T 1673

MIT/LCS/TR-340

REASONING ABOUT PREFERENCE MODELS

Michael Paul Wellman

This blank page was inserted to preserve pagination.

Reasoning About Preference Models

by

Michael Paul Wellman

May 1985

© Massachusetts Institute of Technology 1985

This research was supported (in part) by the National Institutes of Health Grant No. NIH-5-R24-RR01320-02.

Reasoning About Preference Models

by
Michael Paul Wellman

This report is a modified version of a thesis submitted to the Department of Electrical Engineering and Computer Science on May 21, 1985 in partial fulfillment of the requirements for the Degree of Master of Science

Abstract

Programs that make decisions need mechanisms for representing and reasoning about the desirability of the possible consequences of their choices. This work is an exploration of *preference models* based on utility theory. The framework presented is distinguished by a *qualitative* view of preferences and a *knowledge-based* approach to the application of utility theory. The design for a comprehensive preference modeler is implemented in part by the *U*tility *R*easoning *P*ackage (URP), a collection of facilities for constructing and analyzing preference models.

Qualitative mathematical reasoning techniques are employed to develop *partial* specifications of single-attribute utility functions from qualitative preference assertions. Functions are described in terms of gross behaviors, symbolic forms, and parametric constraints. Appropriate dominance-testing algorithms are chosen from a knowledge base of stochastic dominance routines based on qualitative properties of the utility function. URP constructs multiattribute utility functions from a set of independence conditions by applying proof rules from a knowledge base containing the important decomposition theorems from the literature. Proof rules describe the logical relations among independence conditions and functional forms. Hierarchical decompositions are structured automatically.

Flexible model construction provides the potential for interpreting preference choices under a procedure that does not depend on the underlying utility model. Model-independent interpretation enables assessment under a wide range of descriptive theories of preference choice. Domain-specific preference knowledge is incorporated in URP by tying domain concepts to the modeler's *technical vocabulary*. Development of a health preference knowledge base illustrates how preference modeling could be included in a knowledge-based system for a particular application.

Thesis Advisor: Peter Szolovits
Title: Associate Professor of Computer Science

Keywords: utility theory, decision making, mathematical modeling, preferences, qualitative reasoning, knowledge-based systems, multiattribute utility, stochastic dominance, utility assessment, medical decision making

Acknowledgements

Alvin Drake and Ralph Keeney for some early advice and encouragement.

David Bell for much of the same and for looking over part of the knowledge base.

Members, past and present, of the Division of Clinical Decision Making at Tufts-New England Medical Center for two years of lively case discussions and for putting up with my medical ignorance. Special thanks to Van Dunn for the medical example, and to Jim Hollenberg for helping me select and set up the case as well as for many fruitful discussions about AI and Decision Analysis. And to Stephen Pauker for general tutelage, and for keeping me somewhat closer to the real world.

Kathe Unrath for eliminating some of the stylistic monstrosities from this document.

The members of the Clinical Decision Making Group at MIT for lots of good discussions, stimulating ideas, and in general for providing a nice place to work. Particular thanks to Zak Kohane for many heated discussions about this project and for helpful suggestions on several drafts, to Elisha Sacks for draft-reading and for helping me to get the QM package to do what I needed, and to Robert Kunstaetter for tolerating my naive medical questions and some of my more outrageous arguments.

Peter Szolovits, my advisor, for continued support and guidance.

Table of Contents

List of Figures ii

List of Tables ii

Abstract i

Acknowledgements ii

1 Introduction 1
- 1.1 Preferences . . . 1
- 1.2 Preference Models . . . 2
- 1.3 The Basic Approach . . . 3
  - 1.3.1 Qualitative View of Preferences . . . 4
  - 1.3.2 Knowledge-Based Approach to Utility Theory . . . 4
- 1.4 Model Construction and Analysis . . . 5
  - 1.4.1 Model Construction . . . 6
  - 1.4.2 Model Analysis . . . 7
- 1.5 Status of the Implementation . . . 7
- 1.6 Guide to the Thesis . . . 8

