (1959), Chipman (1960), Aumann (1964), and Radner (1964) among others.

Lexicography may be stated formally as follows: Suppose the attributes are ordered so that A_1 is the most important attribute

to the decisionmaker, A_2 is the next most important, and so forth. Then, take

$$\{\mathcal{A}^\dagger\} = \max'_{\{\mathcal{A}^j\}} a_1^j .$$

If this set has a single element, then this element is the most preferred alternative. If there are multiple maximal alternatives, consider

$$\{\mathcal{A}^{\dagger\dagger}\} = \max'_{\{\mathcal{A}^\dagger\}} a_2^j .$$

If this set has a single element, then stop and select this alternative. If not, consider

$$\{\mathcal{A}^{\dagger\dagger\dagger}\} = \max'_{\{\mathcal{A}^{\dagger\dagger}\}} a_3^j .$$

Continue this process until either (a) some $\{\mathcal{A}^t\}$ with only a single element is found which is then the most preferred alternative, or (b) all n attributes have been considered, in which case, if the remaining set contains more than one maximal element, they are considered to be equivalent.

Additive Weighting

In many multiple-attribute choice situations we can think in terms of the relative importance of each attribute. In our example, is accuracy more important than vulnerability, and if so how much more important? Importance judgments, of course, are contingent upon the mission, but we assume that the decisionmaker has these in mind. If the decisionmaker can choose a numerical measure of importance, then he may weight each attribute value by this measure to get a weighted average of the contribution of each alternative. He will then select the alternative with the highest weighted average.

Such a weighting procedure does not disregard any of the original n attributes, because all n attribute values of an alternative are used to form the weighted average. This is the primary difference between the additive weighting approach and those discussed to this point. Because the procedure of weighting uses the regular arithmetical operations of multiplication and addition, however, the attribute values must be both numerical and comparable, as in the form of Table 3. This is clearly a much more stringent requirement than any of those underlying the previous procedures. Further, it is also necessary to find a reasonable basis on which to form the weights reflecting the importance of each of the attributes.

Let us consider the following form of our missile example: Suppose that by introspection and analysis we as decisionmakers can determine the following weights for each of the attributes: range 0.10, delivery time 0.10, total yield 0.20, accuracy 0.25, vulnerability 0.15, and payload delivery flexibility 0.20. In other uses of this procedure, cost or any other attribute could, of course, be included. The weights have been normalized to sum to 1; in this way we can tell at a glance the relative importance of each as a percentage of the total. Multiplying these weights by the corresponding attribute values for each alternative and then summing across alternatives we get the following weighted averages: system X, 6.5; system Y, 4.6; and system Z, 6.7. Using this method we as decisionmakers would choose system Z because it has the highest weighted average.

Let us examine this procedure in the context of the above example to note more clearly the various information requirements and assumptions made. Let us consider the weights first. In the example above, for instance, the weights imply that total yield is twice as important as range. What does it mean to say that one attribute is twice as important as another? Another implication is that accuracy is as important as range and vulnerability combined. Again, the ambiguities of such judgments must be apparent. Although some techniques have been suggested for forming such weighted values, for example Churchman and Ackoff (1954), Eckenrode (1965, and Miller (1966), the decisionmaker must be very careful in his assignment of numerical values.

Even after weights are formed, we must still put the attribute values into a numerical and comparable form so that they can be multiplied by the weights, and then summed into a weighted average. They must be comparable because we are going to combine across attributes; a "high" value for one attribute must receive approximately the same numerical value as "high" values of other attributes. If this is not the case, if for example, a "high" value for range is 9000 while yield is 90, then the normalized weights are meaningless since the range value will surely dominate the choice. In such a case either the weights or the attribute values must be rescaled. It is usually easier to effect the latter, because rescaling the weights for such a purpose would remove their relative importance in the scale.

When weights are assigned and attribute values are numerical and comparable, some arbitrary assumptions still remain. Note that 10 multiplied by 0.10 and 4 multiplied by 0.25 both yield the same product: 1.0. If we interpret 10 as being an exceptional attribute value and 4.0 as being a below average attribute value, then this identity implies that an exceptional delivery time and a somewhat below average accuracy just offset each other. By "offset each other" we mean that both make the same contribution to the weighted average. Thus, there exist some difficulties in interpreting the output of the multiplication of attribute values by weights. At this point it might also be necessary to reconsider the weights and attribute values assigned: Does an exceptional delivery time exactly offset a below average accuracy? In fact, can such judgments be made?

Let us suppose that at this stage in the analysis the decision-maker is satisfied with the weighted values of each alternative on each attribute. Even then, addition across attributes may require some unrealistic assumptions. Because we assigned the attribute values on each attribute separately, the operation of adding the weighted attribute values for a given alternative assumes that there are no complementarities (or "spill-overs" or "spin-offs") in such problems. Thus a high total yield, for example, is valued in itself, irrespective of the accuracy, delivery time, or any other attribute. A simple weighting can give misleading results if, for example, a high total yield is

of little value unless the accuracy is at least average and the delivery time is not too long.

By using the weighting procedure described and by restricting our focus to the two attributes of total yield and accuracy, we get the following results:

<u>Attributes</u>	<u>Weighted Average</u>
Very high total yield, very low accuracy	$1.80 + 0.25 = 2.05$
Very low total yield, very high accuracy	$0.20 + 2.25 = 2.45$
Very high total yield, very high accuracy	$1.80 + 2.25 = 4.05$
Very low total yield, very low accuracy	$0.20 + 0.25 = 0.45$

What serious meaning can we attach to these numbers?

Using the scaling values of Fig. 2 and the weights given above,* a very high total yield has a weighted value of 1.80 (i.e., 9.0×0.20), a very high accuracy has a weighted value of 2.25 (i.e., 9.0×0.25), a very low total yield has a weighted value of 0.20 (i.e., 1.0×0.20), while a very low accuracy has a weighted value of 0.25 (i.e., 1.0×0.25).

A system with a very low accuracy may provide very little chance of destroying the target, even if the total yield is very high; a system with a very low total yield may result in minor destruction of the target even if the accuracy is very high. What sense does it make to say that a system with both very high yield and very high accuracy is only 98%, or $(4.05-2.05)/2.05 \times 100\%$, better than the very low accuracy-very high yield system, and only 65%, or $(4.05-2.45)/2.45 \times 100\%$, better than the very low yield-very high accuracy system? It may well be that the first two systems have little overall value since both fail to produce target destruction, while the third system (with both high yield and high accuracy) has a very high value. The point of this example emphasizes that attributes cannot often be considered

* Range 0.10, delivery time 0.10, total yield 0.2, accuracy 0.25, vulnerability 0.15, and payload delivery flexibility 0.20.

separately and then added together. Further, as long as the attributes are considered separately, this example cannot be modified to yield reasonable results by some new rescaling.

Because of the complementarities between the various attributes, the approach of weighted averages may give misleading results. But when the attributes can in fact be considered separately (i.e., when there are essentially no important complementarities), the additive weighting procedure described can be a very powerful approach to multiple-attribute decisionmaking. Since a single number is arrived at for each alternative, and since these numbers will usually be different, this procedure often leads to a unique choice. For this reason, and because it has some intuitive appeal, it is frequently used. The additive weighting procedure has been used for choices as diverse as choosing jobs (Fishburn, 1964; Miller, 1966), making business investment decisions (Churchman, Ackoff and Arnoff, 1957), choosing computing equipment (Miller, 1966), selecting products (Terry, 1963), selecting military hardware (Aumann and Kruskal, 1958); Bryan, 1964) and choosing disarmament strategies (Aumann and Mashler, 1966). Some sort of weighting procedure is used in almost all multiple-attribute decisionmaking. Unfortunately, the conditions--that is meaningful weights, numerical and comparable attribute values, and independence of attributes--necessary for its successful application are seldom checked out, although they may hold in some cases. For a further discussion of tests of necessary conditions see, for example, Fishburn (1964).

