

Decisions with Multiple Objectives

Preferences and
Value Tradeoffs

Ralph L. Keeney
Howard Raiffa



Decisions with Multiple Objectives

Preferences and Value Tradeoffs

RALPH L. KEENEY

*Institute of Safety and Systems Management
University of Southern California*

and

HOWARD RAIFFA

*Frank P. Ramsey Chair in Managerial Economics
Graduate School of Business and Kennedy School of Government
Harvard University*



CAMBRIDGE
UNIVERSITY PRESS

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Victoria 3166, Australia

© Cambridge University Press 1993

First published in 1976 by John Wiley & Sons, Inc.
First published by Cambridge University Press 1993

Printed in Canada

Library of Congress Cataloging-in-Publication Data

Keeney, Ralph L., 1944–

Decisions with multiple objectives : preferences and value
tradeoffs / Ralph L. Keeney and Howard Raiffa.

p. cm.

Includes bibliographical references and index.

ISBN 0-521-44185-4 – ISBN 0-521-43883-7 (pbk.)

I. Multiple criteria decision making. I. Raiffa, Howard, 1924–
II. Title.

T57.95.K43 1993

92–31874

658.4'03–dc20

CIP

A catalog record for this book is available from the British Library.

ISBN 0-521-44185-4 hardback

ISBN 0-521-43883-7 paperback

Decisions with
Multiple Objectives

To
Tjalling C. Koopmans,
whose standards of excellence are
an inspiration to all who know him.

Preface to the Cambridge University Press Edition

Way back in the mid-seventies when we wrote *Decisions with Multiple Objectives*, our objectives were clear. We aspired to develop a sound, practical theory for analyzing decisions with multiple, competing objectives; to make this theory and framework of thought accessible to a diverse audience of scholars and practitioners of the art of decision analysis; and to promote applications of multiattribute analysis. Today, our objectives are exactly the same. In the intervening sixteen years, the theory most relevant to applications has been augmented a little and the application of the theory has greatly expanded.

Decision analysis is widely recognized as a sound prescriptive theory. When a decision involves multiple objectives – and this is almost always the case with important problems – multiattribute utility theory forms the basic foundation for applying decision analysis. The theory explicitly addresses the value tradeoffs and uncertainties that are invariably at the center of multiple-objective decisions. The experience of numerous applications indicates that the theory available in 1976 is still the most relevant theory available today. The important addition since 1976 concerns value functions that address strength of preferences between pairs of consequences (see Dyer and Sarin, 1979; Bell and Raiffa, 1988).

The art of applying multiattribute utility has expanded in many directions since 1976. Collectively, decision scientists know much more today about identifying and structuring objectives and about measuring their achievement in terms of objectives. We are more skilled at assessing utility functions, and computer programs greatly assist analyses and sensitivity analysis. In addition, the use of structured values, both qualitative and quantitative, to help define decision situations, create alternatives, and communicate about decisions has been developed.

The identification and structuring of objectives essentially frames the decision being addressed. It sets the stage for all that follows. Recent work has developed techniques to involve relevant stakeholders in the process constructively, and to probe more deeply into the meaning, relevance, and relationships of objectives (see von Winterfeldt, 1987; Keeney, 1992). The computer program Logical Deci-

sion (Smith, 1989) is helpful in assessing utility functions. Given the utilities of possible consequences, Logical Decision, IDEA (Whitfield et al., 1989), Super-tree (McNamee and Selona, 1990), Decision Programming Language (Call and Miller, 1990), and three programs by Kirkwood (1991) can analyze alternatives effectively.

Applications of decision analyses with multiple objectives have been summarized in several publications. Corner and Kirkwood (1991) has a comprehensive review of more than one hundred applications covering a broad range of fields including energy, manufacturing and services, public policy, and health care. A few earlier applications are discussed in depth in Bell, Keeney, and Raiffa (1977) and more recent applications, including personal decisions, are found in Keeney (1992). In the past decade, decision analysis involving multiple objectives has expanded to address the substance of important classes of problems more specifically. Examples of these are books that focus in-depth on siting major facilities (Keeney, 1980), the art and science of negotiation (Raiffa, 1982), and medical decisions (Sox, Blatt, Higgins, and Marton, 1988).