2 A Medical Decision Example 11
- 2.1 The Decision Problem . . . 11
  - 2.1.1 Case Description . . . 11
  - 2.1.2 The Analysis . . . 12
- 2.2 Qualitative Single-Attribute Utility . . . 13
- 2.3 Dominance Testing . . . 15
- 2.4 Multiattribute Utility . . . 18

3 Reasoning About Qualitative Behaviors 23
- 3.1 Functions in URP . . . 23
  - 3.1.1 Functional Forms . . . 24
  - 3.1.2 Evaluations . . . 24
  - 3.1.3 Qualitative Behaviors . . . 25
- 3.2 Manipulating Qualitative Behaviors . . . 25
  - 3.2.1 Qualitative Behavior Assertions . . . 25
  - 3.2.2 Qualitative Behaviors and Function Components . . . 26
- 3.3 Qualitative Mathematics . . . 29
  - 3.3.1 Functions in QM . . . 29

3.3.2 Partially Specified Functions 30
3.3.3 Encoding Function Behaviors 30
3.3.4 Enforcing a Functional Form 31
3.3.5 Risk Functions 32
3.3.6 Describing Behaviors 32
3.4 Qualitative ⇒ Intuitive? 32

4 Stochastic Dominance **34**
4.1 Dominance 34
4.1.1 The General Case 34
4.1.2 The Case of Known Functional Forms 34
4.2 Stochastic Dominance 36
4.3 Encoding Dominance Routines in URP 37
4.4 Multiattribute Dominance 38

5 Reasoning with Utility Theorems **40**
5.1 Representing Independence Axioms 40
5.1.1 Axiom Schemata 40
5.1.2 Axioms as RUP Propositions 41
5.2 Representing Theorems 42
5.3 Goal-directed Inference 43
5.3.1 Interpreting Theorems 44
5.3.2 Set-Membership Constraint Propagation 44
5.3.3 An Example Theorem Interpretation 47
5.4 Control of Reasoning 50
5.4.1 Default Control Mechanisms 50
5.4.2 Hierarchical Decomposition 51
5.4.3 Heuristic Control 52

6 The Decomposition Knowledge Base **53**
6.1 Formalization 53
6.2 The Extent of URP's Knowledge 55
6.2.1 The Decomposition Task 55
6.2.2 Coverage 56
6.3 Attributes 57
6.4 Independence Axioms 58
6.4.1 special-case-of 59
6.4.2 rel-prop 59
6.5 Theorems 59
6.6 Multiattribute Functional Forms 60

6.7 Preferential and Utility Independence 60
6.8 Value Independence 61
6.9 Joint Independence 62
6.10 Conditional Independence 62
6.11 Weaker Axioms 63
6.12 Interpolation Independence 64
6.13 Aggregate Conditions 64
6.14 Axioms not Implemented 65
6.15 Assessing the Decomposed Function 65
6.15.1 Assessing Scaling Constants and Conditional Functions . . . 66
6.15.2 Relation of Axioms to Conditional Functions 66
6.15.3 Encoding of Assessment Algorithms 66

7 Assessment with Preference Choices **68**
7.1 Model-Independent Interpretation 68
7.2 Inequalities vs. Indifference Points 69
7.3 Constraints on Utility Functions 69
7.4 Alternate Descriptive Theories 71
7.4.1 Descriptive Theories of Preference Choice 72
7.4.2 Interpreting Choices with Respect to a Theory 75

8 Domain Preference Knowledge **78**
8.1 Modeling and the Ground Domain 78
8.1.1 Modeling as Abstraction 78
8.1.2 Technical Vocabularies 79
8.1.3 Expressing Domain Concepts in a Technical Vocabulary . . . 80
8.2 A Domain Preference Knowledge Base 81
8.3 Built-in Attributes 81
8.3.1 Specifying Attributes in URP 81
8.3.2 Built-in Attributes for Medical Decision Making 84
8.4 Predefined Qualitative Properties 86
8.4.1 Monotonicity 86
8.4.2 Risk Aversion 88
8.4.3 Other Qualitative Properties 92
8.5 Predefined Independence Conditions 92
8.6 Special Functional Forms 93
8.6.1 DEALE 94
8.6.2 Quality-Adjusted Life Years 95
8.6.3 Time Streams 96
8.6.4 Value of Life 97