In addition to its wide use in practice, additive weighting procedures are receiving renewed theoretical emphasis. In psychology, linear models have received attention in studies of both choice behavior (e.g., Adams and Fagot, 1959) and conjoint measurement (Luce and Tukey, 1963; Luce, 1966). Even in economics, where additive models have long been explicitly rejected because of the unrealistic non-complementarity assumption, there have been indications of new explorations (for example, Strotz, 1957; Houthakker, 1960). There are numerous related theoretical developments in operations research--after all, linear programming is such a model--even in those that directly consider multiple-attribute choices (for example, Pfanzgal, 1959; Aumann, 1964; Fishburn, 1964, 1965, 1966).

The additive weighting procedure described above can be formally stated as follows: Suppose we have a weighting function (w_i) mapping from ΠA_i to the real half line. Then the most preferred alternative is the A^{\dagger} , that is, the

$$\max_{\{A^j\}} \sum_i w_i \cdot a_i^j .$$

Note that while the weighting procedure is linear in the attribute scores \bar{a}_i^j , it is not necessarily linear in the basic attribute values, a_i^j .

Effectiveness Index

There is no requirement, of course, that the functional form of a combined attribute approach must be additive. A more general weighting model could entail multiplications, exponentiations, or any type of mathematical operation. For example, we could form something like

$$\text{exp} - \left[\frac{\sqrt{\text{Vulnerability} \times \text{Time}}}{\text{Accuracy} \times \text{Log}(\text{Yield}) + \text{Log}(\text{Range} \times \text{Flexibility})} \right]$$

for our missile example, where this number could be interpreted as a type of probability index of expected target destruction. Clearly, in a real example the relation could be much more complex and would certainly contain numerous parameters having significance for the problem.

This procedure has very strict information requirements, similar to the additive weighting model discussed above. While no explicit weights need to be assigned over attributes, the form of the functional relation (and particularly its parameters) play the same role in an implicit way. The attribute values must be assigned cardinal values and the comparability can either be done directly as in the preceding method, or it can be done implicitly in the functional form itself. Thus, all the reservations about additive weighting apply to this procedure when it is used directly, that is, when decisionmakers are asked to use their judgment in scoring the attribute values and setting up the overall

function. In fact, this procedure is much more demanding since it is much easier, in general, to put meaningful numbers into an additive weighting than into one in which more complicated expressions are being used. An example of a multiplicative procedure applied to research project evaluation is given by Mottley and Newton (1959).

The general weighting models obviously have a greater potential than additive ones for capturing the important interdependencies in a decisionmaking situation. The arbitrariness of the functional form allows a representation of all types of complementarity relationships. The great difficulty is in specifying the function that is to capture these relations. In some cases a multiplicative relation is quite natural (e.g., speed and distance), and perhaps through training decision-makers can be taught to think in terms of more complicated expressions. At the present time, though, additive weighting is much easier to comprehend and enjoys much greater popularity.

The major use of the general weighting procedures is not in the systemization of decisionmaking judgment; they are primarily used in systems analysis work where logical relations among the attributes exist. The task then is to formulate the appropriate functional form to capture these logical relations. It should be clear at this point that this is not the problem we are addressing. In a sense we are considering the problem at the next level, that is, how do we decide among alternatives using the inputs from systems analysis, but also considering important intangibles or other attributes that could not feasibly be considered in the logical formulation.

The form of the general expression in such a logical analysis is given various labels. Some common ones are "mission-related measure of effectiveness," "effectiveness index," "rating index," and "figure of merit" (see Quade, 1964 and Charnes, Cooper and Thompson, 1965). When cost is included in the consideration, along with the performance measures, this is often called a cost-effectiveness study. As noted above, then, these effectiveness measures and the costs, along with intangibles and other attributes not subjected to such analysis, would form the attribute inputs to the decision situations we are studying here.

The formal statement of these general weighting or effectiveness index models is a slight generalization of the preceding case. The most preferred alternative is the \mathcal{A}^\dagger that is the

$$\max_{\{\mathcal{A}^j\}} f(A_1, A_2, \dots, A_n)^j,$$

where the range of the function is the real half-line. Obviously, for such a general functional form the cross partial derivatives need not be zero (as they are for additive weighting models) and hence complementarities can be defined.

Utility Theory

In decision problems where there is a high degree of uncertainty about the outcomes of the various alternatives, it is more useful to consider the scope of the possible outcomes than the attributes of the alternatives. Thus, the multidimensional aspects of this formulation come, not from multiple attributes as such, but from the multiple events that can occur and can yield different outcomes for any given alternative. Any decision problem that is considered in the form of information about the multiple attributes of each alternative must be recast if it is to be amenable to an (axiomatic) utility analysis.

Instead of considering range, yield, vulnerability, and other attributes of weapon systems, we may use the information on the attributes to assign a utility function to outcomes of various uncertain events that may impinge on the performance of the weapons system. For example, one uncertain event might be whether or not the enemy can establish a defensive system of a particular level of effectiveness. For this event we then consider the effect on each of the alternatives (as a whole) we are considering--that is, the effect of this event with regard to a particular system's range, attack capability, etc.--and then assign a utility function or value to the particular alternative for the uncertain event considered. Such an assignment would

be carried out for each uncertain event. If there is more than one event and its complement to be considered, we would form joint events. Then probabilities of the occurrence of each particular (joint) event would be formed. These would be multiplied separately by the corresponding utilities for each alternative to yield an expected utility function for each alternative. The alternative with the largest expected utility would then be chosen.

The utility theory described here can be traced back to Daniel Bernoulli (1738), but it was its axiomatization by von Neumann and Morgenstern (1947) that provided the basis for applying it in decision situations. The basic idea is that the decisionmaker, if he obeys certain axioms, can assign numbers to uncertain outcomes--numbers that are measureable on an interval scale--that is one with an arbitrary origin and scale unit, such as temperature. The uncertainty is handled by assigning probabilities to the (sometimes unique) uncertain events. This probability does not have to be a relative frequency or in any other way be "objective," but is personal to the particular decisionmaker. As long as the decisionmaker's beliefs obey the mathematical axioms of probability, his personal beliefs may be called "probabilities." The basis for these subjective probabilities was given by Thomas Bayes (1763). The synthesis of utilities and subjective probabilities was first given by Ramsey (1931), with perhaps the clearest current statement due to Savage (1954).

While this approach is especially useful when the main element in a decision problem is uncertainty, when the main element is processing complexity--that is, much information that is reasonably certain for each alternative--the approach described above tends to further complicate the problem. Instead of helping to sort out the information on attributes, it assumes that the decisionmaker can do this processing in his head and come out with a figure for the utility of an alternative with all those attributes for a given uncertain event. The consideration of the uncertain events, which really reflects the purpose to which the alternative will be put, should be an element

considered by the decisionmaker,* but it is perhaps unrealistic to ask him to come up with a utility figure without providing him with some means of sorting out the attribute information for each alternative.