There should be significant interplay between descriptive studies of how people do process information and make decisions and prescriptive decision analysis to help people make decisions that are consistent with their values and understanding of the problem. Recently, there has been a major trend to incorporate the insights from descriptive studies into the prescriptive analyses to render them more effective and useful. Many of the descriptive concepts and ideas now used in prescriptive analysis are discussed in Kahneman, Slovic, and Tversky (1982), von Winterfeldt and Edwards (1986), Bell, Raiffa, and Tversky (1988), and Edwards (1992).

As stated in Section 1.3.3, we assumed in this book that the decision problem had been identified and the alternatives specified. However, we recognized that problem selection and alternative creation are critically important. Since our values indicate the reason we are interested in any decision situation, it seemed reasonable to use values as structured in this book to guide the thought process to select decisions and alternatives. This concept, called value-focused thinking (Keeney, 1992), is developed to create better alternatives for any decision problem that we face and to identify decision opportunities more appealing than the decision problems that confront us.

Teaching from This Book

It may be useful to perspective readers and teachers to suggest how we would revise this book if we spent the substantial time on such a revision, and what supplementary material we would use when teaching from the current version. The responses are essentially the same. Chapter 1 still is a useful introduction. Material in Chapter 2 on structuring objectives would be updated with material

from Keeney (1992, Chapters 3 and 4). A discussion of the theory and results involving strength of preferences would be added to Chapter 3. The main theory in Chapters 4–6 and 9 would be changed only in a minor way. Specifically, we would focus a bit more on the art of assessing utility functions and the use of computers to assist the assessment. A Chapter 11 would be added to indicate how structured values can be used both to create alternatives and decision opportunities using value-focused thinking.

Chapter 10 on the aggregation of individual preferences would be substantially enhanced to account for several developments. The revision would include material on the structuring and combining of objectives in problems involving multiple stakeholders (von Winterfeldt, 1987; Keeney, 1992), the major advances in applying multiattribute utility in negotiation contexts (Raiffa, 1982), and the result that a group von Neumann utility function is consistent with assumptions analogous to Arrow's (Keeney, 1976). In other words, when strength of preferences and interpersonal values are explicitly addressed, there is a logical and consistent way to create a group utility function.

The major revisions of the book would have to do with the multitude of examples that could be presented along with those in Chapters 7 and 8. There are now so many good applications that we would carefully need to find a set that would illustrate the many key insights available from decision analyses with multiple objectives. This is precisely the difficult choice problem that we hoped to create when we originally wrote *Decisions With Multiple Objectives*. Thank you to the many original readers who applied the theory and helped make this problem a reality.

Ralph L. Keeney
Howard Raiffa

REFERENCES

- Bell, D. E., R. L. Keeney, and H. Raiffa, eds. (1977). *Conflicting Objectives in Decisions*, Wiley, New York.
- Bell, D. E., and H. Raiffa (1988). Marginal value and intrinsic risk aversion. In *Decision making: Descriptive, Normative, and Prescriptive Interactions*, D. E. Bell, H. Raiffa, and A. Tversky, eds., Cambridge University Press, Cambridge, pp. 384–397.
- Bell, D. E., H. Raiffa, and A. Tversky, eds. (1988). *Decision Making: Descriptive, Normative, and Prescriptive Interactions*, Cambridge University Press, Cambridge.
- Call, H. J., and W. A. Miller (1990). A comparison of approaches and imple-