8.7 Summary: Using the Knowledge Base 98

9 Conclusion **100**
9.1 Perspectives on this Research 100
9.1.1 Decision Making Programs 100
9.1.2 Utility Analysis 103
9.1.3 Mathematical Modeling 108
9.2 Other Work 111
9.2.1 AI and Decision Analysis 111
9.2.2 Knowledge-Based Approaches to Mathematical Modeling . . 114
9.3 Limitations 116
9.3.1 Preference Modeling Issues 116
9.3.2 The Implementation 117
9.3.3 The Approach 118
9.4 Future Work: Planning with Preferences 119
9.5 Contributions 121
9.5.1 Summary 121
9.5.2 Outlook on Preference Modeling 122

References **125**

A Qualitative Behaviors of Utility Functions **136**
A.1 Behavior Descriptions 136
A.2 Terminological Definitions 137

B Axiom Schemata **138**

C Multiattribute Theorems **142**

D Set Membership Constraints **150**

E Multiattribute Functional Forms **152**
E.1 Two-attribute Forms 152
E.2 n-attribute Forms 154

F Health Preference Knowledge **156**
F.1 Built-in Attributes 156
F.2 Predefined Qualitative Behaviors 157
F.3 Predefined Independence Conditions 157
F.4 Special Functional Forms 158
F.5 Relations to Domain Concepts 158

List of Figures

1.1 A high-level view of the program's behavior 5

2.1 Determining the qualitative behavior of an asserted functional form 16
2.2 Dominance relations in the medical example 18
2.3 X decomposed via the multiplicative functional form 20
2.4 Two possible decompositions for X based on assumptions 3 through 7 21

3.1 Logical relations among qualitative behaviors of functions 26

4.1 A dominance test specification in URP 37

5.1 A sample axiom schema 41
5.2 A sample URP multiattribute utility theorem 43
5.3 Properties enforced by the **union** constraint, Z = Y1 ∪ Y2 46
5.4 Restricting possible bindings using constraint propagation 49

8.1 An example of an attribute definition in URP 82
8.2 Constructing preference models in a rule-based expert system. 99

List of Tables

2.1 Baseline life expectancies for each strategy 12
2.2 Notation for utility attributes in the medical decision example 19

3.1 Qualitative behaviors that imply single-attribute functional forms . . 31

4.1 Four dominance categories with their associated utility function classes 36

8.1 Some useful measurement properties and their permissible transformations (x denotes the measured quantity) 83

1. Introduction

The work described in this thesis is an investigation of an approach toward representing and reasoning about preferences in a machine. These preferences "belong" to some agent (person, institution, society) called the *decision maker*. We wish to encode them in a computer program so that it may be more effective in helping the agent make decisions, or even in making decisions for the agent by proxy.

Decision scientists have long been concerned with formal models of preference, and have assembled over the years an impressive body of theory which bears directly on the enterprise of representing preferences in computers. Attempts to employ these models in decision-making computer programs, however, have met with only limited success, mainly because the implementations have incorporated a highly restricted view of the formal preference theories. Taking such a view is difficult to avoid using conventional programming strategies.

This project arose out of the observation that utility theory (like many mathematical modeling disciplines) is tremendously richer and remarkably more flexible than reflected in any program performing utility analysis. This observation, together with advances in knowledge-based programming technology and qualitative reasoning methods from artificial intelligence, suggests the approach related in this thesis.

1.1 Preferences

People make decisions to influence the course of events in ways that they deem desirable. Different actions make the various possible outcomes more or less likely to obtain; choosing among available acts is an exercise in selecting probable futures. It seems important, then, that computer programs that make decisions or help humans to do so should have some means to reason about the desirability of the possible consequences of these decisions.