In both the procedures described earlier and those to be described later in which the attribute information is treated directly, the purposes to which the alternative will be put are assumed to have been considered in the earlier phases of the problem (see Fig. 1), and thus to have led to the consideration of the particular alternatives with which we are confronted. In the process of forming an expected utility for each alternative, it should be clear that the whole problem is considered and that it is eventually mapped down to a single dimension, that is, a numerical scale, a value of which is the expected utility. Since applying this utility approach would require presenting information in a form different from that of Figs. 1 and 2, we shall not work through an example of utility theory as applied to the weapons system problem. (The brief description in the first paragraph of this section should indicate the type of approach used.) We discuss utility theory here only because it is a very important approach to decision problems--but problems of a form slightly different from those considered in this Memorandum.

INTERMEDIATE DIMENSIONALITY

Intermediate between the two major categories of procedures discussed so far are those procedures that consider more than one but less than the full number of dimensions. Of the two procedures to be discussed here, the first is closely related to the linear weighting and effectiveness index approaches and has been widely used as a normative technique, especially in design, rather than final choice, situations. The second procedure stems from recent developments in psychometrics and has heretofore been considered only descriptively. We shall pursue some of its normative implications.

*Of course, the techniques described throughout this Memorandum have (at least implicitly) embodied consideration of the mission and the environment as we have noted. The present discussion is concerned with explicitly attaching quantitative values to uncertainty about the environment.

Trade-offs

In multiple-attribute situations we are often interested in getting information of this sort: if we can settle for a lower value on one attribute, how much can we increase the value of another attribute? For example, if we are willing to give up 10MT of yield, how much quicker delivery time can we get, other things remaining equal? That is, we are asking to what extent we can "trade" yield for time. In a number of decision situations the consideration of trade-off information allows us to make the alternatives much more comparable than they were initially. Thus, we can consider trade-offs that make alternatives equivalent for several attributes, and then examine more carefully the remaining attributes. In this way we can reduce the dimensionality of the actual problem to something less than the original number of dimensions.

In a decision problem having alternatives characterized by n attributes, we can form $\binom{n}{2}$ trade-offs between pairs of attributes. In our missile example of 6 attributes we can form 15 trade-offs for each alternative. Considered in this way, trade-offs have certainly not reduced the dimensionality of the problem. Not all trade-offs will be relevant, however; in fact, when two attributes are independent, it may not be possible to get a higher value on one even though we are willing to give up a great deal of value in another attribute. In other cases it may not be possible to obtain some of the trade-off information. Even with this in mind, many trade-offs still may remain, and consequently may tend to confuse the decisionmaker.

Compounding these difficulties is the fact that having only one trade-off value for each pair of attributes is a very special case. This case is special because the trade-off ratio is independent of the value of each separate attribute. This is obviously not the usual case. Up to a certain point perhaps we can buy more yield at a constant dollar cost, but after this point, it may become increasingly more expensive to buy a higher yield figure. Such interdependence is clearly true for most pairs of attributes, although a constant trade-off ratio may be a fair approximation over some range of the values of the two attributes. The more trade-off ratios to be

considered for a given pair of attribute values, the more complex the decision problem becomes. Because of the necessity of generating so much new information, we shall not work through the missile example using the trade-off procedure.

Trade-off information is more useful when designing multiple-attribute alternatives than when choosing among final versions of them (Schamberg, 1964). In final choice situations, the multiple-trade-off ratios cannot easily be consolidated into an overall figure. Even after trade-off ratios are formed, they may not do much to reduce the complexity of the original problem, and unless treated carefully may actually increase it. The most useful trade-off information is that involving cost and some other attribute.

Trade-off ratios are often used in economics (Baumol, 1959). An alternative term commonly used for trade-offs is "marginal rates of substitution," i.e., the rate at which one attribute may be substituted for another at the margin. Although not dealing directly with attributes of particular alternatives, when the economist considers various products in a commodity bundle or various factors in a production plan, he is in fact considering attributes of that bundle or plan. The economics of consumer demand depicts these trade-offs by indifference curves reflecting various combinations of the two commodities (i.e., attributes). The slope of, or the tangent to, these indifference curves is the marginal rate of substitution or the trade-off ratio. In a similar manner we could draw two sets of indifference curves for pairs of attributes: one representing pairs of attribute values where the decisionmaker was indifferent to the combinations; and a second one representing curves showing the technological rate at which one attribute may be traded for another. Only when the indifference curves are straight lines do we get a single trade-off ratio, or one that does not depend on the level of attribute values.

Trade-offs may be formally defined as ratios of the partial derivatives of two attributes. They may be obtained either directly as individual ratios or from a function relating the attributes, as for example, $f(A_1, A_2, \dots, A_n)$ where f may be one of the linear or more

general weighting functions previously described. If the point of interest for attribute A_i is a_i , then the trade-off ratio between attributes A_i and A_j is given by

$$\frac{\partial A_i}{\partial A_j} \Big|_{A_1=a_1, \dots, A_n=a_n} = \frac{(\partial f / \partial A_j) \Big|_{A_1=a_1, \dots, A_n=a_n}}{(\partial f / \partial A_i) \Big|_{A_1=a_1, \dots, A_n=a_n}}$$

If the value of this partial derivative is independent of the particular values of the attributes A_i and A_j then the trade-off ratio is constant. Interdependencies between the attributes are shown by the cross partial derivatives, that is, $\partial^2 f / \partial A_i \partial A_j$. Clearly, for the simple additive weighting models all these cross partial derivatives equal zero.

Non-metric Scaling

Several of the procedures discussed to this point have required the decisionmaker to evaluate the relative importance of attributes. In lexicography, for example, the decisionmaker had to judge the single most important attribute and then characterize the alternatives by this single attribute. To the extent that ties occur, of course, the decisionmaker must decide upon the most important remaining attribute; in effect, he must be prepared to give a complete ranking of the importance of all attributes, although this will seldom be the case.

On the other hand, the additive weighting procedure, which represents alternatives by the entire number of attributes, requires assignment of (cardinal) weights to each attribute that will reflect its relative importance in the decision problem. Intermediate between these extremes one can, of course, construct procedures representing alternatives in k dimensions, where $1 < k < n$, using judgments about the relative importance of the attributes. For example, the 3 most important attributes may be considered while the remaining $n - 3$ attributes may be neglected in the choice process. Such a procedure would have to consist of two main parts, the first specifying how to

choose the 3, or more generally k , attributes to be considered, the second specifying how the decisionmaker should choose his alternative from the reduced problem after considering the 3, or k , attributes.

Let us consider one such procedure using both these parts. In order to reduce the number of attributes, let us ask the decisionmaker to judge the similarity of a number of (fictitious) alternatives, given in pairs. By systematically varying the attribute values we can discover which k attributes he seems to attend to. Procedures for obtaining this attention information may also yield a spatial representation of each alternative in the k dimensional space. If this spatial representation does in fact truly characterize each alternative, it may also be used to determine the decisionmaker's most preferred alternative. Suppose we specify an ideal object, one with the most preferred values on each of the attributes. Then after placing this object in the appropriate place in the spatial representation (according to its attribute values), we can determine the distance of each of the other alternatives from this ideal. The alternative that is closest to the ideal object would then be the chosen alternative.