- mentations for automating decision analysis. *Reliability Engineering and Systems Safety*, 30, 115–162.
- Corner, J. L., and C. W. Kirkwood (1991). Decision analysis applications in the operations research literature, 1970–1989. *Operations Research*, 39, 206–219.
- Dyer, J. S., and R. K. Sarin (1979). Measurable multiattribute value functions. *Operations Research*, 27, 810–822.
- Edwards, W., ed. (1992). *Utility Theories: Measurements and Applications*, Kluwer Academic Publishers, Dordrecht, Holland.
- Kahneman, D., P. Slovic, and A. Tversky, eds. (1982). *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.
- Keeney, R. L. (1976). A group preference axiomatization with cardinal utility. *Management Science*, 23, 140–145.
- Keeney, R. L. (1980). *Siting Energy Facilities*, Academic Press, New York.
- Keeney, R. L. (1992). *Value-Focused Thinking*, Harvard University Press, Cambridge, Massachusetts.
- Kirkwood, C. W. (1991). Computer programs for multiattribute evaluation. Technical Report DIS 91/92-7, Department of Decision and Information Systems, Arizona State University, Tempe.
- McNamee, P., and J. Selona (1990). *Decision Analysis with Supertree, 2nd Edition*, Scientific Press, South San Francisco, California.
- Raiffa, H. (1982). *The Art and Science of Negotiation*, Harvard University Press, Cambridge, Massachusetts.
- Smith, G. R. (1989). *Logical Decision: Multi-Measure Decision Analysis Software*, Logical Decisions, Inc., Golden, Colorado.
- Sox, H. C., M. A. Blatt, M. C. Higgins, and K. I. Marton (1988). *Medical Decision Making*, Butterworths, Boston, Massachusetts.
- Whitfield, R. G., C. E. Swietlik, M. F. Warren, R. Fuja, M. J. Jusko, J. P. Peerenboom, and E. C. Johnson (1989). *IDEA-Interactive Decision Analysis: Users' Guide and Tutorial*, Technical Report ANL/EES-TM-378, Energy and Environmental Systems Division, Argonne National Laboratory, Argonne, Illinois.
- von Winterfeldt, D. (1987). Value tree analysis: An introduction and an application to offshore oil drilling. In *Insuring and Managing Hazardous Risks: From Seveso to Bhopal and Beyond*, P. Kleindorfer and H. Kunreuther, eds., Springer, New York, 349–377.
- von Winterfeldt, D., and W. Edwards (1986). *Decision Analysis and Behavioral Research*, Cambridge University Press, Cambridge.

Preface to First Edition

If we wanted the title of this book merely to convey the subject matter, it would be some horrendously complicated concoction such as: "On Cardinal Utility Analysis with Multiple Conflicting Objectives: The Case of Individual Decision making Under Uncertainty from the Prescriptive Point of View—with Special Emphasis on Applications but with a Little Theory Thrown In for Spice."

Our actual title, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* is longer than we think a title should ideally be, but, unfortunately, it is too short to prevent unjustified sales. Even in such a simple case, it is not so easy to *balance* among the conflicting objectives: convey the subject matter, minimize the length, and promote justified sales but prevent unjustified ones.

To an ever growing circle of people, "Decision Analysis" has carved out for itself a niche in the literature of operations research, systems analysis, management sciences, decision and control, and cybernetics. Decision analysis looks at the paradigm in which an individual decision maker (or decision unit) contemplates a choice of action in an uncertain environment. The approach employs systematic analysis, with some number pushing, which helps the decision maker clarify in his own mind which course of action he should choose. In this sense, the approach is not *descriptive*, because most people do not attempt to think systematically about hard choices under uncertainty. It is also not *normative* since it is not an idealized theory designed for the superrational being with an all-powering intellect. It is, instead, a *prescriptive* approach designed for normally intelligent people who want to think hard and systematically about some important real problems.

The theory of decision analysis is designed to help the individual make a choice among a set of *prespecified* alternatives. Of course, decision analysts do admit that an insightful generation of alternatives is of paramount importance, and they also take note of the often-overlooked fact that good analysis of a set of existing alternatives may be suggestive of ways to augment the set of alternatives. But this is a side point that is not suitable for development in a preface.

What is important here is that the usual analysis (after suitable modeling has been done) involves two distinctive features: an uncertainty analysis and a preference (or value or utility) analysis. There has been a great deal that has been written on the uncertainty phase: statistical validation of a model, uses of historical and experimental data for inference, the codification of judgmental estimates by the decision maker and by expert groups, and so on. In comparison with this voluminous literature on the uncertainty side little has been written about the value or preference side of the picture. The present book helps improve the balance.