Mainstream AI research—while enormously concerned with representing beliefs and states of knowledge—has for all practical purposes ignored the issue of preferences. Note that goals as employed by AI planning programs do not function as expressions of preference; at the very best they can be said to mean the most desired outcome. More often, goals are used to describe intentions or desired reasoning

states, which are nothing at all like desired states of the world.[1] Furthermore, in most real decision-making or planning situations it is necessary to consider possible outcomes other than the one most preferred. The program must be able to compare the desirability of various partially satisfied or uncertain goals, which may not be related to the absolute goal by any identifiable distance metric.

1.2 Preference Models

This work is an exploration of the use of *preference models* as a basis for reasoning about desires in decision making. Let us define a *prospect* to be a probability distribution over consequences, or *outcomes*. The underlying assumption of preference modeling is that there exists a *preference relation*, denoted by $\succ$, which orders the set of prospects. A preference model is a structure used to represent this order relation.

Preference models can be used to answer questions of the form "Is x preferred to y?" where x and y are prospects relevant to the decision being considered. The decidability of such questions depends on the degree to which the preference model is specified. In general, a preference model will characterize only a subset of the true preference relation it represents. Thus, any partial description of a preference order may be termed a preference model.

The ultimate goal of this effort is to develop a representation for preference models that will accommodate arbitrary amounts and types of information about the preference relation, and a reasoning capacity that will enable us to reconstruct as much as possible of the relation from whatever information is available. A formal characterization of the sorts of information that may be incorporated in preference models and the language used to describe it may be called a *preference calculus*. While I cannot claim to provide a sufficiently formal or principled development, this work may be viewed as a step toward the definition of a useful preference calculus.

The next major step (which requires a leap of faith, since I do not intend to provide detailed justification here) is that the foundations for this preference calculus are to be found in utility theory. A fundamental result of this theory is that, given a few appealing axioms,[2] there exists a *utility function* which maps prospects to real numbers (called *utilities*) in such a way that the utility of any prospect is equal to

[1] Doyle [23] discusses the importance of distinguishing between desires and intentions in a decision making program.

[2] By appealing, I mean that they generally express properties that one would like to conform to in

the sum of the utility of each certain outcome weighted by its probability in that prospect. Utilities rank the prospects by preference; $u(p_1) > u(p_2)$ if and only if $p_1 \succ p_2$. If we base our preference model on the notion of a utility function, we can exploit a large body of results from utility theory, and more general analytical techniques from the mathematics of functions. This thesis is largely an account of some of these possible exploitations.

Formally, a utility function u maps states of the world (outcomes) x in the *outcome space* $X = X_1 \times \cdots \times X_n$ to real numbers in such a way that the optimal decision is the one that maximizes the expected value of this function. Often x is expressed as a vector of attributes, $(x_1, x_2, \ldots, x_n), x_i \in X_i$, where the values of the x_is totally describe the aspects of the outcome pertinent to the decision.

While a major preference modeling effort is often undertaken for the sake of the *insight* it brings to decision making, I will be primarily concerned with the use of preference models for the analysis of a particular decision.[3] For our purposes, a decision involves a choice among a set of strategies, each of which is associated with some resulting outcome $x \in X$ (or more generally, a prospect $p \in P$, where P, the *decision space* is the set of simple probability distributions over X). To make the decision, we must figure out which strategy has the greatest expected utility.

1.3 The Basic Approach

The main features that distinguish the attitude toward utility modeling developed here from traditional computerized utility modeling are its qualitative view of preference structures and its knowledge-based approach to the encoding of utility-theoretic concepts. Taken together, these features result in a system which emphasizes model structuring over model parameterization. This attitude has been embodied in a collection of programs and overall framework for preference model-

decision making. This is not to say that these axioms are totally accepted. However, it is not my intention to enter the debate at this point. A particularly lucid axiomatic development of expected utility can be found in Savage's important work [92]. Alternative axiomatizations are provided by Fishburn [35]. The original existence proof is due to von Neumann and Morganstern [107].