This scaling procedure suggests several points. For one thing, each pair of alternatives given in the judgments must be ranked in terms of its similarity. This requires quite fine discrimination and may be quite difficult for the decisionmaker. In addition, the distance measure used to form the spatial representation assumes that the attributes are independent (i.e., non-complementary). The attribute information can be in any form. It could, for example, be in the non-numerical, non-comparable form of Table I since the scaling procedure itself produces numerical, comparable values on each of the k dimensions, $k \leq n$. It should also be noted that none of the k dimensions necessarily corresponds with single attributes of the original problem.

The techniques for obtaining such spatial representations come from the field of psychometrics where they were first developed by Shepard (1962), then later extended by Kruskal (1964). These procedures are closely related to multiple factor analysis (see for

example, Torgerson, 1958). Klahr (1967) was the first to apply spatial representation in going from similarity judgments to preferences in a most interesting study of the multiple-attribute decisions made by college admission officers when choosing among student applicants. The work on these techniques has thus far focused solely on descriptive aspects, but there seem to be some normative implications along the lines suggested above worth further development. Because of their descriptive emphasis and their complexity of analysis, however, we shall not apply this scaling procedure to the weapons system decision.

The non-metric scaling procedure may be formally stated as follows. First, in forming the spatial representation we need a ranking of the similarity of all pairs of the q (fictitious) alternatives. Thus if \odot means "is more similar than" we have a chain $A^1 A^2 \odot A^1 A^3 \odot A^2 A^3 \odot \dots$, etc., for all $q(q-1)/2$ pairs. Let x_{A^j} be a point in t dimensional space representing alternative A^j ; then x_{A^j} has coordinates $(x_{A^j}^1, x_{A^j}^2, \dots, x_{A^j}^t)$. The distance between any two points $x_{A^j'}$ and $x_{A^j''}$ is defined to be

$$d_{A^j' A^j''} = \sqrt[r]{\sum_1^t (x_{A^j'} - x_{A^j''})^r}$$

where $r \geq 1$. For $r = 2$ we have the Euclidean metric. We want to construct the smallest space of dimension k , such that the distance rankings will be congruent with the similarity rankings, that is, $d_{A^1 A^2} < d_{A^1 A^3} < d_{A^2 A^3} < \dots$. After constructing such a space, we then want to locate an ideal object $A^\#$ in this k dimensional space, $A^\# = (x_{A^\#}^1, x_{A^\#}^2, \dots, x_{A^\#}^k)$. The most preferred alternative is the A^\dagger such that

$$A^\dagger = \min_j d_{A^\# A^j}$$

Summary of Procedures

In the previous sections we have considered a number of procedures relevant to decisionmaking among alternatives characterized by multiple attributes. We have focused on their normative implications even though some of them had previously been considered only descriptively. In order to remind the reader of the main properties of these procedures we present a summary in Table 4 for each of the procedures discussed. All of the procedures presented in the table have been discussed in the Memorandum. Any new material is self-explanatory. As with any summary, the following table oversimplifies similarities and differences among the approaches. The reader should consult the previous sections for a more careful description.

IV. EXTENSIONS AND COMBINATIONS OF PROCEDURES
FOR MULTIPLE-ATTRIBUTE DECISIONMAKING

UNCERTAINTY IN ATTRIBUTE VALUES

Up to this point we have assumed that each attribute value was known, and that that value was unique. But we recognize that the information available to the decisionmaker is often highly uncertain, especially in research and development decisionmaking (Meckling, 1964; Quade, 1964). There are various ways of representing the decisionmaker's uncertainty. Perhaps the simplest is by using a range of values rather than a point estimate of attribute values. By giving an interval or a range of values, as for example, total yield from 65 MT to 90 MT, instead of a single point estimate, say 80 MT, we assume that the actual value will lie within this interval. Let us consider how the use of such a range rather than a single value would affect the procedures we have discussed.*

The dominance procedure readily extends to the use of a range by representing uncertainty in attribute values. Various possible interpretations of dominance can be made in such situations. In a strong form of dominance an alternative may be considered dominated if, for each attribute, all of its maximum values (that is, the values at the high or most preferred end of the range) are pairwise smaller than the minimum values of some other alternative. A weaker form of dominance would state that an alternative is to be considered dominated if, in a pairwise consideration, its maximum and minimum values are never better (and are sometimes worse) than the maximum and minimum values, respectively, for some other alternative over all attributes of those alternatives. Both the strong and weak forms of the dominance notion can be usefully applied at different times. We should observe, though, that these extensions would tend to make the dominance procedure even less applicable than it is when single values are used, because far fewer alternatives would be inadmissible.

*The reader should note that we are using the term "range" here in the statistical sense of an interval of the highest and lowest values of a particular variable. Earlier "range" was used in the physical sense of the distance that a weapons system could feasibly travel and arrive on target.

The satisficing method can be extended in several ways. One would be to regard an alternative as unsatisfactory only if any one (or more) of its maximum values is lower than the goal value for the corresponding attribute. This particular interpretation would result in a larger number of satisfactory alternatives, because, in effect, the most favorable value, rather than the most likely, of any alternative would pass the satisficing test. If this consequence is undesirable, other extensions of satisficing with different implications can be made. We should note here that if the goal values also are uncertain, if the decisionmaker doesn't want to pin himself down to a single value, then the extension of satisficing to uncertainty becomes less amenable to analysis.

In considering extensions of the maximin and maximax procedures, we are in effect adding another level. Each alternative is now characterized by two endpoint values and we can take either the maximum or minimum end point. For example, we could consider a very conservative "maximinimin" approach. In this case we would use the minimum value for each attribute (that is, the low end of the range); then, in a manner identical with the regular maximin, characterize each alternative by the minimum value for any attribute; and then select the alternative with the largest such minimum value. Maximinimin differs from maximin in the extra step of minimizing across attribute values for each attribute before minimizing across attributes for each alternative. Maximin cannot do this because there is only one such attribute value. In a similar manner a maximinimax, a maximaximin, and a maximaximax procedure could be developed. The maximinimin procedure discussed above is much more conservative than the regular maximin and the maximaximax procedure would be much more optimistic than the ordinary maximax.

Lexicography can be extended in a manner similar to the extension of dominance. After identifying the most important attribute (as in regular lexicography) we can reject any alternative whose largest attribute value is lower than the lowest attribute of some other alternative, for that (most important) attribute. Then considering only the reduced set of alternatives, we can apply the procedure

to the second most important attribute, and continue this process as long as necessary. In general, we would have to consider more attributes with this modified lexicography and probably would need a complete ranking of the relative importance of the attributes. In addition we would probably end up with more than one alternative after considering all n attributes. However, even though this procedure is weaker than the regular lexicography, it is less arbitrary and seems to be a quite reasonable approach.

The additive weighting method cannot be so easily extended to incorporate uncertainty. This method requires a good deal more computation, and presents two problems. In order to characterize each alternative, for example, we may consider taking all combinations of highest and lowest attribute values. This would yield 2^n weighted averages for each alternative instead of the single weighted average used in the regular additive weighting procedure. Which of these weighted averages would then represent the alternative? Further, we could expect that the decisionmaker would be uncertain about the exact weights to be used, because the weighted averages for each range could be expected to overlap considerably. If we assume any uncertainty in the weights of the attributes, this method would clearly become too cumbersome, computationally, to be effective. Using the highest and lowest attribute values, for example, we would have to compute $2^n \times 2^n$ weighted averages.