At present, this gross imbalance is also unfortunately very much in evidence in applications. Several person-years of effort will be utilized developing, modifying, and verifying an elaborate simulation model that outputs the possible levels of several indicators of interest resulting from any particular policy. Perhaps the output is synthesized in terms of a few graphs or tables and a summary report is written for the *decision maker*. This decision maker then struggles for perhaps a week with the implications of the alternatives and then chooses an alternative. The score: person-years on the uncertainty side of the problem, a week on the preference side. We feel that the shifting of a little effort—perhaps only a few person-months—to the preference aspects could lead to significantly improved decision making in many situations. In this book, we suggest how you might constructively use more effort on the preference aspects of analysis.

An illustrative example can help set the stage. A decision-making unit must make a policy choice in a complicated environment. Imagine that the problem is so complicated that a computer-based simulation model is designed such that for each policy choice under review, a scenario can be generated that indicates how the future might unfold in time. Now suppose that the analyst effectively summarizes the relative desirability of any future scenario not by a single number but, let us say, by a dozen well-chosen numbers: some reflecting costs, others reflecting benefits. Since these output performance numbers may simultaneously deal with economic, environmental, social and health concerns, these summarizing indices will, in general, be in incommensurable units. To complicate matters, suppose that stochastic elements are involved in the simulation so that, for a single policy choice being investigated, repeated simulation runs result in different sets of summary performance measures. The joint probability distribution of these performance measures as made manifest through repeated realizations of the simulation will, in general, indicate that these 12 measures are probabilistically dependent. Now assume you are a harassed decision maker sitting in front of an output display device deluged with a mountain of conflicting information. You are confused. What should you do? How can you sort out the issues and start thinking systematically about your choice problem: which policy should you adopt in the real setting? We believe this book addresses your problem and has something constructive to say about it that is not merely platitudinous.

We are convinced of one thing: The decision maker cannot simply plug these

incommensurate output performance measures into an objective formula that someone has proposed *ex ante* without any reference to the real-world meaning of the various measures. Instead, our prescriptions lead us in an opposite direction: we advocate that the responsible decision maker force himself to think hard about various value tradeoffs and about his attitudes toward risky choices and we suggest ways that this process can be systematically examined by dividing his complicated choice problem into a host of simpler choice problems.

The methodology will in a step-by-step fashion force the cooperating decision maker to articulate a rank ordering of all potential outcome vectors—in the illustrative example, an ordering of all 12-tuples. This rank ordering can be thought of as constituting a set of indifference curves plus an orientation in 12-space. But this is not enough since repeated simulations of the same policy will produce, because of stochastic elements, different 12-tuples. Our problem is a familiar one by now, and the utility theory of von Neumann–Morgenstern comes to the rescue. This theory tells us that in order to satisfy certain compelling behavioral desiderata, the decision maker must assign to each 12-tuple a single number, referred to as the utility of that 12-tuple, and this assignment must be such that:

1. The more preferred the 12-tuple the higher is the associated utility.
2. These utilities must be scaled in a way that justifies the maximization of expected utilities.

In order to evaluate the relative desirability of a given policy alternative, we then must (1) generate for each simulation run a set of output 12-tuples, (2) associate to each 12-tuple a utility, and (3) average the sequence of utilities generated by repeated runs for the same policy. Finally, we choose the policy that maximizes the expected utility. Built into the assignment of utilities are all the aspects of risk aversion or proneness that we feel are necessary to include. That this can be done and how it is done is the subject matter of utility theory, reviewed in Chapter 4.

Having stated our general approach, can it actually be done in practice? We argue yes, and we substantiate our case by citing many examples illustrating how it has already been done in practice. It's not easy to do; but what are the alternatives?

Outline of the Book

For conceptual purposes, the material presented can be partitioned into four main categories: (1) the *structuring* of multiple-objective problems: Chapters 1 and 2; (2) the *theory* of quantifying preferences over multiple objectives: Chapters 3 through 6; (3) the *applications* of that theory: Chapters 7 and 8; and (4)

special topics; Chapters 9 and 10. We only briefly elaborate here since a more detailed outline is found in Section 1.6.

Chapter 1 introduces the subject matter of concern more systematically than has been done above. Our basic problem is phrased in terms of the analysis of decision trees instead of in terms of a stochastic simulation model, but the distinction for our purposes does not matter. In Chapter 2, we acknowledge that in a given context the set of objectives and attributes are not *given* for a problem. Some suggestions are made for generating and structuring appropriate sets of objectives.