[3] Although, as will be explained later, I believe that models created using this methodology offer the potential for greater insight due to their deep, explicit qualitative structure. And there is no reason to believe that any of the insight-yielding characteristics that are inherent in traditional utility analysis efforts will not carry over to this methodology. Unfortunately, I do not know how to prove an insight theorem.

ing called URP (for *U*tility *R*easoning *P*ackage). The sections immediately below describe features of the general attitude; descriptions of the actual URP framework and programs make up the remainder of this thesis.

1.3.1 Qualitative View of Preferences

An important novelty of the approach described here is in the qualitative mechanisms used to analyze preference structures. Utility analysis is usually considered a primarily quantitative[4] task, and previous utility modeling aids (basically assessment tools) have reflected that bias. I believe that the emphasis on numerical computation has tended to obscure the more important modeling and analysis issues, just as implementations of mathematical models often hide the central qualitative structure of the problems they are supposed to represent. Recent work in artificial intelligence (that of Kuipers [67], Forbus [39], and Sacks [90], for example) has shown that it is often possible to reason effectively about the qualitative behaviors of mathematical structures without resorting to overly precise descriptions. Such a capability would be enormously useful for utility analysis, because (as with any kind of psychological measurement) virtually any numeric data is uncomfortably precise. A more complete account of the potential advantages is given in section 9.1.2.

This qualitative view meshes nicely with the notion of a preference model as a *partial* description of a preference order. The view is supported by URP in its tendency to focus on propositions about the utility structure rather than on direct specification of components making up the structure. Specific methods contributing to this focus are described below.

1.3.2 Knowledge-Based Approach to Utility Theory

The second major distinction in how URP models preferences is that it relies on utility theory knowledge expressed in a (relatively) declarative format, rather than compiled into analysis procedures. Encoding utility-theoretic concepts in explicit knowledge structures leads to advantages in flexibility, extensibility, and explain-

[4] It is difficult to avoid the confusion inherent in words like *qualitative* and *quantitative*. Naturally, all mathematical models are composed of relationships among quantities, and in that sense they are quantitative. But it is important to distinguish qualitative *descriptions* of quantities from numeric descriptions. A symbolic expression is just one form of qualitative description. The term *qualitative* as used here is almost always meant in contrast to *numeric*, or more generally in opposition to excess precision of any sort.

ability, which will be noted as we get into specifics. For now I will just mention that the process of building URP has resulted in the collection of a large corpus of utility-theoretic knowledge, encoded in a uniform representation, *executable* by URP's interpreter. The creation of this knowledge base (especially that pertaining to multiattribute decomposition) may be a worthwhile exercise in its own right.

1.4 Model Construction and Analysis

A high-level view of the URP preference modeling framework is depicted in figure 1.1. Given a probabilistic model of the alternative strategies, the program will produce decisions based on a specification of preferences. The probabilistic model is given as $(p_1, \ldots, p_m)$, where p_i, the result of the ith strategy, is a prospect in P (that is, a simple probability distribution over the outcome space, X).