Undoubtedly simpler methods could be devised for handling uncertainty in attribute values and relative attribute weights, but any method is going to entail a considerable increase in computational requirements, and will of necessity require arbitrary choice by the decisionmaker.

The similarity between the trade-off approach and the additive weighting approach, including the partial derivatives described in the preceding section, leads us to the belief that allowing uncertainty in attribute values would complicate the information requirements in the trade-off approach as it does in additive weighting.

Utility theory has been developed especially for handling uncertainty--but uncertainty of a different kind; non-metric scaling,

like the weighting methods, would tend to break down under the increased computational demands. We shall not examine extensions of the utility theory and non-metric scaling procedures further because of their special nature.

The additive weighting and trade-off methods may seem appropriate when we consider the simple deterministic version of the multiple-attribute decision problem; but when we consider the more realistic case of uncertainty in the information, these procedures become computationally too cumbersome to be effective. While dominance, satisficing, and lexicography might seem too arbitrary and simple in the deterministic form, their simplicity becomes a virtue when uncertainty is considered. Thus, we conclude that if the problem is realistically represented in deterministic form, it may make sense to attempt to apply the more formal weighting and trade-off methods, assuming the other necessary conditions such as independence are satisfied. If uncertainty in the attribute values or attribute weights is a vital part of the problem, however, then the less formal methods of dominance, satisficing, and lexicography merit serious consideration.

Using a statistical concept of range is not the only way of accounting for uncertainty. Another--and an even more demanding way--is to introduce probability distributions. It might be possible to formulate a marginal probability density function of the attribute values for each attribute describing every alternative. By considering each of the functions separately, we could then analyze the effects of uncertainty. A preliminary attempt at treating this form using the dominance method was given by Fox (1964). A second and similar way of incorporating probabilities assumes that we can get the joint probability densities over all characteristics. The analysis then proceeds by Monte Carlo methods (Quade, 1964). For an interesting exploration of these possibilities see Timson (1968). Probability density functions are worthy of further research, but we shall not address ourselves to them in the remainder of the Memorandum. We note only that the same conclusions reached in the preceding paragraph are likely to hold: The more formal weighting and trade-off methods will

become even more cumbersome, while the dominance, satisficing, and lexicographic procedures will become relatively more reasonable. In the next section we shall consider uncertainty by means of a range of attribute values. We shall use an example which demonstrates the ways in which a combination of the methods already described may prove valuable.

A SUBSYSTEMS EXAMPLE

Let us consider another military example of a multiple-attribute choice problem, this time a subsystems decision that might occur in developmental management. While the example is loosely based on characteristics of the pressure suit for the Apollo mission (Kovit, 1965), the designs described here are hypothetical.

Suppose the decisionmaker must choose one design from among several alternative pressure suit designs for future development. After considering the mission, the attributes chosen as most relevant are mobility, comfort, life support, pressurization, metabolic temperature control, meteoroid protection, thermal protection, and ultraviolet radiation. These attributes are not the only factors that could be considered; others might include, for example, bulkiness, weight, cost, etc. Nevertheless, let us assume that the other attributes are either fixed at one particular level by the mission requirements, are identical for all designs, or are unimportant. For the chosen attributes we have simplified and reduced the amount of information to the minimum necessary for purposes of the example. On radiation we consider only ultraviolet rays, excluding infrared and the intense visible. Such restriction is necessary to keep the example within reasonable limits. The characterization of each suit design is given in Table 5 in terms of each of these attributes.

Following the consideration stressed in the preceding discussion, we assume some uncertainty in the attribute values, and hence, provide a range of values for each attribute. We assume that these ranges bound the expected performance of the suit for that attribute, in that any attempt to improve a characteristic beyond the stated limits would entail a major redesign effort. Note that, other things being equal, for the first six attributes higher values are preferred,

Table 5
CHARACTERIZATION OF FIVE PRESSURE SUIT DESIGNS

Attribute	Alternative Design				
	I	II	III	IV	V
Mobility	average to good	fair to average	poor to fair	fair to good	fair to average
Comfort	good, but perhaps minor waste management problems	good, but perhaps some minor hot and cold spots	good, no problems expected	average, minor pressure points expected	average, minor hot and cold spots expected
Life Support (hr)	1-2	3-4	4-5	2-3	2-3
Pressurization (psi)	3.6-3.8	3.6-3.7	3.8-3.9	3.7-3.8	3.6-3.7
Peak Metabolic Temperature Load (Btu/hr)	1600-1900	1800-2000	1900-2200	2000-2200	1700-2000
Primary flux Meteoroid Protection (km/sec)	25-30	30-35	30-35	27-32	25-32
Maximum Thermal gain (Btu/hr)	180-240	210-230	150-170	230-270	230-260
Total ultraviolet Radiation (%)	5-6	3-5	1-2	4-7	3-7

while for the last two attributes lower values are preferred. With all of these considerations in mind, let us consider how we could apply the procedures discussed in this Memorandum.

When extended to cases involving uncertainty in the attribute values, the dominance procedure has two basic interpretations: strong and weak. Using the strong definition of dominance, we note that we do not obtain any dominated alternatives: There is no single design whose maximum values are no better than the minimum values for some other design over all attributes. (If there had been such a strongly dominated design, it surely would have been removed from consideration before this point in the decisionmaking.) Design II does dominate design V, however, in the weak definition; that is, for all attributes, the maximum and minimum values of design V are never better (and sometimes worse) than the corresponding maximum and minimum values of design II.

In order to use a satisficing approach we would have to formulate a meaningful set of minimal attribute values--or specifications--for each attribute. This might be relatively easy to do in a case such as this where the mission is well defined. For example, we might set forth the following minimal requirements: average mobility, average comfort, 3 hr life support capacity, 3.7 psi pressurization, 1900 Btu/hr peak metabolic heat load, 30 km/sec primary flux meteroid protection, 250 Btu/hr maximum heat gain, and 5 percent total ultraviolet radiation. Applying these requirements would rule out design I (because its maximum life support is at best 2 hours) and design III (because its mobility will not even be average). By applying these goals, though, we should realize that we are ruling out the suit design with the best mobility (design I) and the suit design (III) with the best performance on most of the attributes including comfort, life support, pressurization, thermal protection, and radiation. A recognition of these consequences might stimulate a revision of the requirements.

For the lexicographic procedure we would need information on the relative importance of the attributes. If we know, for example, that the most important attribute is life-support capacity, we would choose

design III, since it has the maximum life-support capacity--assuming we use the weak interpretation of strict preference discussed under dominance above. If we used the strong interpretation of strict preference, we would also have to consider design II since it may achieve 4-hr life-support capacity. We would then have to consider the second most important attribute, say pressurization, and would choose design III because it strongly dominates design II on this attribute. We extend the same definitions from dominance to lexicography, because lexicography is a type of dominance procedure on one dimension at a time. Because mobility is not one of the assumed two more important dimensions, the poor to fair rating of design III would not be considered.

An uncertainty version of the maximin or maximax procedures would require a scaling of attribute values to make the values comparable across attributes. An examination of the information given in Table 5 should demonstrate to the reader that obtaining such comparable values is not a simple task. For this reason, in addition to the general inappropriateness of maximin and maximax types of procedures, we shall not consider them in detail here. Utility theory also will not be applied to the pressure suit design example because of its inappropriateness when dealing with information in the form of Table 5. In addition, non-metric scaling methods will not be applied because of computational demands.