The theory, Chapters 3 through 6, present techniques for quantifying preferences over multiple objectives. In order to obtain a von Neumann–Morgenstern utility function in such cases, one must address two separate issues: value trade-offs among objectives and attitudes toward risk. Chapter 3 looks at value trade-offs under conditions of certainty. Chapter 4 restricts itself to a single objective and introduces concepts and techniques that are needed in quantifying and assessing risk attitudes. This chapter essentially reviews single-attribute (i.e., unidimensional) utility theory. Chapters 5 and 6 consider both of these issues simultaneously; they present multiattribute (i.e., multidimensional) utility theory. Because of its length we divided this material into two segments: two attributes (Chapter 5) and more than two attributes (Chapter 6).

Multiattribute utility is already sufficiently developed to make worthwhile contributions to some important complex problems. Chapters 7 and 8 deal with applications that present support for this claim; many problems are discussed where preferences have been quantified using multiattribute utility. These include structuring corporate objectives, examining operational policies of fire departments, allocating school-system funds, evaluating time-sharing systems, siting nuclear power facilities, treating medical problems such as cleft lip and palate, and so forth. In each case, we describe the problem context in which the preference assessments took place. We want to communicate some of the art as well as the theory and procedures of using multiattribute utility analysis. Chapter 8 discusses the theory and procedures developed in earlier chapters in a major case study: the development of airport facilities for Mexico City.

Chapters 9 and 10 on special topics examine preferences over time and aggregation of individual preferences, respectively. Each of these important problems can be cast and naturally studied in a multiattribute framework. As shown, many of the results of Chapters 3 through 6 are relevant to the *time* and *group* problems. These two problems are often added complicating features in multiple-objective problems.

Our Intended Audience

Decision making is of such a pervasive interest that it is hard for us to exclude any group. This book is certainly relevant to analysts, policy makers, policy

advisors, economists, designers, engineers, and managers, among others. Meaningful and important applications can be found in business, public policy, engineering design, resource management, public health and medicine, educational management, and so on.

This is a big book and not all of it has to be read. There are parts, especially Chapter 6 and the latter part of Chapter 9, where the mathematics will be discouragingly complicated except for the mathematical pros. It would be helpful if the nonmathematical reader were already familiar with the rudiments of decision analysis as explicated by Raiffa (1968) or by books at a similar level such as Schlaifer (1969) and Brown et al. (1974).

Depending on interests, you may wish to read only a selection of the chapters. Chapters 1 and 2 on structuring the multiple-objective problem can be read with no prerequisite. Similarly, if you are willing to accept the abstract formulation of the problem, Chapters 3 through 6 are essentially self contained. Even within this group, the reader with some mathematical background could begin with either value tradeoffs (Chapter 3), unidimensional utility theory (Chapter 4), or multiattribute utility theory (Chapters 5 and 6). For a full understanding of the applications in Chapters 7 and 8, a knowledge of the main theoretical results of the book is required. However, a reader, interested in the domain of applicability of multiattribute utility and a feeling for how it is used in a specific context, could pick them up reading only Chapters 7 and 8. Before reading Chapters 9 or 10, it is advisable in most cases to at least glance through Chapters 3 through 6. However, a reader who feels at ease with the level of mathematics in these chapters could begin with either Chapter 9 or 10 and only refer back to the basic theory chapters when references indicate it may be worthwhile.

To our knowledge, there are no other books that overlap much in content with this one. However, most of the theoretical results have appeared in professional journals. Many of these results come from researchers other than ourselves. We have attempted to reference the original contributions appropriately so that you can easily trace the development of any particular topic. A large bibliography of these works is included following Chapter 10.