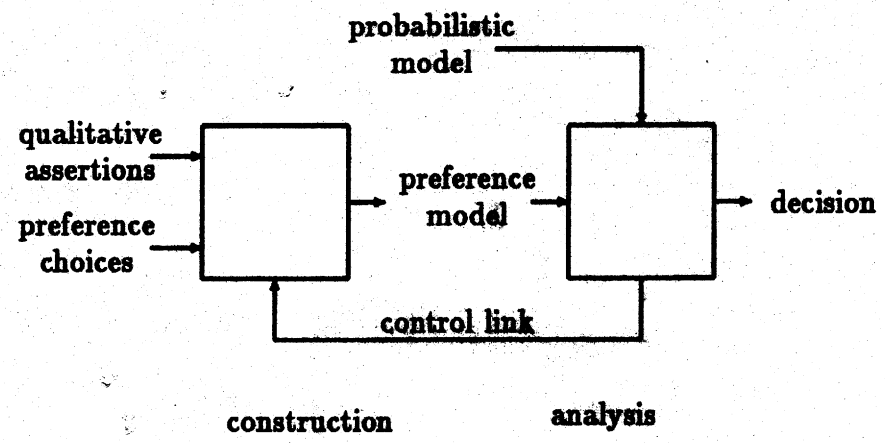

Figure 1.1: A high-level view of the program's behavior

Preference specifications are in terms of the same basic variables as the probabilistic model: the attributes $X_1, \ldots, X_n$. The two main categories of specifications are qualitative utility assertions and preference choices. Qualitative utility assertions are propositions that refer to preference conditions for attributes or sets of attributes. Some illustrative examples, stated in English (rather than URPese[5]), with the js and Ys standing for attributes and attribute sets, respectively:

[5]It will soon be demonstrated that URPese is closely related to the technical language of utility theory.

- Preference for prospects over attributes in Y_1 is independent of the fixed value of the rest of the attributes
- All other things equal, higher levels of (numeric) attribute j_1 are preferred
- The attributes in Y_2 are complementary
- The marginal value of equal increments of attribute j_2 is decreasing with j_2
- An additive form is valid for the utility function over Y_3

The implications of these assertions on the preference model will be made clearer below.

Preference choices are assertions similar to those generated by the standard lottery technique of utility assessment, where the assessor is required to report preferences in hypothetical choices among simplified prospects. The form and use of preference choices is the subject of chapter 7.

The output of the first black box of figure 1.1 is a data structure representing the current preference model. The model is a collection of assertions about the utility structure, containing some combination of functional constraints, algebraic forms, parametric constraints, and other qualitative properties. Because the model is the central object of reasoning activity, it is useful to separate the process into two components: *model construction* and *model analysis*. This is a natural distinction, because the two activities use very different types of knowledge and inference mechanisms. The processes are not totally independent, of course, since we construct models in the first place for purposes of analysis. Therefore, there must be a control link tying the results of the analysis to the goals of the construction.

1.4.1 Model Construction

Several mechanisms are employed to construct preference models from the qualitative properties and preference choices input. Of course, by my definition these raw input assertions alone can constitute a preference model, but we are typically more interested in higher-level inferences we can make from these premises.

URP's most completely developed model construction mechanisms deal with assertions about *independence axioms* and *qualitative behaviors* of single-attribute utility functions. Examples of URP's reasoning in these areas are given in the next chapter; deeper analyses of these mechanisms appear in chapters 3 and 5.

1.4.2 Model Analysis

The goal of the model analysis module is to determine which of the available strategies yields the highest expected utility given a probabilistic model and a preference model. This will not always be possible, either because the expressions for expected utility are underconstrained, or because the analyzer is not powerful enough to recognize the dominance of an expression.[6] If the procedure succeeds in proving that one alternative is optimal, the process is finished. Otherwise, it is necessary to refine the preference model. Examples of refinement goals are to constrain a certain parameter or set of parameters, to determine the applicability of a functional constraint such as monotonicity or concavity, or to find a simpler utility function that has fewer parameters.

1.5 Status of the Implementation

This thesis report describes both a design for a comprehensive preference modeler and an implementation that realizes part of that design. I use "URP" to refer to both of these at times; where it may be unclear I explicitly distinguish the *implemented* URP from the URP *framework*. I have taken the implementation to a point that demonstrates the feasibility and potential of the approach, and at the same time provides some interesting and useful capabilities. Throughout this thesis I point out parts of the implemented URP that need to be extended, and discuss ways in which the program could be improved.

Development of the complete preference modeling system as envisioned here could potentially require an effort similar in size to the CYC[TM] project. Because the system revolves around a task rather than a particular method, a comprehensive program would be composed of a multitude of inference mechanisms and representations. Indeed, even the implemented portion of URP includes several distinct components. Coordinating the use of these separate programs for a central task is a considerable engineering problem in itself.