In attempting to apply the additive weighting procedure, we note the need to elaborate on the information given in several ways. As with the maximin and maximax procedures, we must have all the attribute values on a common scale--and more specifically for additive weighting procedures, this scale must be numerical. So, for example, we have to ask how many points to assign to "average mobility," and given this assignment, how many points to assign to "3-hr life support" to make these scales congruent. After constructing such a comparable numerical scale we still need information on the relative importance of each of the attributes. These weights must also be numerical. Thus, although with lexicography we needed to know that, for example, life support and pressurization were the two most important attributes, with the additive weighting procedure we must determine how much more important

life support is than pressurization. Even after developing the numerical attribute scales and the numerical weights, we still have the difficulty of deciding how to deal with the ranges of attribute values. With design I, for instance, do we take the scale value representing "average mobility," or the scale value representing "good mobility"-- or do we consider both? Clearly, we do not want to compute $256 (=2^8)$ weighted averages for each design, which would be the case if we took all combinations of the endpoints of the range of attribute values. Compounding these difficulties is the requirement for independence of attributes when adding across attributes in forming the weighted averages. In summary, note then that an additive weighting procedure when applied to situations that are not highly abstracted poses some very serious problems.

Related to the additive weighting procedure is the trade-off approach. Trade-off information needed would be of this form: How much improvement in life-support capacity can we get if we give up a particular amount of mobility? Since we are uncertain about the performance parameters (as reflected by the range of attribute values) we may also be uncertain of the trade-offs that can be made. Thus, it might be necessary to express the trade-offs in the form of ranges. When this uncertainty is combined with the uncertain attribute level to which the trade-off is applied, and to the many possible combinations of trade-offs, it is clear that trade-off information must be used selectively, if at all, in multiple-attribute choice problems where there is uncertainty in attribute values.

ECLECTIC APPROACHES TO THE MULTIPLE-ATTRIBUTE PROBLEM

Combinations of Approaches

Up to this point we have considered each procedure separately. However, it should be obvious to the reader that some of the procedures are quite complementary to others. In this section, then, we shall consider a combination of three of the procedures: dominance, satisficing, and lexicography. We shall describe this eclectic approach in the context of the pressure-suit decision problem. This approach

is, of course, only one of a large number of possible combinations, but seems to be both less arbitrary and more computationally feasible than most other combinations.

As a first step the decisionmaker would compare the alternatives with his minimal attribute requirements. Let us use the same requirements: average mobility, average comfort, 3-hr life support capacity, 3.7 psi pressurization, 1900 Btu/hr peak metabolic heat load, 30 km/sec primary flux meteoroid protection, 250 Btu/hr maximum heat gain, and 5% total ultraviolet radiation. The application of the satisficing procedure would rule out design I (because of unsatisfactory life support) and design III (because of unsatisfactory mobility). To make sure that alternatives have not been excluded on arbitrary grounds, two factors should be determined: 1) can the life support capacity of design I or the mobility of design III be improved without making any other attribute values of these alternatives unsatisfactory, or 2) can the minimal requirements given above be weakened on life support or mobility and still be consistent with the mission (Quade, 1964).

In making the first determination--that is, improving the excluded alternatives--trade-off information, if available, could be quite valuable. (As was noted when the trade-off procedure was initially discussed, it is most useful when dealing with design and redesign matters). Thus, for example, if we could slightly modify design III to reduce some of the life support equipment (but not below 3-hr capacity) and in the process improve mobility, then this modified design III would be satisfactory. In making the second type of determination--that is, weakening the requirements--we would ask, for example, if the mission could be accomplished with a suit design having only (at best) fair mobility. Let us suppose, to illustrate various combinations of the multiple-attribute procedure, that we do indeed need 3-hr life support and at least average mobility, and that design I and design III cannot be modified to yield these values. We have thus reduced the number of alternatives from five to three by removing designs I and III from further consideration.

In the next step in the eclectic approach we apply the dominance procedure. As noted in the preceding section, design V is (weakly)

dominated by design II because, for each attribute, the maximum attribute value of II is at least as good as the corresponding minimum attribute value of V and in some cases the attribute values of II are better than the corresponding values of V. This suggests, then, that there is no reason for choosing design V when design II is available, and that we should exclude design V from further consideration.

As the next step we now apply the lexicographic procedure. Note that one of our previous concerns about lexicography--that it may only consider a small number of attributes, maybe only one (unless the alternatives have many common values on the most important attributes)--has been somewhat allayed by first applying satisficing and dominance. Because satisficing and dominance are both procedures that consider all n dimensions, the possible arbitrariness of considering only a small number of attributes using lexicography has been considerably lessened. Since we know that life support is still the most important attribute, the use of lexicography (in the strong preference sense) would suggest the choice of design II over design IV because the life-support capacity of design II is at worst 3 hr while the life-support capacity of design IV is at best 3 hr.

The eclectic approach described in this application leads to a final choice of design II. It is obvious that different minimal requirements and uses of strong and weak preference might have yielded another final choice even if the same general approach of satisficing-dominance-lexicography were applied. Different combinations of procedures, of course, could yield quite different results.

The eclectic procedure did not require that the attribute values be made numerical; in fact, it did not even require that the values for different attributes be made comparable. Values for one attribute of a particular alternative were compared only with values of the same attribute for other alternatives. Thus, no arbitrary scaling methods were necessary. The eclectic approach did require information on minimum attribute values, but some kinds of specifications are usually available if the mission is clear enough to be seriously considered.

In addition, we needed to know the attribute that was considered most important by the decisionmaker. We did not, however, have to get

a numerical weight for the importance of each attribute to the decision-maker. If such a reliable weighting was available, we would, of course, use it, but experience and introspection both suggest that meaningful numerical weights are very difficult to obtain. By not using a procedure that combines values across attributes (such as the additive weighting procedure), we did not have to assume the independence of the attributes.

Clearly, the approach described in this section has two virtues: requiring a minimum amount of hard-to-get information, and making relatively few arbitrary assumptions. Using combinations of procedures seems to be more reasonable than using any of the procedures separately. For example, the ways in which lexicography complements satisficing and dominance were noted above. The combined procedure considered here could also be used in decision problems where the alternatives are uncovered in a sequential fashion rather than simultaneously, as was the case in Table 5. By first applying satisficing, the decisionmaker forms a standard by which to evaluate single alternatives. The accumulated satisfactory alternatives can be compared with each other by the dominance procedure. Only then do we require information about the relative importance of attributes in order to choose among the non-dominated, satisfactory alternatives.

Current Uses of Multiple-Attribute Methods

Some authors have applied a combination of methods to multiple-attribute decision problems, although they do not usually identify individual methods used. For example, Briskin (1966) uses first dominance and then trade-offs in a deterministic scheduling problem, while Pinkel (1967) uses an especially interesting combination of procedures on a weapons system choice problem quite similar to the one we used earlier in this paper.

Pinkel suggests particular attributes (different from ours) that contain time, cost, and performance factors. In a sense this is more general than our first example, dealing primarily with performance, but as we emphasized when the example was introduced, we are interested in

describing the techniques we surveyed, and these techniques are not contingent on any particular attributes.

Given a group of attributes like those in Table 1, Pinkel attempts to reduce the amount of information to those factors that are truly significant to the decision at hand. He accomplishes this by first reducing the number of alternatives and the number of attributes for the alternatives that remain, and then applying some rules of judgment to arrive at a final choice.