Ralph L. Keeney
Howard Raiffa

Laxenburg, Austria
May 1975

Contents

Preface to the Cambridge University Press Edition	xi
Preface to First Edition	xv
1. THE PROBLEM	1
1.1 Sketches of Motivating Examples	1
1.2 Paradigm of Decision Analysis	5
1.3 Comments About the Paradigm	7
1.4 Complex Value Problems	14
1.5 Classes of Examples and Methodological Nightmares	19
1.6 Organization of Remaining Chapters	26
2. THE STRUCTURING OF OBJECTIVES	31
2.1 Objectives and Attributes	32
2.2 Generation of Objectives and Attributes	34
2.3 Hierarchical Nature of Objectives	41
2.4 Sets of Objectives and Attributes	49
2.5 Proxy Attributes and Direct Preference Measurements	55
2.6 Summary and Perspective on the Practical Aspects of Specifying Objectives and Attributes	64
3. TRADEOFFS UNDER CERTAINTY	66
3.1 The Multiattribute Value Problem	66
3.2 Choice Procedures That Do Not Formalize Preference Structures	69
3.3 Structuring of Preferences and Value Functions	77
3.4 Preference Structures and Value Functions for Two Attributes	82

3.5	The Case of Three Attributes	100
3.6	The Case of More Than Three Attributes	108
3.7	Assessment of an Additive Value Function: An Abstract Hypothetical Example	117
3.8	Willingness to Pay	125
3.9	Brief Summary and Guide to Some Relevant Literature	129
4.	UNIDIMENSIONAL UTILITY THEORY	131
4.1	The Motivation for Utility Theory	131
4.2	Direct Assessment of Utilities for Consequences	140
4.3	Unidimensional Utility Functions	141
4.4	Risk Aversion	148
4.5	A Measure of Risk Aversion	159
4.6	Constant, Decreasing, and Increasing Risk Aversion	165
4.7	Proportional Risk Aversion	174
4.8	Monotonically Decreasing and Nonmonotonic Utility Functions	179
4.9	A Procedure for Assessing Utility Functions	188
4.10	Illustrations of the Assessment of Unidimensional Utility Functions	203
4.11	Explicating a Single Attribute by Means of Multiple Attributes	212
4.12	Conditional Unidimensional Utility Theory	214
4.13	Where We Stand	218
5.	MULTIATTRIBUTE PREFERENCES UNDER UNCERTAINTY: THE TWO-ATTRIBUTE CASE	219
5.1	Approaches for Multiattribute Assessment	219
5.2	Utility Independence	224
5.3	Additive Independence and the Additive Utility Function	229
5.4	The Implications of Mutual Utility Independence	232
5.5	Use of Certainty Equivalents	242
5.6	Utility Functions with One Utility-Independent Attribute	243
5.7	What to Do if No Independence Properties Hold	255
5.8	Assessment Procedure for Multittribute Utility Functions	261
5.9	Interpreting the Scaling Constants	271
5.10	The Assessment of a Utility Function in a Hospital Blood Bank	273

6. MULTIATTRIBUTE PREFERENCES UNDER UNCERTAINTY: MORE THAN TWO ATTRIBUTES	282
6.1 Overview of the Chapter	282
6.2 Utility Functions with Three Attributes	286
6.3 The Multiplicative Utility Function	288
6.4 The Multilinear Utility Function	293
6.5 The Additive Utility Function	295
6.6 Assessing Multiattribute Utility Functions	297
6.7 A Fundamental Relationship Between Preferential Independence and Utility Independence	310
6.8 Relationships Among Utility Independence Assumptions	316
6.9 Decomposition of Multiattribute Utility Functions	319
6.10 Additional Representation Theorems	324
6.11 Hierarchical Structures and Conditional Preferences	332
6.12 Summary	344
Appendix 6A Generalization of Preferential Independence and Utility Independence	345
Appendix 6B Evaluating the Scaling Constant in the Multiplicative Utility Function	347
Appendix 6C An Interactive Computer Program for Assessing and Using Multiattribute Utility Functions	349
 7. ILLUSTRATIVE APPLICATIONS	 354
7.1 Air Pollution Control	355
7.2 Preference Tradeoffs among Instructional Programs	365
7.3 Fire Department Operations	377
7.4 Structuring Corporate Preferences for Multiple Objectives	390
7.5 Evaluating Computer Systems	408
7.6 Siting and Licensing of Nuclear Power Facilities	412
7.7 Other Applications	417
 8. AIRPORT DEVELOPMENT FOR MEXICO CITY: A CASE STUDY	 436
8.1 The Problem	436
8.2 Background Information	437
8.3 Evolution of the Analysis	441
8.4 The Static Model	442