The usefulness of the implemented part is illustrated by the example of chapter 2. There it is shown that the program is able to reduce the set of admissible options from those considered in an actual medical decision based on weak, qualitative assumptions about the patient's preference structure. Various kinds of reasoning

[6] One expression *dominates* another if its value is greater for any allowable assignment of the parameters.

about the qualitative properties themselves and how they relate to the utility function are also demonstrated. Though the program cannot always take advantage of the inferences it makes about the function for decision making, a decision analyst would find these facilities useful for developing utility models that could then be used in conventional decision analysis tools (along the lines of the program described by Pauker and Kassirer [82], in routine use at Tufts-New England Medical Center).

The facilities implemented for multiattribute decomposition are limited to model-structuring. To use a multiattribute model generated by URP for an actual analysis, one would have to transfer the model to a package that performs multiattribute assessment and evaluation.[7] Utility theorists may also find the program useful for exploring novel decompositions based on different combinations of axioms. One can test the implications of new results in multiattribute utility by encoding them as URP independence concepts and theorems, adding them to the existing knowledge base.

The main objective of this work has been to demonstrate that the overall approach to preference modeling is feasible and worthwhile. An assessment of how this project has satisfied that goal is deferred to the concluding section.

1.6 Guide to the Thesis

This introductory chapter has served as an abstract description of the preference modeling task, and as an overview of the approach to preference modeling embodied by this project. The rest of this thesis describes the URP preference modeling framework in greater detail: including working representations and mechanisms as well as unimplemented components that have been worked out to some degree.

Chapter 2 provides the setting for the rest of the thesis by describing a medical decision example and illustrating some preference modeling capabilities of URP. The case exemplifies the kind of decision problem of interest: a situation where there are several feasible strategies, none of which is obviously the best course of action. The example also serves to provide a context for the more detailed descriptions of URP's mechanisms that follow.

[7] For example, Keeney and Sicherman's MUFCAP [62] can accept as input many of the models that might be constructed by URP. Though using a conventional utility analysis package sacrifices some of the advantages of the URP framework (the ability to reason about partially specified functions, for example), there is still substantial benefit to structuring the model using URP.

The next chapter describes the first of two major preference modeling mechanisms implemented in URP: reasoning about the qualitative behavior of single-attribute utility functions. URP's mechanisms for relating these behaviors to the mathematical structure of the utility function is discussed, with emphasis on qualitative mathematical reasoning methods employed. The utility-theoretic content of URP's single-attribute knowledge is delimited. Deriving decisions based on URP's representation of single-attribute functions is the subject of chapter 4, which also includes a discussion of dominance-testing procedures incorporated into the URP framework.

The following two chapters deal with URP's facilities for multiattribute decomposition. Chapter 5 describes the program's use of independence conditions to derive multiattribute functional forms, based on proofs of their validity generated from a knowledge base of decomposition theorems. This chapter describes the theorem language as well as the interpreter that uses the proof rules to derive the functional forms. The knowledge base itself is described in chapter 6, an in-depth discussion of the multiattribute utility theory known to URP. There I describe the concepts appearing in the knowledge base, and provide rationale for their inclusion. I also argue that the knowledge base covers a substantial portion of the important decomposition results developed by utility theorists.

In chapter 7 I discuss how assessment based on hypothetical preference choices—the mainstay of traditional utility analysis programs—may be incorporated in the URP framework. The major point of the chapter is how URP's flexible model structuring capabilities and capacity for reasoning about incompletely specified models leads to less restrictive assessment. Of particular interest are the ability to work with ad hoc constraints, and the intriguing notion of interpretation under descriptive theories of preference choice.

Representation of domain-specific preference knowledge is the subject of chapter 8. There I argue that concepts from a domain-independent modeling system must be encoded in a technical vocabulary, and therefore are not necessarily intuitive for modeling in particular application domains. This problem is remedied by the construction of a domain/modeling concept mapping, in the form of a knowl-