Let us first consider his reduction process. The first stage involves removing, for any given attribute, all the values of that attribute that do not differ by more than the precision of the estimate of them. As a procedure for attempting to deal with uncertainty in the attribute values, this resembles the procedures discussed earlier in Sec. IV. Pinkel's second stage reduction involves a type of satisficing procedure where any alternative that has one (or more) unsatisfactory attribute value is eliminated. At the next stage the remaining attribute values of the remaining alternatives are compared and any dominated alternatives are removed. Hopefully, at this point a number of alternatives have been eliminated, and further, even for the alternatives that remain, a number of attribute values should have been suppressed.

After this reduction process, Pinkel then applies some rules of judgment. Here we find a lexicographic flavor (in the recommendations of things to look for first), and also the implicit suggestion of trade-offs. It can be noted, then, that Pinkel's approach is a good example of what we are advocating.

Other combinations of approaches could be suggested; however, at this point we are particularly interested in having the reader think in terms of methods that would be appropriate for his own decision problems. Even though the approach applied to the space suit decision and Pinkel's procedure on the weapons system decision seem quite reasonable, the reader should keep in mind the characteristics of the individual methods used. Thus, we see the earlier discussion of characteristics and the summary in Table IV as providing a solid basis both for evaluating eclectic methods proposed by others and for building up a new combination of methods.

BIBLIOGRAPHY

- Ackoff, Russell L. (1962): Scientific Method: Optimizing Applied Research Decisions, John Wiley & Sons, New York.
- Adams, E. and R. Fagot (1959): "A Model of Riskless Choice," Behavioral Sciences, 4, No. 1, 1-10.
- Aumann, R. J. (1959): "Assigning Quantitative Values to Qualitative Factors in the Naval Electronics Problem," Naval Research Logistics Quarterly, 6, 1-16.
- and J. B. Kruskal (1958): "The Coefficients in an Allocation Problem," Naval Research Logistics Quarterly, 5, 111-123.
- (1964): "Subjective Programming," in M. W. Shelly and G. L. Bryan (eds.) Human Judgments and Optimality, Wiley, New York 1964.
- and M. Maschler (1966): "Independent Preferences: An Area of Applicability of Utility Theory to Disarmament Problems," in Robert J. et al, Development of Utility Theory for Arms Control and Disarmament, Arms Control and Disarmament Agency Report.
- Baumol, W. (1959): Economic Dynamics, 2nd ed., Macmillan, New York.
- Bayes, T. (1763): "An Essay Towards Solving a Problem in the Doctrine of Chances," The Philosophical Transactions, 53, 370-418, reprinted in Biometrika, 45 (1958), 296-315.
- Bernoulli, D. (1738): "Specimen Theoriae Novae de Mensura Sortis," Commentarii Academiae Scientiarum Imperialis Petropolitanae, 5, 175-192, translated as "Exposition of a New Theory on the Measurement of Risk," Econometrica, 22 (1954), 23-36.
- Blackwell, D. and M. A. Girshick (1954): Theory of Games and Statistical Decisions, Wiley, New York.
- Briskin, L. E. (1966): "A Method of Unifying Multiple Objective Functions," Management Science, 12, No. 10, B-406 - B-416.
- Bryan, S. E. (1964): "TFX--A Case in Policy Level Decision Making," Academy of Management Journal, 54-70.
- Charnes, A., W. W. Cooper, and G. L. Thompson, (1965): Chance Constrained Programming and Related Approaches to Cost Effectiveness, Carnegie Institute of Technology, Pittsburgh, Pa., Management Science Research Report No. 39.
- Chipman, J. S. (1960): "The Foundations of Utility," Econometrica, 28, 193-224.
- Churchman, C. W. and R. L. Ackoff (1954): "An Approximate Measure of Value," Operations Research, 2, 172-180.
- , R. L. Ackoff, and E. L. Arnoff (1957): Introduction to Operations Research, Wiley, New York.
- Cyert, R. M. and K. R. MacCrimmon (1967): "Organizations," in G. Lindzey and E. Aronson (eds.), Handbook of Social Psychology, 2d ed., Addison Wesley, Reading, Mass.

- and James G. March (1963): A Behavioral Theory of the Firm, Prentice-Hall, Englewood Cliffs, N. J.
- Debreu, G. (1954): "Representation of a Preference Ordering by a Numerical Function," in R. M. Thrall, C. H. Coombs and R. L. Davis (eds.), Decision Processes, Wiley, New York.
- (1960): "Topological Methods in Cardinal Utility Theory," in K. J. Arrow, S. Karlin and P. Suppes (eds.), Mathematical Methods in the Social Sciences, Stanford Univ. Press, Stanford, Calif., 16-26.
- Dewey, John (1933): How We Think, Heath, New York.
- Eckenrode, R. T. (1965): "Weighting Multiple Criteria," Management Science, 12, No. 3, 180-192.
- Encarnacion, Jose, Jr. (1964): "A Note on Lexicographic Preference," Econometrica, 32, No. 1-2, 215-217.
- Festinger, Leon (1964): Conflict, Decision, and Dissonance, Stanford Univ. Press, Stanford, Calif.
- Fields, D. S. (1966): "Cost/Effectiveness Analysis: Its Tasks and Their Interrelation," Operations Research, 14, No. 3, 515-527.
- Firstman, S. I. and D. S. Stoller (1966): "Establishing Objectives, Measures, and Criteria for Multi-phase Complementary Activities," Operations Research, 14, No. 1, 84-99.
- Fishburn, Peter C. (1965a): "Independence in Utility Theory with Whole Product Sets," Operations Research, 13, No. 1, 28-45.
- (1965b): "Independence, Trade-offs and Transformations in Bivariate Utility Functions," Management Science, 11, No. 9, 752-801.
- (1966): "A Note on Recent Developments in Additive Utility Theories for Multiple-Factor Situations," Operations Research, 14, No. 4, 1143-1148.
- (1966): "Additivity in Utility Theory with Denumerable Product Sets," Econometrica, 34, No. 2, 500-503.
- (1964): Decision and Value Theory, Wiley, New York.
- Fisher, G. H. (1966): The Analytical Bases of Systems Analysis, The RAND Corporation, Santa Monica, Calif., P-3363.
- Fox, Peter D. (1964): "A Theory of Cost-Effectiveness for Military Systems Analysis," Operations Research, 13, No. 1, 191-201.
- Franklin, Benjamin (1945): "How to Make a Decision," in N. Goodman (ed.), A Benjamin Franklin Reader, Crowell, New York, 786.
- Frisch, R. (1966): Maxima and Minima, Rand McNally, Chicago.
- Gulliksen, Harold (1956): "Measurement of Subjective Values," Psychometrika, 21, No. 3, 229-244.
- (1964): "The Structure of Individual Differences in Optimality Judgments," in M. W. Shelly and G. L. Bryan (eds.), Human Judgments and Optimality, Wiley, New York, 72-84.

- and S. Messick (eds.) (1960): Psychological Scaling, Wiley, New York.
- Hausner, Melvin (1954): "Multi-dimensional Utilities," in R. W. Thrall, C. H. Coombs, and R. L. Davis (eds.), Decision Processes, Wiley, New York.
- Heuston, M. C. and G. Ogawa (1966): "Observations on the Theoretical Basis of Cost-Effectiveness," Operations Research, 14, No. 2, 242-266.
- Hitch, C. J. (1964): "Analysis for Air Force Decisions," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387-PR, 13-23.
- and R. N. McKean (1960): The Economics of Defense in the Nuclear Age, The RAND Corporation, Santa Monica, Calif., R-346; also Harvard Univ. Press, Cambridge, Mass., 1960.
- Houthakker, H. S. (1960): "Additive Preferences," Econometrica, 28, No. 2, 244-257.
- Klahr, C. N. (1958): "Multiple Objectives in Mathematical Programming," Operations Research, 6, No. 6, 849-855.
- Klahr, David (1967): Decisionmaking in a Complex Environment, unpublished Ph.D. dissertation, Carnegie Institute of Technology, Pittsburgh, Pa.
- Kovit, Bernard (1965): "The Space Suit," Space Aeronautics, 43, No. 6, 43-49.
- Kruskal, J. B., "Multidimensional Scaling by Optimizing Goodness of Fit to a Non-metric Hypothesis," Psychometrika, 29, 1-27.
- Luce, R. D. (1966): "Two Extensions of Conjoint Measurement," Journal of Mathematical Psychology, 3, No. 2, 343-370.
- and J. W. Tukey (1964): "Simultaneous Conjoint Measurement: A New Type of Fundamental Measurement," Journal of Mathematical Psychology, 1, 1-27.
- MacCrimmon, K. R. (1967): "Some Normative and Descriptive Implications of the Decision Theory Postulates," in Karl Borch (ed.), Risk and Uncertainty, MacMillan, New York.
- May, K. O. (1954): "Transitivity, Utility, and Aggregation in Preference Patterns," Econometrica, 22, 1-13.
- McKean, R. N. (1964): "Criteria" in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387, 81-91.
- McNamara, Robert S. (1966): "New Perspectives in International Security," Atlantic Community Quarterly, 4, No. 2, 229ff.
- Meckling, W. H. (1964): "Strategies for Development," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387, 227-236.
- Mellon, W. G. (1962): "An Approach to a General Theory of Priorities," Princeton Univ. Press, Princeton, N. J., Econometric Research Program Memorandum 42.

- Miller, George A. (1956): "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," Psychological Review, 63, No. 2, 81-97.
- Miller, J. R. (1966): "The Assessment of Worth: A Systematic Procedure and its Experimental Validation," unpublished Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Mottley and Newton (1959): "The Selection of Projects for Industrial Research," Operations Research, 1, No. 6, 740-751.
- Newell, A., J. C. Shaw, and H. A. Simon (1958): "The Elements of a Theory of Human Problem Solving," Psychological Review, 65, 151-166.
- Papandreou, A. et al (1957): "A Test of a Stochastic Theory of Choice," University of California Publications in Economics, 16, 1-18.
- Pardee, F. S. (1965): State-of-the-Art Projection and Long-Range Planning of Applied Research, The RAND Corporation, Santa Monica, California, P-3181.
- Pfanzgal, J. (1959): "A General Theory of Measurement--Applications to Utility," Naval Research Logistics Quarterly, 6, No. 4, 283-294.
- Pinkel, B. (1965): The Impact of the High Development Cost of Advanced Flight Propulsion Systems on Development Policy, The RAND Corporation, Santa Monica, Calif., RM-4560-PR
- (1967): On the Decision Matrix and the Judgment Process: A Developmental Decision Example, The RAND Corporation, Santa Monica, Calif., P-3620.
- Quade, E. S. (1964): "Introduction," "Methods and Procedures," "Pitfalls in Systems Analysis," and "A Missile Comparison," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387-PR; also Rand McNally & Co., Chicago, Ill., 1964.
- Radner, R. (1964): "Mathematical Specification of Goals for Decision Problems," in M. W. Shelly and G. L. Bryan (eds.), Human Judgments and Optimality, Wiley, New York.
- Ramsey, F. P. (1931): "Truth and Probability," The Foundations of Mathematics, Kegan Paul, London.
- Roberts, H. V. (1965): "Probabilistic Prediction," Journal of the American Statistical Association, 60, No. 309, 50-62.
- Savage, L. J. (1954): Foundations of Statistics, Wiley, New York.
- Schamberg, R. (1964): "Technological Considerations," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387-PR, 179-198.
- Schelling, T. C. (1964): "Assumptions about Enemy Behavior," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387-PR, 199-216.
- Seiler, K. (1966): "A Cost Effectiveness Comparison Involving a Trade-off of Performance, Cost, and Obtainability," Operations Research, 14, No. 3, 528-531.

- Shephard, R. N. (1962): "The Analysis of Proximities: Multi-dimensional Scaling with an Unknown Distance Function, I and II," Psychometrika, 27, 125-140, 219-264.
- (1964): "On Subjectively Optimum Selection Among Multi-attribute Alternatives," in M. W. Shelly and G. L. Bryan (eds.), Human Judgments and Optimality, Wiley, New York.
- Sigford, J. V. and R. H. Parvin (1965): "Project PATTERN: A Methodology for Determining Relevance in Complex Decision-Making," IEEE Transactions on Engineering Management, EM-12, No. 1, 9-13.
- Simon, H. A. (1955): "A Behavioral Model of Rational Choice," Quarterly Journal of Economics, 69, 99-118; also reprinted in Simon (1957)
- (1957): Models of Man, Wiley, New York.
- (1960): The New Science of Management Decision, Harper, New York.
- (1964): "On the Concept of Organizational Goal," Administrative Science Quarterly, 1-22.
- Smith, Jack W. (1956): "A Plan to Allocate and Procure Electronic Sets by the Use of Linear Programming Techniques and Analytical Methods of Assigning Values to Qualitative Factors," Naval Research Logistics Quarterly, 3, No. 3, 151-162.
- Stevens, S. S. (1959): "Measurement, Psychophysics, and Utility," in C. W. Churchman and P. Ratoosh (eds.), Measurement: Definitions of Theories, Wiley, New York.
- Strotz, Robert H. (1957): "The Empirical Implications of a Utility Tree," Econometrica, 25, No. 2, 269-280.
- Terry, H. (1963): "Comparative Evaluation of Performance Using Multiple Criteria," Management Science, 9, No. 3, 431-442.
- Thrall, Robert M. (1954): "Applications of Multidimensional Utility Theory," in R. M. Thrall, C. H. Coombs and R. L. Davis (eds.), Decision Processes, Wiley, New York, 181-186.
- Timson, F. S. (1968): Measurement of Technical Performance in Weapon System Development: A Subjective Probability Approach, The RAND Corporation, Santa Monica, Calif., RM-5207-ARPA.
- Torgerson, W. S. (1958): Theory and Methods of Scaling, Wiley, New York.
- Trull, S. G. (1966): "Some Factors Involved in Determining Total Decision Success," Management Science, 12, No. 6, B270-B280.
- Wald, A. (1950): Statistical Decision Functions, Wiley, New York.
- Wohlstetter, A. (1964): "Analysis and Design of Conflict Systems," in E. S. Quade (ed.), Analysis for Military Decisions, The RAND Corporation, Santa Monica, Calif., R-387, 103-148.
- Yntema, D. B. and W. S. Torgerson (1961): "Man-Computer Cooperation in Decisions Requiring Common Sense," IRE Transactions on Human Factors in Electronics, 2, No. 1, 20-26.